Towards Data-Driven Internet Routing Security

by

Cecilia Andrea Testart Pacheco

B.Eng.Sc., Universidad de Chile (2010)
M.Eng., École Centrale Paris (2010)
S.M., Massachusetts Institute of Technology (2016)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2021

Author . . . .
Department of Electrical Engineering and Computer Science
August 27, 2021

Certified by . . . .
David D. Clark
Senior Research Scientist
Computer Science Artificial Intelligence Lab
Thesis Supervisor

Accepted by . . . .
Leslie A. Kolodziejski
Professor of Electrical Engineering and Computer Science
Chair, Department Committee on Graduate Students

Towards Data-Driven Internet Routing Security

by

Cecilia Andrea Testart Pacheco

Submitted to the Department of Electrical Engineering and Computer Science
on August 27, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

The Internet infrastructure is critical for the security and reliability of online daily life. The Border Gateway Protocol (BGP), the defacto global routing protocol, was not designed to cope with untrustworthy parties, making BGP vulnerable to misconfigurations and attacks from anywhere in the network. Recently, unintended large-scale misconfigurations caused significant amount of Internet traffic towards major providers to be dropped for hours, and through BGP attacks, perpetrators have stolen millions in fraudulent transactions. Nonetheless, little has changed in operational environments despite the many proposals to increase security by the research, standardization and industry communities. The problem space is complex: it involves multiple stakeholders, with different interests and available resources, and increasingly, geopolitical challenges. Yet, these stakeholders ultimately need to cooperate and coordinate their efforts to improve security. This dissertation proposes a holistic approach to study routing security. It includes the assessment of barriers of adoption of technical proposals to secure BGP, the empirical analysis of exploitations and misconfiguration due to BGP design flaws, as well as the empirical study of the mitigation strategies deployment and benefits. This analysis reveals the extent of misbehavior and misconfiguration in the use of BGP, and the benefit that operational security practices provide. It also discusses this new evidence in the context of tradeoff that have prevented the adoption of routing security. Finally, it provides a set of actions, which could be orchestrated by a bottom-up industry effort or top-down by governments, and directions for future technical work that would encourage collective adoption of security in BGP.

Thesis Supervisor: David D. Clark
Title: Senior Research Scientist
Computer Science Artificial Intelligence Lab

Acknowledgments

My PhD years have certainly been quite the (fun) ride. They shaped me academically and personally, I am definitely not the same person that entered the program. These years have been intense, amazing and full of surprises on the personal and academic level. They have even included a global pandemic. I can't say all this has always been easy to navigate. But a PhD takes a village, and I've been lucky to have two (and more!) villages of people supporting and guiding me through the journey. I am deeply grateful to all of you. It has definitely been worth the ride and the prize!

David Clark, my advisor, stands out as an amazing mentor and advisor. He has not only introduced me to research and the academic world, and guided me in finding topics I'm passionate about that are worth spending time on, he also helped and supported me in finding balance between work and family. I could never say thank you enough for the support, patience, and the fond memories until the very end of this journey.

I'm also deeply thankful to all my collaborators. Philipp Richter, Alberto Dainotti, Arthur Berger, I have greatly enjoyed working with you. Thank you for all the hard pushes before deadlines and making them fun! I've learned so much from each of you. You have definitely shaped me as a researcher and I hope to continue to collaborate with you. I'd also like to thank all the members of the Advance Network Architecture (ANA) group—Karen, Steve, Bill, John, it's always a joy to share research with you. Thank you also to Josephine Wolff, Alex Gamero-Garrido, Georgios Smaragdakis, Volker Stocker, Brandon Karpf, Zane Markel, Nathaniel Fruchter, Sam DeLaughter and all that have spend time in ANA bringing companionship through the years. Thank you Danny Weitzner and the leadership of the Internet Policy Research Initiative at CSAIL for supporting my work and all IPRI folks for constantly expanding my understanding of the relationship between computing and society. And thank you to kc Claffy and Nickolai Zeldovich. You have provided invaluable feedback to make this dissertation and my research better.

Until fairly recently, I didn't quite get how relevant mentors were in our life and careers. Nonetheless, I've had mentors support me at various stages of my career and turning points in my life, and I am immensely grateful for that. Thank you Rafael Correa, Alejandro Jofré, and Claude Puech, my early mentors who gave me the confidence I would be successful in an academic career and supported me in figuring out the area I liked. Special thanks to Jo Piquer, who one day after working a few years together explained to me what a PhD was all about and why I should pursue one as soon as possible. And since then, he has always kept an eye on me, my family and my career. Jo, it is always a pleasure to share life with you. Thank you Michael Sipser and Leslie Kolodziejski for helping me navigate MIT and EECS all these years. Thank you Radhika Nagpal, Jennifer Rexford, kc Claffy and Justine Sherry for showing me that I have a place in academia. Thank you David Choffnes and David Nino for your support in navigating the academic job market in such an unexpected year.

In this journey, I've met some of the most inspiring and kind people that I've ever known. They have been excellent friends and Sofia's family in the US. Thank you

to the fantastic TPP crowd: Tiziana, Angi, Stacey, Corinne, Lauren, Julius, Anne, Alex. You are so fun and stimulating, I would always put things aside spend time with you. Thank you Mike for sharing the ride since the early year and continuing now to the next steps! Thank you Leilani for all the hours spent together, navigating this journey so closely. Thank you Clement and Reyu for sharing with me your openness, kindness and perspective. Thank you Yuly, Ricardo, Antonia, Dani and the Chilean crowd that has spend time with us in Boston area. Thanks to my mom/dad-friends that gracefully share their take on the parenting ride. Thank you also to all my friends scattered in Chile, France and around the world, that have always been a Whatsapp or FaceTime call away.

And of course, I am also deeply grateful for the unconditional family village that I have. Thank you to my husband Ricardo, who pushed me to apply to MIT and has also supported me in every possible way to make it until the end. Thank you for being a dedicated father to Sofia and soon to our baby boy. Thank you for being my other half, you are exactly the kind of father our kids need. And thank you Sofia for filling my life with surprises and awe about our world even in the craziest moments. Thank you for letting me be your mother and hold your hand in this life. And thank you baby boy, for choosing me as your future mom, writing this dissertation with me and supporting me with your kicks during my defense. I'm also terribly lucky to have very supportive grand parents, parent, additional parents, uncles, aunts and cousins. Thank you to my *abuelos* Alfonso and Cecilia, how have never failed to keep in touch and follow my journey. Thank you to my mom Patricia (or Nounou), for coming to see us often, keeping us connected to our roots and being a space-travelling grandma for Sofia. Thank you Dad, Javier, Kike, Horacio, Consuelo for your love and support. Thank you Pipa for all the hugs and love, and being the best godmother Sofia could have. Thank you Pola, Alfredo, Paula, Roro, and the ones that left too early and are watching us from the sky. Thank you to my in-laws, Ricardo, Sandra, Gabriel, Nathalie, Hernán, for always keeping in touch and making sure Sofia spends (virtual) time with her family.

Finally, I'd like to eagerly thank all the early childhood educators that have provided loving care to Sofia, allowing to work with a peace of mind. Thank you Debra, Latoya, Karen, Carolyn and Muriel for taking such good care of my baby girl and teaching me so many parenting tricks!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The Internet is made of interconnected networks that allows computers or other devices in one network to communicate with computers and devices in other networks by exchanging data packets. These packets travel according to the Internet Protocol (IP) suite. When the computer devices wanting to communicate are not in the same network, packets need to travel through a path of different networks to reach their destination. Networks use a routing protocol called Border Gateway Protocol (BGP) to select routes to the different IP address destination reachable in the Internet. In other words, BGP is the global routing protocol that is used by routers to receive, select and propagate the information of available paths to reach IP addresses. These paths are then used to send IP packets towards their destination. Therefore, BGP plays a critical role in the correct delivery of packets travelling between networks in the Internet, impacting the security and reliability of online daily life.

The functional goal of the Internet is to deliver (well-formed) packets to their destination as expeditiously as possible, making availability a key property and expectation. There are three core systems, part of the initial Internet protocol suite, that allow packets to reach their destination: (1) IP addressing, which distributes unique identifiers to device interfaces; (2) the Domain Name System (DNS), which translates the hierachically-distributed name space into IP addresses; and (3) routing, which provides the information to select paths to IP addresses in use. In addition, the Certificate Authority (CA) system was designed and deployed to provide integrity and confidentiality to connections using end-to-end encryption based on digital certificates issued and distributed through the CA system. These four systems have all been designed—and modified—to encourage availability. However, routing plays a major role in making communication between host in different networks available. In addition, routing can disrupt the correct operation of the DNS and CA system when wrong path are selected, putting confidentiality and integrity at risk . If routers selects routes with incorrect networks in the path, traffic many not be delivered to the proper destination, potentially breaking the availability, integrity and confidentiality of communications over the Internet. For these reasons, the correct operation of BGP is a fundamental building block in ensuring the correct operation of the Internet.

Unfortunately, BGP was not designed to cope with untrustworthy parties and is vulnerable to misconfigurations and attacks from anywhere in the network. Intentionally or unintentionally, any network in BGP can modify the path towards IP addresses in the Internet, causing traffic to shift unexpectedly. Traffic to addresses impacted by the change may be dropped in the wrong destination, but the network causing the path change can also inspect or manipulate that traffic, or send traffic from those addresses. As an example, unintended misconfigurations have caused significant amount of Internet traffic towards major providers including Google, Facebook and Amazon to be dropped for hours. Moreover, attackers have stolen millions in fraudulent transactions by stealing user's credential through BGP attacks. And even governments have (mis)used BGP to dropp and discard traffic to social platforms during protests, making this platforms unavailable for communication. These problems are longstanding and increasingly pressing.

There have been many proposals to secure BGP by the the research, standardization and industry communities, to ensure routes in BGP go to the proper destination through acceptable paths. These solutions intend to prevent third parties from tampering with BGP communication and networks changing path to IP addresses they do not own or violating routing policies from business agreements. Some of them modify the protocol, others create authoritative database to check information received in BGP, use overlay networks or add optional cryptographic signatures to BGP messages. Nonetheless, despite the many alternatives—some dating from the 1990s—only a few of them have been implemented and no solution currently has major deployment and use.

The problem space is complex: it involves multiple stakeholders, with different interests and resources available, and increasingly, geopolitical challenges. Networks in the Internet are of diverse nature, administered by different types of organizations from around the globe, with different operational environments, goals, business models and incentives. Additionally, for most solutions to have impact, many networks need to deploy and use them, *i.e.,* it requires coordination and cooperation between these heterogeneous organizations. Furthermore, there are few direct incentives for networks to provide additional security, since it protects resources from other networks and not their own. Indeed, networks can only prevent the spread of incorrect routing information they receive and forward, but they cannot directly block another interconnected network from erroneously or maliciously changing routes they do not legitimately own or administer. With all the challenges to deploy better routing security, currently there are not many barriers for networks that misbehave, and there is no associated penalty other than peer pressure among network operators.

Goal of this dissertation

The goal of this dissertation is to develop methods based on empirical data to improve our understanding of the systemic impact of BGP design flaws and provide key insights on routing security protocol and policy design. It explores the reasons that have prevented security improvements in BGP, assesses the pervasiveness of malicious activity and the spread of misconfigurations and evaluates the impact of security frameworks being deployed by scrutinizing routing data.

This dissertation considers security issues that impair the proper function of BGP, *i.e.,* the selection of paths that allows packets to be delivered to their destination without interference from third parties. This includes intentional BGP attacks, as well as unintentional misconfigurations that lead to paths along under-provisioned links which are likely to impact the delivery of traffic along those paths, undermining availability. As mentioned before, availability is key to the proper functioning of applications running on top of the Internet.

This work brings light into the struggle of improving routing security and the need for a deeper understanding of the systemic impact of security issues based on empirical evidence. The empirical studies reveal the the extent of malicious behavior and misconfigurations allowed by BGP flaws. They also evidence how implementation and operationalization decisions as well as other non-technical aspects influence the outcome of security solutions. The results and discussion can aid the design of security protocols and the implementation and operationalization process to get solutions into production, providing insights into how to overcome barriers to the adoption of security. Finally, the proposed future work and actions to support routing security can guide policymakers, industry and other related organizations' efforts in the area.

The remainder of this chapter is structured as follows: Section 1 introduces BGP, its role in the Internet Procotol suite, how it works and the ecosystem of networks; Section 2 describes the main BGP flaw; Section 3 describes the current state of BGP security; and Section 4 provides a roadmap to the contents of this dissertation.

1 The Border Gateway Protocol

The Border Gateway Protocol (BGP) was developed in the late 1980s for selecting the paths packets would use to travel between networks in the Internet. In BGP, networks' border routers (gateways) exchange information about IP address blocks in use and available paths, and select the preferred path to those blocks. The border routers then send data packets towards their destination by selecting the route based on the packet IP address destination. The first Internet Engineering Task Force (IETF) [1] Request For Comment (RFC) formalizing this routing standard was published in June 1989 [3]. The current version of the protocol is called BGP-4 and was first standardized in 1994 and later updated in 1995 and 2006 [4–6]. The rest of this section provides an overview of the Internet Protocol suite and the role of routing, how BGP works, the relationships between networks and general characteristics of networks in today's Internet.

1.1 The core Internet Protocols

The Internet is a global network of over 70,000 independent networks that use the Internet Protocol Suite to communicate between devices. The goal of this protocol suite

[1] The IETF is an open standards organization that develops Internet standards and documents methods, behavior and practices through Request For Comments (RFCs) authored by networks operators, engineers and computer scientists [2].

is to properly deliver packets from the source host to the destination host, whether in the same or different network. Each network is a set of linked computers and other devices that share resources and can support a large range of applications. Indeed, networks in the Internet are very heterogeneous, with different purpose, administered by different types of organizations, with different business models, in different geographies and using different types of networking technologies. However, they all implement the needed parts of the Internet Protocol Suite for their devices to communicate as needed to support their applications.

There are three core systems in the Internet Protocol suite that enable the delivery of packets in the Internet, and more recently an additional system was deployed to protect packets integrity and confidentiality:

1. **The IP addressing system:** this system allocates IP addresses that identify hosts (*i.e.,* devices interfaces) in a network, and provides the location of the host in the network. Thus, IP addresses within a network establish the paths to reach the hosts.

2. **The Domain Name System (DNS):** this system consists of two part: a hierarchical naming system that distributes names to services and resources connected to the Internet, and a database to translate between names and the IP address space of the host where the service or resource is located.

3. **The routing system:** this system allows routers to select the preferred path, between those available, to IP addresses in use by other networks. Routing provides the paths for IP packets to reach their destinations.

4. **The Certificate Authority (CA) and Transport Level Security (TLS) system:** this system provides confidentiality and integrity between communicating applications and devices. It consists of two parts: the authoritative databases of records delegated by CAs, and the transport level security handshake protocols used to verify records and distribute the keys for end-to-end encryption of communication.

Although each of the core Internet protocol plays a distinct role, the proper functioning of the routing system is crucial for availability. Indeed, to deliver packets in different networks, the addressing system requires the routes provided by the routing protocol. Routing problems can thus impair the availability of Internet communication.

In addition, routing can impair the proper functioning of the DNS and CA system, putting the integrity and confidentiality of communications at risk. For names to be translated to IP addresses, queries are sent to Domain Names Servers and answers are sent back to host following paths provided by the routing system. Routing also plays a similar role in the CA system. Some protocols also use Internet packet delivery as a way of bootstrapping trust. For instance, Let's Encrypt, a popular certificate authority, uses the Automatic Certificate Management Environment (ACME) protocol [7] to authenticate domain owners. ACME authenticates domain owners by looking up the domain's DNS record and sending an HTTP challenge to the IP address returned by DNS. This step assumes that the Internet correctly routes the DNS query and the HTTP challenge to the appropriate destination. Thus, an adversary that can subvert where packets are routed can not only affect the availability of Internet services, but also undermine the confidentiality and integrity of applications, by subverting how trust is bootstrapped to begin with. Therefore, the correct operation of BGP is essential in ensuring the correct operation of the Internet. The next sections explain how routing works and its security flaw.

1.2 How BGP works

The Border Gateway Protocol (BGP) is the inter-domain routing protocol, providing the necessary information for networks to choose preferred routes, among those available, to IP addresses used by other networks. These routes are then used to send packets to their destination. BGP was designed to allow for dynamic changes in routes.

In BGP, each independent network operated as a single administrative domain in the Internet is called an *Autonomous System (AS)*, and uses an Autonomous System Number (ASN) to identify itself. ASes establish BGP sessions with their neighboring ASes and exchange information about routes to IP addresses in use by networks. Then, through the BGP route-selection process, ASes independently decide on which path to use for each set of IP addresses. ASes also choose which IP addresses are reachable through them to each of their neighboring networks based on their business relationship. BGP-selected routes feed the forwarding table of ASes' border routers.[2] Border routers use the forwarding tables to decide where each incoming IP packet is sent. IP packets are forwarded to the border router of the next-hop AS in the selected route matching packets' IP address destination.

Routes in BGP have two critical types of information: *IP prefixes* and *AS paths*[3].

- IP prefixes: These IP address blocks are sets of Internet Protocol addresses with a common prefix of a given length. For example prefix 18.0.0.0/16 represents all IP addresses that share the first 16 bits. To forward packets, routers look for the longest prefix match between prefix entries in the routing table and the destination IP address, thus choosing the route towards the smallest block of IP addresses including the destination address.

- AS paths: These ordered lists of ASNs represent the network-level paths to reach given IP prefixes. For example, a path of length 2 going from AS A to AS B and to AS C is A, B, C. C is the AS where the IP addresses in the prefix are located, *i.e.,* it is the *origin* of the path.

[2] BGP requires that the route announced to an AS neighbors is the same route that is effectively installed in the forwarding tables of that AS, the BGP Additional Paths feature is implemented and negotiated between BGP peers at the start of a session.

[3] Also referred as *BGP Path*.

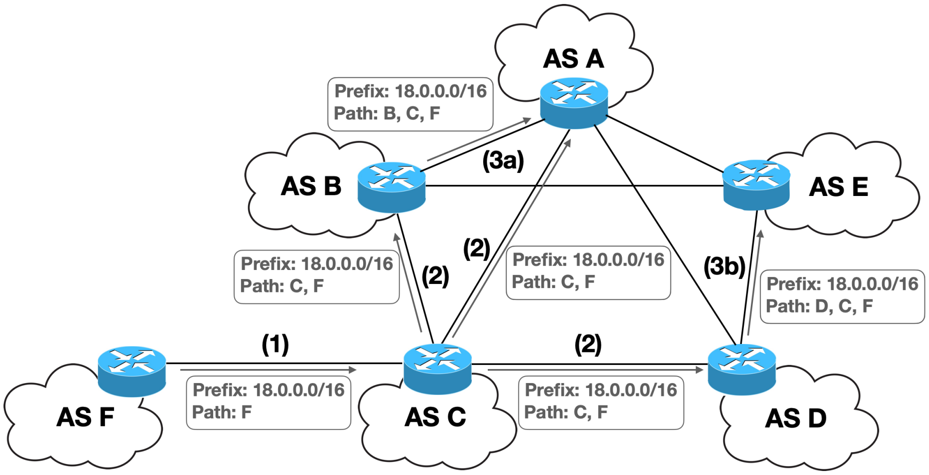

Figure 1-1: How BGP works: in this example, Autonomous System (AS) F *originates* the IP address block 18.0.0.0/16 in a BGP announcement sent to AS C (step (1)). Then AS C forwards that announcement to neighboring network A, B and D (step (2)). The ASes receiving the announcement repeat the process, choosing to which neighbor to forward the announcement and adding their AS on the path. ASes are identified by numbers in BGP, letter are used in this sketch for convenience.

BGP messages containing information about IP prefixes are called BGP updates. Updates can *advertise/announce* a route to a prefix, signaling a reachable route to the prefix through the sender AS. Updates also can *withdraw* a route to a prefix through the sender AS, signaling a previously sent route became unreachable.

When an AS *originates* a BGP announcement for a given IP prefix, *i.e.,* it puts itself as the *origin ASN* in the AS path, it signals that hosts in its networks are using IP addresses from that prefix. Then ASes that receiving the announcement and forwarding it will add their ASN at the beginning of the path, thus the origin ASN is the left-most AS in a path. Figure 1-1 shows how BGP works in a simple topology of ASes. AS F sends an announcement originating prefix 18.0.0.0/16 (step (1)). Then, AS C forwards that announcement to ASes A, B and D adding itself in the path (step (2)). Afterwards, AS B decides forward that announcement to AS A but not AS E (step (3a)) and AS D forwards the announcement to AS D but not AS A (step (3b)), both ASes adding their ASN to the AS path. Given the border router in AS A receives multiple route to the same prefix, using BGP decision making process, it will chose only one of those routes for its routing table.

Prefixes and origin ASes are particularly relevant in BGP. First, when an AS originates a prefix, *i.e.,* claims that hosts in its network use IP addresses within that prefix, other networks expect that AS to legitimately hold the allocation of that block of IP addresses and the authorization to originate it in BGP. IP address delegation process has changed over time as the Internet has grown and currently is handled by the five Regional Internet Registries, each representing a geographical zone.[4,5]

[4] The five RIRs are: RIPE NCC for Europe and Middle East, ARIN for North America, APNIC for the Asia-Pacific region, LACNIC for Latin America and the Caribbean, and AfriNIC for Africa.

[5] Many IP addresses were in use by organization before the creation of RIRs, they are called

Second, the route selection process in border routers is done per IP prefix, even if one IP prefix may be part of a larger IP address block also announced in BGP. For instance if both 18.0.0.0/16 and 18.0.0.0/18 are announced in BGP, ASes select a route to each. When an IP packet is forwarded towards its destination, the longest prefix match rule is used, which means that IP addresses covered by the longest prefix (*e.g.,* 18.0.0.0/18 in the example above), will follow the route of the longest prefix. This rule is used by networks operators for traffic engineering and for reducing the size of the routing table (using only one prefixes for a very large block of addresses). However, it is exploited by attackers and contributes to the quick spread of certain misconfigurations in BGP.

Each AS can decide on its criteria for route selection. Many times the route with the shortest path to a given prefix is chosen but the relationships of the AS choosing the route and its neighbors can modify that choice by giving routes different levels of preference.

1.3 AS relationships

ASes decide whether to forward or not a BGP announcement and the level of priority of routes based on the AS relationship they have with their neighboring ASes. These relationships depend on business agreements reached between networks with respect to the financial conditions of the traffic between them. There are two basic types of relationships [9]:

- Customer-provider relationship: If AS F pays AS C to have access to the broader interconnected network (as it is depicted in figure 1-1), AS F is a *customer* of the *transit provider* AS C and they establish a customer-provider relationship. Providers usually forward announcements from their customers to all its neighbors, as providers indeed benefit from and have a business obligation to send traffic to its customer. Similarly, most providers share with their customer their whole routing table, forwarding all their announcements.

- Peer-to-peer relationship: If AS C and AS D decide to exchange their own traffic freely (*i.e.,* neither of them pays the other to exchange traffic), AS C and AS D established a peer-to-peer[6] relationship. ASes usually only forward their customers and own routes to ASes they have a peer-to-peer relationship with and do not forward routes from their providers.

There are other types of relationships between ASes. Some ASes establish *complex relationships* that depend on the interconnection location and type of traffic. For instance, AS C and AS D could be customer-provider for some interconnections or some type of traffic and AS C would pay AS D for that, but also be peer-to-peer for

legacy addresses. Some of the first RFCs kept track of IP addresses (and other) delegations in the early days of the Internet [8]

[6] ASes' neighbors are also known as *peers*, which can bee confusing when talking about relationships given an AS might not establish peer-to-peer with all its BGP peers (neighbors).

other interconnections or type of traffic, and neither of them would pay the other for that. In addition, many organization have multiple ASes and use BGP to interconnect them. Two ASes from the same organization are know as *sibling ASes* [10].

Routing route-selection process takes into account AS relationships to comply with business agreements and also because AS relationship impact the provisioning of links between networks. If traffic is sent through under-provisioned links, the availability of Internet communication is impaired.

1.4 Networks in BGP

Currently, there are over 70,000 ASes exchanging information in BGP, and there are over 1 million IPv4 and IPv6 prefixes originated by these ASes. The size of ASes in terms of customer networks and address space they originate varies significantly, creating a very diverse ecosystem of networks.

The largest ASes have thousands of customer networks and exchange traffic between them for free. These ASes, known as *Tier 1 ASes*, also originate thousands of prefixes each. Examples of these ASes are AT&T (AS7018), Lumen[7] (AS3356) and Telia (AS1299). There are also content provider ASes that originates vast amounts of address space but that only provide transit to a few ASes, usually from the same organizations. Examples of these ASes are Amazon (AS16509) and Google (AS15169).

However, most ASes the Internet are small. About 60,000 ASes have no customer, and over 20,000 of those ASes originate only one prefix. ASes with no customers are known as *stub ASes* (AS F in figure 1-1 is a stub AS). Stub ASes can have one or multiple providers. If they have more than one provider, they are known as *multihomed ASes*.

2 Problems with BGP

The critical design flaw in BGP is that there is no validation of any information in BGP messages, impacting the correct delivery of IP packets to their destination. If no additional security measures are used, information in BGP updates is directly considered in the route selection process and then forwarded. Therefore, ASes in the Internet can intentionally or unintentionally announce incorrect and false information. In other words, any AS can originate any given block of IP addresses or modify the AS path when forwarding an announcement (*e.g.,* by claiming to be in the path or shortening the path). Given that there is no verification of the information in BGP announcements, wrong routing information can then spread through the network, modifying border router's forwarding table and therefore where packet are sent across the Internet.

Figure 1-2 illustrates a case of how wrong information can spread in BGP, where an AS originates an announcement for a prefix it has not been authorized to use. AS X wrongly originates prefix 18.0.0.0/16. When AS D receives this announcement, it must pick between two competing announcements for the same prefix: the one from

[7] Formely CenturyLink and Level 3.

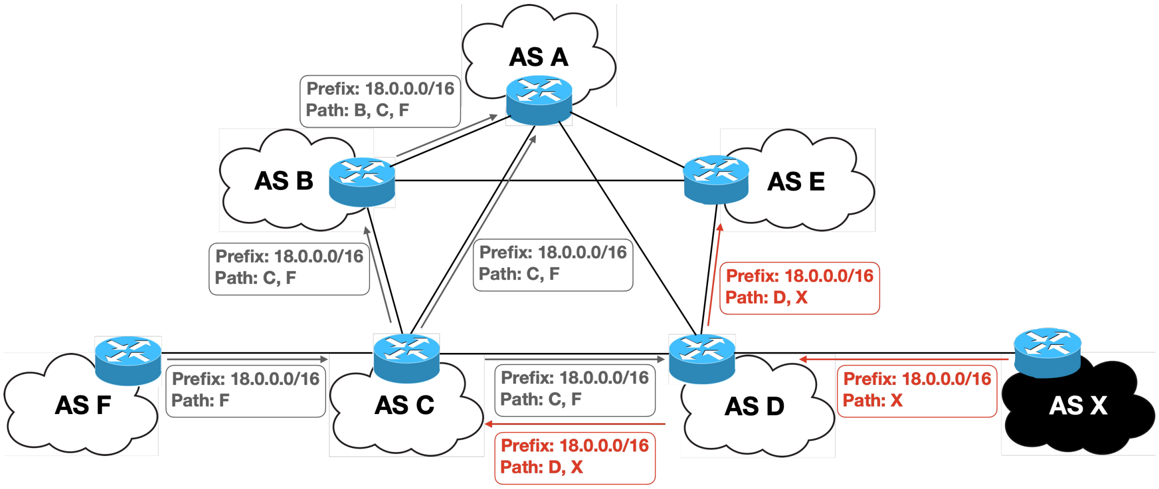

Figure 1-2: Prefix Hijack: AS X attempts to hijack address block 18.0.0.0/16 by originating an announcement for that prefix, thus falsely asserting that it has been allocated those addresses. When AS D receives this false announcement, it must pick whether to use and forward this one or the one it received from AS C. Since the announcement from AS X is shorter, it might choose to forward that one. Then AS D and any AS that accepts the announcement originated by AS X will forward traffic intended for 18.0.0.0/16 to AS X as opposed to AS F.

its neighbor AS X, or the longer one from AS F via AS C. If it picks the one with the shortest path and forwards it, a number of ASes may choose and continue to spread the wrong information. AS a result, many networks end up sending traffic for 18.0.0.0/16 to AS X.

An AS is said to do a *BGP hijack* when it *intentionally* introduces incorrect information in a BGP message. In a *prefix hijack*, an ASN appears as the origin of an IP prefix which that ASN is not authorized to use in BGP. Figure 1-2 represents the case of a prefix hijack by AS X. In a *path hijack*, an ASN appears in the path to an IP prefix when it should not appear, potentially having modified part of that path. For instance, in figure 1-2, AS X could send an announcement for prefix 18.0.0.0/16 with path X,F even if there is no link between AS X and AS F. In the false path X,F, the origin is the legitimate owner of the prefix, but the second hop is not correct. This wrong BGP announcement can still spread through BGP even though it is not shorter than the correct announcement. For instance, if AS X is a customer of AS D and AS C is a peer or provider of AS D, AS D will likely choose the route to prefix 18.0.0.0/16 through its customer AS X, and spread that wrong information to its neighbors. In both prefix and path hijack, there is nothing that the victim (AS F in the example in figure 1-2) can do to stop the wrong announcement, beyond reaching out to AS X and other ASes and asking them to stop sending or forwarding the wrong announcement.

Hijacks are used by malicious actors for different purposes. The simplest harm that can result from incorrect information in BGP is that traffic goes to the wrong part of the Internet, where it is then discarded. This outcome leads to a loss of availability between the intended communicating parties. As a recent example at the time of writing, on February 5th 2021, the main network provider from Myanmar hijacked Twitter address space in BGP, following a request by the Myanmar government to prevent people from accessing the service in the days following the military coup. In addition, hijackers can also fake an end point or abuse addresses from the hijacked prefixes. For example, on April 24th, 2018, a small network from Ohio hijacked Amazon address space by advertising blocks of IP addresses from Amazon with AS10297 as the origin. This prefix hijack lasted about two hours, during which crypto-currencies from a crypto-wallet housed in Amazon Route 53 Cloud Services were stolen [11–13]. Similarly, BGP hijacks have been used to abuse IP addresses as part of an online advertisement fraud scheme [14] and to send spam campaigns [15].

BGP's lack of information validation makes it challenging to limit the spread of misconfigurations. Sometimes, an AS unintentionally sends information to its BGP neighbors that it should not have shared (*e.g.,* it was intended for its internal use only) given the business agreement and AS relationship with its neighbors. This may put this ASN as the origin or in the path to routes that should not include that ASN and which end up modifying packet forwarding in the Internet, in the same way that intentional hijacks do. These events are known as *route leaks.* As an example, on June 24th, 2019, a major route leak happened when a small network from Pennsylvania leaked over 65,000 internal routes to its provider, who then forwarded them to Verizon, quickly reaching most of the Internet. As a consequence, traffic to core Internet services (*e.g.,* Domain Name System servers) as well as major content providers such as Amazon, took a detour via the smaller networks causing severe congestion in their links, impairing service availability [16]. Similar route leak events have taken large portion of European traffic through China for hours before reaching their final destination in the US or back in Europe, and have also sent traffic for major US providers such as Google and Akamai through Russia or Brazil [17, 18].

This critical security flaw with BGP has been known for a long time. Indeed, BGP inherited this flaw from the precursor protocol called Exterior Gateway Protocol (EGP). Already in 1982, it was mentioned in RFC 827 that false information sent in EGP could have great impacts on traffic by sending it the wrong way [19]. Nonetheless, at the time BGP was designed, there were only a few hundred networks [20] and most of the people behind them probably knew each other. However, this design flaw has been considered a critical vulnerability that needs to be addressed for a long time.

3 Overview of BGP Security

This section provides a high level overview of BGP security, giving pointers to parts of upcoming chapters where specific areas of BGP security are discussed in more detail.

There have been many proposals to secure BGP from the IETF, industry and academic communities [21–38]. To secure BGP, some solutions propose changes to the protocol, other add (optional) components or build on an overlay network. Most security proposals include a verification step to validate (parts of) the information sent in BGP announcements using cryptographic records. However, solutions use different methods to build and distribute authoritative routing information. Chapter 2 discusses the main proposals, their differences and similarities, the ones that have been implemented and their deployment and use (if any).

Given the number of different proposals, previous work has also focused on evaluating different aspects of security proposals such as performance and efficiency [39, 40], limitations and advantages concerning security guarantees [41], techniques used to secure BGP [42], dynamics of their architecture [43], and the incremental security benefit solutions bring in partial deployment and fully deployed. [44, 45] . Additionally, in [46] the authors study in detail the work of the Secure Inter-Domain Routing working group (SIDR-WG) at the IETF,[8] which has developed many standard track RFCs describing the parts of frameworks to secure BGP and their use. Although most proposals to secure BGP have been published for many years and dynamics impairing their deployment have been studied, it is still surprising that no clear direction has emerged for BGP security. Indeed, even though network operators do consider BGP hijacking to be a concern [47], little has changed in operational environments.

Other bodies of work in BGP security focuses specifically on monitoring and detection of hijacks and misconfigurations from BGP announcements [48–61]. Some monitoring solutions have been implemented and produce feeds of suspicious activity in BGP based on public and private collection of BGP updates [55, 62]. There are works that study behaviors seen in BGP related to specific types of networks (*e.g.,* spammers [15] or Bulletproof hosting ASes[9] [63]), of resources [64] , or of BGP updates from given peers [65]. The literature related to BGP hijacking is discussed in Chapter 3 and the literature related to detecting route leak misconfigurations is discussed in Chapter 4.

In 2012, the IETF standardized the Resource Public Key Infrastructure (RPKI), a framework developed by the SIDR working group. The RPKI is currently the most deployed and used framework to secure BGP, although it has taken a long time for ISPs to adopt its use. Increasingly, many work study different aspects of the RPKI. This scheme and the related literature is discussed in Chapter 5.

Despite the previous work studying different issues related to BGP security, the systemic impact of BGP design flaws at scale remains ambiguous. Indeed, Internet routing seems to “just work" most of the time. The approximate volume of hijacks and misconfigurations in BGP and how these volumes have evolved over time are open questions. Similarly, the benefits of security proposals in production and the impact of the related operational practices are unknown. The lack of understanding of these issues at scale has made it hard for consensus on a direction to emerge and to incentivize any effort to improve BGP security at all levels —from operators to policymakers. This dissertation seeks to start answering these questions and provide clarity about the systemic impact of BGP security issues and solutions.

[8] `https://datatracker.ietf.org/wg/sidr/about/`

[9] ASes dedicated to hosting cyber-criminals

4 Roadmap

This section provides more details about specific studies carried in each chapter and the main contributions.

Chapter 2: Historical Review of BGP Security

Chapter 2 offers a historical review of the different ideas put forward to secure interdomain routing. It examines where the ideas came from, implicit trust choices, the required infrastructure and the residual vulnerabilities of proposals. Reviewing these aspects gives insight into why adoptions rate are so limited. Many solutions use similar cryptographic techniques to authenticate and verify data. Nonetheless, there is a remarkable lack of consensus on what needs to be secured or validated, and the approach to be taken to build and distribute the authoritative database. These disagreements have prevented solutions to get critical support for deployment.

Chapter 3: Empirical Analysis of Malicious Behavior

Chapter 3 focuses on systemic malicious behavior coming from networks that repeatedly perform hijacks over time, with the goal of evaluating the pervasiveness of malicious activity in BGP and the types of hijacking malicious actors are perpetrating. Out of the more than 70,000 networks in the Internet, this work finds about 800 networks with suspicious behavior. Further scrutinizing these networks, it finds significant evidence of malicious behavior, as well as misconfigurations and false positives linked to benign forms of hijacking. The findings reveal the existence of BGP *serial hijackers*—networks that persistently perform hijacks in BGP, showing that there are few barriers to performing even the basic forms of routing attacks, with almost no consequences to operators. Finally, this work narrows the focus on a small set of suspicious networks to the point that fully automated detection of serial hijackers and network reputation scoring systems can be envisioned in the future.

Chapter 4: Empirical Analysis of Misconfigurations

Chapter 4 aims to characterize the prevalence of route leak events at the edge of the Internet that shift core traffic, creating bottlenecks and impairing the availability of IP prefixes involved in those events. It builds a tool to detect and monitor route leaks over time based on AS path characteristics, using network centrality metrics to identify paths that violate business structure. The findings illuminate the characteristics of identified route leak activity with regards to timing, size and spread and the dynamics of route oscillation. The results provide an extended route leak dataset to study the impact of specific protocol configurations in the spread of such misconfigurations, the main victims and involved networks.

Chapter 5: Empirical Analysis of Defenses

Chapter 5 focuses on the use of defense mechanisms to stop the spread of hijacks and misconfigurations in BGP and their effectiveness. Finding that the adoption of the Resource Public Key Infrastructure (RPKI) framework standardized by the IETF in 2012 finally gained traction in 2019 and 2020, this work centers on this framework. It first measures the changes in the amount of invalid BGP information forwarded by networks using RPKI to validate routing information they receive in BGP. Then, it evaluates the impact this adoption has on the overall spread of incorrect routing information. It finds that even when less than 10% of networks have adopted this practice, the spread of invalid and potentially illicit announcements in BGP is reduced by 10-15%. As many network operators and researchers have expressed that benefits from RPKI require almost full adoption, this result—based on real-world evidence—is game changing for understanding and advocating for RPKI and ROAs adoption. This work also increases the incentive of providers to deploy operational security by providing a method to passively track operator RPKI usage over time, able to identify different settings and problems with this practice. Finally, it examines technical and non-technical barriers that influence the operational practices and the outcome of the use of RPKI information for routing decision.

Chapter 6: Insights to Secure Internet Routing

Chapter 6 examines technical and non-technical barriers that have hindered the adoption of security proposals, using the insights extracted from the empirical evidence of the previous chapters. It also discusses empirical evidence that brings to light how these barriers play out at scale, and proposes future work that can help overcome the barriers. Finally, it discusses a set of actions that would facilitate the emergence of a common direction and encourage collective adoption of security in BGP.

Chapter 2

Historical Review of BGP Security

This chapter presents a survey of IETF Request for Commment (RFC) documents formalizing the Border Gateway Protocol (BGP) and BGP additional features and BGP security proposal put forward by industry, academia and standardization organizations. The objective is to identify what has stalled progress in the adoption of BGP security.

Previous work has evaluated different aspects of security proposals, including performance and efficiency [39, 40], security guarantees [41], techniques used to secure BGP [42], dynamics of their architecture [43], and their security benefit in partial and full deployment [44, 45]. Building on this previous work, this study analyzes security proposals with the goal of identifying the main barriers and disagreements preventing their implementation and adoption. It finds long-lasting disagreements and lack of systemic review of non-technical aspects of proposals such as the choice of trust and the management and operation of support infrastructure. This pervasive lack of consensus has not allowed a common direction for routing security to emerge, holding back the full implementation and deployment of BGP security proposals. An original version of this work appeared in [66].

This chapter starts by describing the methodology and data used for the analysis in Section 1. Then, Section 2 discusses the evolution of BGP and awareness of its flaws over time. Section 3 reviews proposals to secure BGP from the IETF, academia and industry. Section 4 compares and discusses the proposals life-cycle, the main differences between proposal and the disagreement that have hold back consensus. Section 5 dives into the challenges of validating information in BGP. Finally, Section 6 summarize the conclusions and contributions.

1 Methodology

The goal of this study is to identify the main barriers and disagreements stalling the adoption of better BGP security. The work is divided in two steps. First, it aims to better understand the evolution of BGP and awareness of its security issues. Thus, it studies the main purpose, motivation and discussion of security issues in IETF Request For Comment (RFC) documents describing all version for BGP as well as optional additional features of the protocol.

Second, this work aims to identify disagreements between BGP security proposals and related framing or discussion in further research works through the review of security proposals and related surveys. The security proposals considered are the ones that had some traction and are specifically focused on improving BGP, not any generic routing algorithm (although they may use aspects from those works). For IETF proposals, the traction meant that RFCs were updated, mentioned in BGP protocol updates or discussed in operational practice documents. For industrial and academic proposals, the traction was assessed from consideration in BGP security literature surveys [41–43, 46] as well as other works studying specific aspects of BGP security proposals such as performance and efficiency [39, 40], security guarantees [41], techniques used to secure BGP [42, 43], and their security benefit in partial and full deployment [44, 45]. The analysis infers the different aspects, goals and motivations of proposals from the main documents, complemented with information from accompanying documents and surveys.

2 BGP evolution and vulnerabilities

The function of the routing protocol in Internet communication, how BGP works and its main design flaws are described in Chapter 1 Sections 1 and 2. This section describes the evolution of the protocol as evidenced by RFCs related to BGP over the years and the vulnerabilities discussed in RFCs.

2.1 BGP evolution

Since the first RFC describing BGP authored in 1989, there have been four versions of BGP, with BGP-4 having had two major updates [3–6,67,68]. From the beginning, the main goal of BGP has been to exchange network reachability information between ASes. Figure 2-1 is a timeline of RFC documents with the main BGP protocol documents on the left side and the extensions and their respective updates on the right side. As can be seen, ever since the first formalization, BGP has been in constant evolution. Either the main protocol itself or the many extensions have been modified or created to accommodate the evolving usage of BGP and the evolution of the routing ecosystem of the public Internet.

A major change to the protocol was made in version 4. In previous versions, networks were advertised according to the hierarchical class system. Instead, BGP-4 supports classless inter-domain routing, which means available networks are identified by an IP address prefix and a prefix length. It also meant that longest prefix match —the more specific network match— became the base behavior for forwarding data packets.

Additionally, the first versions of BGP (BGP-1, BGP-2 and BGP-3) had the option of including authentication data in messages. However, no specific authentication mechanism for that option was ever formalized in RFCs. In BGP-4, the authentication option was deprecated as it was not being used.

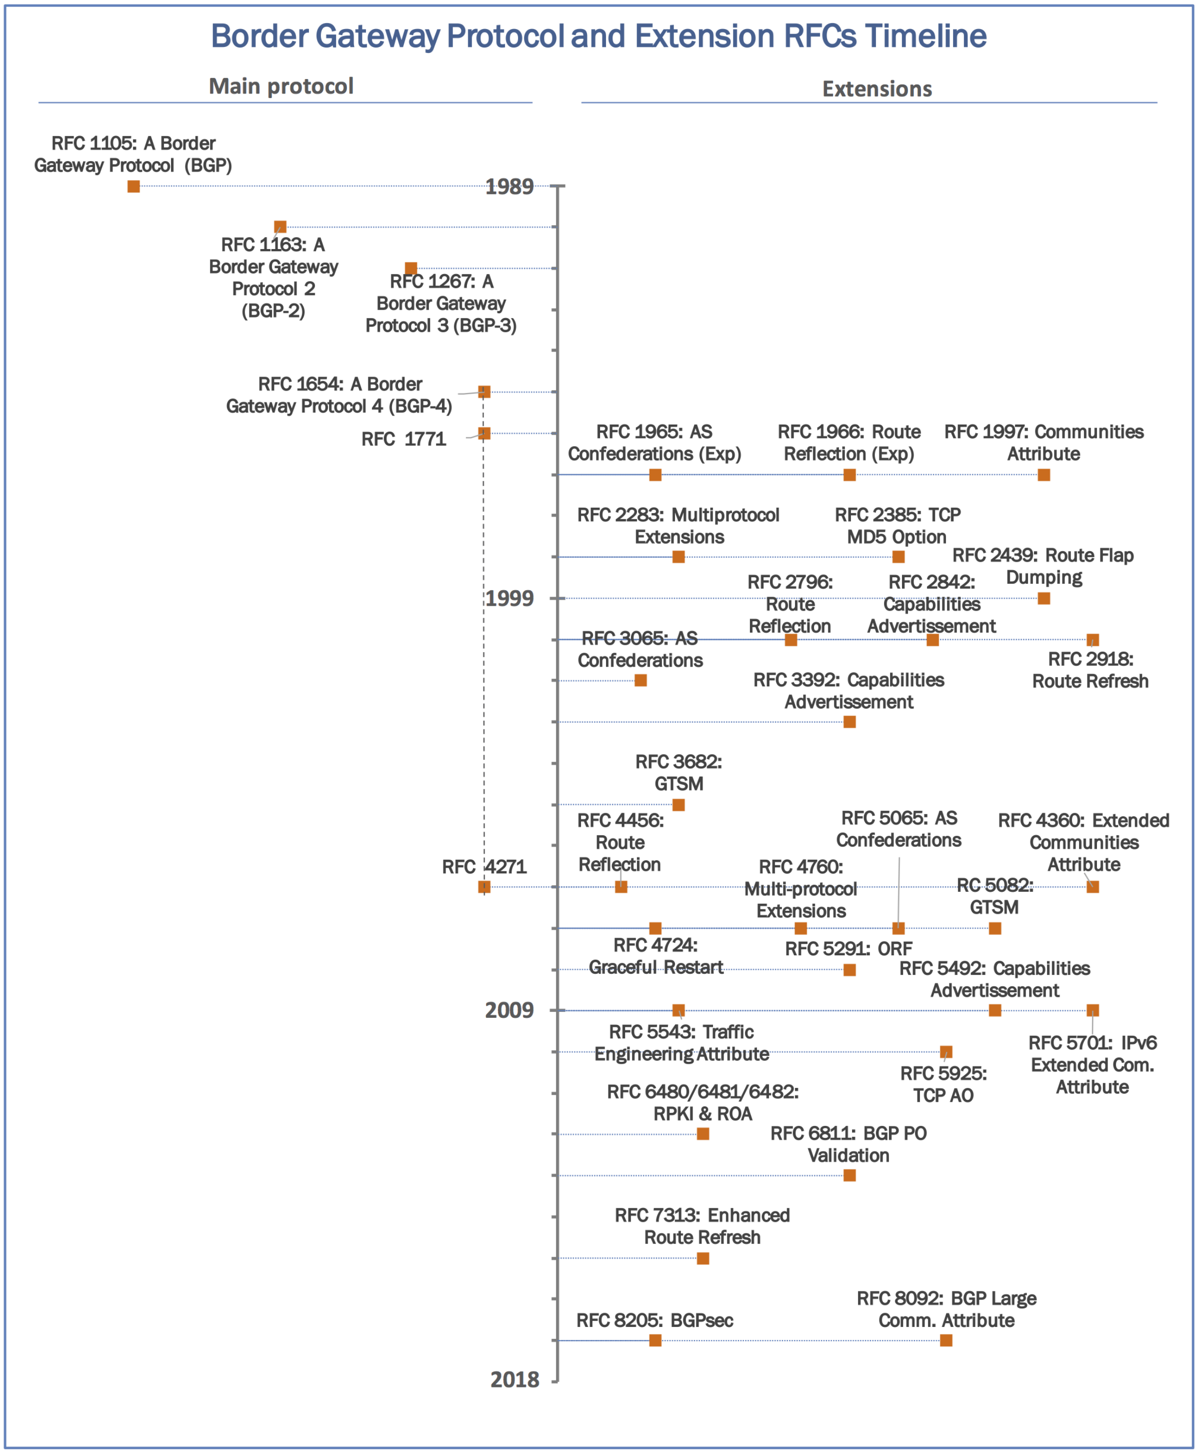
Border Gateway Protocol and Extension RFCs Timeline
Main protocol
Extensions
RFC 1105: A Border Gateway Protocol (BGP)
1989
RFC 1163: A Border Gateway Protocol 2 (BGP-2)
RFC 1267: A Border Gateway Protocol 3 (BGP-3)
RFC 1654: A Border Gateway Protocol 4 (BGP-4)
RFC 1771
RFC 1965: AS Confederations (Exp)
RFC 1966: Route Reflection (Exp)
RFC 1997: Communities Attribute
RFC 2283: Multiprotocol Extensions
RFC 2385: TCP MD5 Option
RFC 2439: Route Flap Dumping
1999
RFC 2796: Route Reflection
RFC 2842: Capabilities Advertisement
RFC 3065: AS Confederations
RFC 2918: Route Refresh
RFC 3392: Capabilities Advertisement
RFC 3682: GTSM
RFC 4456: Route Reflection
RFC 5065: AS Confederations
RFC 4360: Extended Communities Attribute
RFC 4271
RFC 4760: Multi-protocol Extensions
RC 5082: GTSM
RFC 4724: Graceful Restart
RFC 5291: ORF
RFC 5492: Capabilities Advertisement
2009
RFC 5543: Traffic Engineering Attribute
RFC 5701: IPv6 Extended Com. Attribute
RFC 5925: TCP AO
RFC 6480/6481/6482: RPKI & ROA
RFC 6811: BGP PO Validation
RFC 7313: Enhanced Route Refresh
RFC 8092: BGP Large Comm. Attribute
RFC 8205: BGPsec
2018

Figure 2-1: Timeline of standard track IETF RFCs of the Border Gateway Protocol (on the left side) and extension (on the right side), noting the ones considered experimental (exp). Horizontal lines indicate RFC documents year of publication.

According to BGP RFCs, the fundamental priorities for the development of the main BGP protocol are:

- Ability to enforce destination based AS-level policies: each AS needs to be able to enforce the policy of choice concerning destinations reachable through its network.

- Scalability and efficiency: BGP needs to be able to handle the increasing number of prefixes advertised in the public Internet without using too much traffic.

- Dynamic routing while limiting convergence time: BGP needs to accommodate frequent changes to reachability information without taking too long in its route selection process.

- Identification of routing loops: BGP needs to provide a mean to prune routes from the route selection process if a loop is identified to prevent endless looping of data packets.

- Limiting manual configuration of routing policies: BGP needs to allow for automated decisions based on general policies configurations.

- Flexibility for complex and creative routing policies: BGP has to be flexible to accommodate new developments to support policy routing.

When the fourth version of BGP was introduced in 1994, BGP had been deployed in different networking environments and many independent interoperable implementations existed. All inter-domain routing has been done using BGP since the early 1990s.

Since BGP-4 was introduced, more than 15 extensions have been developed and standardized in IETF RFC documents. Extensions are depicted in the right side of figure 2-1. These extensions add new capabilities to the base protocol. Most BGP extensions were designed to improve and extend BGP operation, to facilitate policy management, or to provide some level of security to the protocol.

2.2 BGP design flaws awareness

As experience with BGP-4 accumulated, awareness and efforts to overcome of its shortcomings and vulnerabilities developed. The main aspect of security discussed and addressed in BGP-related RFCs is availability. There is an understanding that BGP has to provide as much availability as possible, including in the case of an attack or unintended failure. Consequently, to increase availability, most BGP extensions aim at reducing management complexity and manual configuration of BGP speakers —which are prone to unintended failures— and routing instability, given the increasing complexity of network topologies and routing policies.

Concerning the correctness of information sent in BGP, there are two main categories of security flaws:[1]BGP's lack of ability to guarantee the integrity of information

[1] This flaws were initially mentioned in relationship to a precursor protocol [19] and were later thoroughly described in a BGP vulnerability analysis RFC [69].

sent between peers, and BGP's lack of ability to guarantee the correctness of routing information exchanged between peers. BGP runs on top of TCP/IP connections, which are not secured by default. Therefore, a third party could interfere with the communication between two BGP routers. Moreover, BGP lacks any mechanism to validate the information sent in BGP messages, which means that neither the IP prefix, nor any AS in the path nor any optional attribute included in BGP announcements is verified before routers integrate that information in their route selection process. BGP inherited the lack of validation from the precursor protocol called Exterior Gateway Protocol (EGP) and the risk were described in EGP's RFC [19]. In addition, as noted earlier in Section 2.1 the first three versions of BGP included an authentication mechanism to prevent a third party to act as a legitimate BGP speaking router and inject false information. However, even using the authentication mechanism for BGP communcation, given that BGP tuns on top of TCP, attackers could still interfere with the communication. In addition, as no mechanism was formalized to do the authentication and networks operators were not using this option, the authentication option was removed in the last version of BGP.

3 Efforts to secure BGP

This section summarize security additions to BGP considered in this study.

3.1 IETF security proposals

At the IETF, there have been many efforts to improve BGP operation, including two working groups addressing routing security, the Routing Protocol Security Requirements (2002-2009) and the Secure Inter-Domain Routing (2006-2018) working groups. The work that was documented in RFCs is summarized below.

3.1.1 BGP extensions for communication security

The first BGP extensions to address routing information security focused on securing the communication between BGP peers setting up BGP sessions. At a high level, the threat model these security extensions consider is the ability of an attacker to interfere in the communication between peers. However, they imply different attack vectors and capabilities. Three BGP extensions described in RFCs fall in this category:

- The TCP-MD5 option was introduced in 1998. It does not involve changes to BGP, but rather uses a TCP option for carrying an MD5 digest that is used to verify TCP packets integrity using a password known to both ends of a BGP session.

- The Generalized Time-To-Live (TTL) Security Mechanism (GTSM) was introduced in 2004 as simple hack to increase BGP transport level security without the burden of TCP-MD5 configuration. Again, this extension does not modify the BGP protocol itself but rather uses TCP features to allow BGP speakers to verify that a received BGP message was sent by a router a hop away.

- The TCP Authentication Option (TCP-AO) obsoleted the TCP-MD5 option in 2010, as the MD5 cryptographic algorithm was considered a weak mechanism. Nowadays, all BGP implementations are required to support TCP-AO [36]. TCP-AO uses a scheme similar to TCP-MD5 but with a stronger message authentication mechanism and a re-keying option to update secret keys without manual configuration.

All of these extensions are currently implemented and deployed. However, even if deployed into border routers, most ISPs do not use these optional extensions. In particular, GTSM use is limited to simple topologies. In 2015, even though TCP-MD5 had been obsoleted and replaced by TCP-AO, it was more widely deployed and used than TCP-AO [70], and many ISPs had never changed TCP-MD5 password since they started to use it years ago [71]. The most recent Best Current Practice for BGP security [70] recommends the use of GTSM in direct peering links and that TCP-AO should be preferred to TCP-MD5 when implemented. However, the document also recommends operators to consider the operational burden and computational cost of using TCP-AO as it might not be suited for all operational environments and could significantly impair performance. .

3.1.2 Validation of information in BGP

In 2006, the IETF started the Secure Inter-Domain Routing (SIDR) working group [72] to address routing information vulnerabilities in BGP. The group developed two solutions for validating routing information sent in BGP announcements: the Resource Public Key Infrastructure (RPKI) with Route Origin Authorizations (ROAs) for validating IP prefixes and origin ASes in BGP announcements, and BGPsec, a mechanism for providing path security. The threat model for both proposal is a network sending incorrect information in BGP. RPKI and ROAs prevent the announcement of an IP prefix with an origin AS that has not been authorized to announce such prefix in BGP. BGPsec prevents the announcement of an IP prefix with a path that was not one explicitly authorized at each hop. The next paragraphs describe in more detail the solutions standardized by the IETF.

RPKI and ROAs for prefix origin validation in BGP

The infrastructure developed by the IETF to validate IP prefix and AS origin in BGP announcements is part of the Resource Public Key Infrastructure (RPKI) and was published in 2012 [37, 73, 74]. Using the RPKI, organizations that have been delegated IP address blocks by Regional Internet Registries (RIRs) can issue assertions stating the Autonomous System Number (ASN) of networks authorized to originate those address blocks in BGP.[2,3]

This IETF framework to secure BGP is based on a Public Key Infrastructure (PKI) and provides cryptographically verifiable attestations of number resources al-

[2] There are many IP addresses that were in use by organizations before the creation [8].

[3] See Chapter 1 Section 1.2 for more details on how BGP works, prefixes, origins and paths.

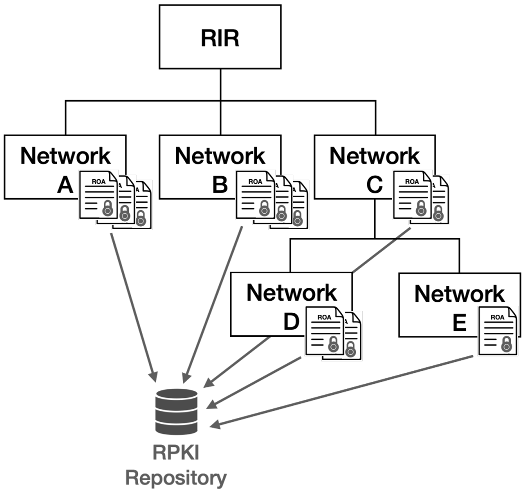

Figure 2-2: Example of hierarchical delegation of IP address blocks and Route Origin Authorizations. The RIR is the root of trust and delegates IP address space to networks A, B and C, which can then issue Route Origin Authorizations (ROAs) for their respective address space. Network C further delegates parts of its address space to networks D and E. ROAs are then accessible through the RIR's RPKI repository.

locations [37], *i.e.,* the allocations of IP address blocks and AS numbers (ASNs). The five RIRs are the *roots of trust* and manage the RPKI repository of records for their geographical zone.

To use the RPKI to validate IP prefix and origin AS in BGP announcements, two distinct step are needed: (1) organizations need to issue Route Origin Authorizations for their delegated address space indicating the ASN authorized to originate that address space in BGP (Figure 2-2); and (2) border routers in networks need to have access to a cache of ROAs to validate data in BGP announcements.

When an RIR delegates an IP address block to an organization, it can also give the organization the authorization to sign Route Origin Authorizations (ROAs) that authorize (parts of) the IP address block to be originated by the specified ASN. The ROAs are added to the respective distributed repository system. Other networks can fetch those records and cryptographically verify the ROAs following their delegation tree until the root of trust (on of the five RIRs).

Figure 2-2 depicts an example of an RPKI delegation tree representing the hierarchical delegation of IP address blocks and the issuance of ROAs for the delegate address space. The RIR is the root of trust and delegates IP address space to networks A, B and C. These networks can then issue ROAs for all or part of their delegated address space to indicate the ASN authorized to originate that address space in BGP. Network C further delegates parts of its address space to networks D and E and thus authorize those networks to issue their ROAs for their address space. ROAs are then accessible through the RIR's RPKI repository.[4]

To validate BGP information using RPKI ROAs, networks need to fetch ROAs

[4] ROAs can be stored in RIRs' RPKI repository or stored by networks and indexed and linked to from the RPKI repository

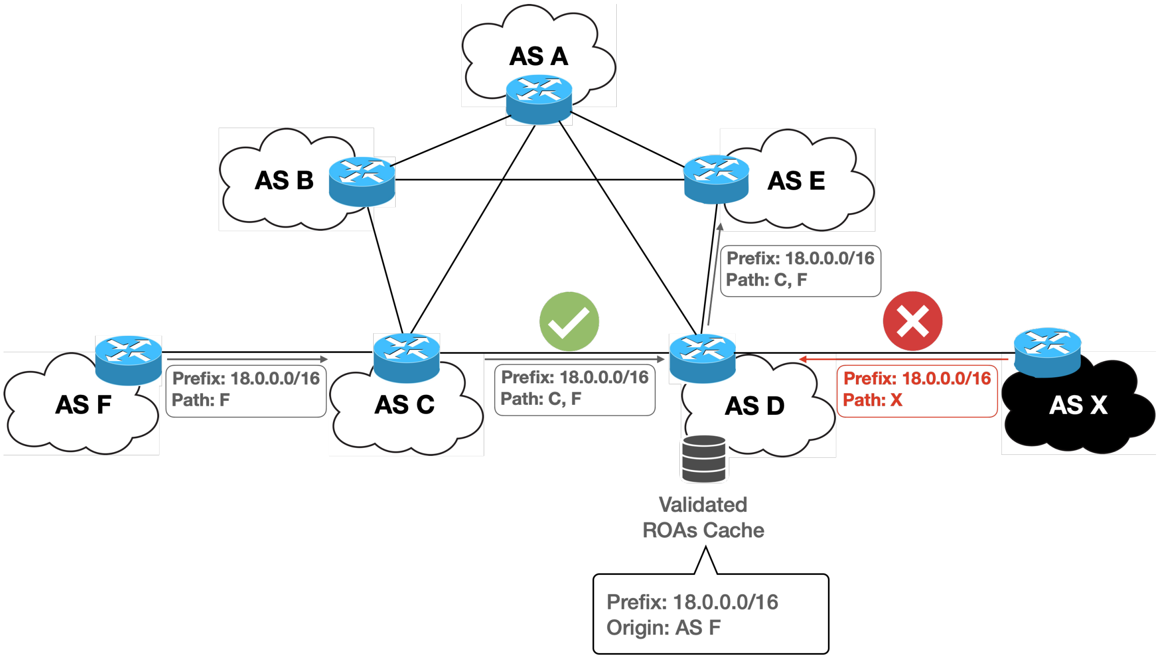

Figure 2-3: Example of Route Origin Validation. Network D has an accessible cache of validated ROAs including the ROA for prefix 18.0.0.0/16. AS X is not authorized to originate 18.0.0.0/16 and thus the announcement from AS X fails the validation by AS D and is dropped. In contrast, the announcement from AS C to AS D for 18.0.0.0/16 is validated as F is the authorized origin AS.

from RPKI repositories and cryptographically verify them following the delegation tree. This process creates a cache of validated ROAs that networks need to make accessible to their border routers. Then, using the information from validated ROAs, routers can perform the Route Origin Validation (ROV) by checking that the AS at the origin of a BGP announcement appears in a ROA for the related address space and is thus authorized to originate it in BGP.

Figure 2-3 illustrate the ROV process where network D validates BGP announcements it receives from IP address block 18.0.0.0/16. The border router in AS D has access to a cache of ROA previously cryptographically verified, which includes the ROA for prefix 18.0.0.0/16. AS F is authorized to originate 18.0.0.0/16 in BGP. Thus, the announcement from AS X is invalid since AS X appears to be the origin. The border router drops that announcement, preventing its further spread. In contrast, the announcement form AS C is valid because AS F originated it. AS D then forwards that advertisement to AS E.

The RPKI infrastrucure and ROAs do not directly introduce any change in BGP, the prefix origin validation is taken into account through router configuration. However, networks need to fetch, verify and make accessible validates ROAs to their border routers.

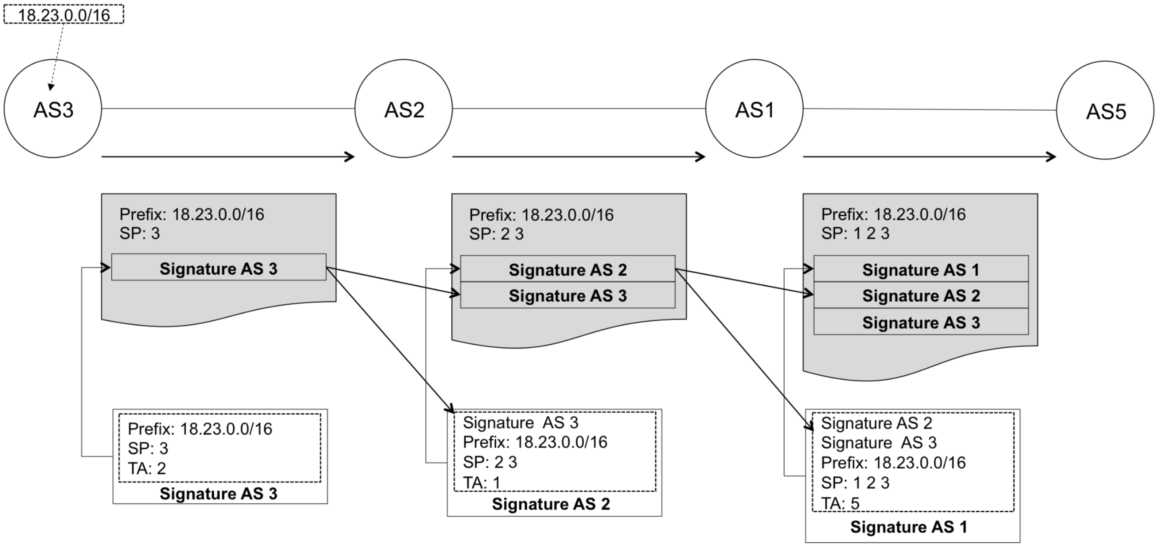

Figure 2-4: Simplified representation of BGPsec announcements of a route being forwarded through three ASes. Each AS adds its signature of the Secure Path attribute. The signature is computed using all previous signatures, the Secure Path attribute itself and the ASN of the AS that will receive the BGPsec message.

BGPsec for AS path validation

To validate AS path information in BGP announcements, the IETF SIDR working group developed BGPsec. It was published as a standard in September 2017 [38,75]. BGPsec was designed to give assurance to ASes receiving a BGPsec message and verifying all cryptographic signatures, that the ASes listed in the BGPsec path have explicitly authorized the advertisement of the route to the subsequent AS in the path. The BGPsec protocol and its operation relies on the RPKI infrastructure for the storing and distribution of signed objects used in the validation of the AS path. Using specific BGPsec records fetched from the RPKI databases, ASes receiving a BGPsec message can validate all signatures and thus the path from the latest AS until the origin of the path.

BGPsec introduces a new path attribute to include the digital signature of all ASes in the path, binding the prefix and BGPsec path along the way. Figure 2-4 shows a simplified representation of a BGPsec announcement for prefix 18.23.0.0/16 originated by AS 3 and forwarded to AS 2 and then re-advertised to AS 1, and then AS 5. Each AS in the path prepends its AS number in the new Secure Path attribute along with its signature. ASes compute the signature using all previous signatures, the Secure Path attribute and the target ASN–the ASN to which the BGPsec message will be sent.

For validating a BGPsec message with routing information, all included signatures need to be validated. This can use significant computation resources, especially for long AS paths. Border routers implementing BGPsec can perform the signatures validation by themselves or use an offline validation mechanism similar to RPKI with ROAs [76]. Equality important, an AS using BGPsec can advertise route announcements even if the verification fails. After validating the BGPsec Update message, it is expected that the BGPsec router will include this result in the route selection process. However, this is left as a matter of local AS policy for each AS.

3.1.3 Secure operational practices

BGP security related work at the IETF also includes the formalization of operational practices that increase BGP security by limiting the spread of incorrect information. In February 2015, the IETF published a set of best current practices (BCP) for BGP operations and security, BCP 194 [70], which included two operational practices to validate routing information sent in BGP: Route filtering and the Internet Routing Registries.

Route filters are mechanisms used to discard routes either received from or to be forwarded to another AS based on the IP prefix address, the AS path or other attribute of the routes. They are widely use for enforcing routing policies (*e.g.,* to forward BGP announcements to customers but not provider ASes). ASes also use filters to drop BGP announcements containing known false information (*e.g.,* reserved address space) and their prevent the propagation.

The Internet Routing Registries (IRRs) are a distributed public registry of routing information for network operators. Some of the registries in the IRRs are managed by RIRs but most are managed by private entities. Networks can add objects in the IRRs to document the prefixes they originate in BGP and routing policies between with other networks among others. The IRRs use a Routing Policy Specification Language (RPSL). Even though the IRR is compose of about a dozen independent registries, RPSL allows to globally assemble the objects in the different registries in a single routing registry [77]. However, the organization managing the IRRs do not validate the information in their registry and different registries can have conflicting records [78].

Route filters and the IRR are old ideas. Policy filtering has been in place since the NSFNET was the backbone of the Internet [79] and their use can be extended to reduce the spread of incorrect information in BGP. The RIPE community [80] had been using registries for routing information from networks operators before the creation of the IRRs, and their experience was the base for the design of IRRs using RPSL. However, some network operators oppose this practice because of its management overhead and the risk of mistakes disrupting traffic. . Nonetheless, some ISPs, especially in Europe, require their peers and customers to have their information up-to-date in an IRR registry.More recently, in 2014, a group of network operators signed a document called *Mutually Agreed Norms for Routing Security (MANRS)* describing actions that ISPs should take to increase routing security and the use of route filters and the IRR are two of them [**?**]. Still, filtering and the use of the IRR was inconsistent among ISPs [70], even though it has been shown that using prefix filtering with origin validation techniques provides comparable security to origin and AS path validation while AS path validation mechanisms are still in deployment [45].

3.2 Academia and industry security proposals

There has been significant work looking to secure BGP outside the IETF. Based on the security proposals considered in different surveys and evaluation work, the list below describes their main aspects.

- Securing BGP AS Path with Predecessor information: In 1996, Smith and Garcia-Luna-Aceves published one of the first set of measure to protect BGP [21]. Their proposal is based on cryptographic mechanisms to provide confidentiality and integrity to BGP communication. It also includes changes to BGP messages to allow AS to validate IP prefixes and AS paths of BGP announcements.

- Secure Border Gateway Protocol (S-BGP): S-BGP was developed by a group at BBN Technologies and was published in 2000 [22]. It uses the IPsec cryptographic mechanism [81] to protect BGP messages integrity and confidentiality. Additionally, it relies on a Public Key Infrastructure (PKI) where RIRs are the roots of trust. It also adds a new path attribute to send AS signature in route announcements that can be validated by each AS along the path.

- Hop Integrity for routing security: In the early 2000's, the Hop Integrity protocol was developed by Gouda *et al.* in collaboration with IBM research Labs [23]. It uses a cryptographic mechanism through the support of two new protocol layers that need to be added in the protocol stack of routers to provide integrity check and exchange new secret keys smoothly.

- Origin lists for false origin detection: In 2002, Zhao *et al.* proposed the use of a list with IP prefixes and the ASNs authorized to advertise them in BGP. This list would be sent along announcements using of an existing BGP optional attribute. Then, similar to the RPKI and ROAs, network operator should validate information in BGP using the records in the list.

- Secure Origin BGP (soBGP): In the early 2000s, the secure origin BGP (soBGP) protocol was developed by a group mostly within Cisco System and was published in 2003 [25]. The soBGP protocol considers the use of cryptographic mechanisms to secure BGP communication. It also relies on a PKI infrastructure, but the root of trust for IP allocation is a "small number of well-known entities" that would then issue certificates to other ISPs and organizations, forming a "web of trust" for allocating IP prefixes [25]. soBGP also requires AS to publish a list of BGP peers and the set of policies that the origin AS would like to apply to the route announcement of an IP prefix. It also includes a new type of BGP announcements to distribute certificates among ASes.

- The Interdomain Route Validation (IRV) Protocol: Also in the early 2000's, a group working for the ATT research lab developed the IRV Protocol and published it in 2003 [26]. It is based on a decentralized query systems that connects ASes and is used to verify routing information from BGP. ASes may host or designate an IRV database to speak authoritatively about their network status and routing information.

- The Secure Path Vector (SPV) Protocol: Researchers from UC Berkeley and Carnegie Mellon University published the Secure Path Vector (SPV) Protocol in 2004 [27]. SPV protocol is based on a series of cryptographic mechanisms that authenticate ASes, allows ASes to authorize route announcements and ensures message freshness through certificates expiration. SVP also relies on hierarchical certificate structure equivalent to the Public Key Infrastructure used in S-BGP [22] for allocating IP prefixes.

- Listen and Whisper: In 2004, researchers at UC Berkeley presented the two mechanisms *Listen* and *Whisper* to improve BGP security [28]. The Whisper protocol is a monitor system based on a signature scheme that is included in BGP announcements. If a BGP speaker receives two routes to the same IP prefix with the origin, it can verify if the signatures are consistent and choose the route coming from the AS with the lower penalty metric.

- Pretty secure BGP (ps-BGP): In 2004, the Pretty Secure BGP (psBGP) was presented by researchers from Carleton University based on the analysis of S-BGP and soBGP [29]. psBGP also uses IPsec [81]. To authenticate AS numbers, psBGP relies on a centralized PKI like S-BGP [22]. To validate the allocation of IP prefix, each AS creates a list of the prefixes it originates and additional lists for the prefixes its peers originate. psBGP also requires each AS in the path to append its signature of the route information to be sent in the route announcement.

- External Security Monitors (ESM) to secure BGP: In 2006 researchers from Cornell University published a mechanism to use an overlay network of ESMs to monitor and secure BGP traffic, verifying the correct modification of the AS path at each hop and check origin authentication certificates [30]. The authors propose the use of a decentralized PKI infrastructure called Grassroots [82], where ASes are able to directly issue their certificates for their prefixes. Certificates would be send over the ESM network.

- In 2006, Qiu and Gao [31] published Hi-BGP after studying previous proposals to secure BGP. Like soBGP [25], Hi-BGP relies on a "web-of-trust" PKI infrastructure to issue IP prefix ownership certificates, requires the use of transport-level security or encryption, and introduces a new type of BGP message to send certificates. However, Hi-BGP asks AS to publish full and accurate routing information including prefix ownership, AS links and AS relationships to verify routing information in BGP.

4 Life-cycle of ideas to secure BGP

Looking chronologically at proposals to secure BGP allows us to follow the life-cycle of ideas and identify cross-pollination between proposals. This process illuminates areas of consensus and areas of disagreements related to BGP security. This section summarizes the findings first with respect to IETF proposals, and then with respect to proposals coming from academia and industry. It includes a summary of the main distinctions and disagreement between all proposals at the end.

4.1 Life-cycle of IETF proposals

The RPKI, ROAs and BGPsec standards were clearly influenced by security proposals outside the IETF. In particular, ROAs and BGPsec draw aspects from S-BGP [22]. In fact, the authors of S-BGP actively participate in the SIDR working group [72]. S-BGP was the first proposal to present a Public Key Infrastructure with the RIR as the root of trust. Many other proposals also used a PKI system to store and distribute assertions to validate information in BGP, although not all PKI structure proposed had the the RIR as roots of trust. S-BGP was also the first proposal to include signatures from border routers using BGP in the different ASes in the path of a route announcement, and include it in the announcement itself. Other proposals also include signatures from AS or their border routers along the path of a BGP announcement with varying schemes.

The SIDR solutions had to comply with a set of requirements, [46] reviews the requirements in detail. One of the most relevant requirements was that SIDR solutions had to be backward compatible with BGP-4. One of the consequences is that in the use of ROAs and BGP to validate information in BGP, the outcome of the validation process does not immediately create changes in the route selection process and its priority is left as a matter of AS local policies.

The RPKI framework and ROAs are currently under deployment, with the fraction of IP prefixes used in BGP covered by a ROA slowly increasing over the years. Currently ROAs cover over 30% of prefixes in BGP, and also over 30% of all the address space advertised in BGP [83]. To put this in perspective, in 2013, 87% of BGP announcements were originated by ASes that had registered their IP address blocks in one registry part of the IRRs [78]. Nonetheless, the use of IRRs has reduced over time and many records in IRRs are currently out of date. Recent work has extended the RPSL language to include RPKI object, adding a new channel for the distribution of ROAs [84], converging both solutions.

BGPsec is earlier in the cycle of deployment, with the standard only published in September 2017, although it was much talked about during its development and the concept has been discussed for more than a decade [85]. In addition, the designers of BGPsec recognize that it may be a long time before BGPsec is widely adopted, in particular because of its CPU and memory requirement [86]. There are efforts underway to optimize BGPsec performance [87]. Nonetheless, BGPsec has strong opponents who consider unacceptable the trade-offs of implementing it and disagree with the need to secure AS path as defined in BGPsec [88]. More recently, another IETF effort focuses on the development of the Autonomous System Provider Authorization (ASPA) a standard that would allow partial path validation with less processing and communication burden [89].

4.2 Life-cycle of academia and industry proposals

The main motivations that guided the work of BGP security proposals from academia and industry are varied. From the need to protect BGP communication to developing from scratch secure monitors to verify the correct operation of BGP in border routers, the reasons why protocol designers considered their solution to be a good one are very different. Indeed, most of this security proposals do not share the same threat model and for instance they do not agree about what needs to be secure. And even when solutions are tackling similar vulnerabilities (*e.g.,* validation of IP prefix and the AS originating a BGP announcement) they consider different options on how it should be secure, leading to different trade-offs and what is considered acceptable by the authors. As an example, some solutions such as Listen and Whisper and the use of origin lists only require BGP information validation when there are conflicting BGP announcements, while other solutions validate information for all BGP announcements.

In addition, there is a clear influence of the earlier proposals in latter ones. Many proposals take elements of the early proposals, either by improving a part of the mechanism or by clearly opposing some principle and provide an alternative. For instance, S-BGP proposed a hierarchical PKI infrastructure to validate IP prefix and ASes originating BGP announcements, which other solutions later also included. However, other proposals such as soBGP proposed a more decentralized PKI structure, instead of only having RIRs as roots of trust. soBGP proposes a PKI with a small group of trusted network operators at the top. Furthermore, ESM proposed an even more decentralized PKI structure where each AS would issue its own certificates and only in case of conflict would ASes verify the related certificates and attestations.

Furthermore, usually the authors motivation guided the development of one specific feature of their proposal and then borrowed aspects from other proposals to cover broader threat models. For example, SPV designers focused on developing a more efficient mechanism to validate AS Path than the one from S-BGP, but relies on a PKI infrastructure with the 5 RIRs as roots of trust like S-BGP for validating IP prefix and AS origin information.

None of the security proposals by academia or industry were fully implemented and only few of them are still discussed in BGP security works. However, many of these works influenced the SIDR developments described in section 3.1.2 and some can potentially have influenced current BGP monitoring services, in deployment either by ISPs or other entities.

Finally, although there are many distinctions between proposals, there is significant consensus between them about the techniques to use to validate different pieces of information in BGP announcements. Whenever validation is involved, most solution proposed the use of similar cryptographic records and signatures. Many rely of PKIs to issue the related attestations and make them available to other networks.

4.3 Why insecurity is persistent

Following along the life-cycle of ideas and main motivations behind these proposals surfaced a number of specific challenges involving disagreements, incentives and governance, as reasons to either modify or start from scratch the development of another proposal. These challenges have impeded the selection or consensus on a scheme from all the proposed options. The following list summarizes the main challenges:

- Persistent disagreement as to which BGP vulnerabilities are the most important and should be prioritized. There are several points in the design of BGP that represent potential security vulnerabilities.
- Lack of agreement as to which proposals are actually practical, taking into account the issues of deployment and operation.
- Lack of framework that drives toward a consensus. Reaching agreement requires advocates to give ground, and there is no reward for doing so.
- Misaligned incentives for actors. ISPs will almost certainly bear the major cost and complexity of deploying a change to BGP, but they are not the beneficiaries of the changes. It is primarily the end points that benefit, not the ISPs, that benefit from reduced hijacks.
- First-mover disadvantage. The first ISPs to deploy a mechanism may see no real benefit, either to themselves or their customers. The investment in the mechanism increases their costs, making them less competitive.
- This global problem is not easily be shaped by domestic regulation.

In addition, the next list summarizes the main differences between BGP security proposals:

- **The threat model:** Solutions are focused on protecting from distinct threat models and as such they differ in the changes proposed to BGP or BGP operation.
- **Authoritative source:** Some proposals have network providers (or a subset) as the roots of trust and authoritative source of assertions to validate information in BGP, whereas other proposals have the 5 RIRs as roots of trust.Some propose to use historical BGP data.
- **Information to validate:** Some proposals aim to validate IP address blocks and the AS that originated them; others look to validate the whole AS path reported in BGP messages.
- **Validation strategy:** One approach is that every router should always validate each BGP announcement; another is that a router need only validate an announcement if there is a conflicting announcement. As described in Chapter 1 Section 2, hijacks of IP address blocks that are not supposed to be routed can still be harmful.

As seen from this list, there are a number of challenges that relate to the validation of routing information in BGP announcements. Indeed, the mechanisms, supporting infrastructures and choices of trust implied in how BGP announcements are validated against other source of information are all slightly different from one proposal to the other. Section 5 below describes in more details these challenges.

5 The challenge of validating information in BGP

There is a general observation that informs all the approaches requiring the validation of informaiton in BGP: *validating routing information is contrary to the primary goal of the Internet*, which is availability—to deliver data. Validating mechanisms require an additional step before sharing routing data that can prevent routing data from being shared, so by definition they interfere with availability. The goal of validating information in BGP is to be able to block the spread of information that appears to be invalid. This means that some routing information that could potentially be valid (although deemed invalid at the moment of the check) will not be taken into account to decide where and how to route data packets between networks. Therefore, packets may end up being dropped, impairing the availability of some service.

Creating a database of records that specifies what information in BGP announcements should be valid. An AS that receives an announcement would check the announcement against the database and reject it if it does not match. This idea is fine in principle but raises several critical issues in practice: there needs to be an authoritative source of the information in the database, records need to be stored and kept up to date, the operation of the database needs to be secured, and operators need to be willing to share information about their internet resources. The next paragraphs further describe each of these specific challenges.

5.1 Defining the authoritative source:

Since the era of the telephone system, designers have incorporated databases that are supposed to be the source of authority as to the state of the network, and the results have usually been problematic. The insight that emerged is that as the elements of the network become more sophisticated and able to communicate about their state, the network itself should be the ground truth about the network.[5]

The idea that the network itself is the authoritative record of what the network is (and should be) is actually very empowering and can reduce management costs and complexity for network operators and for the organization that would managing the authoritative database. In a dynamic and decentralized network such as the current

[5] As a simple example, in the era of telephone service based on copper pairs, the phone companies tried to maintain a database of which pairs were in service, but workers in the field would make changes and not update the database, which then led to future confusion and failed installs. There is no way to “query" a copper pair and ask what its state is. However, with current access technology such as fiber to the home or hybrid fiber coax, the elements in the network are active, and so it is possible (for example) to send a control message down a fiber to confirm what is at the other end and whether it is working. In this way, the network itself can be the authoritative record.

Internet, it is hard for any one authority to know about all resources that are active and where they are currently located. Even for network operators, it is hard to keep track of all states of its network elements (*e.g.,* IP address blocks in use). If there was a way for the different parts of the network to self-report their status, no additional steps to update a database would be required.

Unfortunately, self-reporting opens the door to malice, because elements can now lie about their state. Malicious actors could get control of a network element and by having it lie about its state and resources, they would be able to use IP addresses they do not own or fake endpoints, as explained in Chapter 1 Section 2. Indeed, false routing information has been found in IRRs, where network provider self-report their routing records with little oversight [90].

If the network itself cannot be the authoritative source of what should go in BGP, then the question is what should be the source. One option is to have network operators (or a selected group of them) be the authoritative source. Still, operators could lie about their resources, either to incorporate new resources or again with malicious intention. Another option would be to have Regional Internet Registries (RIRs), which are already in charge of regional IP address block delegations, be the authoritative source. In this case, the trust is transferred to RIRs and their processes. However, RIRs only know high-level information about IP address delegations. RIRs do not know how organizations distribute their resources (*e.g.,* in smaller address blocks) or which networks (ASes) are allowed to originate a given IP address block in BGP. Therefore, the responsibility of filling the database would have to be shared: RIRs control the range of IP address blocks for which an organization may issue records and then network operators issue the records in agreement with RIR's authorization. This option nonetheless still relies heavily on RIRs processes, which has been a concern for opponents of this approach [91].

5.2 Having access to up to date records:

A critical challenge for a validation database to work for Internet routing, is to keep it up to date with all dynamic changes occurring in all parts of the network. Owners of addresses make a change to how they are using them and forget to update the database. Then other networks on the Internet may end up dropping legitimate BGP advertisements because they do not match the information in the database, and the resulting losses of connectivity are hard to debug: this forgotten update does not immediately break communication; it only makes some addresses unavailable, and it can take some time for operators to realize this is happening and figure out how to fix it.

One of the key concerns of transit provider when implementing the validation of BGP records is whether to drop or block BGP announcements from their customer, even if these records are not updated or are malformatted. For instance, AT&T decided to continue to accept invalid BGP announcements from their customers after implementing BGP information validation in 2019 [92]. This highlights the tension between information validation and the prime mission of network operators, which is to provide connectivity (and therefore availability) to their customers and users.

Another challenge is that not only do records need to be up to date, but network elements need to have access to updated records in as close to real-time as possible. This means that eventually, when such a database gains full adoption, the infrastructure supporting it would need to be provisioned so that all networks in the Internet can quickly access it to fetch all the records. This has already proven to be a difficulty in the use of BGP validation. In [93] the authors measure the connections of networks (ASes) to RPKI databases and find that many of them were unable to reach all delegated publication points of such databases.

5.3 Protecting records from attacks:

A database used for the validation of global routing information in itself becomes a point of control and an attractive target for malicious attack. There are multiple ways routing can be disrupted because of the database system. An attacker can make (part of) the database system unavailable to the rest of the Internet through a denial-of-service attack, which would allow the temporary spread of invalid advertisements. An attacker can also seek to change the database, and in this case the scope of mischief is broad but the attack can be very hard to identify in a large and distributed database. As mentioned earlier, in IRRs database, malicious actors have created fake records to make their BGP hijacks more stealthy by having records for the address block they hijack in a routing database system [90]. Similarly, the role of RIR in the RPKI has raised concerns about their potential of misbehavior as RPKI authorities [91].

Another risk of having a centralized database system is that governments may seek to use the database to regulate online activity in accordance to their specific laws and regulation, which may not be shared by the rest of the global Internet. As an example, in 2011, the US Federal Bureau of Investigation (FBI) asked the Dutch police to enforce an order issued by a US court related to record in a database of RIPE NCC, the Regional Internet Registry covering Europe [94]. Researcher have proposed ways to limit the power of entities in charge of critical databases [95].

5.4 Forcing transparency of network provisioning:

For information validation to provide security, database records need to be specific enough to match routing operations, which ends up revealing business relationship between the owner of an address block and the network authorized to originate it in BGP. For instance, if a network has an agreement with another network for back-up or for Denial of Service (DoS) attack protection, it needs to issue records letting either the back-up network or the DoS protection network advertise its IP address blocks in BGP. Thus, the records in the validation database give a thorough view of routing business relationship and some operators have voiced concerns about issuing records for this reason [96].

6 Conclusion

This chapter reviewed proposals to secure BGP, studying the life-cycle of the main ideas and motivation of the proposals to identify and frame similarities and differences in proposals' implied choice of trust, residual vulnerabilities, mechanism used and infrastructure requirements.

The findings comprise long-lasting disagreements concerning many aspects of how to secure BGP and BGP information, paired with lack of systemic review and framing of non-technical aspects of proposal such as the choice of which supporting infrastructure and related organizations to trust. Indeed, most proposals to secure BGP require the validation of routing information in BGP announcements, which presents many challenging aspects to operationalize, mostly linked to the creation and distribution/accessibility of authoritative records to compare with information received in BGP. The pervasive lack of consensus on how to validate BGP information and other aspects of security proposals has held back the implementation and deployment of BGP security proposals.

Chapter 3

Empirical Analysis of Malicious Behavior

The previous chapter reviewed BGP security proposals. It finds that they consider different threats models of how BGP design flaws would be exploited by attackers in BGP. This chapter focuses on studying systemic malicious behavior coming from networks that repeatedly perform hijacks over time. The goal is to evaluate the pervasiveness of malicious activity in BGP and the types of hijacking malicious actors are perpetrating by characterizing harmful network behavior in BGP.

Other than anecdotal evidence, little is known about the frequency and harm of malicious network behavior. Many works on hijack detection have proposed heuristics and methods to identify illegitimate route announcements [48–55, 97]. However, identifying hijacks based on single incidents is fundamentally limited in accuracy. This study instead focuses on systemic malicious behavior coming from malicious networks that repeatedly perform hijacks over time, which can be detected with higher accuracy.

This analysis scrutinizes network behavior using 5 years of BGP data from over 1,400 networks and a manually constructed ground truth, to identify and characterize the routing-level behavior of *BGP serial hijackers*—networks that persistently perform hijacks in BGP. It analyzes how the behavior of serial hijackers differs from that of honest operators to uncover key distinctions between malicious and benign network activity and the variability in behavior in both groups.

This work also involves the training of a machine learning classifier to identify networks in the Internet with behavior similar to serial hijackers. Out of the more than 70,000 networks in the Internet, the classifier finds about 900 networks with such behavior. Of this group, about 400 of those networks were also identified as malicious by services monitoring malicious activity in the Internet such as Spamhaus, because of criminal behavior or repeated spam campaigns. Other networks are related to common misconfigurations in BGP, and a few are found to be false positives linked to benign forms of hijacking (*e.g.,* used for DDoS mitigation).

This study is the first to reveal the existence of BGP serial hijackers, which persistently perform hijacks in BGP, showing that there are few barriers to performing even the basic forms of routing attacks, with almost no consequences to operators. It illuminates key characteristics of their behavior and how it differs from benign BGP behavior. It shows that, through analysis of readily available public BGP data—without leveraging blacklists or other indicators—it is possible to identify dominant patterns of serial hijackers. The analysis provides a state of affairs of networks' BGP behavior that can be reproduced by operators to assess the full extent of networks' hijacking activity to the point that fully automated detection of serial hijackers and network reputation scoring systems can be envisioned in the future. This findings have thus relevance for the operator community, increasing transparency of networks' behavior and supporting network operators to identify suspicious ASes *a priori*, potentially allowing for preventive defense. Finally, this work narrows the focus on a small set of suspicious networks that can be further studied to better understand attackers capabilities and behaviors.

An original version of this work appeared in [98], and datasets and results are publicly available.[1]

1 Background

BGP's lack of route authentication and validation remains a pressing problem in today's Internet. The lack of deployment of basic origin validation of route announcements in BGP not only makes the Internet more susceptible to connectivity issues due to misconfigurations, but also opens the door for malicious actors. While a long-standing problem, its severity becomes clear in numerous recent reports of widespread connectivity issues due to BGP misconfiguration [99], as well as hijacking events of popular destinations in the Internet [100]. Episodes range from simpler attacks with the goal of using blocks to send spam emails [64] to more sophisticated misuse of BGP to intercept traffic or steal crypto currencies [101].

While the operator and research communities have devoted substantial resources to improve the state-of-the-art of BGP security (*i.e.,* the RPKI [102]), little has changed in production environments. Today, operators can use monitoring services [62] to automatically detect potential hijacks of their prefix announcements. Current hijack detection systems typically rely on assumptions of prefix ownership and track origin changes in the global routing table. If an event is detected, the victim network can react and attempt to get in contact with the perpetrator or its upstream networks to solve the problem. However, many times this contact is not fruitful or not even possible. At that point, victims of hijacks are only left with the option of publicly disclosing the event in network operator mailing lists in the hope that peer pressure and manual interventions by other networks, such as filtering announcements or refusing to provide transit, will remediate the situation.

1.1 Related Work

There have been many efforts in the research community to characterize BGP hijacking events [64,103] and to develop hijack detection systems using different approaches,

[1] Auxiliary material can be found at `https://github.com/ctestart/BGP-SerialHijackers`.

metrics, and vantage points [48–55]. What most BGP hijack detection systems have in common is that *(i)* they are *reactive* in nature, *i.e.,* they identify hijacking events only after they occurred, and *(ii)* they are event-based, *i.e.,* they track individual hijacking events. While most systems focus on detecting individual BGP hijacking events, some attempt to identify the source of the cause and a few even tackle mitigation and remediation [97] in a case by case nature.

However, malicious BGP behavior by an actor is sometimes consistent over time, creating opportunities for methods based on longitudinal analysis, potentially informing *proactive* approaches (*e.g.,* scoring systems) and providing situational awareness. We indeed find that many hijacking events disclosed in operator mailing lists and network security blogs involve malicious Autonomous Systems (ASes) that repeatedly hijack prefixes, *i.e.,* originate prefixes allocated to and routed by other networks. In fact, some of these ASes show malicious activity in the global routing table for *multiple years*, and we refer to networks of this type as *serial hijackers*. Serial hijackers pose an ongoing threat, yet they have received surprisingly little attention in terms of empirical assessment.

Thus, in contrast to most earlier works on BGP hijacks, the approach of this works starts by profiling the network-wide BGP prefix origination behavior of ASes, leveraging the repetitiveness of their actions. Few previous works study network-wide behavior of malicious actors. In [15], the authors study BGP announcements dynamics of prefixes found in email spam blacklists. They find that some spammers use short-lived (a few minutes long) BGP route announcements of large address blocks to send spam from IP addresses scattered throughout the advertised prefix. In [104], the authors study ASes that are over-represented in blacklists of phishing, scam, spam, malware and exploited hosts. Analyzing a month of BGP data, they find that these ASes are more likely to become unreachable and that they have more changes in their connectivity than most ASes in the Internet. Konte et al. [63] developed a system to identify bulletproof hosting ASes, leveraging features such as frequent rewiring of transit interconnections. This work is complementary in that it focuses on a specific group of malicious ASes, *serial hijackers*, on behavioral characteristics related to their BGP origination patterns (*i.e.,* it does not leverage any data other than BGP for the classification), and specifically study *long-term* behavior of networks.

1.2 Introducing Serial Hijackers

To bootstrap our analysis, we first introduce the *serial hijacker* network type, and illustrate some of its pertinent characteristics by example. We review related work in the field of hijack detection and network profiling, and present a roadmap for this paper.

Since as of today, no reliable and widely deployed system to automatically discard illegitimate BGP route announcements exists, the network operator community frequently relies on mailing lists (*e.g.,* NANOG [105]) to exchange information about illegitimate BGP announcements and to coordinate efforts to limit their propagation and impact by blocking announcements from networks originating such prefixes.

The key observation that motivates this work came from studying 5 years of

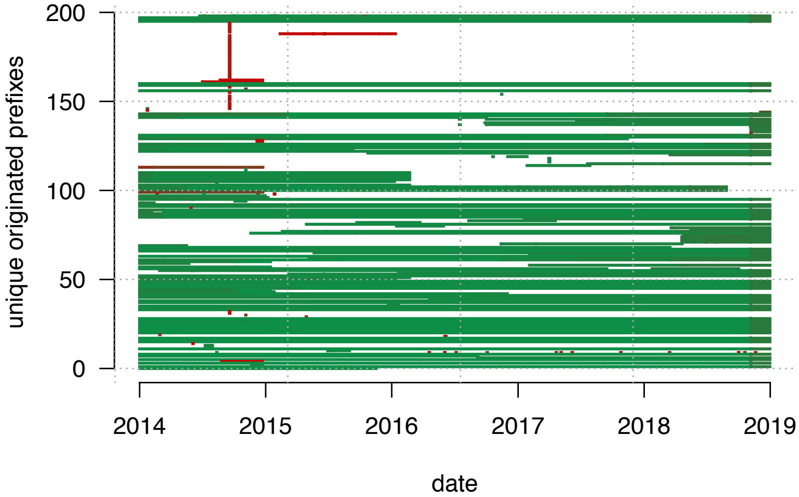

(a) Legitimate AS: Prefix origination of AS5400 (British Telecom) over the course of 5 years. This AS originates prefixes consistently over long time periods.

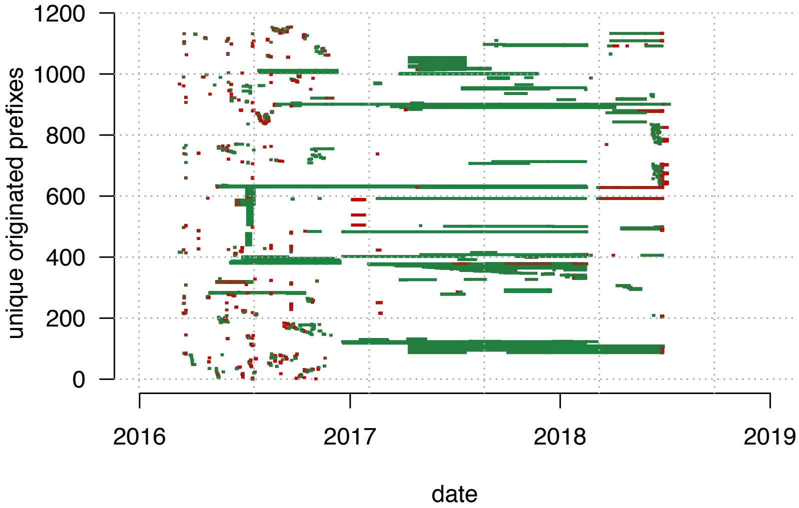

(b) Serial Hijacker AS: Prefix origination of AS3266 (Bitcanal) over the course of 3 years. This AS announces a large number of prefixes over short time periods.

Figure 3-1: Long-term prefix announcement behavior for a regular AS, and a serial hijacker AS. Each originated prefix is a row on the y-axis. Prefixes are colored in red if their normalized visibility in the global routing table is less than 15%. We sort prefixes numerically and show time (3-5 years) on the y-axis.

threads from operator mailing lists: many reported hijacks are not "one-off" events, where a previously unknown AS number starts to advertise prefixes. Instead, we often find reports of the very same ASes repeatedly carrying out prefix hijacks. In fact, some of these networks continue to hijack different prefixes over the course of multiple *years*.

Figure 3-1b shows a visualization of the origination activity of AS3266, a network that was repeatedly reported to hijack address space. Over the course of 3 years, this AS originated almost 1,200 unique prefixes, and we observe a highly irregular pattern of short-lived origination of disparate address blocks. To put this behavior in contrast, Figure 3-1a shows the origination activity of AS5400 (British Telecom). This network, a large British residential and mobile ISP, shows a much more steady pattern, longer prefix announcement times, and an overall constant, and monotonically increasing number of advertised prefixes. However, legitimate ASes can also exhibit irregular patterns (see the white space between lines indicating a prefix was not originated at that time), often due to configuration issues of the network in question or of third-party ASes. Thus, metrics and systems attempting to isolate ASes with potentially malicious behavior must be chosen and evaluated carefully to allow for robustness. From Figure 4-1 it becomes clear that these two networks show wildly different long-term behavior in the global routing table. The goal of this study is to identify and scrutinize the dominant prefix origination characteristics of this important class of networks: serial hijackers.

1.3 Roadmap

The rest of this chapter is organized as follows: In § 2 we first describe how we build a ground-truth dataset of serial hijacker ASes, as well as a control set of legitimate ASes. We also introduce our longitudinal dataset that covers 5 years of BGP activity at a 5-minute granularity. We introduce necessary data cleaning and preprocessing steps in § 3. In § 4, we first introduce a set of behavioral characteristics and pose hypotheses on how the behavior of serial hijacker ASes might differ from legitimate ASes. For each category, we introduce different metrics to capture AS behavior and study in detail how serial hijackers' BGP origination behavior differs from legitimate ASes in our ground-truth dataset and how our metrics capture these differences. With our metrics in hand, in § 5 we proceed and train a machine-learning model to identify networks in the global routing table exhibiting similar behavior to serial hijacker ASes. In § 6, we present a broad and detailed study of the ≈ 900 networks flagged by our classifier "in the wild". Finally, we feature three networks in case studies in § 7, and discuss implications and limitations of our work as well as avenues for future work in § 8.

2 Datasets

This section first describes the datasets we leverage for identifying serial hijackers and a control group of legitimate ASes. Then it introduces our longitudinal BGP dataset.

Start date	Jan 1, 2014 00:00:00 UTC
End date	Dec 31, 2018 23:55:00 UTC
Snapshot files	525,888
Unique prefixes	6,044,333
Unique ASNs	76,769
Prefix-origin pairs	7,351,829

Table 3.1: Raw dataset properties.

2.1 Legitimate ASes and Serial Hijackers

Legitimate ASes: We start our selection of legitimate ASes using the participants to the Mutually Agreed Norms for Routing Security (MANRS) initiative [106]. MANRS is a global initiative started by network operators and supported by the Internet Society, which proposes a set of actions, such as filtering and global validation of Internet resources, that network operators can implement to foster routing security. Since MANRS participants voluntarily agree to implement a set of proactive security measures in BGP, it is unlikely that they would repeatedly—and willingly—engage in BGP misbehavior or malicious activities. 272 ASes[2] are part of the MANRS initiative. Additionally, we manually select 35 ASes that represent the full spectrum of routed ASes: major end-user ISPs, enterprise networks, content/cloud providers, and academic networks. For these ASes, we are reasonably certain that the administrators do not willingly engage in repeated hostile activity.

Serial Hijacker ASes: Finding ground truth on serial hijacker ASes is a more difficult task: we process 5 years worth of email threads on the NANOG [105] mailing list and extract 23 AS numbers for which network operators repeatedly disclosed hijacking events. We note that for each of these ASes the email threads included several address blocks that had recently been (or were being) hijacked. Furthermore, in 4 cases, hijacker ASes were mentioned in connection to hijacking events spanning multiple years.

In the remainder of this study, we use the set of *Legitimate ASes* and *Serial Hijacker ASes* to first study the dominant characteristics of serial hijackers in § 4, and to later train a classifier to identify these characteristics in the larger AS population in § 5.

2.2 Longitudinal BGP Dataset

We base our study on snapshots taken from the global routing table computed every 5 minutes over a time period of 5 years, leveraging historical BGP data from all available RIPE and RouteViews collectors. Starting on January 1st, 2014 and ending

[2] Later in § 5 we only leverage MANRS ASes that have originated at least 10 prefixes in the 5 years considered in our study.

	IPv4	IPv6
Snapshot files	524,556	524,290
Unique prefixes	1,907,397	196,136
Unique ASNs	75,261	22,248
Prefix-origin pairs	2,317,168	196,137

Table 3.2: Dataset properties after removal of incomplete snapshots and very low visibility prefix-origin pairs.

in December 31, 2018, we build an individual routing table for each peer (network that feeds into any of the collectors) of each collector every 5 minutes using RIB dumps and BGP updates received over the respective peer-collector BGP sessions. For each of these routing tables, we extract prefix and origin AS numbers to generate 5 minute snapshots listing prefix-origin AS pairs (*prefix-origins* in the following) together with the count of peers observing them. Each snapshot file contains between 560,000 and 1,240,000 prefix-origin pairs. We obtain 288 files per day, 525,888 snapshot files in total. Across the entirety of our dataset covering 5 years, we find 7,370,019 unique prefix-origins to be advertised by at least one peer. We find a total of 76,769 unique ASes and 6,044,333 unique prefixes. Table 3.1 summarizes the main properties of the dataset.

3 Data Preprocessing

In this section, we describe the necessary steps to de-noise our dataset, and to convert individual snapshots into aggregated prefix-origin timelines for further analysis.

3.1 Dataset De-Noising

Variability of BGP peer availability: We leverage the count of peers that see and propagate an individual prefix-origin pair as a proxy for the prefix-origin visibility in the global routing table. Figure 3-2 shows the maximum visibility of IPv4 and IPv6 prefix-origin pairs in each snapshot file, *i.e.,* the maximum number of peers that reported the same prefix-origin pair to any of the RIPE or RouteViews collectors. Over the course of 5 years, the maximum visibility increases from the 250-300 range for IPv4 and 160-210 range for IPv6 in 2014 to 400-500 (IPv4) and 300-400 (IPv6) in 2018, mainly a result of increasing participation of networks in the BGP collection infrastructure. However, we see constant variability, *e.g.,* caused by lost BGP sessions between peers and collectors, or outages of individual collectors. Indeed, we find a number of episodes of significant reduction in the number of peers with active connections to collectors. During the 5 year period, the lowest maximum peer count is 83 for IPv4 and 102 for IPv6. In order to reduce the impact of significant peer disconnections and other BGP collector infrastructure problems, for IPv4 and IPv6, we do not consider a snapshot file if the maximum peer count drops below 20% of the

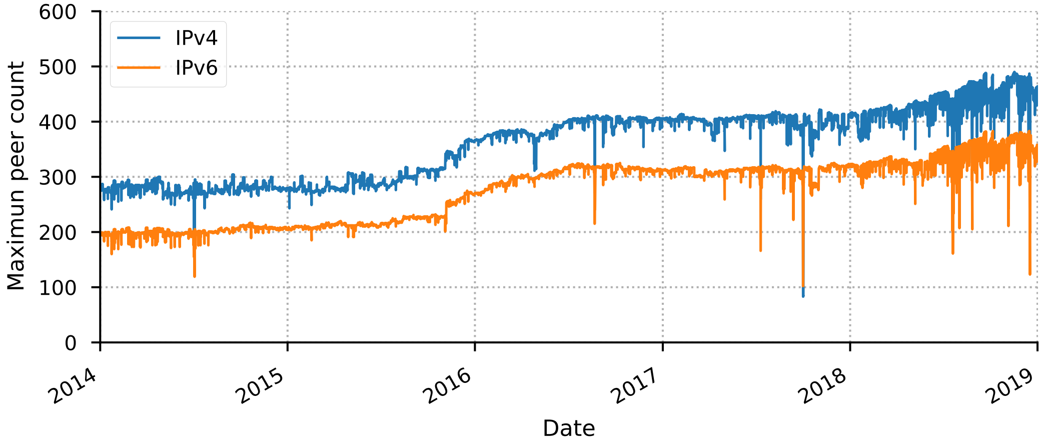

Figure 3-2: Variability and growth of the maximum visibility (max. number of peers) in the collector infrastructure of RouteViews and RIPE RIS combined.

median maximum peer count of the previous week for the same protocol. In total, for the 5 year period, we ignore 1332 (for IPv4) and 1598 (for IPv6) snapshot files, representing 0.25% and 0.30% of all available files respectively.

Highly localized BGP advertisements: In every snapshot file, we find prefix-origin pairs with very low visibility. These BGP advertisements can either be the result of highly localized traffic engineering efforts or related to misconfigurations and errors of the collector infrastructure itself or of a single, or a few, of their connected peers (recall that the total number of peers ranges between 300 and 500 for IPv4 during our measurement period). We remove prefix-origin pairs that were seen by 5 or fewer peers. While we specifically track both low-visibility and high-visibility prefix advertisements in this work, these cases of very low visibility are unlikely to represent actual routing events of interest for this study. We find that, on average, of all prefix-origin pairs of a snapshot file, less than 20% of IPv4 and 15% of IPv6 prefix-origin pairs are seen by 5 or fewer peers, but point out that they represent only 0.09% of IPv4 and 0.1% of IPv6 prefix-origins found in the routing tables of BGP collectors' peers at the time of the snapshot. Two thirds of the low-visibility IPv4 prefix-origins are announcements more specific than /24, and three quarters of IPv6 prefix-origins more specific than /48. Table 3.2 summarizes the properties of the cleaned routing dataset for IPv4 and IPv6. We note that although filtering very low visibility prefix-origins reduces the overall number of prefix-origin pairs from some 7.4M to 2.5M, it only represents ≈ 0.1% of all BGP collectors' peers routing table data during the time of the study.

3.2 Aggregating Snapshots to Timelines

Our methodology to go from individual snapshot files to a suitable data representation for longitudinal analysis of prefix-origin characteristics consists of 3 steps:

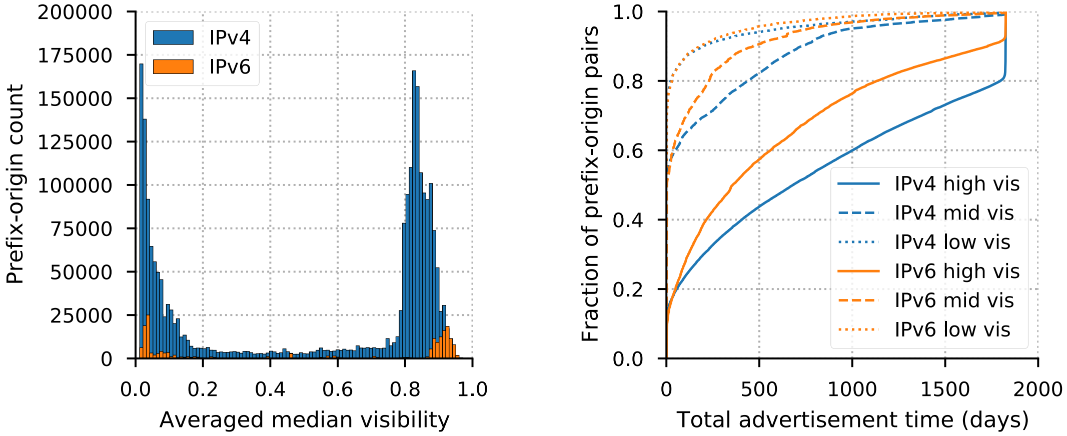

(a) Distribution of prefix-origin averaged median visibility. Most prefixes have either high (> 0.75) or low (< 0.15) visibility.

(b) Prefix-origin total advertisement time for different visibility levels for IPv4 and IPv6.

Figure 3-3: Visibility of prefix-origin pairs in the global routing table.

***(i)* Normalizing visibility:** To deal with absolute changes in peer count when evaluating prefix-origin visibility, we normalize the raw prefix-origin peer count from *each snapshot* by dividing the absolute visibility of a prefix-origin pair by the maximum peer count seen in each snapshot for the respective protocol (IPv4 or IPv6). Our normalized visibility thus is in the $(0, 1]$ interval for each prefix-origin pair.

***(ii)* Building prefix-origin timelines:** We next create *timelines* for each prefix-origin aggregating the 5-minutes-apart snapshot files, requiring *(i)* constant existence of the prefix-origin pair in consecutive snapshot files,[3] and *(ii)* a steady level of visibility of the prefix-origin pair. We find that prefix-origin visibility is overall relatively stable, but we want to capture significant changes. For each prefix-origin timeline, we require that the visibility range (maximum visibility minus minimum visibility) of the prefix-origin pair in all contained snapshots does not exceed 0.1, that is 10%.[4]

***(iii)* Classifying prefix-origin pairs by visibility level:** We next tag each prefix-origin pair with its aggregated visibility, *i.e.,* the median visibility of all contained timelines, weighted by their duration. Figure 3-3a shows a histogram of the visibility for all prefix-origin pairs. Here, we observe a bi-modal behavior: for IPv4,

[3] Since some snapshot files are not considered due to low BGP peer availability (see § 3.1), consecutive files can be more than 5 minutes apart.

[4] We note that for a single snapshot file, visibility of prefix-origins is strictly bi-modal, *i.e.,* visibility is either close to 1 or close to 0. Our threshold of 0.1 thus works well to capture significant changes.

65.3% of prefix-origin pairs show visibility greater than 0.75, while 26.1% show visibility lower than 0.25 (55.9% and 32.6% for IPv6 respectively). To better understand the relationship of prefix-origin visibility and the total time they are originated by an AS, we leverage this bi-modal behavior of visibility and classify prefix-origins according to 3 levels of visibility as follows:

- Low visibility: prefix-origin pairs with an averaged median visibility of less than 15% of active peers.

- Medium visibility: prefix-origin pairs with an averaged median visibility of less than 75% but more than 15% of active peers.

- High visibility: prefix-origin pairs with an averaged median visibility of 75% of active peers.

Figure 3-3b shows the total time that prefix-origin pairs are visible in the global routing table for high, mid and low visibility, for IPv4 and IPv6. We note that, generally, high visibility prefix-origins are present in the global routing table for longer time periods when compared to medium visibility prefix-origins, and low visibility prefix-origins. Note that in Figure 3-3b, the maximum duration is naturally constrained by our measurement window of 5 years.

In the next section, we leverage our generated prefix-origin timelines from step *(ii)* and the visibility and total advertisement distribution from step *(iii)* to compute features at the prefix-origin and AS level to scrutinize the prefix origination behavior of serial hijackers in the global routing table.

4 Dominant Origin AS Characteristics

Since little is known about BGP behavior of serial hijacker ASes other than the anecdotal evidence that these networks are repeatedly involved in BGP hijacks, we start with a mental exercise of describing how origination behavior of a network dedicated to malicious activity might look like in our BGP data. We identify five main characteristics:

- **Intermittent AS presence:** BGP activity of hijackers might be intermittent. We expect some serial hijackers to have offline periods, during which they do not originate any prefix and are thus not present in the global routing table.

- **Volatile prefix origination behavior:** We expect hijackers to show higher variability in terms of the number of originated prefixes over time than legitimate ASes. Further, we expect serial hijackers to change prefixes more frequently, resulting in a higher number of *unique* prefixes originated by serial hijackers when compared to the average number of originated prefixes.

- **Short prefix origination duration:** We expect that serial hijackers originate prefixes for shorter time periods than legitimate ASes. However, we also expect to see short-term origination of prefixes from legitimate ASes due to misconfigurations (*cf.,* Figure 3-1a). We expect that different visibility levels of such events might help to disambiguate hijacks from misconfiguration events.

- **Fragmentation of originated address space:** We expect that serial hijackers originate prefixes allocated to different RIRs (Regional Internet Registries), whereas most legitimate ASes originate prefixes allocated to a single RIR, reflecting geographic boundaries of ASes. Further, we expect that some serial hijackers originate unassigned address space.

- **Multi-Origin conflicts (MOAS) of originated prefixes:** Since hijackers originate address space routed by other ASes, we expect to see a significantly higher share of MOAS conflicts for prefixes originated by hijackers, when compared to legitimate ASes. We note, however, that there are also benign cases of MOAS conflicts that are not indicative of hijacks. We take the behavioral characteristics, *i.e.,* duration and frequency, of MOAS conflicts into account to disambiguate such cases.

In the remainder of this section, we elaborate and test each of these assumptions, introduce metrics that can capture these behavioral patterns, and contrast the behavior of our ground truth *serial hijackers* against our manually selected 35 *legitimate ASes* (*cf.,* § 2.1). We test the relevance of our metrics using the broader set of ground truth ASes in § 5 using a machine-learning classification algorithm. The features used to train the algorithm are based on the properties described in this section.[5]

4.1 Inconsistency and Volatility of AS Activity

To exemplify differences in AS activity, Figures 3-4a and 3-4b show the number of originated IPv4 and IPv6 prefixes over time for a legitimate AS (AS7922, top), and a serial hijacker AS (AS133955, bottom). Here, we see a strong contrast: while the legitimate AS is present in the global routing table 100% of the time, we see that the serial hijacker AS showed activity in 2015, no activity in 2016, and then again higher levels of activity starting in mid-2017. Although the number of prefixes originated by both ASes varies over time, the legitimate AS shows an overall much more stable origination pattern. We note, however, that also legitimate ASes can show high levels of short-term variability, as evidenced in Figure 3-4a. This peak is the result of AS7922 de-aggregating large prefixes for localized traffic engineering purposes to handle an infrastructure problem in 2015.[6]

Intermittency of AS presence: To investigate the length and frequency of AS offline periods, we compute two metrics: the number of times an AS stops originating prefixes (offline drop count), and the percentage of time an AS originates prefixes during its entire lifetime (active time), where the active time is the range between the

[5] The full feature list can be found at `https://github.com/ctestart/BGP-SerialHijackers`.

[6] A contact in AS7922 confirmed this incident.

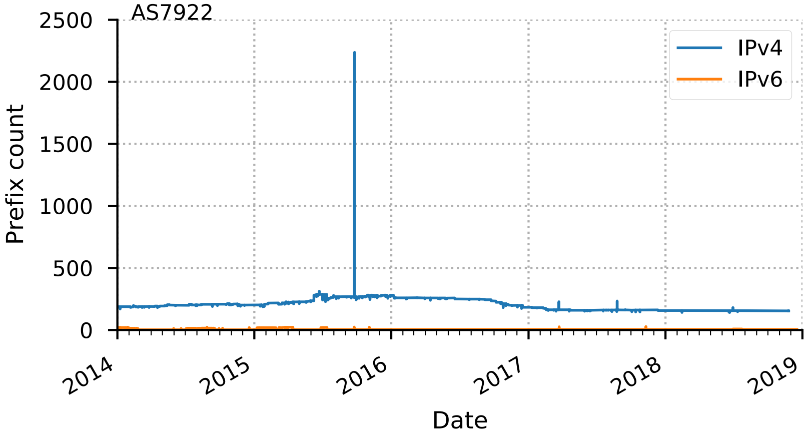

(a) Prefixes originated over time by a legitimate AS (AS7922).

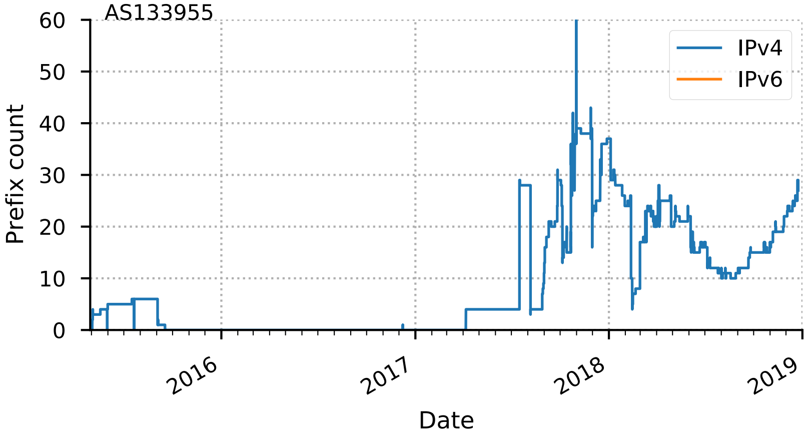

(b) Prefixes originated over time by a hijacker AS (AS133955).

Figure 3-4: Example of changes in prefix origination over time.

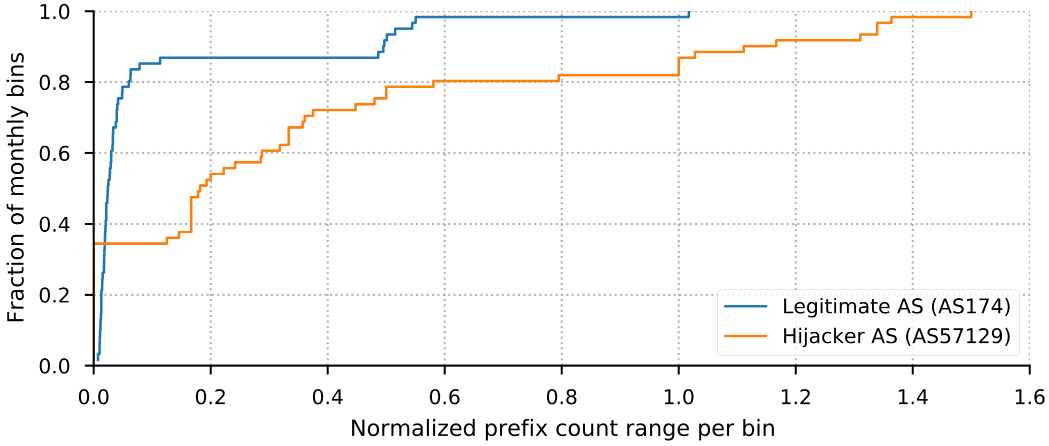

Figure 3-5: Example ASes: Monthly prefix count range normalized by median prefix count. The hijacker AS shows higher volatility in the number of advertised prefixes resulting in larger prefix count range values.

first and the last visible prefix origination of an AS. Figure 3-6a shows the distribution of these two metrics for legitimate and hijacker ASes. We find that all legitimate ASes cluster in the lower right corner, *i.e.,* once they start originating prefixes they are almost always seen originating prefixes, being active close to 100% of the time. In contrast, a large share of the serial hijacker ASes have lower overall activity times and we see multiple offline drops, *i.e.,* instances where an AS ceased to originate any prefix.

We also compute these metric for ASes originating IPv6 and obtain similar results (not shown). However, we find a few legitimate ASes that show a low activity-time percentage and high count of offline drops. Possible explanations include the fact that some networks may have originated IPv6 prefixes for testing purposes (recall that we cover a period of 5 years) before starting to steadily announce IPv6 prefixes and thus have experienced offline periods in IPv6.

Volatility in the number of originated prefixes: To quantify volatility in the number of originated prefixes over time (*e.g.,* as shown in Figure 3-4b), we partition our dataset into different time bins: one day, one week and one month. Then, for each AS and bin we compute statistics over the number of originated prefixes: range, median, and the absolute number of prefix changes. We normalize both the range and the number of prefix changes by the median number of advertised prefixes. This is to allow for more variability for large ASes, as compared to small ones. Figure 3-5 shows the distribution of the normalized range of originated prefixes for monthly bins for a legitimate AS (AS174) and a serial hijacker (AS57129). In a legitimate AS (AS174 in Figure 3-5), we see that their normalized range is small for most time bins, since the number of prefixes originated during a typical month does not vary much. AS57129, a serial hijacker, on the other hand, shows a higher number of bins with higher normalized ranges.

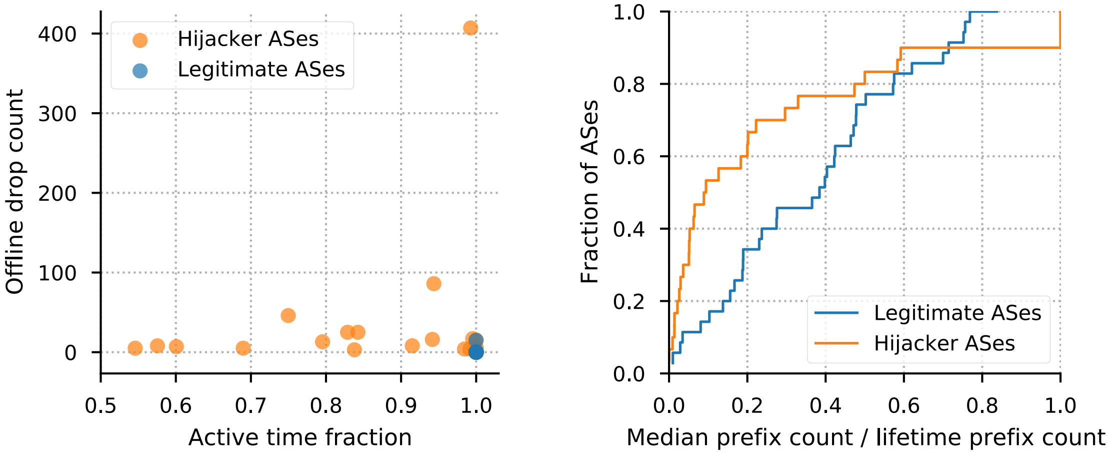

(a) Fraction of active time and offline drop count per AS. Many hijacker ASes are only intermittently visible in the global routing table, resulting in an active time < 1 and multiple instances of offline drops.

(b) Median prefix count divided by lifetime prefix count per AS. Hijacker ASes originate a smaller share of their lifetime prefixes at a given time, i.e., they have a higher turnover rate of prefixes.

Figure 3-6: Volatility metrics of prefix origination behavior for serial hijackers and legitimate ASes.

Volatility in the set of originated prefixes: So far, we developed metrics that can capture volatility in the number of originated prefixes over time. Next, we are interested in the stability of the *set* of originated prefixes. In particular, we want to capture if an AS typically advertises a fixed set of prefixes (the legitimate case) or if it "hops" through a large number of unique prefixes. To this end, we compute the median number of originated prefixes per AS, and we divide this median by the total number of unique prefixes this AS ever announced over the course of 5 years. The distribution of this ratio for legitimate and hijacker ASes (Figure 3-6b) suggests that serial hijackers tend to show a lower ratio compared to legitimate ASes, which indicates that they have a higher turnover of prefixes. Note however, that some legitimate ASes also show a low ratio, if, *e.g.,* a network had a route leak or misconfiguration problem that significantly increased the number of prefixes it advertised for a short period of time. Nonetheless, these types of events do not occur frequently in our set of legitimate ASes and our metric separates our two classes well.

4.2 Prefix-origin Longevity and Visibility

In this section, we study the dynamics of individual prefixes originated by ASes, in particular how hijackers' prefix total duration and visibility in the global routing table differ from prefixes originated by legitimate ASes.

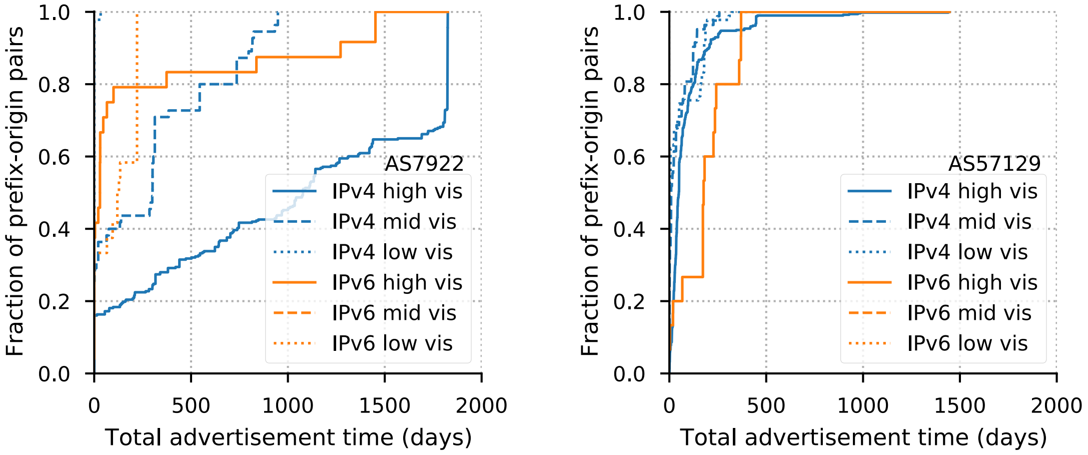

(a) Legitimate AS example: Total prefix advertisement time. Over 50% of prefixes are originated for more than 1,000 days.

(b) Hijacker AS example: Total prefix advertisement time. Over 50% of prefixes are originated for less than 50 days total.

Figure 3-7: Advertisement longevity of prefixes originated by legitimate and serial hijacker ASes.

Longevity of prefix announcements: Our hypothesis is that hijackers originate prefixes for a shorter period of time than legitimate ASes. While we find this clear distinction when looking at aggregate data, *i.e.,* hijackers' median prefix-origin duration is 27.25 days v.s. 264.17 days for legitimate ASes, we found it challenging to identify a threshold that separates short-term and long-term prefixes and hence separates our two categories of ASes well. To sharpen the picture, we next take the visibility of announcements into account.

Longevity vs. visibility level: Figures 3-7a and 3-7b show the distributions of the total advertisement time of prefix-origin pairs, for different levels of visibility, for a legitimate AS and a serial hijacker AS. AS7922, a legitimate AS, has a large fraction of long-term originated prefixes, *i.e.,* more than 50% of high visibility IPv4 prefixes it originates are advertised for over 1,000 days. On the other hand, the lower the visibility the larger the share of short-term prefixes. We notice that most of the low visibility prefixes that AS7922 originates have a very short total advertisement time. Indeed, a large share of the prefixes advertised by AS7922 for only a short period of time come from highly localized traffic engineering efforts used to handle infrastructure problems and hence have very limited visibility in the global routing table (*cf.,* § 4.1). AS57129, a serial hijacker, however, shows vastly different behavior: some 50% of high visibility IPv4 prefixes originated by AS57129 have less than 50 days of total advertisement time, and the share of short and long-term prefixes it originates is very similar for all levels of visibility.

When plotting ASes by median prefix visibility and total advertisement time (3rd

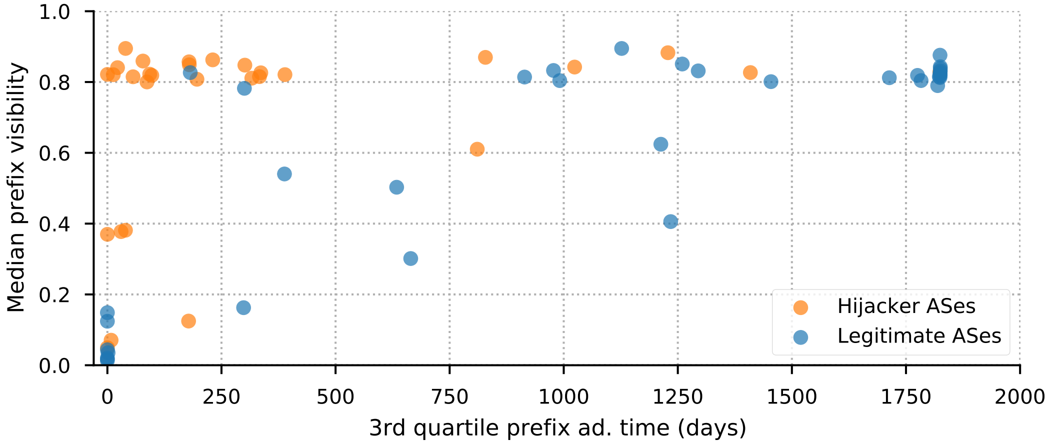

Figure 3-8: Advertisement time and visibility per AS. Hijacker ASes show shorter, high-visibility announcements.

quartile shown, Figure 3-8), a large portion of serial hijacker ASes cluster in the high visibility, low advertisement time corner (upper left). In contrast, legitimate ASes are spread out and high visibility is correlated with longer advertisement time for these networks. Thus, we find that the longevity of prefix origination can only be meaningfully leveraged to separate our two classes of ASes when qualified by their visibility level.

4.3 Address Space Properties

In this section, we study different properties of the IP addresses that ASes originate. We take into account the Regional Internet Registry (RIR) that assigned originated IP addresses, whether ASes originate bogon or unassigned IP space, and if originated prefixes were originated by other ASes at the same time (MOAS conflicts).

Address space fragmentation: Our hypothesis is that legitimate ASes only originate address blocks that were allocated to them by a respective Regional Internet Registry (RIR). Since most networks are limited in geographic scope, and individual RIRs cover individual geographic regions, we expect most legitimate ASes to either originate addresses from a single RIR, or, if they originate prefixes from different RIRs, they would still be concentrated in one of them. Since we do not expect serial hijackers to originate address space allocated to them, nor respect RIR boundaries, we expect them to originate prefixes from multiple RIRs, and show much less concentration on any particular RIR. To express concentration of originated address space across RIRs, we compute the Gini coefficient of ASes' RIR distribution using the percentage of prefixes ASes originate from each of the five RIRs. A Gini of 0.8 means all IP resources come from one RIR, whereas a Gini index closer to 0 means resources are uniformly distributed across the 5 RIRs. Figure 3-9a depicts the distribution of serial hijackers and legitimate ASes with respect to the Gini coefficient over the

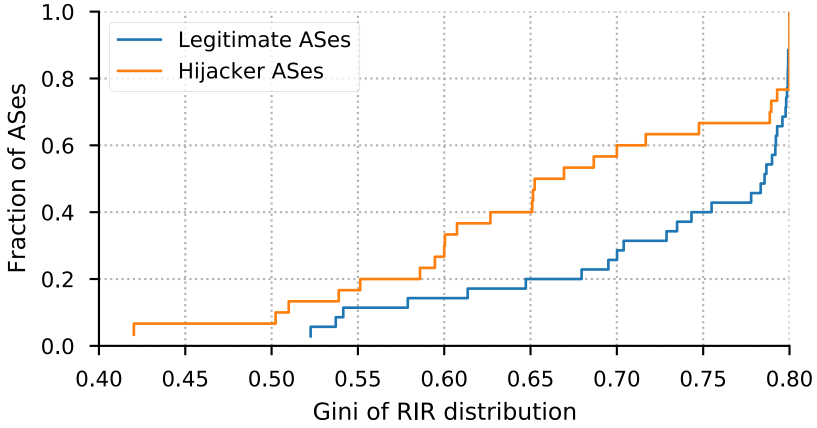

(a) Gini coefficient of originated prefix RIR concentration per AS. Serial hijackers' prefixes are more spread out over different RIRs when compared to legitimate ASes.

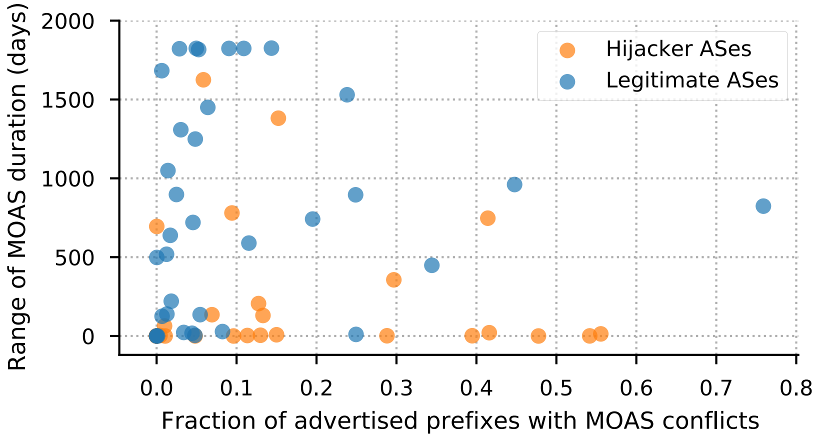

(b) Fraction of prefixes with MOAS conflicts and range of MOAS duration per AS. Some hijacker ASes show a higher fraction of prefixes with MOAS conflicts with a low duration range of MOAS conflicts.

Figure 3-9: Specific address space characteristics example for legitimate and serial hijacker ASes.

RIR distribution. We observe that many serial hijackers show a lower Gini coefficient compared to legitimate ASes, meaning that the prefixes they originate are comparably more uniformly distributed among RIRs. This is in contrast to legitimate ASes, which typically show high RIR concentration.

Multiple Origin AS prefixes: We compute the number of prefixes and the share of address space an AS originates that is also originated by another AS at the same time, *i.e.,* the prefix has Multiple Origin ASes (MOAS) in the global routing table. Figure 3-9b shows per AS the fraction of advertised prefixes with MOAS conflicts (x-axis) and the range of the duration of the MOAS announcements (y-axis). We chose to show the range of the MOAS duration, since we found that serial hijackers have almost exclusively short-term MOAS announcements, resulting in a small MOAS duration range, whereas legitimate ASes show variable MOAS durations, with many short-term and long-term prefix originations with MOAS conflicts, resulting in a large MOAS duration range. Many serial hijacker ASes have a very short range of MOAS duration and a significant share of the address space they originate are MOAS prefixes, which is what we would expect for illegitimate MOAS events (*e.g.,* replaced by new ones as they are detected). We note that, as expected, some legitimate ASes show MOAS conflicts, but that these MOAS events typically last much longer than those of serial hijackers.

5 Towards Scalable Classification of BGP Misbehavior

Next, we describe how we build a classifier to identify more ASes in the global routing table that exhibit a prefix origination behavior similar to serial hijackers. We start by explaining the main challenges faced when training a model with our dataset, and elaborate on our resulting choices for our model and its main parameters. We then discuss the features we use, their importance, and present the final ensemble classifier and its accuracy metrics. We present the results of the classification based on our trained classifier in § 6.

5.1 Challenges Faced

We face three main challenges when applying machine learning algorithms to classify whether ASes show behavioral patterns of serial hijackers: *(i)* heavy-tailed and skewed data, *(ii)* limited ground truth, and *(iii)* class imbalance.

Heavy-tailed and skewed data: The routing data on which our analysis is based is extremely heterogeneous. In almost all dimensions, individual prefixes and ASes are heavily concentrated at some level but then there is a long tail of outliers, making the data difficult to normalize. In addition, some of our features range from zero to one (*e.g.,* the Gini coefficient expressing concentration of address space across RIRs described in § 4.3), while other features, such as the total advertisement duration (described in § 4.2) ranges from 5 minutes to 5 years.

Small ground truth: As discussed in § 2.1, building a ground truth dataset including serial hijackers and legitimate ASes is challenging. In total, our ground truth dataset consists of 230 labeled ASes. We only select ASes originating at least 10 prefixes in the 5-year dataset. This includes all hijackers but only 217 ASes from our legitimate AS group described in § 2.1. Therefore, we must carefully select a model to avoid *overfitting*.

Class imbalance: We do not expect that a large share of routed ASes exhibiting serial hijackers' behavior. The true share of such ASes is unknown, and if we were to make an educated guess, we would only expect to find this behavior for a small number of ASes, *i.e.,* less than 1% of routed ASes (over 75,000 ASes are routed in our dataset in the 5-year period). Class imbalance is also present in our ground truth dataset: we only have 23 serial hijacker ASes vs. 217 ASes in the legitimate group of our labeled ground truth.

5.2 Our Classifier

Choice of Classifier: We choose a tree-based classifier since decision trees do not require normalized data and work well with large dimensions and heavy-tailed data such as the features we built to capture different aspect of BGP origination behavior. More specifically, we use Extremely Randomized Trees (Extra-Trees) classifiers [107]. An Extra-Trees classifier is an ensemble (forest) of decision trees that picks feature thresholds to split nodes at random, instead of fitting the threshold to the training data like in a common random forest classifier. This added randomness greatly reduces overfitting, another of our main challenges as discussed in § 5.1.

Model accuracy for parameter selection: To properly select model parameters (sampling methods, forest size, feature selection) without reducing the training data by doing an n-fold cross-validation, we use bootstrapping samples (subset samples) in the training phase of the individual trees and compute the classifier Out-Of-Bag (OOB) error estimate. OOB error estimation is a method to measure the prediction error of random forests, where a lower OOB error indicates higher accuracy of the model. The OOB error estimate is the average error for each data point p in the training sample computed averaging the prediction of trees trained on a bootstrapping sample (bag) not including p [108]. The OOB score has been shown to converge almost identically as the n-fold cross-validation test error and is an established method to validate random forest classifiers [109].

Sampling techniques: To address class imbalance, we try different under- and over-sampling methods to create balanced training sets for our classifier, by either

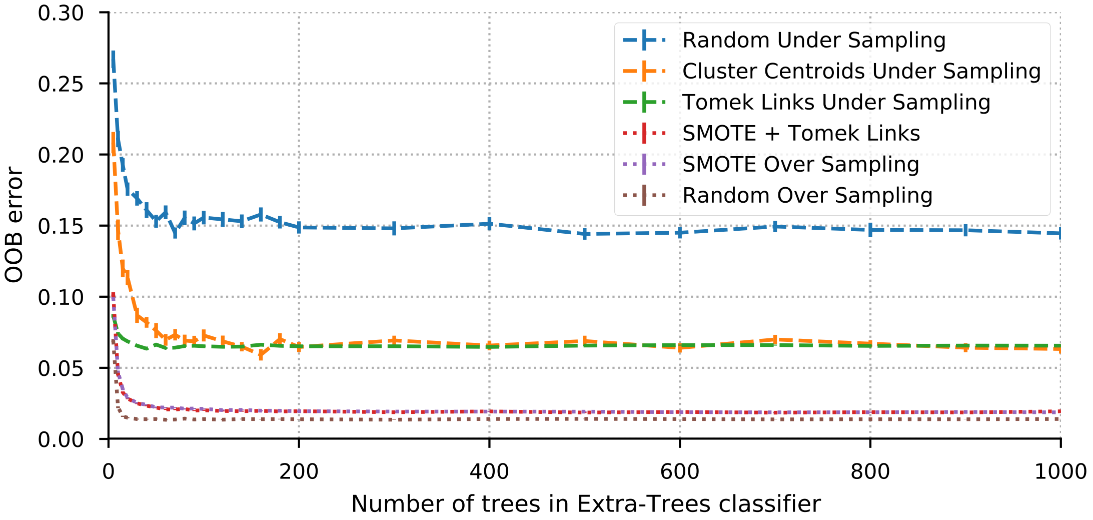

Figure 3-10: Mean Out-of-bag accuracy scores and error bars of sets of 100 Extra-Trees classifiers trained using different sampling techniques for increasing forest sizes.

under-sampling the majority class (selecting only a few legitimate ASes) or over-sampling the minority class (artificially expanding the set of serial hijackers) in our original ground truth. Figure 3-10 shows the mean OOB scores (and error bars) of sets of 100 Extra-Trees classifiers trained using 6 different sampling technique for different forest sizes. We observe that techniques that are purely based on under-sampling perform worse than techniques that include an over-sampling step. In addition, over-sampling techniques use different rules and randomness to expand the serial hijacker set and thus no two synthetic training sets are equal. We therefore decide to use a mixture of over-sampling techniques for the training of our classifier, so that it leverages the different distributions of misclassified points to improve its generalization ability [110].

Feature selection and importance: Based on the extensive manual analysis described in § 4, we select 52 features that capture BGP behavior according to 8 categories: ASN presence in the global routing table, prefix origination behavior, longevity of individual prefix advertisements, prefix visibility, longevity vs. visibility level, prefix set stability, address space fragmentation, and MOAS statistics. The features capture different characteristics and statistical behavior of the properties discussed in § 4, such as the median origination time of high visibility prefixes and 90^{th} percentile of the distribution of daily changes in prefix origination.

To assess feature importance, we compute the *drop column feature importance* for each feature.[7] The drop column importance captures how the classifier accuracy

[7] Given most of our features are computed from the same raw BGP data, selecting features by usual random forest feature importance ranking or information gain is not adequate [111–113].

actually varies when a feature is not considered in the training phase [114]. We learn that all categories have positive median drop column importance, *i.e.,* they all add to the accuracy of the model. We thus proceed to feed all 52 features to train our final classifier.

The trained classifier: Our final ensemble classifier is based on the vote of 34 Extra-Trees classifiers of 500 extremely randomized trees each, and each trained on a different balanced synthetic training set computed using one of the 3 over-sampling algorithms we selected. The model OOB error estimate is 2.5%. We program our classifier using the `sklearn` and `imblearn` libraries [115] in Python, which have the Extra-Trees classifiers and sampling algorithms pre-programmed.

False positives from the training set: Using the OOB predictions for the training set, the ensemble classifier precision and recall are 79.3% and 100% respectively. Although our serial hijacker set is small, the high recall rate supports our hypothesis that our small group of serial hijacker have distinctive characteristics in their BGP prefix origination behavior. We note however that the classifier precision is only about 80% — a strong reminder that the behavior of ASes selected by the classifier is not *necessarily illegitimate.* Even in our legitimate group, there are a few ASes that present similar characteristics to serial hijackers. Indeed, throughout all the different classifiers we tested, there are 6 ASes in our legitimate group that get consistently misclassified. Looking in more detail at these ASes, we find that two of them are from Verisign, an organization that offers DDoS protection, and are hence benign cases of serial hijackers, which we discuss in § 6.3. Two other ASes have only originated prefixes for a short period of time and are not currently being routed, which could have adversely affected our metrics and classification. The last two ASes are hosting organizations showing irregular BGP behavior of which the cause is unclear to us.

6 Investigating BGP Misbehavior in the Wild

In this section, we describe the output from our ensemble classifier. We feed the classifier with features based on IPv4 prefix-origin routing data of ASes that originate at least 10 prefixes in the 5 years of our dataset. Of the 19,103 ASes in our prediction set, our ensemble classifier finds 934 ASes having similar behavior to serial hijackers, we refer to them as *flagged ASes.* We note that the group of flagged ASes is fairly consistent across classifiers trained using different combinations of sampling methods and forest sizes. For models with an OOB error score of 4% at most, at least 95% of the ASes flagged by that classifier where also flagged by the final classifier. In the next sections, we first describe general characteristics of flagged ASes and compare them to *non-flagged* ASes. Then, we further scrutinize flagged ASes, breaking them into sub-categories.

	Flagged ASes			**Non-flagged ASes**		
	1^{st} qrt.	median	3^{st} qrt.	1^{st} qrt.	median	3^{st} qrt.
Count		934			18,169	
Prefix count	18	41	101	14.0	23.0	53.0
Active time	65.9%	99.2%	100%	99.9%	100%	100%
Prefix origination median time (days)	1.8	48.2	176.9	144.6	598.0	1,217.9
Prefix-origin median visibility (%)	51.1%	80.8%	84.2%	79.7%	82.9%	85.3%
Median origination high vis. pfxs (days)	3.4	79.4	227.2	289.7	754.2	1,386.0
Originated/unique prefixes	0.017	0.089	0.222	0.213	0.435	0.684
Gini index from RIR add. concentration	0.575	0.675	0.743	0.80	0.80	0.80
MOAS prefix share	6.7%	22.9%	52.7%	0.00%	6.9%	24.0%

Table 3.3: Summary statistics of selected metrics for ASes *flagged* as having similar BGP origination behavior to serial hijackers ASes and *non-flagged* ASes. For each metric, we show the median value across ASes in each group, as well as the 1^{st} and 3^{rd} quartiles (qrt). Only ASes originating 10 or more prefixes in our dataset (N=19,103) are fed into our classifier.

6.1 Behaviors Captured by the Classifier

Table 3.3 provides summary statistics of some representative metrics for the two classes of ASes identified by the ensemble classifier: ASes flagged as having similar BGP origination behavior to serial hijackers and non-flagged ASes. For each metric, its distribution in flagged ASes is considerably different from its distribution in non-flagged ASes.

Volatile overall BGP behavior: The ASes flagged as having similar behavior to serial hijackers show more sporadic and volatile BGP activity: the 1^{st} quartile of ASN active time is 65.9%, compared to 99.9% for non-flagged ASes. Most prefixes originated by flagged ASes are shorter-lived than those of non-flagged ASes—50% of flagged ASes have a median prefix-origin duration of less than 48.2 days *vs.* only 17.9% of non-flagged ASes.

Large ASes: On average, ASes flagged by our classifier originate more prefixes than the rest—with a median prefix count of 41 compared to 23 for non-flagged ASes. Furthermore, 34 flagged ASes have originated over a thousand prefixes, representing 3.64% in the group, compared to only 1.37% of networks in the Internet announcing more than a thousand prefixes.

Diverse IP sources: ASes flagged by our classifier use IP space spread out across the RIRs—with a median RIR Gini index of 0.675 compared to 0.8 for non-flagged ASes (an RIR Gini index of 0.8 means all prefixes originated by that AS come from only one of the five RIRs). Flagged ASes also exhibit a larger share of MOAS address space than non-flagged ASes, resulting in a median MOAS prefix share of 22.9% *vs.* 6.9%, respectively.

6.2 Indications of Misconfiguration

We find that some ASes were likely flagged as a result of misconfiguration issues in BGP.

Private AS numbers: Per RFC 6996 [116], ASNs $[64512, 65534]$ are reserved for private use. In the group of flagged ASes, we found 114 private ASNs that appear to have very volatile prefix origination behavior with relatively low visibility. A possible explanation is that due to router misconfiguration, these AS numbers appear at the origin of BGP AS-paths. As many ASes filter out prefixes originated by known reserved AS numbers, the spread and visibility of these misconfigurations is often limited. Some of the serial hijackers in our ground truth dataset exhibit lower visibility too, which is likely why these behavior got captured by the classifier.

Fat finger errors: Our classifier flagged all of the single-digit AS numbers. Indeed, the origination behavior of these ASes appears to be extremely volatile using the longitudinal routing data. We note however, that apparent origination of prefixes by theses ASes does not necessarily reflect actual routing decisions by the owner or network with given AS number. The prefix originations by these single digit ASes are likely mere results of misconfigurations, where an origin network accidentally adds an additional AS number (behind its own) to its BGP advertisements. These so-called "fat finger errors" [117] commonly occur when configuring a router to perform AS path prepending, a traffic engineering technique that artificially lengthens the AS path in order to make the advertised path less desirable in the BGP decision process [118]. A notable example of an AS flagged by our classifier is AS5, an AS whose registered company went out of business 20 years ago, periodically revived through router misconfiguration.

Removing private and single digit ASes from our group of flagged ASes, 811 remain.

6.3 Benign Serial Hijackers

In our dataset, we find prefixes originated by 29 DDoS protection networks (*e.g.,* DDoSGuard).[8] 18 of these ASes are flagged by our classifier. We find that a significant share of the address space originated by these networks has MOAS conflicts, representing over 30% of the prefixes they originate in most cases. The DDoS mitigation they perform includes originating prefixes of their customers when a DDoS attack is detected, in order to attract all the traffic destined to the network under attack, "scrub" it (to remove DDoS traffic), and tunnel it to the intended final destination [119]. Thus, DDoS protection networks present a case of "legitimate", or benign, serial hijacking behavior.

[8] Our list of AS numbers of DDoS protection services is manually compiled and hence not necessarily complete.

6.4 Indications of Malicious Behavior

After removing private AS numbers, single digit ASes, and DDoS protection ones, a total of 793 publicly routable ASes flagged by our classifier remain. Next, we assess if our identified ASes show indications of malicious behavior, *e.g.,* spam or probing activity.

Flagged ASes in Spamhaus DROP list: First, we leverage snapshots of the Spamhaus *Don't Route Or Peer* (DROP) ASN list [120], a list of ASes controlled by "spammers, cyber criminals, and hijackers". We have access to 6 snapshots taken between January 1^{st} 2017 until early 2019, containing a total of 451 unique ASes, and we note that 266 of these ASes appear in all snapshots. We compared the ASes flagged by our classifier with those listed in any of the 6 snapshots of the Spamhaus DROP list we have available, finding that 84 (10.6%) of our flagged ASes are present in the Spamhaus DROP list. For comparison, we find only 206 (1.1%) ASes from the non-flagged group are present in at least one snapshot of the blacklist. Thus, flagged ASes are almost 10 times more likely to be in this list of spammers, hijackers and cyber criminals. Of the 266 ASes that are blacklisted in *all* snapshots of the Spamhaus DROP list, 133 originate more than 10 prefixes during our measurement window, and are thus in the set of ASes we classified. Our classifier flags 50 of them as exhibiting serial hijacker characteristics. In other words, based on our feature set, our classifier detects some 38% of all the ASes with enough BGP activity that repeatedly appear on this blacklist, an indicator of persistent malicious activity in this group of ASes.

Spam activity of flagged ASes: We also check for indications of spam activity in our group of flagged ASes. To this end, we leverage 2.5 years of snapshots taken 4 times a day from the UCEPROTECT [121] Level 2 spam blacklist. Attributing prefix ranges from the UCEPROTECT blacklist to ASes is challenging in our case, since our identified ASes are by definition highly volatile and might only temporarily originate prefixes that are otherwise routed by different ASes. We first load all prefixes and their origination time ranges into a prefix trie. We then process the blacklist snapshots, where we *(i)* perform a lookup in our trie to see if the particular blacklisted address block was ever originated by one (or multiple) flagged AS(es), and *(ii)* tag a given prefix-origin as blacklisted, if the prefix was originated by the respective AS at the time it appeared in the blacklist.[9]

We find indication of spam activity for more than a third of ASes flagged by our classifier. Specifically, for 38.3% of our flagged ASes, we find at least one address block originated and simultaneously blacklisted. Note that while ASes that are *victims* of hijacking for spamming purposes might also appear in spam blacklists, we do not expect them to consistently appear in multiple blacklist snapshots. Indeed, We find that when blacklisted, prefixes originated by flagged ASes tend to be blacklisted for a larger share of their advertisement time, *i.e.,* 27% are blacklisted during more than 50% of their advertisement time, compared to 12% for prefixes originated by ASes

[9] We allow for 24 hours leeway before and after prefix origination.

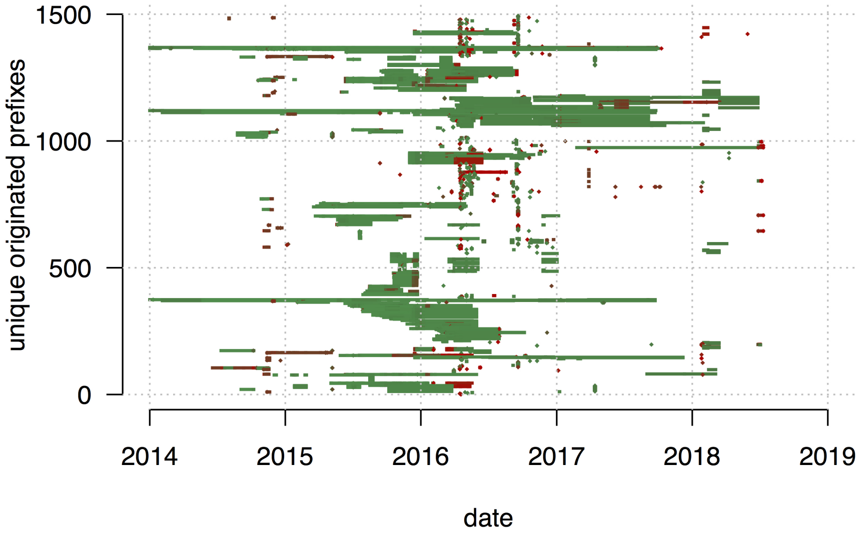

Figure 3-11: AS197426, a known serial hijacker, part of our ground truth dataset.

not flagged by our classifier.

6.5 Big Players

To find possible false positives, we inspect large ASes flagged by our classifier. Using data from CAIDA AS-Rank [9, 122], we find that 4 flagged ASes are in the top 500 ASes by customer cone size, and 21 ASes are in the top 1000. Since it is unlikely that a large prominent transit provider performs serial hijacking, these are probably false positives. Nonetheless, the BGP origination behavior of these large ASes appears to be highly volatile, similar to false positives from the training sample (cf. § 5). As an example, the median of these ASes' median prefix-origin duration is only 69 days compared to 411 for large non-flagged ASes, and they show higher levels of prefix changes—the rate of normalized monthly prefix changes is 1.0 for large flagged ASes vs. only 0.35 for large non-flagged ASes.

7 Case Studies

In this section, we illustrate three cases of ASes actually misbehaving, two of which are not in our ground truth dataset but are instead in the group of ASes identified by our classifier. We picked: AS197426, a serial hijacker from our ground truth dataset that was essentially "kicked off the Internet" in July 2018 because of their repeated malicious behavior [123]; AS19529, an AS flagged by our classifier for which we subsequently found hijacking complaints in a RIPE forum; AS134190, another

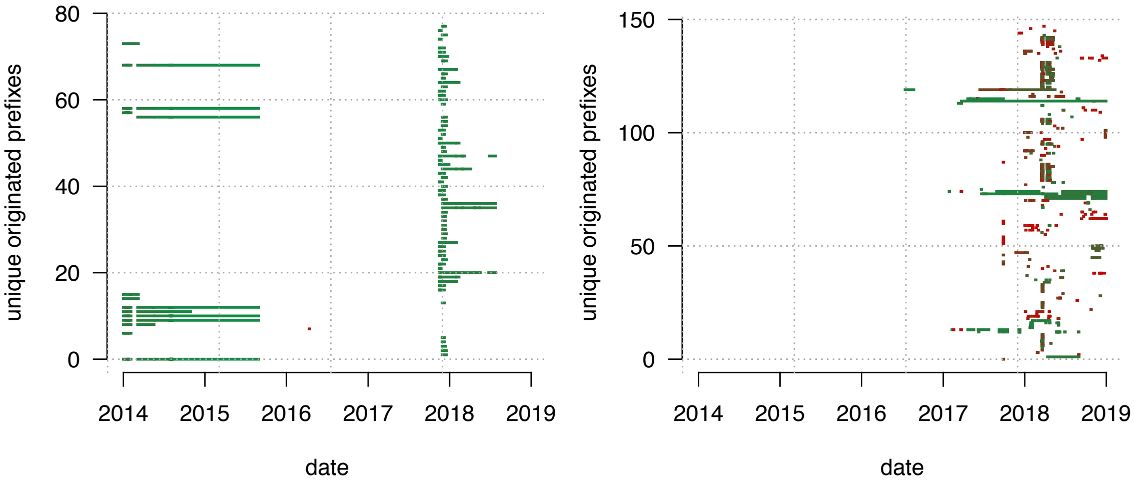

(a) AS19529, a hijacker identified by our classifier for which found corroborating evidence of hijacking activity.

(b) AS134190, the most recent detected case of a potential serial hijacker.

Figure 3-12: Prefix origination behavior for our selected case studies of flagged ASes.

flagged AS, which only recently started to show characteristics of a potential serial hijacker.

7.1 The Quintessential Serial Hijacker

Bitcanal, the "hijack factory", a Portuguese Web hosting firm, has been featured in several blog posts [124–126], since it represents a glaring case of serial hijacking, and one of the few cases in which prolonged coordinated action among network operators, ISPs, and IXPs, finally resulted in complete disconnection of the company's ASes. Bitcanal leveraged several ASNs: in this case study we focus on AS197426, the most active ASN used by Bitcanal.[10] While multiple incidents of hijacks carried out by Bitcanal were featured in numerous blog posts [124–126], we provide a first comprehensive data-driven assessment of their long-term behavior in the global routing table, revealing the full extent of persistent hijacking activity of this network, *i.e.,* an upwards of 1,500 originated prefixes over the course of 4 years.

Figure 3-11 provides a graphical representation of their prefix origination activity, each row represents a different prefix that AS197426 has originated. In the first snapshot file of our dataset in January 2014, AS197426 originates only 4 prefixes, but its origination activity soon ramps up. Already in February 2014, the same AS starts originating 15 prefixes and by October 2014 it originates almost 50 prefixes. The first post about hijacking activity by AS197426 appeared as early as September 2014 stating that it originated unrouted IP addresses that were allocated to a diverse set of organizations [124]. And yet, this was only the start of their serial hijacking spree. Starting in early 2015, we see AS197426 progressively increasing the number of prefixes it originates, and in January 2015, another blog post described recent hijacks

[10] Figure 3-1b features another Bitcanal AS.

by AS197426. Origination activity peaks at ≈ 300 prefixes in the second trimester of 2016, see vertical structures in late 2016 in Figure 3-11. During this time, this AS makes an average of 2.5 changes per day in the set of prefixes they originate. Sometime in 2017, AS197426 was expelled from the German IXP DE-CIX because of their bad behavior. DE-CIX collected and analyzed evidence before contacting the AS and finally suspending their services [127,128]. On June 25, 2018, a detailed email thread on the NANOG mailing list described multiple hijacks carried out by AS197426 and explicitly called out Cogent, GTT, and Level3 to act, since they provided transit to AS197426 [129]. Reportedly, GTT and Cogent quickly suspended their services to Bitcanal. Then, early in July 2018, Bitcanal appeared using other European transit providers (see sporadic activity in 2018 in Figure 3-11), who terminated their relationship with Bitcanal only a few days later. Bitcanal was also present in other European IXPs, including the large LINX and AMS-IX, who terminated services with Bitcanal shortly after. The last transit provider disconnected Bitcanal on July 9, 2018. AS197426 has not been visible in the global routing table since that day.

From 2014 until its disconnection in 2018, our data shows AS197426 originating a total of 1,495 different prefixes. While hijacking activity was reported as early as September 2014, coordinated measures only showed effect and resulted in eventual disconnection in 2018.

7.2 A Recent Hijacker

AS19529, originates about a dozen prefixes in our first snapshot in 2014. As Figure 3-12a shows, 7 of these prefixes were steadily originated for over a year. In April 2016, we see AS19529 withdrawing these prefixes and disappearing from (our proxy for) the global routing table (white gap in Figure 3-12a). Although the ARIN WHOIS record [130] for AS19529 has not been updated since 2012, our dataset shows it returns originating prefixes (31 this time) in November 2017. Then, AS19529 quickly increases the number of prefixes it originates, reaching almost 60 prefixes by the end of 2017. This spike in activity is clearly visible in Figure 3-12a. During these months, new RIPE RIR entries appeared, listing AS19529 as origin of IPv4 blocks owned by a different institution and registered in the ARIN region. At the same time, the legitimate owner of these prefixes raised complaints in a RIPE forum, stating that such RIPE RIR records were incorrect and that the respective address blocks were hijacked [129]. The complaints continued until April 2018 and the result, as of today, is unclear. In our data, we see AS19529 stopping to originate prefixes in July 2018: in its last 9 months of activity, it originates a total of 63 different prefixes, 20 of which are MOAS.

7.3 An Ongoing Potential Hijacker

We see AS134190, for the first time in our data on July 14, 2016, originating only a single prefix for about a month, after which it disappears from the global routing table. In early 2017, AS134190 starts repeatedly originating different prefixes for very short time periods (about a day). Starting in July 2017, AS134190 originates a few prefixes on and off—the small dots in Figure 3-12b—with some burst of activity reaching over 30 prefixes being simultaneously originated. In this period, AS134190 averages almost 10 changes per day in terms of originated prefixes. In November 2018, BGPmon[11], a widely known BGP hijack detection system [62], detected a potential hijack from AS134190 and 10 additional potential hijacks in early 2019. As of today, we have not found further evidence in the form of public complaints about potential hijacks carried out by AS134190.

8 Discussion

Our study was motivated by repeated complaints in the operational community about reiterated, even persistent, prefix hijacking activities carried out by certain ASes. On the one hand, BGP's native lack of validation mechanisms exposes it not just to one-off or stealthy attacks but also to routinely executed, in-the-open, forms of abuse. On the other hand, BGP's inherent transparency, combined with the availability of pervasive and "public" BGP measurement infrastructure (*e.g.,* RouteViews, RIPE RIS) provides the opportunity to uncover systematic malicious behavior, also through the application of automated methods.

In this work, we analyzed the origination behavior of a small set of manually identified serial hijacker ASes, finding that they show distinct origination patterns, separating them from most benign ASes. We further showed that, in spite of limited ground truth and severe class imbalance, it is possible to train a machine-learning classifier that effectively narrows our focus to a set of networks exhibiting similar behavior to serial hijackers: this set accounts for 5.5% ($\approx$ 900) of the examined ASes, 1.4% of all ASes visible in IPv4 BGP. Our analysis also reveals clear potential and specific directions to further reduce this set, to the point that fully automated detection approaches and scoring systems can be envisioned in the future.

Practical relevance

To the best of our knowledge, this is the first work that examines the BGP origination behavior of serial hijackers, a category of networks that has received surprisingly little attention in terms of broad and detailed empirical assessment. We argue that serial hijacking behavior needs attention from both operators and the broader research community to allow for faster mitigation or even prevention of hijacking events.

While, as expected, not all ASes flagged by our classifier are serial hijackers, we note that all such networks do show a highly distinctive origination pattern. Scrutinizing these networks, we found widespread indications of malicious behavior, with flagged ASes being more likely to be in blacklists associated with malicious behavior, as well as different indicators of misconfiguration. Since our system is *orthogonal* to commonly deployed reputation systems (*e.g.,* event-based hijack detection), and works out-of-the-box using readily available public BGP data, we believe that, after

[11] BGPmon was acquired by Cisco and is now merged with Cisco BGPstream service [62]

refinement, the output of our classifier might be used to provide additional scoring data, *e.g.,* in scoring-based reputation systems.

Even after disclosure, hijack reports and discussions on mailing lists typically focus on isolated incidents (*i.e.,* usually the prefixes of the network operator issuing the complaint), and the case of Bitcanal shows that it took *years* to effectively cap hijacking activity and disconnect Bitcanal. Our metrics can compactly, and yet comprehensively, capture the dominant origination characteristics of misbehaving networks. Thus, even after initial disclosure on mailing lists, our metrics and analysis provide an instant picture of the Internet-wide "state-of-affairs" of the networks in question, which can help operators to readily assess the full extent of hijacking activity, and thus inform the process of coordinated mitigation.

Limitations

We note that our classifier is solely based on the routing activity of ASes. We focus on identifying routing characteristics of serial hijackers, which present one particular case of hijacking activity. Our detection mechanism does, naturally, not cover the space of hijacking activity exhaustively. While we find that serial hijackers do show distinct announcement patterns, our classifier does falsely tag some legitimate ASes as having BGP behavior similar to serial hijackers, as reflected in the precision of our classifier of $\approx 80\%$. We hence want to stress that our classifier, while effective in narrowing down the set of flagged ASes to ≈ 900 ASes, can and should *not* be deployed, as is, to generate, *e.g.,* filtering rules. Furthermore, if deployed at any point in the future, there is a potential risk that hijackers could craft their BGP announcements to not exhibit the characteristics captured by our classifier and thus evade detection. Another limitation of our work is that we focus solely on distinct features of the BGP *origination* patterns of networks and therefore on BGP origin hijacks. Hijacks which modify the AS path leaving the legitimate origin AS unaltered are therefore not captured in our data. Our work constitutes an initial view into the properties of serial hijackers with much future work to be done.

Future work

In the future, we plan to extend the features we leverage for classification. Potential additional features include more BGP-derived properties, such as AS-path characteristics of hijacked prefixes, as well as sub- and super-MOAS events. We believe that such features could not only further improve separation of ASes, but also shed light on topological properties of hijackers, *e.g.,* upstream networks and peering facilities leveraged by serial hijackers. We further plan to cross-evaluate our findings with other external datasets. In a first step, we correlated our identified ASes against blacklists, finding indications of persistent malicious behavior.

Our work is based on 5 years of historic BGP routing data, and we point out that some of the dominant characteristics of serial hijackers only become visible when studying routing data at longer timescales. We note, however, that our features to capture advertisement volatility are in fact computed over much shorter timescales *i.e.,* bins of weeks and months, and our address space features might well yield distinctive results when applied to shorter timescales. This suggests that early detection of systematic misbehavior might be indeed possible. We plan to further study the time-sensitivity of our approach to assess closer-to-real-time detection possibilities.

9 Conclusion

This work presents a first in-depth study of the characteristics of serial hijacker ASes. We identify a set of dominant and distinct origination patterns of a set of ground truth ASes. Based on our observation, we train a machine learning model capable of identifying ASes with similar characteristics, *misbehaving ASes*, in the global routing table. Our classifier identifies some ≈ 900 misbehaving networks in the Internet. We find a a wide range of indicators both for misconfigurations, as well as for malicious activity of these networks. Our work presents a solid first step towards gaining a better understanding of malicious behavior in BGP, and use it towards automatically identifying suspicious ASes in the Internet or as input in network reputation scoring systems in the future. Indeed, some networks providers and researchers have already used these metrics and the serial hijacker list to evaluate network behavior and reputation [131].

This study reveals the existence of BGP serial hijackers—networks that persistently perform hijacks in BGP, and for which hijacking represents a significant share of their activity. In other words, in the current state of BGP security, there are very few barriers to performing even the basic forms of routing attacks and networks are able to engage in repetitive malicious behavior with little consequences.

Chapter 4

Empirical Analysis of Misconfigurations

The previous chapter studied systemic malicious behavior in BGP. Unfortunately, BGP is vulnerable not only to attacks but also to unintended misconfigurations. This chapter focuses on misconfigurations that impair availability in the Internet. It studies route leaks: misconfigurations caused by an AS unintentionally sending information to its BGP neighbors that it should not have shared given the business agreement and AS relationship with its neighbors. The goal of this work is to characterize the prevalence of harmful route leak events and thus centers on route leaks that occur at the edge of the Internet and shift core traffic through under-provisioned link and creating bottlenecks.

AS relationship—the business agreements that networks establish with other networks—influence route selection and industry structure, as described in Chapter 1 Sections 1.3 and 1.4. Not all links between networks can handle large amounts of traffic and route leaks making the traffic go on smaller links impact the availability and reliability of Internet services, as mentioned in section 2.

Research on this topic has focused on inferring peering relationships between networks to find route leaks when network paths in BGP do not match the inferred relationships [58–61]. Thousands of route leaks can be detected daily using AS relationships [62] but there is still a high rate of false positives and not all harm availability. Indeed, relationships between networks are increasingly complex and not always deterministic: two networks can peer for some traffic and have a customer-provider relationship for other traffic [132].

This study proposes a new method to detect and monitor route leaks based on their inherent path characteristics and independent of individual network relationship. The detection mechanism is based on the position of networks along the path relative to the core and edge of the Internet, *i.e.,* the centrality of network in the Internet. It focuses on route leaks that result in a small non-transit network accidentally becoming a transit provider for large neighboring networks, making traffic from the core of the Internet go to this small edge network and then boomerang back to the core, *i.e.,* doing a U-turn before reaching its final destination. In the edge network, traffic usually reaches levels of congestion that disrupt Internet service to the impacted IP address blocks. This study leverages the fact that this U-turn pattern is not usually seen in BGP routing announcements.

Using this method, this study builds a dataset of U-turn route leaks based on 15 months of routing data. Then, further scrutinizing the route leak activity, it uncovers the dynamics of route oscillation during the events. Then, leveraging the fined-grained data, it assess the impact of specific protocol configurations in route leak propagation and mitigation, as well as the main victims and involved networks. It finds that large Content Distribution Networks (CDNs) such as Akamai and Amazon are frequent victims of these misconfigurations and that specific BGP configuration limit the spread of most events found in wild. This findings have thus relevance for the operator community, illuminating the harms of these events and helpful practices to limit their spread.

1 Background

The Border Gateway Protocol's (BGP's) inherent lack of security mechanisms continues to pose major challenges to operators around the globe. Reachability information propagated via BGP is typically not verified and can thus be erroneous, either caused by deliberately crafted announcements (*e.g.,* route hijacks), or accidentally, due to misconfigurations. While route hijacks with malicious intent caught the broad attention of the research and operator community [53–55, 117, 133, 134], *route leaks* caused by unintended misconfigurations have received little empirical assessment.

Prime examples of harmful route leaks include scenarios in which a non-transit network exports some or all of the routes it learned from one of its transit providers to its neighboring networks, hence unwillingly becoming a transit provider. Route leaks happen frequently, with new reports of widespread connectivity disruptions appearing regularly [135, 136].

1.1 Introducing U-turn Paths

This work specifically focuses on *U-turn* route leak events: instances in which a network closer to the edge of the Internet unintentionally becomes a transit provider for address space it learned from its own providers. On June 24^{th}, 2019, a major route leak exhibiting these characteristics affected Amazon [1, 16]. Figure 4-1 illustrates BGP path changes during this event for a /22 prefix from Amazon (AS16509) and a leaked /23 subprefix of that block, as seen via AS701 (Verizon). Under normal conditions, traffic from Verizon flows to Amazon via a direct peering link. During the route leak, however, the traffic takes a detour via a more-specific prefix propagated through AS33154 and AS396531 (both small networks that usually do not provide transit for Amazon), back to Level3 (AS3356), and eventually reaching Amazons AS. The sketched route leak affected over 30,000 paths (across multiple prefixes), and Internet traffic made a *U-turn* from the core to the edge and then back to the core. This major route leak instance resulted in widespread congestion, impacting core services of the Internet (*e.g.,* DNS), as well as services provided by Amazon [1].

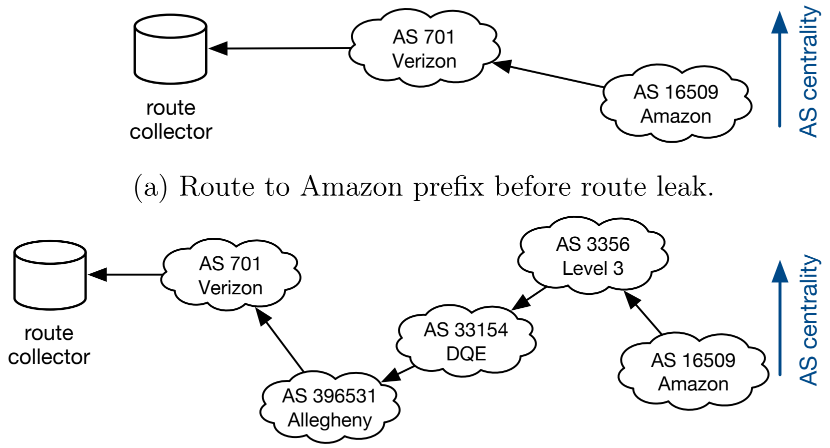

(a) Route to Amazon prefix before route leak.

(b) Route to Amazon prefix during route leak.

Figure 4-1: Recent route leak of Amazon's address space [1] caused traffic to take a detour passing through otherwise unrelated networks as seen by Verizon (AS701).

U-turn route leaks typically do *not* modify the origin AS of the erroneously propagated route, and are the result of unintended misconfigurations by otherwise legitimate networks. These characteristics rule out many well-known hijack detection mechanisms, including systems that detect origin changes [55, 62], bogons and IP addresses squatting [137], as well as reputation-based systems that identify likely malicious networks and actors [98,120]. However, leveraging the unusual U-turn pattern of these route leaks, our approach is based on translating AS numbers in BGP paths to centrality metrics that capture the position of ASes relative to the core and edge of the Internet.

In a nutshell, our approach uses AS centrality metrics to identify BGP paths that traverse an AS closer to the core of the Internet, reach an AS closer to the edge, and boomerang back to an AS closer to the core. We call them U-turn paths.

It is well-known that the AS-level graph of the Internet has long moved away from a strictly hierarchical shape to a flatter, more complex interconnection structure, fueled by increasing interconnection between ASes at IXPs as well as the rollout of wide-area networks of major content providers [138–141]. Consequently, we do not attempt to strictly partition the Internet into an edge or a core, nor infer or rely on individual AS relationships. Instead, we use centrality measures of networks to identify macroscopic discrepancies simultaneously happening in many AS paths that are unlikely to be the result of peering relationships, and are hence indicative of route leak events.

1.2 Capturing AS Centrality

To capture AS centrality in the Internet we use two metrics:

- **AS Customer Cone Size:**This metric is the number of ASes in the cone formed by customers and customers of customers of an AS [9] and is therefore

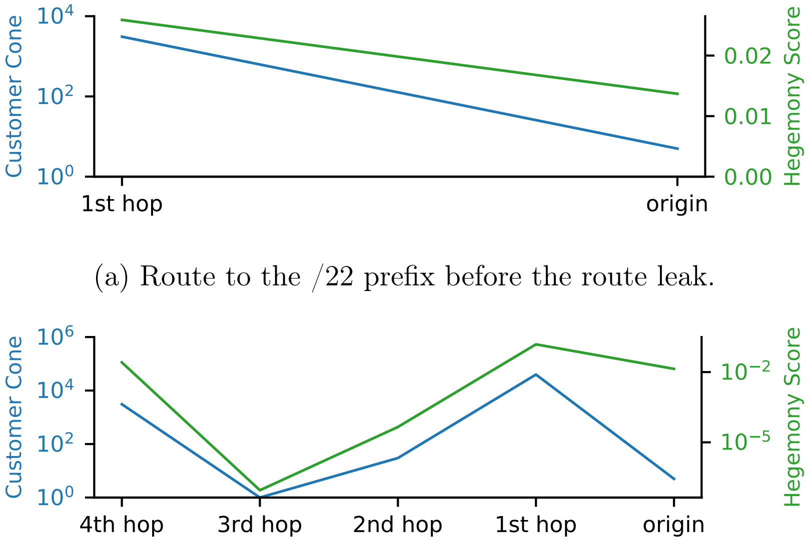

(a) Route to the /22 prefix before the route leak.

(b) Route to the /23 sub-prefix during the route leak.

Figure 4-2: Example of customer cone size and hegemony score changes along path during recent leakage of Amazon address space as seen by Verizon (AS701)

based on AS-relationship assumptions. For most ASes, the customer cone size is fairly stable, slowly increasing or decreasing over many months. AS customer cone size varies between 1, for stub ASes, and almost 40,000 for the largest Tier-1 providers.

- **AS hegemony score:** This metric represents the portion of IPv4 address space that passes through an AS according to control-plane data from unbiased vantage points that are neither too close nor too far from the AS in question [142, 143]. Hegemony scores range between $(0, 1)$. Stub ASes score in the order of 10^{-12}-10^{-5}, depending on the size of the IPv4 address space the AS originates, and large transit providers reach scores as high as 0.1

Leveraging both approaches to capture centrality: While customer cone size and hegemony score are correlated for many ASes, they do characterize AS centrality differently. The hegemony score does not rely on a customer-provider inferred AS hierarchy, but rather views the AS graph as a flat mesh. More importantly, the hegemony score considers the size of the routed address space (*i.e.,* number of routed IP addresses) that either transits or is originated by an AS. The difference between these metrics can be exemplified by Amazon, AS16509, which only has a customer cone size of 5, but reaches a relatively high hegemony score of 0.013 (in the top 25): this network does not provide transit to many other ASes, but originates a large chunk of IP address space.

1.3 U-turn Route Leak Manifestation

We use ASes' customer cone size and hegemony score to study the relative centrality of ASes along the paths seen before and during the route leak of June 24th, 2019 described in § 1.1. Figure 4-2 depicts customer cone size and hegemony score along sample paths toward an Amazon /22 prefix before the route leak, and toward a related (/23) subprefix during the leak. Before the route leak, the path to the /22 prefix originates from AS (Amazon AS16509), which has a lower customer cone size and hegemony score and hence is close to the edge, and reaches the first hop (Verizon AS701). In contrast, during the route leak, there is a U-turn in the route that appears for the /23 prefix: the U-turn is between the first hop AS (AS3356), which has a large customer cone (39,642) and hegemony score (0.1527), the 2nd and 3rd hop, which have very low metrics (customer cone of 30 and 1, and hegemony score of 10^{-5} and 10^{-7}, respectively), and the 4th hop, which again has a large customer cone (3,082) and high hegemony score (0.0259).

In this leak, the minimum difference between the bottom of the U-turn and the two tops of the "U" is 3,081 in customer cone size and 0.0259 in hegemony score. We call this the *depth* of the U-turn. In this specific leak, the bottom of the U-turn in the paths is a stub AS, making it a very deep U-turn, unlike paths that are usually seen in BGP.

1.4 Related Work

Early work on route leaks started by characterizing BGP misconfigurations and specific route leak events. Mahajan et al. in 2002 found that almost 75% of all new prefix advertisements were the result of BGP misconfiguration, though only one in 25 affected connectivity [144]. Hiran et al. studied the China Telecom Route Leak of April 2010 using BGP data [145].

Researchers have also proposed other techniques to detect route leaks, which potentially could complement ours, using BGP data in various ways. Song et al. propose a method based on detecting routing loops [56]. Su et al. detect route leaks by comparing long-lived paths with anomalous ones [59]. Siddiqui et al. propose a local technique whereby an AS uses the RIBs at their own the border routers, and data plane information to infer if a BGP update from a neighbor corresponds to a route leak [57, 58]. Mauch's website shows detected route-leaks [60], which infers a route leak when three "major" networks are all on the same AS-path. Other works on route leak detection use pairwise AS relationships to find route leaks. [61] proposes a set of heuristics based on AS relationships to identify route leaks. In [132], the authors infer probabilistic AS relationships based on BGP data for route leak detection.

In contrast, our work focuses on macroscopic-level anomalies in path that capture violations of industry structure, independent of single AS relationships. In [117] the authors use the changes over time in one of the centrality metrics we use (hegemony scores) to characterize and classify hijacks, using in particular the number of "valleys" in hegemony scores along a BGP path during a hijack as a classification feature. In this work, we use the median AS hegemony of the past 30 days as one way to characterize the centrality of an AS related to the core of the Internet and use it to characterize BGP paths with the objective of detecting U-turn route leaks. In [65], the author define baselines of announcements volume to detect anomalous routing events, including hijacks, infrastructure problems and route leaks. In contrast, in our work we focus only on route leaks making traffic shift from the core of the Internet towards the edge and back. Finally, Sermpezis et al. introduce ARTEMIS, which automatically and quickly detects and mitigates BGP hijacking and unintended misconfigurations. ARTEMIS is designed for the individual AS [55] to leverage its own internal network information to detect anomalies. In contrast, our detection method does not require prior information of prefix and AS relations and only uses publicly available routing metrics to evaluate ASes in BGP paths.

1.5 Roadmap

The rest of this chapter is organized as follows: We describe the datasets we use in § 2. We study the characteristics of a set of route leaks for which we have ground truth in § 3. Leveraging our observations, we introduce our detection mechanism in § 4 and present characteristics of a series of route leaks we identify in 15 months of BGP updates in § 5. Finally we discuss our results, limitations, and implications in § 6.

2 Datasets

In this section we describe the different datasets we collect and use for our study.

Route leak ground truth: We collect ground truth data on 12 major route leak events that did not modify the origin AS of BGP routes during the events[1], which happened in the last 3 years and were publicized in news articles, blog posts, operator mailing lists, or tweets [1, 17, 136, 146–151].

BGP Updates: We collect 15 months (from February 2019 until end of April 2020) of BGP updates from all RIPE RIS and RouteViews collectors. These collectors, 44 in total have about 320 direct peers that provide a full feed of the global routing table to the collectors. Additionally, we collect the BGP updates from the days of the route leak events in our ground truth if they happened outside our measurement window.

BGP RIBs: We collect RIB records from RIPE RIS and RouteViews collectors twice per month in days where there is no known route leak event from February 2019 until end of April 2020. In addition, we collect RIB records from the 24h hours prior

[1] We collected 9 other major events that were reported as route leaks, however in these events the origin AS was changed, which is also a characteristic of a route hijack, and can be detected using readily available tools, e.g., [62]. We thus chose to exclude these events from our ground truth dataset.

to route leaks in our ground truth. On average, we collect some 365M RIB dump records per day.

AS Customer Cone data: We collect all customer cone data available from 2017 to 2020. This metric is computed and made publicly available by CAIDA once a month [9], using 5 days of BGP data from all RIPE and RouteViews collectors. The customer cone tends to be fairly stable for most ASes. However, in the case of particular routing events in the BGP data used to compute the customer cone size, the size might not be representative of the actual customer cone size. To account for this, we compare the last 3 months of customer cone data for each AS and in the case of a variation over 2 orders of magnitude of customer cone size compared to the previous months, we consider the data to be missing (happening in less than 0.02% cases).

AS Hegemony data: We collect all AS Hegemony data available from 2017 to 2020. This metric is computed every 15 minutes using BGP data from all RIPE and RouteViews collectors and made publicly available by the Internet Health Report (IHR) project from Internet Initiative Japan (IIJ). The hegemony score of an AS also tends to be fairly stable but it is known to change in ASes impacted by routing events [117]. To use this metric to evaluate the centrality of an AS, we use the median AS hegemony score of an AS calculated over the previous 30-day sliding window.

3 Anatomy of U-Turn Route Leaks

In this section, we take a deep dive into the properties of a set of ground truth route leaks. We extract the BGP update messages that correspond to the route leaks and apply and test metrics to study the dominant characteristics of these events. This step forms the basis for understanding the prevalence of U-turn path in certain route leaks.

3.1 Distilling Route Leak BGP Update Messages

Public information about route leaks is not only rare but is also sparse and does not typically include the full set of leaked prefixes. Thus, to collect BGP data for the analysis of these route leaks, we need to infer the leaked prefixes and associated updates. Leveraging publicly available BGP data (RIB dumps and updates), we isolate BGP updates related to our ground truth route leaks in a two-step process:

(i) **Selection of route leak candidate updates:** Using the reported time of the route leak, a time window starting 2 hours before the reported time and using the reported ASN of the network responsible for the route leak, we select all BGP updates whose path contains the route-leaking ASN.

		visibility	**related**	**updates by prefix type**			**unique**		**CC U-turn depth**			**Hegemony U-turn depth**		
date	**leak AS**	**% peers**	**updates**	match	sub	new	**prefixes**	**paths**	25^{th}	50^{th}	75^{th}	25^{th}	50^{th}	75^{th}
April 5, 2020	AS7552	100%	459K	75.5%	18.0%	6.5%	9,557	16,251	22	22	51	2.741	2.982	2.982
March 31, 2020	AS50048	99.7%	403K	53.9%	44.2%	1.9%	2,043	33,565	2,934	2,934	2,984	981	981	1,245
February 7, 2020	AS8359	100%	24M	17.3%	81.7%	1.0%	21,028	50,324	0	0	0.02	0	0	0.92
June 24, 2019	AS33154	95.9%	2.2M	36.9%	62.9%	0.2%	65,858	47,735	3,081	6,575	15,755	287,457	287,457	646,448
June 6, 2019	AS21217	98.1%	2.1M	14.5%	80.7%	4.8%	47,663	41,576	17	124	3168	18.48	937.47	14292.01
April 4, 2019	AS60280	94.2%	1.1M	96.1%	0.5%	3.4%	15,713	34,700	5	170	534	3.13	15.26	107.80
November 12, 2018	AS37282	49.0%	33.3K	93.5%	5.1%	1.4%	351	502	0	0	7	0	0	6.68
May 18, 2018	AS263086	99.6%	2.2M	40.4%	54.4%	5.2%	10,771	99,775	166.00	195.00	195.00	271.21	2,066.25	2,066.25
January 24, 2018	AS196737	100%	4.0M	58.9%	17.9%	23.2%	306,471	150,643	226	845	845	721.72	721.72	721.72
December 30, 2017	AS55644	100%	4.4M	31.4%	8.4%	60.2%	397,276	109,966	1	1	13	0	0	0
December 27, 2017	AS4788	100%	1.2M	87.2%	3.9%	8.9%	4,051	50,011	0	7	53	0.57	5.60	43.59
October 21, 2017	AS263361	100%	7.0M	48.4%	49.8%	1.8%	56,019	601,766	1	15	106	0	0	0.0001

Table 4.1: Manifestation of ground truth route leak events in the global routing table.

***(ii)* Removal of updates unrelated to route leak event:** We leverage a RIB dump taken from the day prior to the reported route leak. We then check for each candidate update from the first step, whether *(i)* the prefix announced in the candidate update is present in the respective per-peer RIB of the previous day and *(ii)* whether the reported leaking ASN is found on the path of the entry of the previous day. If both conditions are fulfilled, we remove the candidate update from our analysis, since it is unlikely to be related to the route leak event, but part of the regular operation of the reported ASN.

This methodology leaves us with a set of updates that *(i)* involve the leaking AS in the reported timeframe, and *(ii)* announce prefixes that were not routed via the leaking AS before the route leak event. Hence, these updates are very likely to be directly connected to the route leak event. Using this set of BGP updates, we next characterize how the reported ground-truth leaks manifest in BGP.

3.2 Characteristics of U-turn Route Leaks

In this section we study in detail the visibility, volume, and prefix characteristics of the 12 route leaks included in our ground truth. Then, we scrutinize the centrality of ASes along BGP paths relevant to the leaked prefixes before and during the route leaks, using both customer cone size and hegemony score. Table 4.1 provides a summary characterization of the route leaks included in the ground truth. For each event, we note the reported AS number ("leak AS") and the corresponding date.

High visibility: We express the spread/visibility of route leaks by calculating the percentage of direct peers of our route collectors that saw paths affected by the route leak. Stunningly, we find that all but one route leak reach visibility higher than 90%: route leaks spread far, and hence affect large swaths of the Internet.

Update dynamics: Route leaks result in a large number of updates, 9 out of 12 of our identified route leaks result in millions of updates, most impacting thousands of prefixes and paths. These BGP messages, if not filtered, cause redirection of traffic on the Internet, sometimes to links without the adequate capacity, which in turn causes significant congestion or even loss of connectivity, see e.g. [135].

Comparing the prefixes of individual update messages against the prefixes present in the RIB dumps of BGP collectors from the previous day (recall § 3.1), we find that most route leak updates either replace already existing prefixes ("match") or inject subprefixes into the routing system (as in our previously featured Amazon example). Only rarely do route leaks result in the announcement of previously unrouted prefixes.

U-turns in BGP paths: For all updates, we extract the corresponding AS paths and apply our two metrics of customer cone (CC) and hegemony score to identify U-turns. In particular, we transform the AS path into a sequence of CC/hegemony scores and then identify the existence of U-turns by looking for "dips", *i.e.,* the path goes from an AS with a high score to an AS with a lower score, and eventually back to an AS with a higher score. The AS (or multiple ASes) with a lower score constitute the bottom of the identified U-turn portion of the route. We then assess the *depth* of the U-turn by comparing centrality values for the bottom and top of the U-turn section of the path. Specifically, the *depth* is the ratio of the maximum (CC/hegemony) value on each "side" of the "U" divided by the smallest (CC/hegemony) value of the AS at the bottom of the U-turn. Note also that there can be more than three ASes involved in a U-turn.

Table 4.1 shows, for each route leak and all identified paths, the 25th, median, and 75th percentile of the depth of the identified U-turns. The table reveals a variety of cases, highlighting differences in *(i)* overall depth, *(ii)* number of "deep" U-turn paths, and *(iii)* centrality metric. The depth of U-turns varies dramatically across different route leaks, since it is dependent not just on the leaking AS (*i.e.,* the AS at the bottom of a U-turn), but also on the involved core ASes on the path. In the route leak from June 24, 2019, the U-turn is very deep for both the CC size and hegemony score, while the route leak from October 21, 2017, is much less.

For many route leaks, the CC size and hegemony scores are high for only a subset of the paths: *e.g.,* in the route leak from December 27, 2017, 25% of the paths had a CC U-turn depth of at least 53, but another 25% of the paths did not show a U-turn in either chosen metric. A reminder that any detection approach based on the existence of U-turn paths will only be capable of detecting a subset of the paths of a route leak.

Finally, some route leak U-turns are deeper in customer cone size compared to hegemony score (see the route leak from December 30^{th} 2017) and some the other way round (see the route leak from April 5^{th} 2020), reflecting the different way these metrics express AS centrality.

The obvious diversity of the depths of the U-turns in Table 4.1 is put in an illuminating context when compared with the depths of paths a normal RIB, which is done in Section 4.1 and in particular Figure 4-5.

3.3 Route Leak Dynamics in Detail

So far, we only considered update messages that carry U-turns in their AS path—the key property of the route leaks under study. However, route leaks are a transient phenomenon, and we are interested in understanding how a route leak event unfolds over time. That includes the first BGP update messages reflecting the initial start of the route leak event, the spread of these updates, as well as corrective responses of the routing system, and the eventual end of the route leak event. Autonomous systems react differently and independently to ongoing routing events, causing update bursts to have uneven impact in the overall routing system. In this section, we look in detail to all movements in the routing systems caused by U-turn updates from the ground truth route leak of June 24^{th}, 2019 in Table 4.1.

Figure 4-3 shows the evolution of the number of prefixes involved, the visibility, the count of active (prefix,peer) pairs and the volume of BGP activity over time for the route leak of June, 24^{th},

Varying impact over time: The top graph in Figure 4-3 (blue) depicts the count of active prefixes —prefixes for which U-turn updates get in peers routing table—in the route leak. The larger count of active prefixes happens in the first 40 minutes of this route leak and the count gets close to 0 before going back to higher levels multiple times during the later part of the route leak (see dips in the blue line). Although leaked prefixes come and go during the route leak, they reach to most peers throughout the event. The second graph in Fig.4-3 (orange) shows the number of peers forwarding any route leak prefixes over time. The peer count is fairly stable as most impacted peers have leaked prefixes until the very end of the event. However, as leaked prefixes come and go, the overall impact of the route leak in the routing system is significantly reduced 45 minutes in the route leak. The third graph in 4-3 (red) has the count of (prefix, peer) pairs, *i.e.,* the count of how many leaked prefixes are in each peer's routing table across all peers. This metric clearly shows that at the beginning of the route leak, there were more leaked prefixes in peers routing table than later on, even though most peers had some leaked prefixes until the very end of the event.

Quick spread of U-turn updates: Announcements with U-turns get into peer routing tables very fast. In this route leak, in just about 80 seconds, 250 ASNs peering with the BGP collectors (75% of active peers) forward leaked routes involving over 65,000 prefixes, reaching the maximum spread with over $160,000$ (prefix, peer) pairs. At many points in this route leak, the volume of prefixes in peers' routing table drops significantly (dips in red line) but quickly new U-turn update appear and many leaked prefixes are back in routing tables. This movement only stabilizes at the end of the route leak.

Fast initial correcting reaction, but also oscillations: In this route leak, the first BGP message correcting a U-turn update takes less than 8 seconds to appear, and afterwards there is a constant flow of such messages. Indeed, throughout this

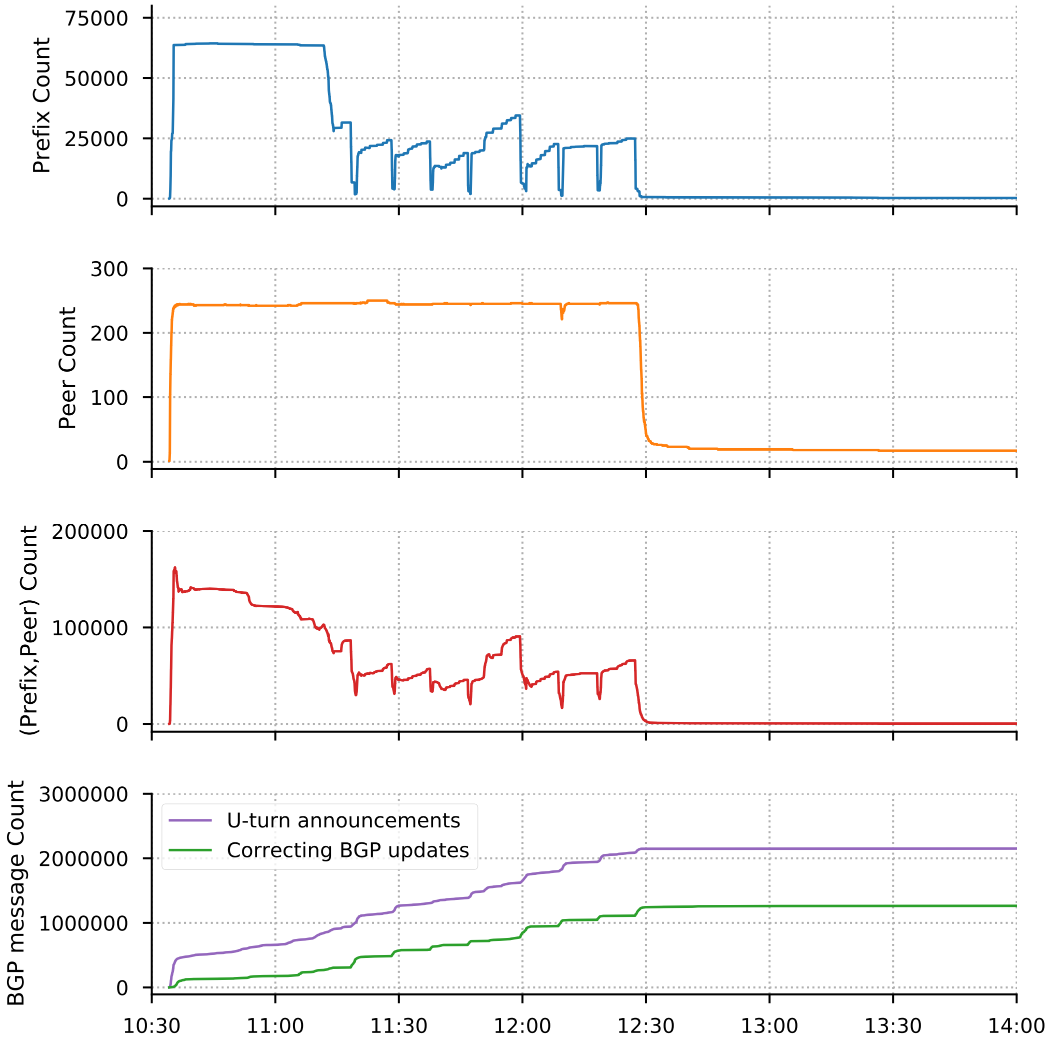

Figure 4-3: Active prefixes, peers, (prefix, peers) pairs, and total number of BGP messages over time for the route leak on June 24, 2019.

event, there is a constant flapping of BGP routes between routes with U-turn in them and routes without U-turns or withdrawals for the involved prefixes. The bottom graph in Fig.4-3 depicts the cumulative count of U-turn updates (purple line) and of correcting BGP updates (green line). In this graph, the sharp increases (the jumps) in the count of correcting updates coincide, as one would expect, with sharp decreases in the number of prefixes (blue line in top graph) and the number of (prefix, peer) pairs. Moreover, the number of prefixes (blue line) drops almost to zero. However, perhaps surprisingly, these jumps in correcting messages are closely followed with similar jumps in U-turn updates (purple line, bottom graph) that in turn are reflected in jumps up in the number leaked prefixes in peers' routing tables.

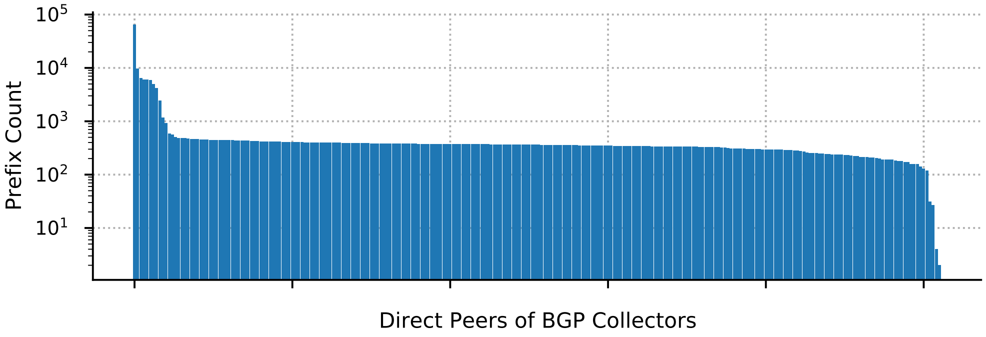

Figure 4-4: Count of leaked prefixes forwarded by BGP collector's peers. Only one peer has over 60,000 prefixes. Most peers have close to a thousand prefixes.

Unequal impact in peers: Not all ASNs that are reached by the route leak have the same impact, and that impact changes throughout the event. Figure 4-4 show the count of leaked prefixes forwarded to BGP collectors by the peers that see the route leak. Only one peer includes more than 60,000 leaked prefixes in its routing table, whereas most peers have only about a thousand prefixes and a few are only lightly impacted.

The end of the route leak: After 1 hour and 54 minutes, 90% of the leaked prefixes in routing tables, i.e. (prefix, peer) pairs, have been removed, putting the count below 10% of its maximum. A minute later, the (prefix, peer) pair count goes below 1%. However, thereafter, the (prefix, peer) pair count decreases very slowly, reaching, after 14.5 hours, a minimum of 261, accounting for 224 different prefixes distributed among 15 peers. We do not detect any BGP messages that further remove the few prefixes left in the next 12 hours. This is likely due to withdrawals that were not seen by the BGP collectors, or routes stuck in routers of ASes along the path before the U-turn ASN [152]. We study properties of bursts of U-turn updates in § 4.3 and devise a method to distill route leak events from these update bursts in § 4.4

4 Route Leaks in the Wild

With our insights on route leaks in hand, we next devise a method to detect such events in the wild. To this end, we first assess the overall prevalence of U-turn paths in the routing system in § 4.1. We then introduce our thresholds for identifying BGP updates with U-turn paths in § 4.2. We then study the properties of identified U-turn bursts in § 4.3 and finally devise a method to distill route leak events from U-turn bursts in § 4.4.

4.1 Prevalence of U-turn Paths

The key feature that triggers our route leak detection mechanism are U-turn paths in BGP update messages, under the assumption that such paths are exceptional cases that do not regularly occur—otherwise, our method would erroneously detect a large number of potential route leak events. In this section, we assess how prevalent U-turn paths are in the general routing system, i.e., we contrast our findings from the previous sections with properties of "regular" BGP paths.

Our observations serve as the basis to devise a system to automatically detect relevant U-turn route leaks in BGP data.

Rarity of U-turns in regular paths: Using the more than 30 days of RIB dump records from all RIPE RIS and RouteViews collectors from days without a known routing event, we study paths that we expect to be more regular, knowing that unknown events might still be impacting the data. On average, after cleaning AS-Sets, loops, private ASNs, and prepending, we find about 32M unique paths in RIBs (28.5M for IPv4 prefixes, 4.7M for IPv6). Of these paths, between 4-6% per day have a U-turn for either customer cone or hegemony score.

Shallow U-turns : Of the few U-turns in RIB dump paths, most U-turns are "shallow" and usually less than 1% are deeper than a full order of magnitude, *i.e.,* the centrality of the bottom AS in the U-turn is more than 10 times smaller than the centrality of ASes at the top of the U-turn, according to either customer cone or hegemony score.

Comparing RIBs and ground truth route leak paths: The contrast between U-turn paths in route leaks versus regular paths becomes very clear when we compare the likelihood of such paths in RIBs and updates from route leaks. All route leaks have a much larger fraction of related AS paths with deeper U-turns than the few U-turns found in RIBs. About half the route leaks have over 50% of the paths with U-turns with a depth ratio of 100 (100 times difference in centrality), which instead happens in less than 0.1% of the paths found in RIBs. Figure 4-5 depicts the distribution of paths by U-turn depth for our ground-truth route leaks, as well as for a RIB dump containing a full set of BGP paths to any destinations, irrespective of the presence of a route leak. Fig. 4-5a shows the distributions for customer cone size and Fig. 4-5b for hegemony score. The distributions of paths by U-turn depth are represented in blue (full RIB), grey (*individual* ground-truth route leaks), and orange (*all* ground-truth route leaks).

4.2 Finding U-turn Announcements

Next, we look for announcements with U-turns in 15 months of BGP updates from all RIPE and RouteViews collectors. From February 2019 until end of April 2020, we process over 317 billion updates.

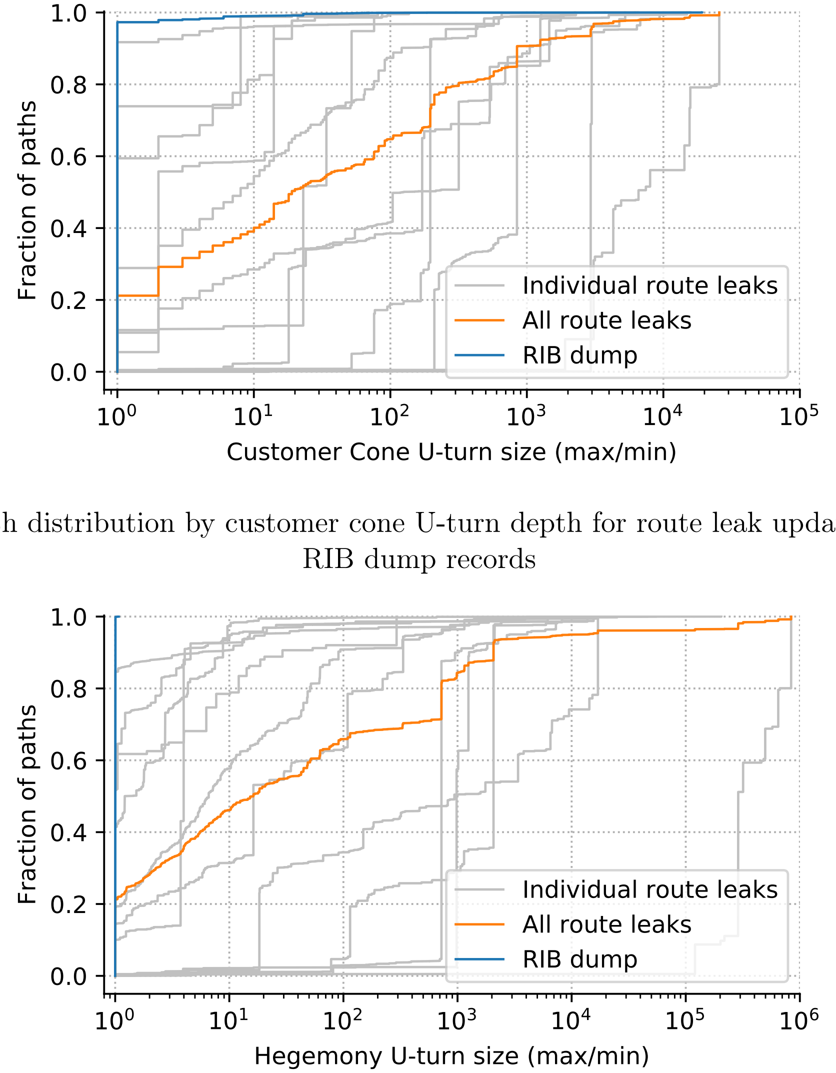

(a) Path distribution by customer cone U-turn depth for route leak updates and RIB dump records

(b) Path distribution by hegemony score U-turn depth for updates from route leaks and RIB dump records.

Figure 4-5: Distribution of route leak and RIB dump paths according to U-turn size measured in customer cone and hegemony score.

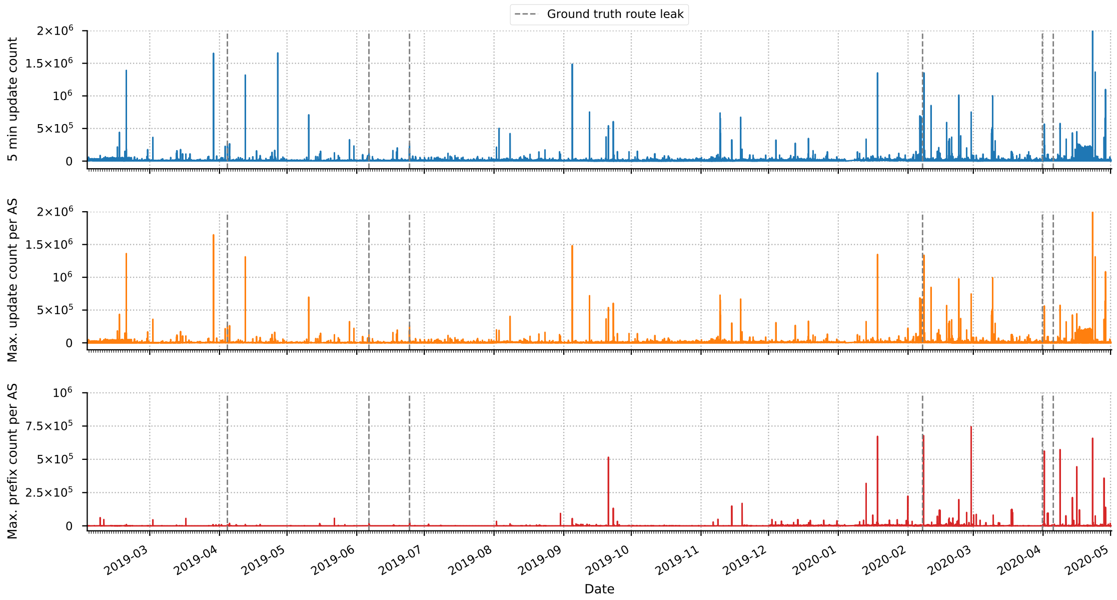

Figure 4-6: Time series for detecting route leaks: count of updates per 5 min bin with U-turn deeper than 20 times in customer cone size and 15 times in hegemony score (top), and maximum update count per U-turn AS (bottom).

Deep U-turns: Our depth metric of customer cone and hegemony score evidently separates many of the paths seen in our ground-truth route leaks from regular BGP paths. However, as visible in Figure 4-5, not all route leak paths make a U-turn, and some of the U-turns are very shallow. Setting our parameters to tag announcements as route leak candidates is hence a trade-off between reducing the number of falsely tagged paths and achieving detection completeness. Based on the analysis of U-turns in paths of announcements from ground truth of route leaks as well as in path from RIBs, we consider a U-turn to be deep when the ratio between the maximum and minimum customer cone is 20 and the ratio between the maximum and minimum hegemony score is 15. These thresholds capture paths in all ground truth route leaks, though only a very small fraction of them are captured for 3 of them. More than 50% of all ground truth route leak paths have U-turns deeper than these thresholds, whereas less than 0.5% of paths in our RIBs satisfy these criteria.

Deep U-turns in BGP updates: Using the thresholds for customer cone size and hegemony, we count the number of announcements with such U-turns per 5 minutes buckets. The top graph in Figure 4-6 shows the count of U-turn announcements, with U-turns deeper than our thresholds for both metrics, in 5-minute bins. Spikes of unusually large numbers of announcements with U-turns deeper than our thresholds are clearly visible. In Figure 4-6, the vertical, green dotted lines indicate the reported start time of route leaks in our ground-truth set: even though these events made the news, there are even larger spikes of unusual update activity.

The middle graph in Fig. 4-6 represents the maximum count of U-turn announce-

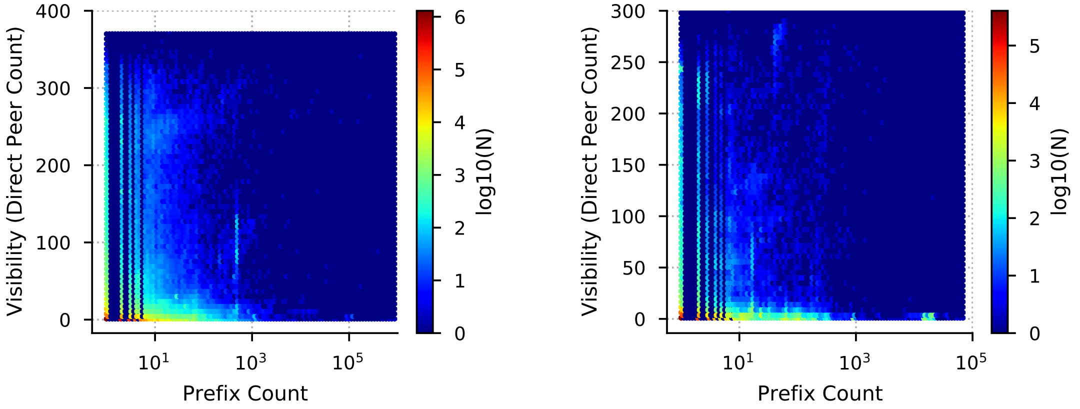

(a) IPv4

(b) IPv6

Figure 4-7: Heatmaps of U-turn announcements bursts by prefix count and visibility (% of active peers) for IPv4 (top) and IPv6 (bottom).

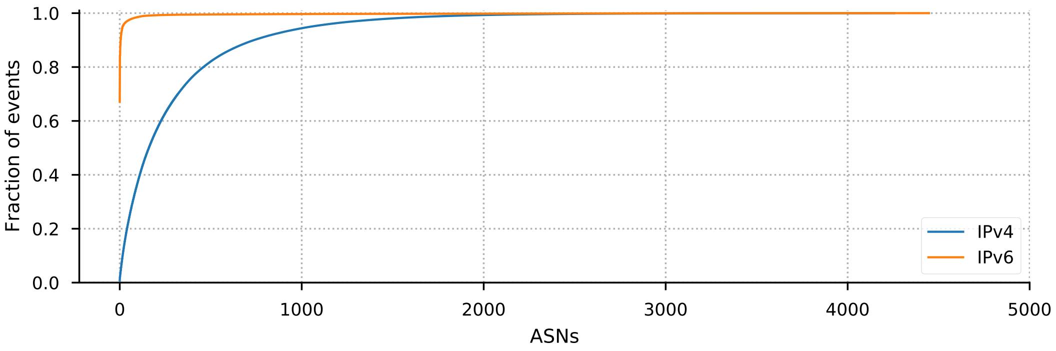

Figure 4-8: Cumulative sum of IPv4 and IPv6 U-turn announcements bursts by ASNs ordered by decreasing count of events.

ments meeting the customer cone and hegemony thresholds with the same AS at the bottom of the U-turn per 5 minute bucket. This graph clearly shows that most—but not all—spikes from the graph counting all U-turn announcements, share the same section of the U-turn in their BGP path and are hence indicative of large-scale route leak events. However, the number of announcements is not directly related to the number of prefixes involved. The bottom graph in Fig. 4-6 shows the prefix count from the announcements with the same AS at the bottom that generated the maximum count (the announcements in the orange graph above).

4.3 Bursts of Consecutive U-turn Announcements

Using the ASN at the bottom of U-turns as identifier, we define a "U-turn burst," or simply "burst" as a sequence of BGP announcements with a U-turn with the same identifier ASN whose inter-epoch times are within 5 minutes. (Note: a burst can be as small as one.) In 15 months of BGP announcements, we find 3.3 million U-turn bursts (2.7M for IPv4 and 750K for IPv6, only 77K events involve IPv4 and v6 prefixes at the same time). For each of these U-turn bursts, we compute the total number of prefixes involved, the visibility it reached in the BGP collector infrastructure and the duration of the burst timestamps. The heatmaps in Figure 4-7a and 4-7b show the distribution of U-turn bursts by prefix count (x-axis) and visibility (y-axis) for IPv4 and IPv6. Note that the metrics we use to evaluate an AS centrality are based on IPv4 routes, thus their relevance for IPv6 is unclear. However, our findings for IPv6 are very similar to our findings for IPv4.

Most bursts involve a small number of prefixes: Over 85% of the U-turn bursts we detect involve 5 prefixes or less. Indeed, we only find 78 events involving more than 15000 prefixes, which is the median size of route leaks in our manually collected ground truth. Nonetheless, the largest event in IPv4 involves over 720,000 prefixes, which is close to the size of the full IPv4 routing table. In IPv6, the largest event involves almost 60,000 prefixes, which is close to 20% of the IPv6 routing table.

Most bursts have limited reach: 82% of the U-turn bursts are visible by 3 or less ASNs that peer with the collectors and only 1.2% are visible by 50% or more peers. Even though most of these U-turn bursts have limited reach, these episodes can still involve many prefixes and potentially have great impact in the AS neighborhood where the events happen. Indeed, there are about 200 bursts with very low visibility (1-5 peers) that involve more than a thousand prefixes, some involving over 100,000 prefixes in total. Indeed, if a route leak happens and is contained in an AS neighborhood where only a few ASNs peer with the collectors, its visibility will be very limited in the infrastructure. Of the more than 1,400 peers to the collector infrastructure, only between 320-350 peers send the equivalent of the full routing table to collectors. These full-feeder ASNs are the ones that allow to evaluate the visibility and reachability of a BGP route.

Most bursts are short-lived: The time interval between the first and the last announcement in a U-turn burst is usually very short. 76% of events last less than 5 min, and only 5% last more than 20 min. Assuming many of these bursts are unintentional misconfigutations or mistakes, it is expected that most get quickly corrected once they are noticed. However, the longest bursts last many days and even weeks. Indeed, the longest burst lasted 230 days. This indicates that some of these announcements are not mistakes. The presence of U-turns in their AS path translates complex relationships between ASes of different size, both in terms of customer and share of address space that they route.

Frequent U-turn bottom ASNs: In the 15 months of announcements, about 5,000 (out of the more than 65,000 ASNs we found in BGP paths) are at the bottom of U-turn paths. In the case of IPv4, 500 ASNs (15%) are responsible for 80% of these events. In the case of IPv6, the count of events by ASNs is even more skewed with one ASN responsible for 67% of ASNs. Figure 4-8 has the cumulative sum of events by ASes ordered in decreasing count of events where they are the bottom of the U-turn.

Even though BGP announcements with U-turn paths are not usual (see 4.1), low counts of U-turn announcements are captured almost everyday by BGP collectors. And most of these announcements come from a set of ASes that frequently appear at the bottom of U-turns in path. Indeed, 92% of U-turn bursts come from ASNs that are detected over 500 U-turn bursts in the 15 months of data. Many of these bursts are probably not misconfigurations or mistakes but they translate complex relationships between ASes in the path. However, U-turn bursts involving many prefixes and reaching higher levels of visibility are unusual. To infer route leak events, in the next section, we consider U-turn bursts that involve more than 100 prefixes and reach at least 10% visibility.

4.4 From Bursts to Events

In this section, we identify route leak events based on the U-turn bursts we detected in Section 4.3. Starting from U-turn bursts that involve more than 100 prefixes and reach at least 10% visibility in the BGP collector's infrastructure, we first remove U-turn bursts from ASNs with a high number of such bursts and then fetch all BGP update messages (announcements and withdrawals) related to the U-turn announcements. This leaves us with $3,015$ U-turn bursts for consideration.

Removal of repetitive and persistent bottom ASes: In this work, we are particularly interested in route leak events that are the result of one-off misconfiguration events, just like the route leak events studied in Section 3. In order to isolate such events, we first remove repetitive, persistent, instances of U-turn bursts in our dataset. In the $3,015$ U-turn bursts, we find that 8 ASes are responsible for 2/3 of the bursts. Even more, one AS is at the bottom of $1,182$ (37%) of the U-turn bursts. Manual investigation of the U-turn bursts related to these 8 ASes suggests that they are likely due to persistent routing misconfigurations spanning for weeks and months that are out of scope of our study. Therefore, in the following, we omit these bursts from our analysis. We filter out all bursts coming from ASNs that are at the bottom in the U-turn more than 10 times in 15 months of data, allowing up to 5 events in one day and no consecutive days with events, leaving 372 U-turn bursts as potential route leaks of interest.

Determining the end of route leak events: The next step to confirm that the selected U-turn bursts correspond to route leaks is to find that the BGP activity related to the bursts induces a reaction in the routing system and that we find an end for the event. With that goal in mind, for each of the 372 potential U-turn route leaks, we fetch all the BGP updates (announcements and withdrawals) that relate to the event by checking whether they correspond to active (prefix, peer) pairs found in the U-turn burst.

	min.	10^{th}percentile	1^{st} quartile	median	3^{rd} quartile	90^{th}percentile	max.
Max. active prefixes	28	105	169	347	750	2,738	707,735
Max. active peers (visibility %)	19 (6%)	44 (13%)	78 (23%)	190 (56%)	266 (78%)	292 (86%)	337(99%)
Max. (prefix,pair) count	152	833	2,204	6,723.5	21,132	48,491	713,523
Duration (sec)	10	129	347.5	1396.34	4458.5	12,519	74,676
Time to max. activity (sec)	0.21	19.64	73.39	362.93	1,197.49	3,768	35,647
Time to 1^{st} correcting update (sec)	0	0	0	0	1.0	7.0	486.0
Total U-turn announcements	310	3,203	8,796	34,857	127,520	379,099	4,334,757
Total BGP correcting updates	327	15,079	39,448	118,419	357,260	1,633,924	48,103,529
Correcting/U-turn announcements	0.24	0.63	1.12	2.19	8.7	60.68	999.3

Table 4.2: Summary statistics of selected route leak metrics. For each metric, we show the minimum, median and maximum, value as well as the 1^{st} and 3^{rd} quartiles and 10^{th} and 90^{th} percentiles.

Based on the study of BGP activity of the ground-truth route leak described in § 3.3, we determine the end of a route leak when the number of active (prefix, peer) pairs gets below 10% of the maximum active (prefix, peer) count of the route leak event. We further require that the (prefix, peer) count does not increase above this threshold in the following 12 hours. Effectively, our approach requires that most of the paths with U-turns that got into the peers' routing tables got removed (either by a withdrawal or by an announcement with a path without the U-turn). Using this definition, we are left with 275 out of 372 potential route leaks (74%). Of the 97 U-turn bursts for which a significant number of U-turn announcements are not modified by later BGP messages, 55 correspond to bursts from ASes with multiple bursts (4-10 bursts). This indicates again that certain ASes, given their routing relationships, are frequently in the bottom of U-turns in BGP paths. The next section studies the 275 route leak events started by U-turn bursts and with a defined end.

5 Drilling into Detected Leaks

Table 4.2 provides summary statistics of representative characteristics of the 275 route leaks we detect, started by U-turn announcements and ended with 90% of leaked prefixes are no longer in the direct peers' routing tables.

5.1 Prefixes and Visibility

Most route leaks involve less than 1,000 prefixes: Figure 4-9 shows a scatter plot of the detected route leaks by the maximum active prefix count and visibility (direct peer count). Even though we find some very large route leaks involving more than 100,000 prefixes, most route leaks have less than 1000 active prefixes at any point in time. This could be the result of mechanisms that either limit the numbers of prefix changes from a peer (the BGP maximum prefix feature [153]) or limit route oscillations (route flap damping [154]). As described in § 3.3, there is a constant route flap between U-turn announcements and correcting updates that might get detected by route flap dampening mechanisms, temporarily limiting the spread of a route leak.

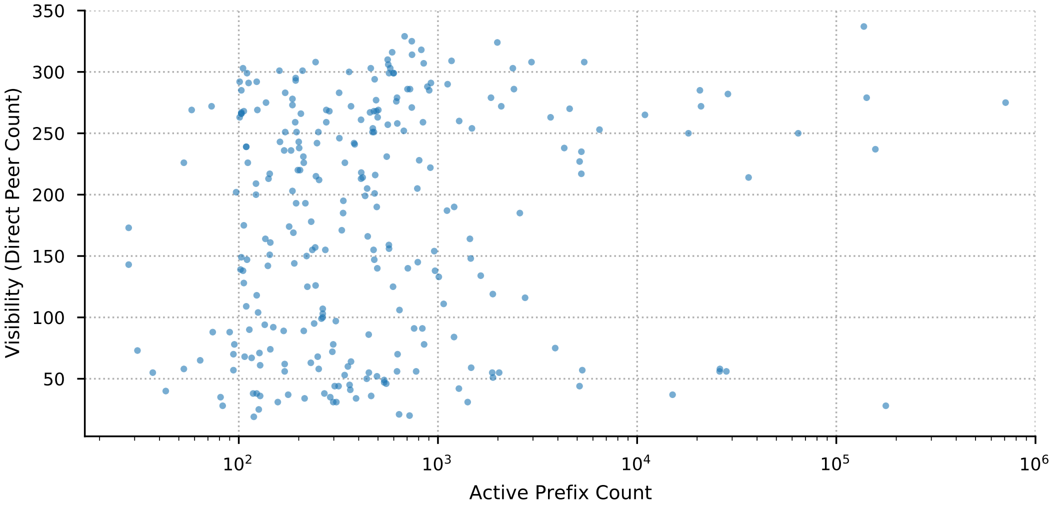

Figure 4-9: Scatter plot of detected route leak events by maximum active prefixes and visibility (direct peer count).

Visibility: The route leaks we detect have varying levels of visibility. Although we started from bursts of U-turn announcements that overall had at least 10% visibility (32-35 direct peers), when incorporating BGP updates that change the state of the involved prefixes in the peers' routing tables, some of these route leaks have lower visibility, as not all the impacted peers have leaked prefixes in their routing table at the same time. Still most detected route leaks reach more than 50% visibility.

Diverse impact in peers: Although most detected route leaks reach over 50% visibility, the actual number of prefixes that get into the peers' routing tables varies significantly across peers, with usually only a few of them getting most of the prefixes. Figure 4-10a shows the cumulative distribution of direct peers by the share of leaked prefixes that get into their routing table during all detected route leaks. 50% of the peers see less than 5% of all prefixes of a route leak. Only 10% of the peers have 80% or more of leaked prefixes in their routing table.

5.2 Timing and Reactions

Quick reaction but long duration: In 70% of the route leaks we detect, the first correcting update takes less than a second to appear and in 93% of them it takes less than 10 seconds. The cumulative distribution of detected route leaks by the time to the first correcting update is shown in Figure 4-10b (green line). The longest time we measure to the first correcting update is slightly over 8 minutes. Even though there is a fast reaction of the routing system, route leaks keep lingering and the total duration of events is much longer. Figure 4-10b shows the distance between the cumulative distribution of route leaks by duration (red line) and the cumulative distribution of the time to the first correcting update (green line). Most route leaks

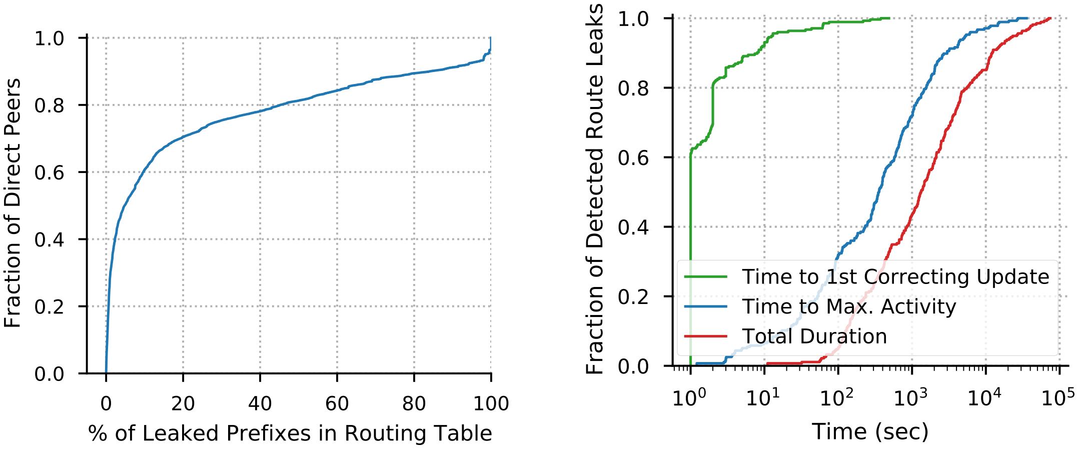

(a) CDF: Percentage of leaked prefixes seen by individual direct peers across all identified route leaks.

(b) CDFs: Time to the first correcting update after a route leak starts, time to maximum activity, and total duration of the event.

Figure 4-10: Characteristics of detected route leaks.

last a few minutes (50% lasts 6 minutes or less). However, 51 (18%) route leaks last 2 or more hours and the longest route leak we find lasts almost 21 hours (20h44m).

Maximum activity early in route leaks: In almost 60% of detected route leaks, the maximum activity in terms of leaked prefixes in peers' routing table (*i.e.,* (prefix,pair)) happens in the first half of the route leak. In Figure 4-10b, the blue line shows the cumulative distribution of detected route leaks by the time since the start to the maximum activity. We note that the blue line is shifted from the red line, highlighting that while route leak events tend to reach maximum activity rather early, most take significantly longer time to resolve.

High volume of correcting updates: In most detected route leaks, the number of correcting updates (removing U-turn path from peers' routing tables) is at least twice as much as all U-turn announcements forwarded by direct peers. In 66 cases (24%), the correcting updates are 10 times or more than U-turn announcements.

5.3 U-turn ASNs and Victims

Table 4.3 lists the 10 U-turn bottom ASNs contributing to the largest number of route leaks we detect. All these ASes have very small customer cones (3 or less) and 6 of them are stub ASes. Looking at the individual events from these ASes, we find that 60% of them (31 out of 52) have U-turn updates that are detected as route leaks by a rule based approach based on Tier-1 ASes relationships [60].

Table 4.4 lists the 10 origin ASes whose routes were affected by route leaks. In the list of top ten victims of the detected route leaks, we find 3 major CDNs and cloud

U-turn AS	Organization	Route leaks	%
AS724	DoD NIC	7	2.5%
AS37697	Webmasters	6	2.1%
AS53053	Bom Tempo Informática	6	2.1%
AS6643	Jive Communications Inc.	5	1.7%
AS52022	Klimenko A.A. PE	5	1.7%
AS138915	Kaopu Cloud HK	5	1.7%
AS61642	NEXNETT Brasil Telecom	5	1.7%
AS38823	PINC AS	4	1.4%
AS262867	Hoje Sistemas De Informatica	4	1.4%
AS50048	NewReal AS	4	1.4%

Table 4.3: Top 10 U-turn bottom ASes by count of route leaks.

Victim AS	Organization	Route leaks	%
AS20940	Akamai Technologies	82	29.2%
AS13335	Cloudflare	56	20%
AS16625	Akamai Technologies	52	18.5%
AS16509	Amazon	47	16.7%
AS9829	Natl. Internet Backbone India	46	16.4%
AS23969	TOT Public Company Limited	41	14.6%
AS9009	M247	41	14.6%
AS9583	Sify Technologies Limited	40	14.2%
AS45528	Tikona Digital Networks	40	14.2%
AS34164	Akamai Technologies	36	12.8%

Table 4.4: Top 10 victim ASes, origin ASes of leaked prefixes, by count of route leaks.

providers, accounting for 5 to the top 10 victims. These providers are large, with 4 of them originating over 3,000 prefixes each, and all of them are distributed, i.e they announce prefixes in multiple geographic locations to a diverse set of peers, including open peering via Route Servers at many Internet Exchange Points (IXPs) [155]. The announcement of their routes to a large number of networks makes them particularly prone to be exposed to U-turn route leak events.

6 Discussion

Regardless of their sometimes massive connectivity impact, route leaks have so far received comparably little systematic analysis. Our work presents a solid first step towards automating the detection of a critical type of route leaks: U-turn route leaks. In these route leaks, a small AS at the edge of the Internet becomes a transit provider for one or many much larger ASes.

Illuminating route leaks: A common assumption is that BGP route leaks are sudden and short events (e.g., [156]). Our insights paint a more complex picture: route leaks can linger for many hours and often result in continued "stress" in the routing system, with constant route oscillations re-introducing U-turn paths and again correcting updates. Gaining a deeper understanding of these dynamics will help in developing and refining approaches to combat route leaks and limit their spread (*e.g.,* by evaluating which practices would have the best outcomes). Our method provides a tool to build a dataset of route leak events and study the complex interactions visible through BGP updates. For these events, we analyze their time dynamics, the activity they result in BGP, their visibility, the frequent ASes at the bottom of U-turns and the top victims. We plan to make both our code and our data publicly available.

Lightweight detection: Our U-turn path detection approach relies on publicly available data, requires little state, and can be applied out-of-the-box to readily identify incoming announcements that are likely related to route leak events. This approach could aid network operators by providing early detection. Current mechanisms to limit the spread and impact of route leaks, e.g., by limiting the maximum number of prefixes a network accepts from a peer [153] or BGP route flap dampening [154], could be enhanced by leveraging our metric to treat or weigh announcements with U-turns differently from regular BGP updates.

Limitations and opportunities: Our approach focuses on the detection of U-turn paths, limiting our visibility to route leaks that result in traffic boomeranging between the edge and the core. Smaller route leak events might not exhibit this behavior and will thus not be detected by our approach. As we show in Section § 3, only a subset of paths caused by a route leak exhibit a U-turn. Our approach will hence only be able to detect a subset of the affected paths in most events. As our approach can readily identify the bottom AS on U-turn paths, an additional step in future work could scrutinize other, non-U-turn paths that include such AS to extend visibility into additional affected paths. In addition, by refining our thresholds for detecting "deep" U-turns, we might be able to extend our detection coverage, albeit at the risk of increasing false positives.

7 Conclusion

This work provides a new method to detect and monitor over time harmful route leaks, capable of impacting the availability of impacted Internet services. The detection is based on inherent path characteristics denoting a violation of industry structure and independent of individual network relationship. As such these events shift core traffic through smaller links, creating bottlenecks and impairing availability. This methods finds almost 300 such route leaks in 15 months of BGP data, that are further studied to characterize the events and the overall prevalence of this type of route leaks.

In addition, leveraging the fine-grained level of BGP updates from these events, this work also examines in detail the complex routing dynamics that arise. It finds that specific BGP configurations impact the spread of route leaks. Studying the owners of IP address blocks impacted by these events, it finds that large Content Distribution Networks (CDNs), such as Akamai and Amazon, are frequent victims of these misconfigurations. Even these large networks for which availability is critical to properly provide their services, are not shielded from the consequences of BGP design flaws.

Chapter 5

Empirical Analysis of Defenses

Previous chapters revealed the extend of malicious activity and the spread of misconfigurations in BGP. This chapter switches the focus to the defense side, to better understand the usage and effectiveness of defenses in BGP. It centers on the Resource Public Key Infrastructure (RPKI) framework standardized by the IETF in 2012 to support the validation of routing information in BGP. Using the RPKI, networks that have been allocated IP address space by one of the RIRs, can issue Route Origin Assertions (ROAs), which can then be fetched by any AS to validate IP prefix and origin AS in BGP announcements.

As described in Chapter 2, there are many disagreements with respect to which direction to follow to improve BGP security. In particular, network operators and academics have expressed skepticism about the real benefit of RPKI and ROAs, given that RIRs manage the support infrastructure and that it only protects against basic kinds of attacks and misconfigurations. Based on empirical data, this work measures the adoption of the RPKI scheme to validate information in BGP, the operational practices and challenges related to it, and importantly, the benefits of this scheme.

Recent research shows that many networks are starting to take first steps to secure their IP prefixes in BGP by using the RPKI framework [83, 157]. The RPKI is the most deployed framework to secure BGP. However, the actual benefit of this practice depends on whether network providers modify their operational practice and configure their routers to use ROA records to validate prefix announcements in BGP. Only recently has there been anecdotal evidence of this practice [92, 158–162]. This study measures the prevalence and benefit of networks validating information in BGP using RPKI.

Using the publicly available and passively collected BGP and RPKI data, this study first proposes a method to track RPKI validation behavior of networks sharing their BGP data with the public collectors. It measures the changes in the amount of invalid BGP information forwarded by networks using RPKI to validate routing information they receive in BGP. Using this method, it finds that, after struggling for years to gain traction, the adoption RPKI validation of routing information finally took off in 2019 and 2020.

This method is able to capture when networks start using the RPKI ROAs to make routing decisions and is able to differentiate the outcome of different operational practices linked its use. Leveraging that insight, this study also examines technical and non-technical barriers that influence the operational practices in the adoption of RPKI for validating routing information in the routing route selection process. This method thus increases the incentive of network operators to adopt this practice by making more transparent which networks are adopting it and the outcome of their operational decisions (and challenges) in its implementation.

Then, this study measures the impact the adoption of RPKI validation has on the overall spread of incorrect routing information by evaluating how it limits the spread of invalid routing information. It finds that even when less than 10% of networks have adopted this practice, the spread of invalid and potentially illicit announcements in BGP is reduced by 10-15%. As many network operators and researchers have expressed that benefits from RPKI require almost full adoption, this result based on real-world evidence is encouraging for the research, standardization, and operator communities. It bodes well for increasing routing security in the Internet and advocates for RPKI and ROAs adoption, showing its benefits even with limited adoption.

This study is entirely based on publicly available datasets. An original version appeared in [163], and updated results on networks using RPKI to validate information in BGP is made public periodically.[1] .

1 Background

The inter-domain routing system of the Internet continues to suffer from major routing incidents, including accidental route leaks causing widespread disruptions [164], and intentional prefix hijacks for malicious purposes [11, 14, 98]. At the heart of the problem lies BGP's lack of mechanisms for route validation described in Chapter 1 Section 2.

The RPKI [37] represents one of the most recent attempts to increase BGP security, providing networks in the Internet with a trustworthy database of Route Origin Authorization (ROAs) that maps IP prefixes to the Autonomous System (AS) number that is authorized to originate them in BGP. The RPKI is backed by strong cryptography, with the Regional Internet Registries (RIRs) serving as trust anchors. Networks can leverage this data to validate that the IP prefix and origin AS in incoming BGP announcements. The RPKI framework and the issuance and usage of ROAs to validate information in BGP is thoroughly described in Chapter 2 Section 3.1.2.

Recent research shows an encouraging trend of both increasing global registration of prefixes in the RPKI (over 30% of routed prefixes are covered by a ROA registered in the RPKI as of August 2021), as well as increasing data quality of actual RPKI records [157]. The RPKI has thus the potential to finally provide a universally trusted database mapping IP prefixes to origin ASes, a major building block to greatly improve routing security.

The increasing registration of prefixes in the RPKI only represents a first step towards securing BGP. The eventual benefit of RPKI registration depends on whether

[1] Auxiliary material can be found at `https://github.com/ctestart/BGP-RPKI-ROA`.

the networks of the Internet enforce the RPKI's contents, *i.e.,* validate routing information in BGP using ROAs and drop invalid announcements, hence not propagating them to their neighbor ASes. Recently, AT&T, a major transit ISP, publicly announced that they started dropping BGP announcements that are invalid as per the RPKI [92], suggesting increasing acceptance and trust by major transit providers in the RPKI. However, besides such anecdotal evidence, we know little about current levels of RPKI *enforcement* in the Internet and, as of today, have no way to assess the resulting benefits of RPKI registration.

To tackle these questions, we empirically study to what degree networks in the Internet validate BGP announcements based on RPKI data and show to what extent the issuance of ROAs, also called registration in the RPKI, benefits networks in situations in which RPKI is needed the most: instances of conflicting BGP announcements in the global routing table, such as those caused by misconfiguration and prefix hijacking.[2]

1.1 Related Work

The IETF has devoted substantial efforts over the last years to develop and document in detail the RPKI framework and ROAs [37, 74, 76, 165–169]. Recently, the research community started to measure RPKI deployment in the Internet. Chung *et al.* provide both an accessible overview of today's RPKI deployment and an extensive study of RPKI registration and usage patterns. They find increasing registration of prefixes and networks in the RPKI, and overall higher data quality of RPKI records, resulting in lower numbers of RPKI-invalid prefixes caused by misconfiguration by the respective operators [157]. Iamartino *et al.* had previously measured problems with RPKI registered ROAs and the potential impact that validation and filtering (*i.e.,* dropping) of RPKI-invalid announcements could have in production [170].

To the best of our knowledge, only two previous academic studies, using two different methods, touched upon the adoption of RPKI-invalid filtering, *i.e.,* the integration of IP prefix and AS origin validation using ROAs in BGP route selection process, finding only negligible RPKI filtering in 2016 and 2017. Gilad *et al.* analyze a month of BGP RIB dumps from 44 ASes [161]. Their passive approach uses all the ASes but the last hop in the AS path of RPKI-valid and -invalid announcements to identify ASes filtering invalid announcements. They find that, in July 2016, only 3 of the top 100 ASes (by customer cone size) were enforcing RPKI-invalid filtering. Reuter *et al.* instead, actively advertise RPKI-valid and -invalid prefixes of address space under their control [162]. They infer which ASes filter RPKI-invalid announcements based on the propagation path of their announcements, finding only 3 ASes filtering in 2017.

Measuring RPKI filtering also caught attention from the operator community. Cartwright-Cox uses active measurements to infer filtering based on presence or absence of ICMP responses from probed IP addresses in RPKI-valid and -invalid pre-

[2] Chapter 1 Section 2 explain BGP hijacks and misconfigurations and the harms they inflict on Internet users.

fixes [171]. In April 2020, Cloudflare deployed a web-based test for Internet users to learn if their Internet Service Provider (ISP) is validating information in BGP using the RPKI and also tracks the adoption of this practice [172].

Our study complements and extends prior work: our passive method to detect filtering of RPKI-invalid announcement focuses on networks that provide a direct and full feed to BGP collectors, which allows for definitive and detailed assessment of RPKI filtering of these networks. Our study is longitudinal, revealing a strong uptake in RPKI filtering deployment in recent years. Most importantly, however, we present a first-of-its-kind assessment of RPKI enforcement and its actual impact and benefit in situations in which the RPKI is needed the most: instances of conflicting prefix announcements in the global routing table.

2 Datasets and Preprocessing

2.1 RPKI and BGP Datasets

To study the visibility of RPKI-valid and RPKI-invalid announcements in the global routing table, we leverage the following datasets.

Longitudinal BGP dataset: To study long-term trends of RPKI filtering behavior, we download and process—using CAIDA BGPStream [173]—snapshots of the routing tables (RIB dumps) of all RouteViews and RIPE RIS collectors on the first day of each month[3] from April 1, 2017 until January 22, 2020.

Fine-grained BGP dataset: To assess the visibility of RPKI-invalid announcements in detail, we process all the BGP updates generated over the month of September 2019 by RouteViews and RIPE RIS collector peers' and we compute 5-minute snapshots of their routing tables using CAIDA BGPStream [173].

RPKI (ROA) data: We take daily snapshots of validated Route Origin Authorizations (ROAs) for every day in September 2019, made available through the RIPE NCC RPKI validator [174]. For longitudinal analysis, we instead leverage the historical dataset of validated ROAs made publicly available by Chung *et al.* [157], selecting snapshots that align with our BGP dataset. A validated ROA consists of a prefix and the AS number authorized to originate that prefix in BGP according to cryptographically signed records in the RPKI. ROAs may include a *maxLen* attribute specifying up to which prefix length the de-aggregation of the ROA prefix is to be considered valid.

2.2 Preprocessing

From BGP snapshots to prefix-origin pairs: As a first step, we remove *bogon* prefixes from our BGP dataset, these include IETF reserved address space, and por-

[3] Or the closest day for which validated historical RPKI data is available.

tions of address space not allocated by IANA to RIRs [137]. We further remove any IPv4 prefixes more specific than /24 or less specific than /8 (more specific than /64 or less specific than /8 for IPv6). Then we extract, for each BGP snapshot (both RIB dumps and those we derive from updates), all visible prefixes together with the advertised origin AS, obtaining *prefix-origin pairs*.[4] For each prefix-origin pair, we save the set of *feeders*—that is, ASes that directly peer with any of the RouteViews and RIPE RIS route collectors—that have a route to the given prefix-origin in their routing table. In the following, we will leverage the set of feeders to assess filtering and to estimate visibility of prefix-origin pairs in the global routing table.

Tagging prefix-origin pairs: We next tag each individual prefix-origin pair in our dataset with its corresponding RPKI state. For each prefix-origin pair, we find the closest snapshot available of validated ROAs and tag the prefix-origin pair with one of the following states: *(i) unknown*: the prefix is not covered by any prefix of validated ROAs in the RPKI; *(ii) valid*: the prefix is covered by a validated ROA, the AS number in BGP matches the one in the ROA, and the prefix length in BGP is at most the maxLen attribute of the ROA; *(iii) invalid ASN:* the prefix is covered by a validated ROA, but the origin AS in BGP does not match the origin AS in any ROA covering the prefix; *(iv) invalid length:* the prefix is covered by a validated ROA, the origin AS in BGP matches the origin AS in the ROA, but the prefix length in BGP is longer than the maxLen attribute, *i.e.,* the prefix is more specific than what is allowed as per the ROA.

3 To Filter or not to Filter: Longitudinal Study

In this section, we provide a macroscopic perspective on RPKI filtering deployment in today's Internet. In particular, we study to which extent some of the transit networks in the Internet do filter BGP announcements with invalid RPKI state and how this filtering behavior evolved over time.

3.1 Detecting Filtering

While there is no practical way to comprehensively study filtering behavior of all networks, we introduce a method to infer RPKI filtering with high confidence for a small but relevant set of ASes. At a high-level, our method is made of two steps: *(i)* we select *full-feeder* ASes, *i.e.,* ASes that share with BGP collectors a number of routes (and thus prefix-origin pairs) comparable to what is globally visible in BGP—in other words, they tend to share the vast majority of, if not all, their preferred routes; *(ii)* we leverage our set of RPKI-invalid prefix-origin pairs to look for significant presence/absence of them in the data full-feeders share.

The essence of this approach is to look for statistically significant absence of RPKI-invalid prefix-origin pairs: *e.g.,* the absence of a single invalid pair in the routes

[4] Note that a prefix can have multiple origins in the global routing table, in this case we extract multiple prefix-origin pairs.

shared by a full-feeder is not a strong indication of RPKI-based filtering; similarly, the absence of a large number of invalid pairs in a shared routing table that is already missing many other valid routes (*i.e.,* from a *partial-feeder*) is not a strong indication of RPKI-based filtering either. The combination of the two factors instead, provides a high degree of confidence. In § 3.3, we validate our method for a few ASes that have publicly stated when they started applying RPKI-based filtering. In detail, we operate as follows.

***(i)* Selecting full-feeders:** We consider a collector's peer a *full-feeder* if the number of prefix-origin pairs shared by that AS is at least 75% of the maximum prefix-origin pair count sent by all feeders. We perform our analysis for IPv4 and IPv6 independently. In Figure 5-1, the orange line shows this threshold for IPv4 (fig. 5-1a) and IPv6 (fig. 5-1b) in September 2019: out of 578 ASes peering with the collectors, we consider 276 to be full-feeders for IPv4 (232 out of 402 for IPv6). We chose 75%, since it separates recent and historical snapshots well.

***(ii)* Detecting filtering of RPKI-invalid announcements:** From the set of full-feeder ASes, we infer an AS to be filtering RPKI-invalid announcements if the number of RPKI-invalid prefix-origin pairs received from that AS is less than 20% of the maximum number of invalid records sent by all full-feeders. Here, we leave some leeway, since previous research [162] has shown that, even if ASes are filtering *most* RPKI-invalid announcements, they usually never filter *all* invalid announcements due to churn in RPKI records and selective filtering (*cf.,* § 3.3). The green dashed line in Figure 5-1, shows these thresholds for IPv4 and IPv6, we infer 21 and 18 ASes were filtering RPKI-invalids announcements in September 2019 for IPv4 and IPv6 respectively.

The representativeness of our approach is limited by the comparably small number of full-feeder ASes: 290 ASes for IPv4 and 246 ASes for IPv6 in January 2020. However, we find that these networks include many global transit providers and mid-sized networks: 187 transit and access ASes (of which 12 are Tier-1 ASes), 36 content providers, and 47 educational/non-profit networks, according to PeeringDB [175]. In total there are 36 ASes in the top 100 CAIDA AS rank and 93 in the top 1,000. This set of ASes thus provides a reasonable approximation to study macroscopic filtering trends of major networks in the Internet.

3.2 Filtering Networks: Trends and Current Status

With our method in hand, we now present a longitudinal analysis of RPKI-invalid filtering behavior. Figure 5-2 shows the evolution of the fraction of full-feeder ASes that filter RPKI-invalid announcements for IPv4 and IPv6. Both protocols follow a similar trend. Initially, slightly fewer ASes filter RPKI-invalid IPv6 announcements compared to IPv4. However, after ASes initially filtering in IPv4 only also implement filtering for IPv6, the share of ASes that we can measure filtering for IPv6 goes up and is higher than for IPv4.

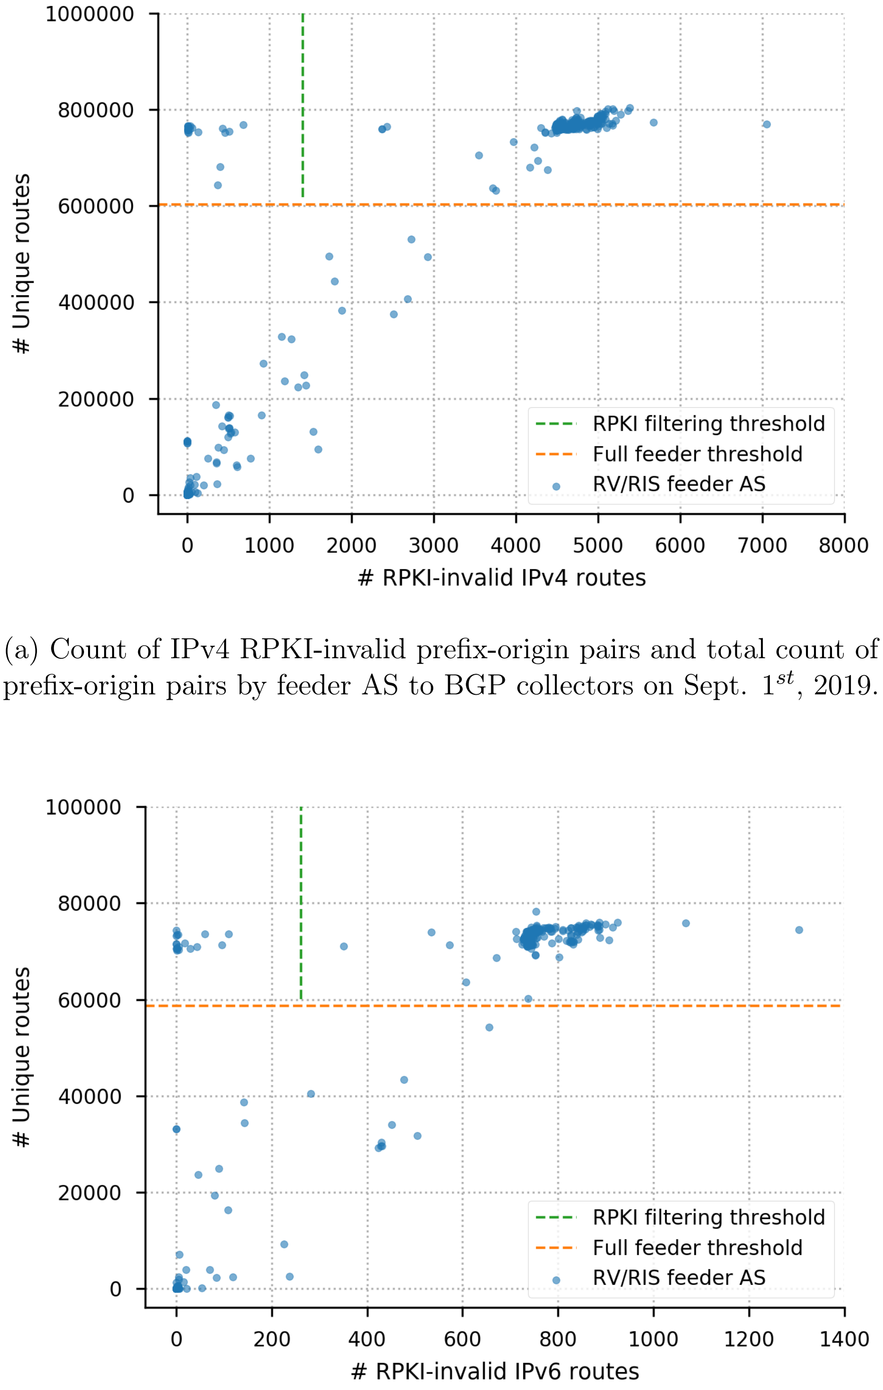

(a) Count of IPv4 RPKI-invalid prefix-origin pairs and total count of prefix-origin pairs by feeder AS to BGP collectors on Sept. 1^{st}, 2019.

(b) Count of IPv6 RPKI-invalid prefix-origin pairs and total count of prefix-origin pairs by feeder AS to BGP collectors on Sept. 1^{st}, 2019.

Figure 5-1: Detection of ASes filtering RPKI-invalid announcements: We infer full-feeder ASes in the groups on the upper left corner are filtering RPKI-invalid announcements for IPv4 and IPv6 independently.

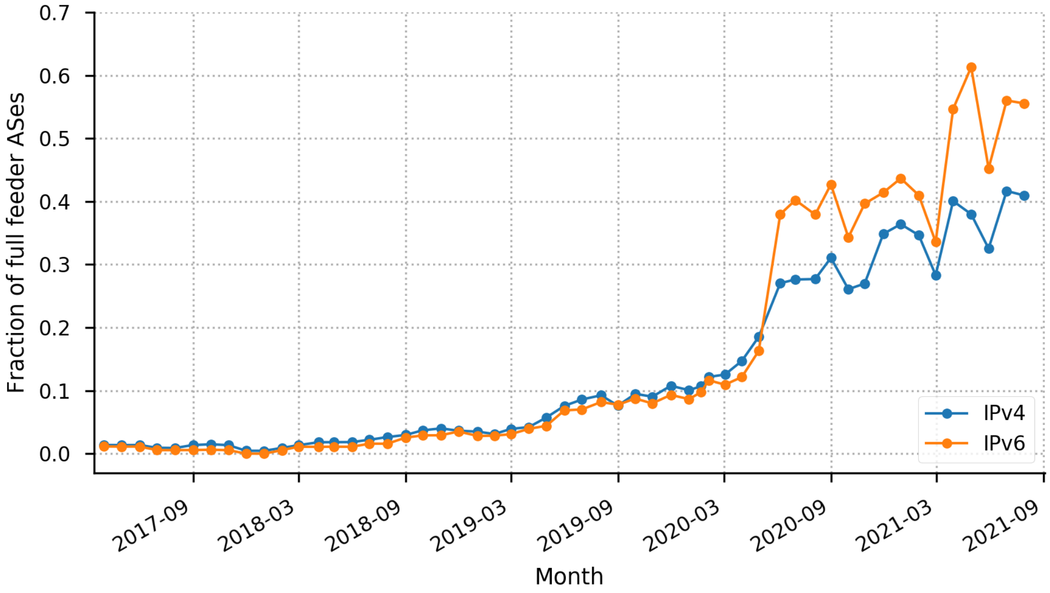

Figure 5-2: Adoption of RPKI filtering overtime: Fraction of RouteViews and RIPE RIS collector full-feeder ASes filtering RPKI-invalid announcements over time. A major increases happen first in mid 2019 and then in mid 2020.

We detect that in April 2017, less than 2% full-feeders were filtering RPKI invalid announcements: 3 out of 219 full-feeder ASes for IPv4 and 2 out of 176 for IPv6. We witness overall low levels of RPKI filtering until April 2018, when a few full-feeder ASes start to filter each month, reaching about 3% one year later in March 2019. From April until August 2019, we see a 3-fold increase in the rate of RPKI filtering adoption. In late January 2020, 11% of full-feeder ASes filter RPKI-invalid announcements in IPv4 and 10% in IPv6, 30 out of 290 and 23 out of 246 respectively. Then, we detect another major increase during 2020. In mid 2021, we measure over 30% of full-feeder ASes filtering in IPv4 and IPv6, 118 out of 335 and 140 out of 294 respectively.

The bulk of the networks filtering RPKI-invalid announcements are either transit or access network providers (31 ASes, 18% of such networks) or educational-research/non-profit networks (12 ASes, 23% of such networks). Early in the adoption of RPKI filtering, we find lower levels of filtering deployment in larger networks: only 2 of the 36 full-feeder ASes in the top 100 CAIDA AS Rank filtered invalid prefix-origins and 10 out of the 93 ASes in the top 1,000 CAIDA AS Rank filtered in early 2020. However, this changed in 2020 and in early 2021 we find 10 out of 36 top 100 ASes filtering RPKI-invalids announcements. In contrast, we only find one out of 36 content providers filtering invalid prefix-origins. RIPE, ARIN and APNIC are the regions with most full-feeder ASes, representing 39%, 20% and 12% of full-feeders ASes from these regions respectively.

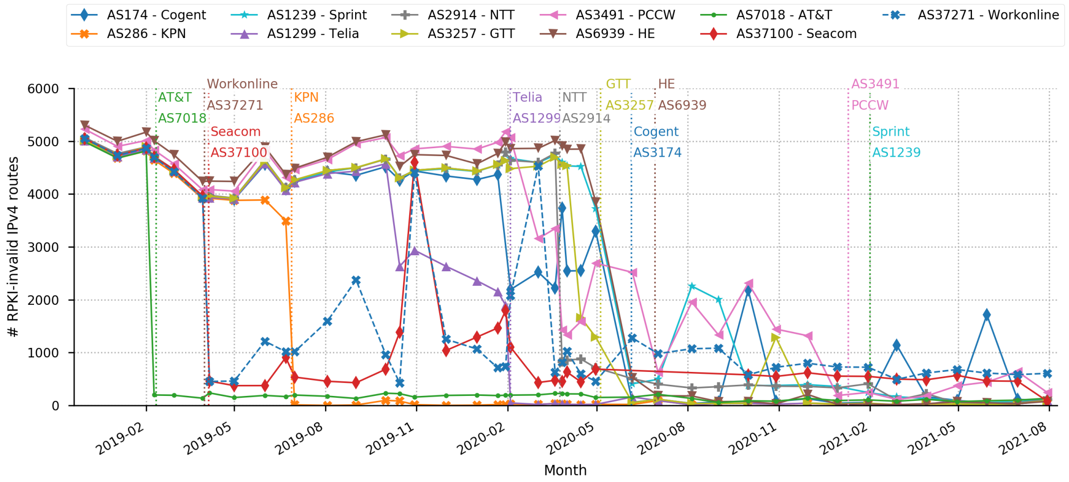

Figure 5-3: RPKI-invalid IPv4 prefix-origin pairs from networks that publicly announced RPKI filtering deployment, vertical lines show the announcement date of deployment completion (dashed) or beginning of deployment (dotted).

3.3 A closer Look at Filtering Networks

We now take look in more detail at which networks appear to be filtering and how the amount of RPKI-invalid prefix-origin pairs varies over time.

3.3.1 Comparison with public announcements of RPKI filtering

Many transit ISPs that provide direct and full BGP feeds to one of our considered collectors have publicly stated that they have deployed or are currently deploying RPKI-invalid filtering: AT&T (AS7018), KPN (AS286), Seacom (AS37100), Workonline Communications (AS37271) and Telia (AS1299), NTT (AS2914), GTT (AS3257), Congent (AS3174), Hurricane Electric (HE, AS6939), PCCW (AS3491) and Sprint (AS1239) [92, 158–160, 172]. Figure 5-3 shows the count of invalid prefix-origin pairs propagated by these ASes from December 2018 until mid 2021, annotated with their public announcement date of filtering implementation.

In our data, we see about 5,000 invalid prefix-origins from all networks in early 2019. In mid February 2019, AT&T publicly stated that they started filtering RPKI-invalid route announcements and afterwards we only detect a few hundred invalid prefix-origins sent to collectors by AT&T. In early April 2019, two major African ISPs, Workonline Communication and Seacom, announced completion of deployment of RPKI filtering, after which we observe only several hundred invalid prefix-origins from these two networks. However, these ASes have encountered operational issues when deploying RPKI filtering and have (partially) stopped filtering for some periods of time, see intermittent upticks [176].

In late June 2019, KPN announced completion of deployment of RPKI-filtering and has only propagated a few dozen invalid prefix-origins to collectors since. Finally, in mid September 2019, Telia announced that it began to deploy RPKI filtering. Shortly after their announcement, we detect a continual decline in the number of invalid prefix-origins forwarded by Telia. Telia completed their deployment in February 2020. Then followed NTT (March 2020), GTT (May 2020), Cogent (June 2020), HE (June 2020), PCCW (January 2021) and Sprint (February 2021). Many of these network followed Telia's gradually approach to deploy RPKI-filtering and we see the number of invalid prefix-origin pairs start to decrease months before their public statement, some with many upticks in between.

3.3.2 Partial RPKI filtering

In our longitudinal study, no full-feeder network ever filters *all* RPKI-invalid announcements. Besides some expected short-term churn, *e.g.,* caused by delays when updating filtering rules, we identified 3 main reasons for persistent partial RPKI filtering:

- **Selective RPKI Trust Anchor (TA) filtering:** we find 6 networks not validating ROAs from the ARIN TA, resulting in a higher share of propagated invalid prefix-origins. Indeed, legal barriers limiting availability of ARIN ROAs have been reported [177].
- **Selective filtering depending on AS relationships:** several network operators announced to implement filtering only for routes received from peers, but not customer networks [92].
- **Operational deployment issues:** network operators reported compatibility issues with RPKI validator implementations and router software, prompting them to deploy RPKI-filtering in a subset of their border routers [176].

4 RPKI to the Rescue: Conflicting Announcements

Our findings of increasing deployment of RPKI filtering in the recent years motivate us to study the effect of filtering in more detail. We first introduce how we process our dataset to allow for analysis of visibility of individual routing events and study the overall visibility of valid/invalid prefixes. Next, we showcase several relevant real-world case studies of conflicting, and hence potentially malicious, prefix announcements. Visibility of a prefix in the global routing table translates directly into its *reachability*, and thus serves as a proxy to study the benefit of RPKI filtering in the wild. In this section, we present our findings for IPv4.

4.1 Tracking Visibility in the Global Routing Table

Aggregating prefix-origin snapshots into *timelines*: To study the visibility of RPKI-registered prefixes, we leverage our fine-grained BGP dataset, consisting of per-feeder snapshots of all prefix-origin pairs every 5 minutes in September 2019 (*cf.,* § 2.1). As a first step, we aggregate adjacent prefix-origin pairs into continuous *timelines.* We require *(i)* that the maximum deviation in visibility within each timeline is

Prefix-origin timelines	**count**	**%**
IPv4 total	883,400	100%
RPKI covered	147,870	16.7%
RPKI-valid	139,537	15.8%
RPKI-invalid ASN	4,203	0.47%
RPKI-invalid length	4,130	0.46%
IPv6 Total	91,313	100%
RPKI covered	19,173	20.1%
RPKI-valid	17,656	19.3%
RPKI-invalid ASN	362	0.40%
RPKI-invalid length	1155	1.26%

Table 5.1: Properties of prefix-origin timelines and their respective RPKI validity states (September 2019).

less than 10%, otherwise we terminate the timeline and start a new one. We express visibility of a prefix-origin pair timeline as the fraction of active feeder ASes that propagate a route to given prefix and origin AS. Secondly, *(ii)* we require consistent RPKI state (valid/invalid ASN/invalid length/unknown) for each prefix-origin timeline.[5] The resulting timelines consist of a tuple of a prefix, an origin AS, a visibility level, its RPKI state, and timestamps. We filter prefix-origin timelines with a private AS number or AS-Set as origin, and prefix-origin timelines with very low visibility, *i.e.,* seen by 3 or fewer peers, since such very low visibility prefixes are unlikely to represent actual events in the global routing table. Table 5.1 shows the properties of our resulting dataset for IPv4 (top) and IPv6 (bottom).

Overall prefix-origin visibility by RPKI state: Figure 5-4 shows CDFs of the visibility of prefix-origin timelines for IPv4 and IPv6, expressed as percentage of active feeder ASes seeing a prefix-origin. Overall, we find that RPKI-valid as well as RPKI-unknown prefix-origins (*i.e.,* prefixes not covered by validated ROAs) show similar visibility levels, with 80% of all prefix-origins seen by 80% or more of feeder ASes (see green and blue lines). RPKI invalid prefix-origins, however, show vastly different visibility: some 20% of these prefix-origins are very localized announcements (seen by less than 5% of feeder ASes, see orange and red lines), and we speculate that these cases are instances of misconfigurations, whether in BGP or RPKI records, which happen to also show up as RPKI-invalid artifacts. More importantly, we find that even invalid prefix-origins with higher visibility show distinctively lower visibility when compared to valid prefix-origins (see concentration of RPKI-invalid at around 70%, compared to over 80% for RPKI-valid in Figure 5-4a). In IPv6, there are even fewer RPKI-valid prefix-origins with low visibility compared to IPv4 (see Figure 5-

[5] For 0.37% IPv4 and 0.13% of IPv6 prefix-origin timelines, the RPKI state changed due to churn in the RPKI database caused by changes of RPKI entries during our measurement window. We remove these instances.

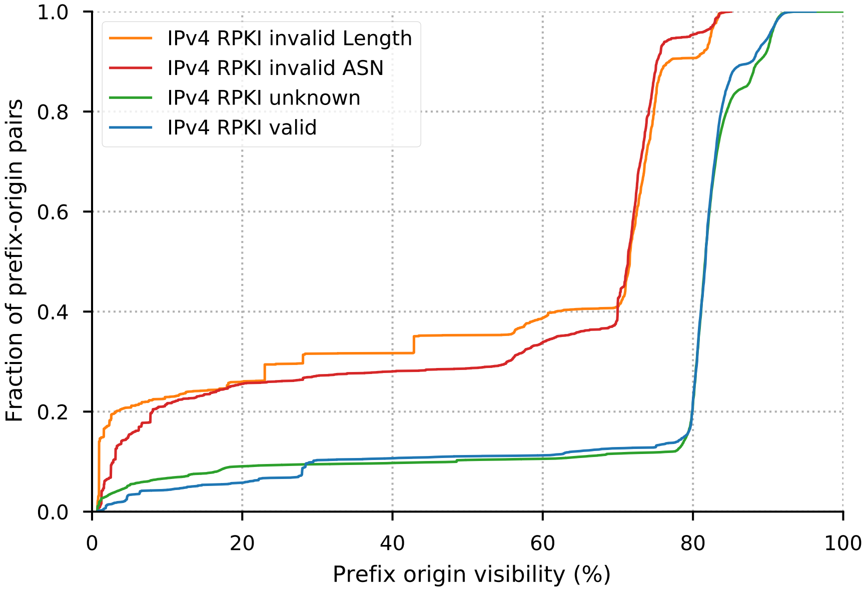

(a) CDF of IPv4 prefix-origin pairs visibility.

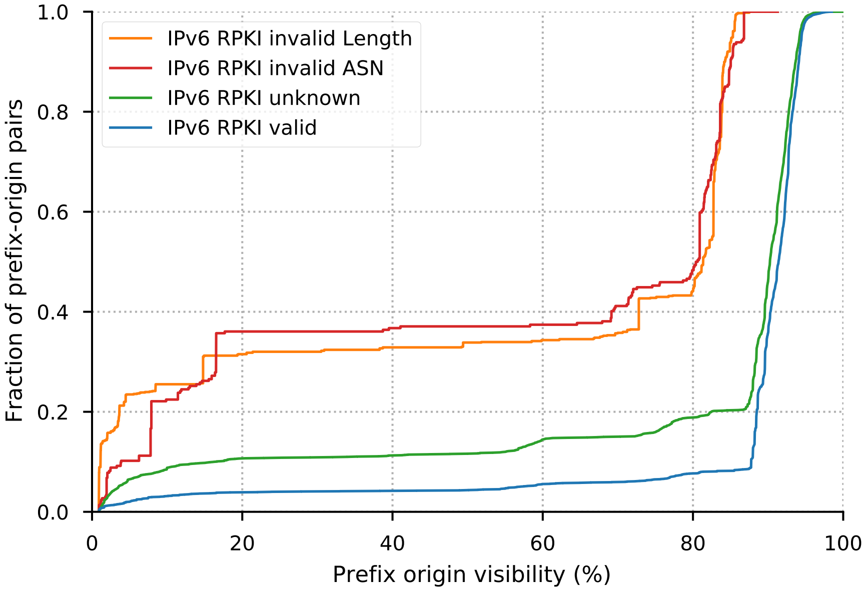

(b) CDF of IPv6 prefix-origin pairs visibility.

Figure 5-4: CDF of IPv4 prefix-origin pairs by visibility during September 2019 for different RPKI state, for IPv4 (top) and IPv6 (bottom).

4b): less than 10% IPv6 prefix-origins have less than 80%visibility compared to 20% for IPv4. This difference in prefix-origin propagation is the direct result of filtering of RPKI-invalid announcements.

4.2 Conflicting Prefix Announcement Scenarios

Next, we study *RPKI in action*, *i.e.,* we want to understand if registration in the RPKI benefits networks in cases of conflicting announcements. In particular, we cover 3 scenarios: *(i)* Multiple Origin AS (MOAS) announcements: instances where two equal prefixes are announced with two different origins, often caused by intentional or unintentional prefix hijacks; *(ii)* subMOAS announcements: instances where an announcement of a more specific prefix points to a different origin AS, also a potential prefix hijack scenario; *(iii)* same-origin subprefixes, instances where a more specific prefix is visible, points to the same origin AS as its parent, but fails RPKI validation due to max length restrictions. This scenario is what we would expect to see in the case of a path hijack, the most advanced form of prefix hijacks [55]. We note that in this work, we do not attempt to classify instances of conflicting prefix announcements into malicious activity vs. misconfigurations. Instead, we base our notion of illicit announcements on the RPKI state of the involved prefixes: if two prefix announcements are in conflict, and only one of them passes RPKI validation, in our analysis we treat the invalid one as if it is an illicit announcement (while it might also be due to incorrect/unissued ROAs). Our argument here is that, irrespective of the root cause of these conflicts, we can study the effectiveness of RPKI filtering under the same conditions that would also hold when a malicious actor injects BGP prefixes to hijack address space.

4.3 Visibility of Multiple Origin AS (MOAS) Prefixes

To study the visibility of prefixes that are concurrently originated by multiple origin ASes, we first isolate our prefix-origin timelines that show *(i)* two origin ASes for the same prefix and *(ii)* one of these prefix-origins is registered in the RPKI and valid. In total, we find about 90,000 instances of MOAS prefix-origin pairs in September 2019 for IPv4, of which some 10% are cases in which at least one prefix-origin is RPKI-valid, while others are not. Of these cases, about 20% (N= 1898) are cases of exactly 2 MOAS prefix-origin pairs one valid and the other invalid according to RPKI records. For IPv6, we find about 41,000 instances of MOAS prefix-origin pairs in September 2019, of which some 133 are cases in which at least one prefix-origin is RPKI-valid while others are not.

Figure 5-5 shows the distribution of the maximum visibility of IPv4 prefix-origin timelines during MOAS conflicts of two prefix-origin pairs, where we partition RPKI-valid and -invalid state, see positive y-dimension in Figure 5-5. Figure 5-6 shows the distribution of the maximum visibility of prefix-origin timelines during MOAS conflicts.

We see a stark difference in the visibility of prefix-origin timelines: RPKI-valid prefixes clearly dominate visibility, with more than 70% of valid prefixes in IPv4

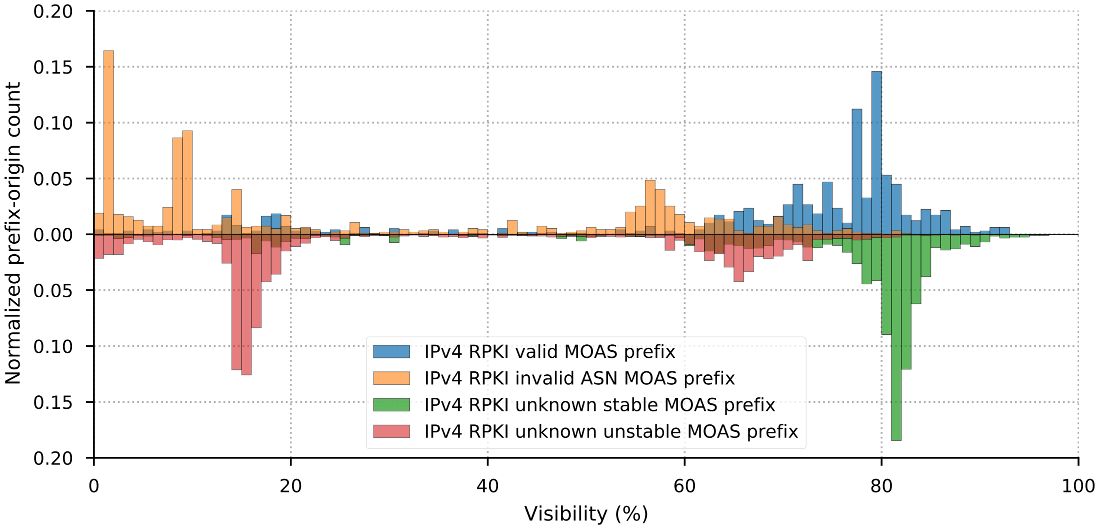

Figure 5-5: Visibility of prefix-origin pairs during MOAS conflicts: RPKI-valid and invalid ASN MOAS prefix pairs in the positive y-dimension, RPKI-unknown MOAS prefix pairs in the negative y-dimension, partitioned as stable/unstable according to total advertisement time during September 2019.

having visibility greater than 70%, and we only see few instances of RPKI-valid prefixes with low visibility (only 12% of instances with less than 30% visibility). IPv6 results follow a similar distribution, with low and high-visibility mode more distanced.

On the other hand, invalid prefix-origin timelines show distinctively lower visibility: In IPv4, some 60% have a visibility level lower than 30%. Some invalid prefixes do reach substantial visibility levels, but we do point out that even those higher-visibility invalid prefixes cluster at around ≈65%, that is, significantly lower when compared to valid prefixes, which cluster at around around ≈80%. In IPv6, only a few prefix-origin timelines reach visibility over 60%. These results are consistent with our expectations: the RPKI benefit should be significant in instances of exact MOAS conflicts, since two prefixes compete for reachability in the global routing table, and even when RPKI filtering is not enforced, some routers still give preference to RPKI-valid announcements over RPKI-invalid announcements as part of the route selection process (discarding an invalid route only if a valid one is available) [178].

To assess the potential benefit of registering a prefix in the RPKI vs. not registering it, we next compare the above studied instances of MOAS conflicts in which the concerning prefix is registered in the RPKI against vanilla cases of MOAS, in which the concerning prefix is not registered, and hence both prefix-origins are of type RPKI-unknown. Here, in the absence of RPKI information, we face the difficult problem of determining which of the conflicting announcements represents the legitimate announcement vs. the illicit one. Taking a pragmatic approach, we leverage stability of announcements as a proxy: In the case of a MOAS conflict where neither prefix-origin is registered in the RPKI, we tag the prefix-origin that was visible for a longer period of time as *stable*, and the conflicting prefix-origin that was visible for a

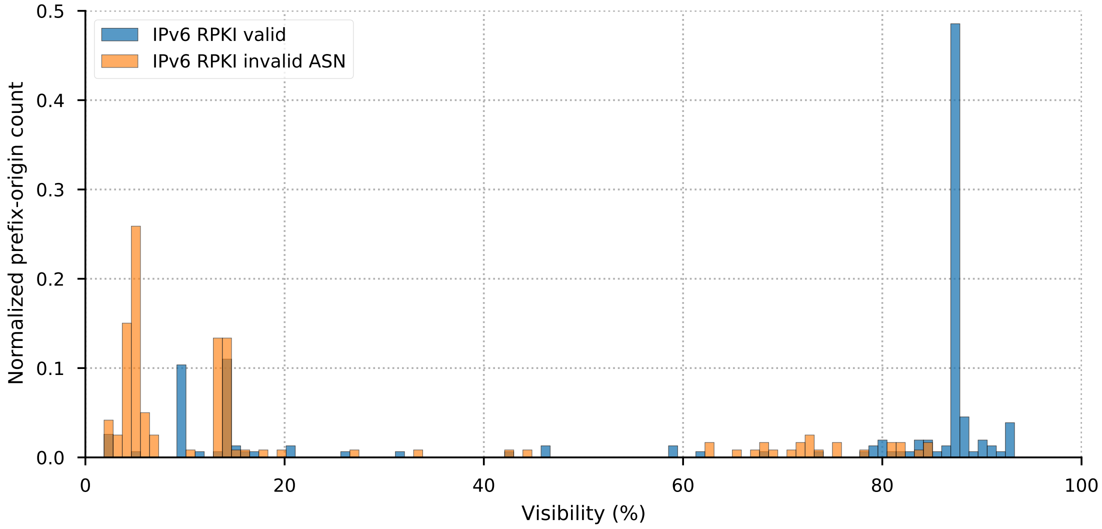

Figure 5-6: Visibility of RPKI covered IPv6 prefix-origin pairs during MOAS conflicts

shorter period of time as *unstable*. We pick only MOAS cases where the stable prefix-origin is announced for a period at least 3 times longer[6] than the unstable prefix-origin counterpart (N=6,374 MOAS events for IPv4). Unfortunately, only about 30 MOAS events show this property and thus we only preform the comparison for IPv4 prefixes. We acknowledge that our heuristic requiring stable prefix-origin be announced at least 3 times longer than the unstable prefix-origin is not a hard-and-fast rule, since there are many potential root causes for unstable announcements (*e.g.,* rewiring, address space transfers, etc.). However, it allows us to present a one-to-one comparison of RPKI vs. non-RPKI scenarios.

We plot the distribution of prefix-origin visibility of RPKI-unknown prefixes in the negative y-dimension in Figure 5-5. We find that, overall, stable prefixes show much greater visibility when a MOAS conflict happens, when compared to their conflicting unstable counterparts. However, contrasting the vanilla case (no RPKI registration, negative y-dimension) against the case in which the prefix is registered in the RPKI (positive y-dimenstion), we see a difference: unstable RPKI-unknown prefixes generally reach higher levels of visibility when compared to RPKI-invalid prefixes. This difference manifests both for very low visibility cases, where RPKI-unknown cluster at around ≈15% visibility, higher than their RPKI-invalid counterparts which cluster at ≈8%, as well as for cases of higher visibility: unstable RPKI-unknown prefixes reach visibility levels of some 70%, while RPKI-invalid cluster below 60%. Indeed, less than 14% of RPKI-invalid MOAS instances reach a visibility over 60% compared to 37% for unstable RPKI-unknown MOAS instances. RPKI registration shows a clear effect on prefix visibility when MOAS conflicts happen.

[6] We tested different thresholds, finding that the modes of the distribution do not change much.

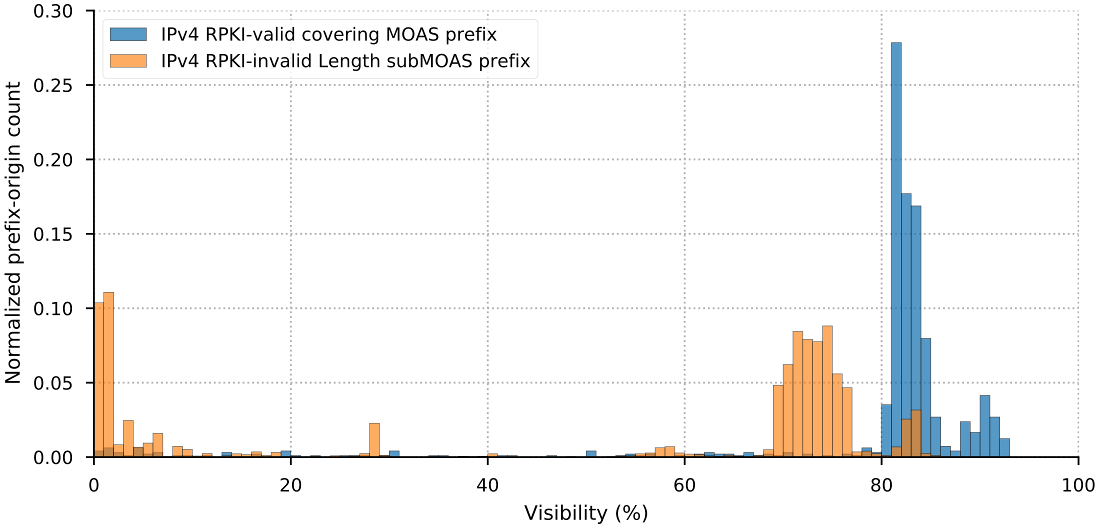

(a) Visibility of RPKI-covered prefix-origins during subMOAS conflicts.

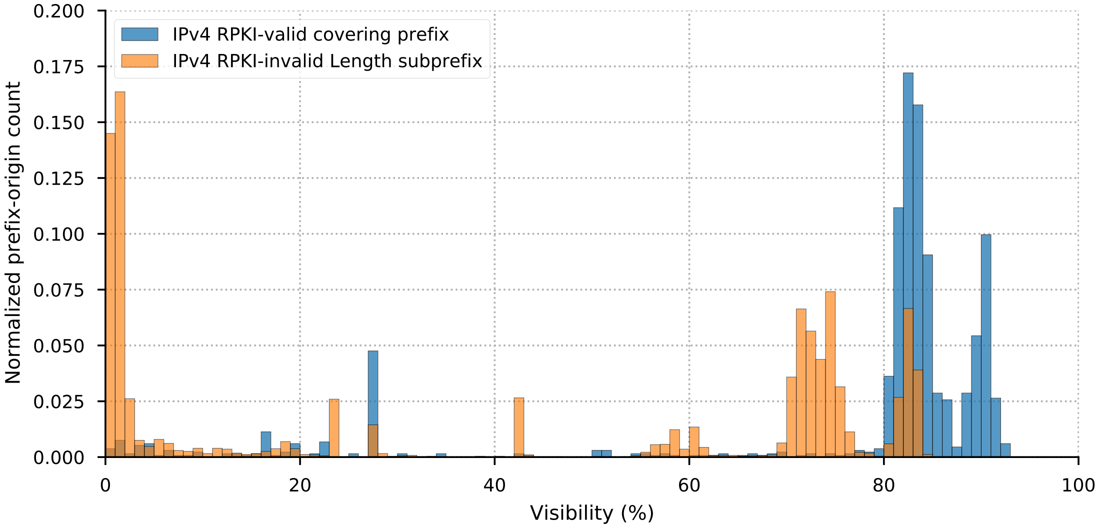

(b) Visibility of RPKI-covered prefix-origins during subprefix conflicts.

Figure 5-7: Impact of RPKI registration in subMOAS and subprefix conflicts for IPv4.

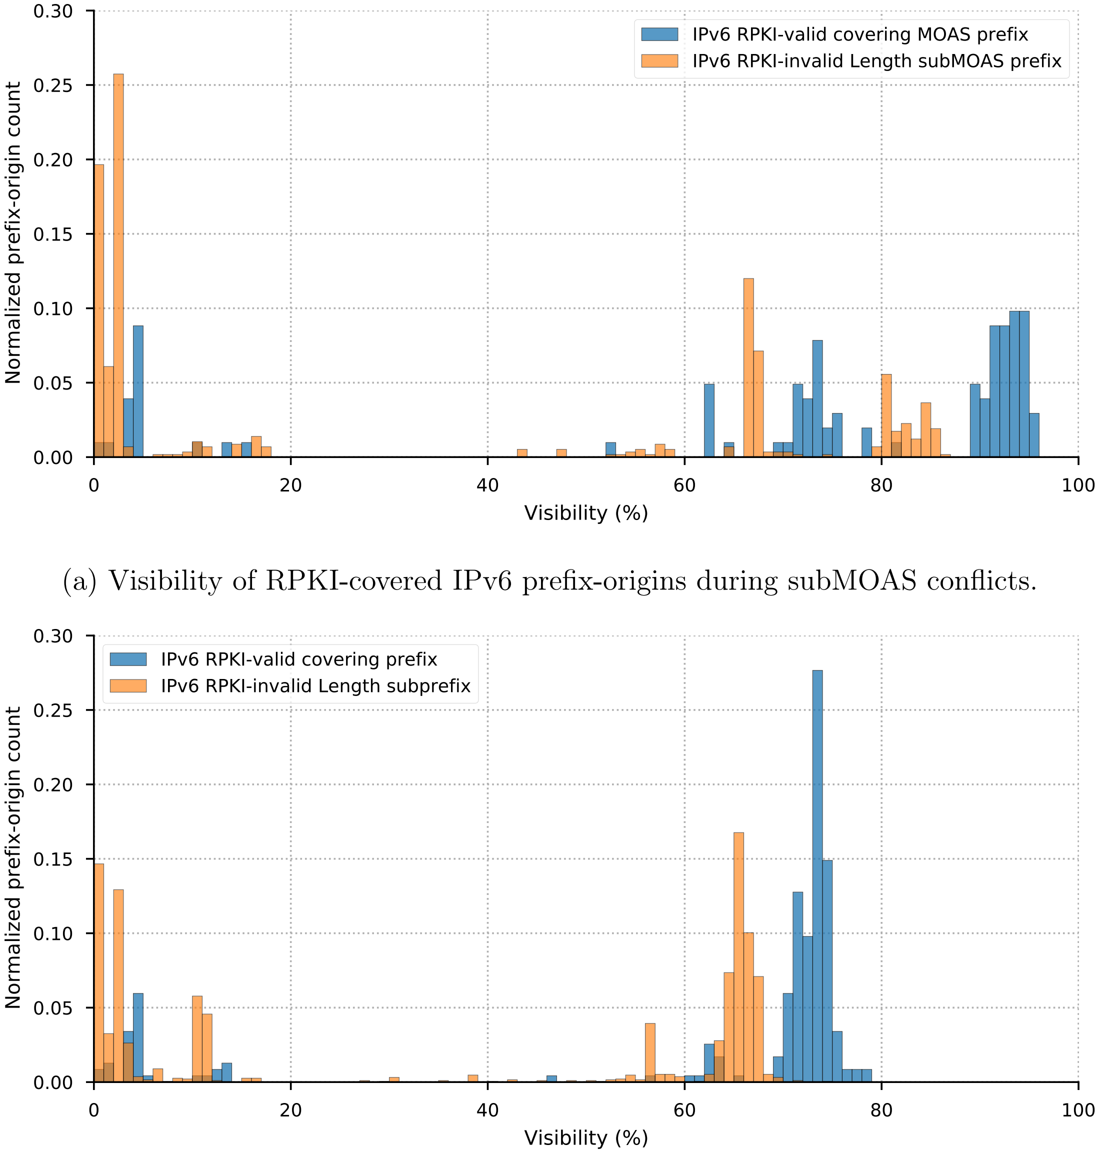

(a) Visibility of RPKI-covered IPv6 prefix-origins during subMOAS conflicts.

(b) Visibility of RPKI-covered IPv6 prefix-origins during subprefix conflicts.

Figure 5-8: Impact of RPKI registration in subMOAS and subprefix conflicts for IPv6.

4.4 Visibility of Subprefix Announcements

We next study instances of subprefix announcements, which instead do not compete with the covering prefix for visibility, since routers use longest-prefix matching, preferring more-specific routes for packet forwarding. For this reason, subprefix announcements can be a powerful and dangerous tool to, *e.g.,* hijack address space and redirect traffic, and their effect has been also evidenced in large-scale routing incidents, including route leaks [164, 179].

To study the impact of RPKI registration on subprefix announcements, we first isolate all incidents of subprefix announcements in our dataset, *i.e.,* we observe a covering (that is, less specific) prefix, covered by a validated ROA in the RPKI, and concurrently a more specific prefix announcement that does not pass RPKI validation—either because of an invalid ASN (subMOAS) or invalid prefix length (subprefix). In total, we find 10,450 instances of RPKI-invalid subprefix and subMOAS announcements in IPv4, conflicting with 2,291 RPKI-valid covering prefixes. Figure 5-7a and Figure 5-8a show the distribution of prefix visibility in the case of subMOAS for IPv4 and IPv6 respectively: if a more-specific prefix announcement fails RPKI validation because it has a different origin AS (N=5,401 subMOAS prefixes, N=966 covering prefixes for IPv4; N=575 subMOAS prefixes, N=102 covering prefixes for IPv6). While the RPKI-valid covering prefixes show high visibility, their invalid counterpart, subMOAS prefixes, show two modes of visibility: some 35% of invalid subMOAS show very low visibility, *i.e.,* lower than 10%. More importantly, however, is the finding that none of the subMOAS prefixes reach the same visibility level as their valid parent: while subMOAS prefixes barely exceed 75% visibility, their valid covering prefixes typically reach some 85% - 90% visibility and 75% reach at least 80% visibility. These observations are consistent with our earlier findings of increasing RPKI filtering, and highlight that RPKI registration also benefits registrants in the case of difficult-to-combat subMOAS situations.

Figure 5-7b shows the visibility for invalid-length subprefix announcements having the same origin AS as their covering RPKI-valid counterpart (N=5,049 subprefix, N=1,325 covering prefixes for IPv4; N=1,903 subprefix, N=235 covering prefixes for IPv6). Recall that the RPKI permits to specify a maxLength attribute, limiting the prefix length of any prefix matching the RPKI record, irrespective of the origin AS. Besides cases of misconfiguration, this scenario also applies in the case of *path hijacks*: instances where an attacker injects a subprefix that allegedly points to the same origin AS as its valid covering prefix, but in fact the attacker redirects traffic to its network. Such attacks can, *e.g.,* be carried out by prepending the valid origin AS at the end of the path after the hijacker's AS number. Such path hijacks present advanced forms of prefix hijacks and are difficult to detect using today's methods [55]. In Figure 5-7b, we see similarly lowered levels of visibility for RPKI-invalid subprefix announcements, even if they point to the registered origin AS. Invalid announcements reach some 70% of visibility, substantially lower when compared to their valid covering prefix. These results show that RPKI registration can benefit networks even in this most advanced case of illicit announcements: subprefix path hijacks.

5 Discussion

Recent research has shown increasing registration of ROAs in the RPKI by networks around the globe. Our work complements these observations, adding an important dimension: RPKI enforcement. We find that a substantial, and growing, number of ISPs in the Internet begin to filter invalid RPKI announcements, including major players such as AT&T. Increasing RPKI enforcement starts to bring direct value to networks, since registration in the RPKI benefits them in real-world scenarios, such as prefix hijacks. Our findings show that already as of today, registration in the RPKI limits the propagation of illicit announcements, in MOAS conflicts as well as in instances of subMOAS and subprefix announcements. While the RPKI protects its registrants in the case of such illicit announcements, we can also expect that increasing RPKI enforcement provides further incentives for networks to keep their RPKI records up-to-date, since stale records and other misconfigurations will have a direct impact on reachability of the respective address blocks. Our method provides a simple way to track current levels of RPKI filtering and to study its impact on illicit prefix announcements.

6 Conclusion

This work provides the first large-scale empirical evidence that validation of information as part of BGP route selection process benefits owners of IP prefixes that register them in authoritative databases such as the RPKI. Illicit announcements of IP prefixes for which networks have properly registered ROAs in the RPKI have limited visibility compared to announcements involving prefixes that are not covered by ROAs. Indeed, we measured that when less than 10% of networks had adopted this practice using the RPKI framework, the spread of invalid and potentially illicit announcements in BGP was reduced by 10-15%. Thus, registering ROAs in the RPKI delivers direct value to networks by effectively reducing the impact of illicit announcements. This empirical result demonstrates that this practice yields significant benefits in the present Internet and these benefits are not something to just hope for in the future. As many network operators and researchers have expressed skepticism about the benefits of RPKI when partially deployed, this result—based on real-world evidence–is game changing for understanding and advocating for RPKI and ROAs adoption.

This work also increases the incentive of providers to deploy operational security using RPKI by providing a method to passively track operator's RPKI-related behavior over time, which can identify different implementation decisions. Continuous monitoring of deployment of RPKI-validation and tracking of implementation decisions and challenges increases the transparency of the adoption and adoption process, providing further incentives for network operators to join the growing group of networks that protect their prefixes by registering ROAs in the RPKI. Indeed, Telia, one of the largest Internet Providers in the world, reproduced our method to monitor theirs and other networks RPKI-validation practices [180].

Chapter 6

Insights and Future Directions

The previous chapters have explored the reasons that have prevented security improvements in BGP, and provided empirical evaluations of the impact of BGP design flaws and solutions at scale and over time. The analysis of the life-cycle of ideas to secure BGP in Chapter 2 surfaced long-lasting disagreements related to the relevant threat model and choices of trust that have prevented consensus for a direction in routing security to emerge. Then the empirical study of malicious behavior in BGP revealed the existence of BGP serial hijackers—networks that persistently perform hijacks in BGP, proving that, in the current state of BGP security, there are very few barriers to performing even the basic forms of routing attacks, and networks are able to engage in repetitive malicious behavior with little consequences. In addition, the empirical study of misconfigurations showed that harmful route leaks frequently happen and that even these large networks for which availability is critical to properly provide their services, are vulnerable to them.

On the defense side, the study of the adoption of routing information validation in BGP using the RPKI and ROAs provided encouraging results. It revealed that this practice is finally gaining traction and that its brings benefits to networks even with limited deployment, contrary to the idea that almost full-deployment is needed for RPKI and ROAs to improve security. It also illustrated how passive monitoring of security practices can make more transparent security decisions of networks and their associated settings.

This chapter takes a step back to discuss barriers to the adoption of routing security measures, their impact and how empirical work can provide much needed evidence to overcome them. The goal is to link together the different parts to provide better evidence to better engage with network operators and all other stakeholders—including policy makers—to improve routing security in a viable way.

After the discussion of the insights, this chapter describes a set of actions that can support the adoption of routing security. These actions cannot be undertaken effectively by a single entity as they must be taken collectively, and could be orchestrated by a bottom-up industry effort, top-down by governments, or by public-private partnerships. Finally, this chapter proposes directions for future work in the area.

1 Insights to secure Internet routing

This section examines the main barriers of adoption of routing security and discusses how empirical evidence can provide a better understanding of trade-offs between different solutions. The insights are organized around three main aspects: trust, threat models, and uncertainty.

Trust relates to the choice by networks adopting routing security practices of BGP security mechanism and the related support infrastructure. BGP security solutions differ in the choice of supporting infrastructure and the organizations involved in the operation. This difference is many times implicit in specifications of security solutions and later under-considered in technical works, but it plays a central role. When adopting a security mechanism, networks decide based on the technical aspects of the mechanism as well as the choice of supporting infrastructure and organization.

The threat models of the different security proposals also differ, particularly when considering bad or wrong behavior on the other end of a BGP session. Most security proposals use cryptographic records and signatures that allowed verification of integrity of the data, *i.e.,* that the information was not changed and it is what the party wanted to say, but that does not prevent a network lying or unintentionally sending wrong information. Depending on the security proposals, the incorrect behavior of other network was considered harmful in different scenarios, but no solution can cover all possibilities of misbehavior in a distributed system such as Internet routing.[1] Thus, choosing a security proposal is also deciding on a threat model.

Finally, given that no solution is perfect, networks adopting routing security are faced with uncertainty about the outcomes of their implementation efforts. Security proposals may interfere with legitimate BGP announcements that happen not to follow security rules but that are not malicious or wrong per se. Proposals might also disrupt other aspect of usual operation. These consequences of security proposals are not usually studied until networks implement a given framework, making harder the initial buy-in from network operators.

1.1 The critical role of trust

The central role of trust in the adoption of routing security became apparent through the review of the many proposals to secure routing, as presented in chapter 2. Indeed, when choosing a security mechanisms, networks decide which infrastructure and related organization to trust. Then, the empirical work of chapters 3, 4 and 5 provided insightful evidence to understand the implications and how to build trust in different aspects of security proposals.

There are two aspects of trust that can influence networks' decision to adopt security proposals: the trustworthiness of supporting infrastructure and related organizations; and the trust that other networks and organization will play they part in routing security. The next paragraphs discuss each aspect in more detail.

[1] This problem is referred as Byzantine failure, where a node in the networks lies but otherwise functions properly [181].

Trustworthiness of supporting infrastructure

BGP security proposals use different mechanisms to support routing security. Each mechanism depends on different organizations as roots of trust to build authoritative databases and manage supporting infrastructure. These organizations decide for instance who can authoritatively claim ownership of a given resource and authorize the resource use in BGP. Indeed, the organizations at the root of trust issue or delegate to others the ability to issue authoritative assertions about routing resources, which significantly influences the content of the authoritative databases. In addition, the operational management of authoritative databases and other supporting infrastructure impacts the decision by networks operators to integrate security proposal into existing systems. For instance, the availability of the data and synchronization frequency of new or modified records can have repercussions in the operational integration and influences the security outcome. Thus, the arguments put forward to prefer one or another type of supporting infrastructure for routing security concern both the risk of misbehavior these organizations, and the operational ability of organization(s) managing the infrastructure.

To build trust in third party infrastructure supporting routing security (*e.g.,* RPKI or Internet Routing Registries' infrastructure), it is relevant to understand and monitor both the behavior and operation of such infrastructure. How networks are implementing and using the scheme can inform about these aspects of trust in supporting infrastructure. As an example informing about trustworthiness in the RPKI framework, chapter 5 presents a technique to measure which networks are integrating RPKI in their routing operations. These networks use RPKI data to validate information in BGP announcement they receive and drop the invalid ones rather than including them in the route selection process. This integration thus impacts routing decisions. Networks that adopt this practice are trusting the RPKI framework and the organizations managing in its infrastructure.[2] The longitudinal analysis reveals the uptake that route validation with RPKI data had in the past few years, providing evidence of an emerging direction for BGP security. It demonstrates that many networks had started to trust RPKI infrastructure to integrate it in their usual operation. In addition, the analysis provides a view of different implementation settings and challenges, relevant to the inner workings and operational readiness of the RPKI framework.

Future work to increase trust in the infrastructure supporting routing security can develop monitoring systems of the operation and use of such infrastructure. Having publicly available data that makes transparent key aspects of RPKI or similar infrastructures' operation and how networks have integrated those mechanisms in their operation gives empirical evidence for network operators to evaluate the trustworthiness of such infrastructures. If network operators do not trust administratively and operationally these infrastructures they will not integrate them in their operation. Knowing other networks use a mechanism and how, and having easy access to relevant operational monitoring can encourage the adoption of secure routing solutions.

[2] The organization managing the RPKI infrastructure and the roots of trust are the Regional Internet Registries (RIRs). See chapter 5 for more detail.

Relying on other parties

A network can only prevent the spread of incorrect routing information it receives and sends. Since it cannot limit incorrect or malicious announcements elsewhere in the Internet, it relies on other networks to prevent the spread of incorrect information related to its resources. The actions of one network impacts the security of resources from other networks, thus improving routing security requires that many networks implement routing security practices. Thus, networks ultimately rely on and have to trust security practices other networks adopt. In other words, the cost networks incur adopting security practices benefits others, creating a *collective action problem.*

A structural solution to overcome the collective action problem is to increase transparency of *who* is doing *what*, involving the different actors and creating a sense of accountability.[3] In the context of routing security, creating the sense of accountability requires more transparent and accessible information about which network is using which security practices and even evaluate the implementation of the practices over time. For instance, through the method developed in chapter 5 to identify networks using RPKI data to validate information in BGP, it is also possible to identify different outcomes coming from different implementation of this practice. Telia, one of the largest transit providers based in Europe, reproduced this method to have an outsider perspective of the outcome of their integration of RPKI validation in BGP route selection, and to compare their efforts with other similar networks [180].

Similarly, studying the dataset of route leak misconfigurations built in chapter 4 revealed that many networks appeared to be using a BGP configuration that limits the size of misconfigurations spreading through the routing system. Tracking these practices and who is appropriately implementing them allows network operators to better understand the security posture as well as the operational competence of other networks, which are necessary for building trust between networks.

Finally, the method presented in chapter 3 to find networks that persistently perform hijacks in BGP provides an overall assessment of the hijacking activity coming from a network. The outcome of this work has been used by networks and other research groups for evaluating network reputation and likelihood of future hijacking activity, other aspects of building trust between networks.

Future work can develop other methods to monitor network's implementation of other security practices and other aspects of routing behavior. For instance, the use of routing information in the Internet Routing Registries (IRRs) or other types of route filtering could be monitored and tracked over time. Similarly, network's overall burstiness behavior of BGP announcements and changes at link level in AS Paths could be evaluated over time to inform about the usual and unusual routing behavior of networks. Additionally, building databases of ground truth routing events is key to identify operational practices and who is enforcing them. Making transparent networks' posture with respect to security practices can develop trust between net-

[3] Transparency does not imply accountability per se. For instance, in a scheme such as the RPKI, there is currently no expectation that networks will implement the integration with the route selection process. There are no consequences for not validating announcements. However, peer pressure and other policy mechanism can create expectations.

works. It becomes visible which networks are playing their part to secure routing, encouraging other networks to act.

Summarizing, to adopt a routing security mechanism, networks need to trust the supporting infrastructure and related organizations, and that other networks will eventually also implement security. Unfortunately, these details were often lost among the technical details and description of mechanics of proposals. The next paragraph discusses this phenomenon.

Choice of trust in security proposals

Documents describing security proposals and later review papers often do not systemically describe and explain the choice of support infrastructure and the involved organizations, as were technical details such as security guarantees, scalability and performance. Different network operators have different postures on trusting other actors and infrastructure. Indeed, currently many partial solutions to improve BGP security co-exist, making it challenging for consensus or clear directions to emerge (see chapter 2 for more discussion of differences between security proposals and trust delegation).

Depending on the security proposals, certain organizations play critical roles in support of routing security. There are risks of misbehaviors in any framework, but discussion around proposals did not contemplate how to build trust in the chosen framework and the trustworthiness of involved organizations. Some proposals concentrate authoritative issuance of routing assertions in a few organizations (*e.g.,* RPKI that has 5 roots of trust) while others are more decentralized (*e.g.,* IRR has more than 20 registries from a mix of for-profit and non-for-profit organizations). Both schemes have different implications related to how trust is distributed in the mechanisms and how network operators assess the trustworthiness of the actors and infrastructure. For instance, in the RPKI framework, Regional Internet Registries (RIRs) manage the *trust anchors* or roots of trust of the RPKI repository for their geographic zone. The policies and practices that RIRs set up for their RPKI repository use and the issuance of authoritative RPKI records considerably influence the adoption and security outcomes of RPKI. As an example, ARIN[4] is the only RIR that requires the signing of a legal agreement before accessing data in ARIN's RPKI repository, resulting in many network deciding not to include ARIN records when they integrate RPKI data to validate routes in BGP routers (for more details in the impact of this practice, see chapter 3.3).

Future work can tackle the question of how to build trust in security frameworks to lower the barriers of adoption. What metrics and behaviors related to security frameworks are relevant for operators to trust the support infrastructure of security proposals? What else could facilitate the process of building trust and the adoption of these proposals? Tackling these questions would help guide discussions about proposals to build the consensus and trust needed to increase adoption, specially at

[4] ARIN (American Registry for Internet Numbers) is the RIR covering North America.

the early stage of deployment.

1.2 Settling on threat models

Different security proposals to secure BGP were designed for different threat models, all based on possible kinds of attacks. Most security proposals agreed on the use of cryptography as a mechanism to verify the integrity of the data sent by a network in BGP, *i.e.,* that the information sent was not changed by a third party and that it corresponds to what the network wanted to say. However, verifying integrity does not prevent a network from lying or unintentionally sending wrong information.

Security proposals considered the incorrect behavior of other networks harmful in different scenarios. For instance, some solutions proposed to verify and validate only part of the routing information sent in BGP, other proposed to verify only in case of conflict with other announcement, and others only when the announcement has not been seen in recent historical data.

Nonetheless, no solution can cover all possible misbehavior in a distributed system. Without additional security frameworks, networks by themselves cannot verify routing information of other networks in the Internet. Thus, a network that lies but otherwise functions properly is hardly distinguishable from a network advertising only legitimate and correct information. Security proposals establish *what* and *when* to verify information in BGP, and the authoritative source. Thus, choosing a security proposal is also deciding on a threat model and the kind of misbehavior of networks that is not tolerated.

Unfortunately, little is known about the pervasiveness and impact of different kinds of attacks in BGP , making it challenging to evaluate the trade-offs between proposals. Indeed, attacks currently in use by malicious actors do not predict how malicious actors will perform once better security is deployed. Bad actors can adapt their behavior and strategy, so blocking one attack vector currently in use will not necessarily reduce and limit overall levels of attacks. Nevertheless, it is unclear that more complex kinds of BGP hijacks (*e.g.,* path hijacks) can spread as easily as the basic hijacks[5] (*e.g.,* prefix hijacks) currently do. Thus, empirical evidence can complement and illuminate relevant aspects of the structural analysis of possible threats.

The empirical evidence presented in chapter 3 demonstrated that hijacks, even in the most basic form, are so pervasive, that there are networks (in large part) dedicated to hijacking. This finding suggest that it is justified to deploy security solutions even if these solutions do not prevent more complex and stealthier kinds of attacks. The bar for attackers is too low and even with repeated evidence of their malicious behavior, there are almost no consequences for perpetrators. In addition, that study provides a list of suspicious networks that can be used to analyze other aspect of malicious behavior, contributing to guide structural analysis.

Future work can develop new uses of automated methods to find different types of malicious behavior and evaluate their pervasiveness at scale and over time. More work is needed to better understand critical vulnerabilities to fix as well as the capabilities

[5] see Chapter 1 Section2 for more details on different kinds of hijacks.

of malicious actors to adapt. In addition, more empirical or simulation studies are needed to understand the viability and reach of different kinds of hijacks, taking into account routing policies and deployment levels of different routing security practices.

1.3 Coping with uncertainty and imperfection

A final challenge to deploying schemes to improve security is the more or less inevitable imprecision of most approaches, and thus the uncertainty regarding the impact to the outcome and operation of (early) adopters. In an ideal world, designers would invent security mechanisms that only prevent *bad* or *unwanted* behavior from happening, but that is an unrealistic expectation. Mechanisms are designed with a set of rules about what constitutes good and usual behavior. However, even if a behavior does not follow those rules, it can still be legitimate. Moreover, there is hardly an agreement about what constitutes bad or unwanted behavior. Indeed, behaviors allowed by security mechanisms can still be malicious or unwanted, and similarly, behaviors blocked by security mechanisms can be part of expected and usual operation.

A fundamental challenge for the design of any security mechanism—including proposals to secure BGP—is the balance between designing a mechanism that undershoots or overshoots. It is very hard to design a mechanism that perfectly draws the line between unwanted behavior and usual operation. As a consequence, proposals either undershoot or overshoot. If they undershoot, they leave options for malicious actors, who are quick to exploit them. If proposals overshoot, they prevent (or inconvenience) too many legitimate efforts. In both cases, it can be argued that the effort to deploy security mechanisms is not worth the effort, based on the mechanism limitations or the harm to operations.

Proposals to improve security are usually evaluated by trying to determine the scope of the bad behaviors they prevent. The estimation of how much inconvenience or disruption to acceptable behavior they trigger is often limited to performance and scalability issues. However, there are many ways in which a security mechanism may inconvenience or disrupt operations, and they are not studied in a systematic a way.

The fact that there might be disruptions to benign operations implies that ongoing measurement of benign activity must be a part of the overall deployment scheme. As an example, software-related issues in the RPKI framework and the potential disruption of customer traffic (*e.g.,* a valid BGP route may be dropped because an authoritative record is out-of-date causing loss of availability to the related IP address block) held back transit providers' adoption of RPKI data for validating BGP information until recently. Operationalizing the security scheme to limit the disruption of usual behavior is critical for reaching higher levels of adoption.

Unfortunately, there currently is no established path to operationalize solutions. Even if a protocol is standardized by the IEFT, there is no ownership and thus no clearly designated party in charge of developing the necessary pieces of software to implement the standard. This effort is currently addressed by many different actors such as router vendors (developing the code to integrate security solution with the routing protocol), administrators of support infrastructure such as the RIRs and independent third parties.

In designing an approach to improve routing security, it is most important to understand the complexities and incentives for deployment, which require empirical evidence to better understand trade-offs. For solutions to have impact, many network operators are required to take action and integrate solutions in their operation, so it is crucial to have a realistic view of how adoption will be encouraged and supported.

Future work can develop infrastructure to support ongoing measurement of both malicious (actual hijacks) and benign activities, to gauge whether the state of security is getting better, and whether collateral impairment of benign operations is being detected and mitigated.

2 Actions to improve routing security

Based on the insights discussed above, this section describes three main actions that would increase the adoption of routing security solutions. These actions can help overcome the fact that routing security cannot be undertaken effectively by a single entity, it requires a *collective effort*. The actions could be orchestrated by a bottom-up effort, top-down by governments, or by public-private partnerships. In particular, governments could encourage collective action by shaping incentives, through regulation, or by direct investment.

- Encourage the validation of information in BGP using authoritative databases. This means (1) encouraging owners of IP address space to register address blocks, whether in use in BGP or not, and the associated networks authorized to originate them in BGP in a publicly available routing database; and (2) encouraging networks to integrate data from these authoritative sources to validate information in BGP announcements before considering the information in the route selection process. Validating information in BGP reduces the impact of incorrect and potentially illicit routing information even in partial deployment. Although only validating IP prefixes and origin ASes still allows other attacks and misocnfiguration to spread, it is a relevant first step that reduces the impact of basic—widely used— forms of attacks.

- Mitigate non-technical barriers to adoption of BGP information validation (*e.g.,* using the RPKI framework or IRRs). These include legal barriers (*e.g.,* contracting terms between address owners and RIRs), and lack of organizational ownership in the software development of key components of security mechanisms (*e.g.,* reaching maturity in key software packages). These barriers indeed make it harder for network operators, and specially early adopters of security practices, to integrate security practices in their usual operation, which can significantly delay—and totally stall—their adoption.

- Organize an ongoing assessment of the overall state of routing security, the steps being undertaken to improve it, and the actions taken by different parties. Providing an ongoing assessment generates incentives to the individual parties to do their part, identifies continuing barriers to progress, and provides information to the larger community about the current levels of security and what barriers remain.

These actions are focused on promoting cooperation between networks and other relevant actors in securing Internet routing. By changing the structure of the dilemma with government creating the necessary incentives or bottom-up private sector initiatives calling for action, and by lowering non-technical barriers of adoption, these actions would move routing towards better security. Having these actions pursuing collective adoption of security practices also reduces the collateral damage of individual organizations deciding to adopt stronger security practices, leveling the playing field. The efforts to track progress and monitor compliance will create a level of accountability for the behaviors and practices of the various actors, which again contributes to the shared sense of trust among those actors.

3 Future directions

Much has changed in the Internet since the first version of the Border Gateway Protocol (BGP) was standardized in 1989. In particular, the scale of the Internet makes it very difficult to detect unwanted and malicious behavior from networks in BGP. To reduce the spread of routing misconfigurations and attacks, network operators need to adopt better security practices. However, although there has been much work on routing security proposal in the last 30 years, little has changed in operational environments. This dissertation provided evidence-based insights on how to improve the adoption of BGP security. In Section 1 of this chapter, many avenues for future work are described in the discussion of barriers to the adoption of routing security. The next paragraphs summarize those ideas.

A key challenge to support the adoption of routing security is to develop network operators' trust in the supporting infrastructure of security mechanisms and their operational management, as well as in the ability of other networks to properly integrate those mechanism in their BGP operation. Therefore, tackling the question of how to build network operators' trust in security frameworks will be critical to reach high-levels of adoptions. What information do network operators need to integrate security mechanisms in their operation? At what granularity? Answering these questions would help design appropriate methods to make transparent and easily available information that can lower the barriers to adoption of routing security practices.

Developing monitoring systems to assess the behavior and operation of infrastructure supporting routing security can contribute to increase trust in the organization managing in the infrastructure. Having publicly available data that makes transparent key aspects of supporting infrastructures' operation and how networks have integrated those mechanisms in their operation, gives empirical evidence for network operators to evaluate the trustworthiness of such infrastructures.

In addition, knowing other networks use a mechanism and how can encourage the adoption of routing security. Thus, building monitoring systems that track all types of practices that impact BGP security can provide information to ease the operationalization and prevent routing disruption in the adoption of BGP security. Making networks' posture with respect to security practices publicly available can develop trust between networks. It becomes visible which networks are playing their part to secure routing, encouraging other networks to act. As an example, the use of route filtering and the validation of BGP data using Internet Routing Registries (IRRs data) should be tracked over time and made publicly available.

Equally important, studying the usual and unusual routing behaviors in BGP, as well as network reputation in general, can help build trust between networks. For instance network's overall burstiness of BGP announcements and changes at link level in AS Paths can inform about networks' usual operational behavior. Additionally, developing new uses of automated methods to find different types of malicious behavior and evaluate their pervasiveness at scale and over time can inform network reputation.

Moreover, building databases of ground truth routing events is key to improve detection of interesting events for network reputation and to identify operational practices and who is enforcing them. External datasets coming from mailing lists, social networks and news can provide a starting point and timeframe of interesting events.

Likewise, more work is needed to better understand critical BGP vulnerabilities to fix and understand the capabilities and costs of malicious actors to adapt. In addition, more empirical or simulation studies are needed to understand the viability and reach of different kinds of hijacks, taking into account routing policies and deployment levels of different routing security practices.

Finally, developing infrastructure to support ongoing measurements of both malicious (actual hijacks) and misconfigurations that spread through the Internet, would allow to gauge whether the state of security is getting better. Then, complementing with measurements of benign activities could help assess collateral impairment of benign operations and whether they are being properly detected and mitigated.

Bibliography

[1] Sebastian Moss. Verizon BGP route leak causes Cloudflare customer outages, AWS issues. https://www.datacenterdynamics.com/en/news/bgp-route-leak-causes-cloudflare-outages-aws-issues/, June 2019.

[2] IETF. Internet engineering task force.

[3] Yakov Rekhter and Kirk. Lougheed. RFC 1105: A Border Gateway Protocol (BGP), Jun 1989.

[4] Yakov Rekhter and Tony Li. RFC 1654: A Border Gateway Protocol 4 (BGP-4), Jul 1994.

[5] Yakov Rekhter and Tony Li. RFC 1771: A Border Gateway Protocol 4 (BGP-4), Mar 1995.

[6] Susan Hares, Yakov Rekhter, and Tony Li. RFC 4271: A Border Gateway Protocol 4 (BGP-4). IEFT RFC 4271, Jan 2006.

[7] Richard Barnes, Jacob Hoffman-Andrews, Daniel McCarney, and James Kasten. RFC 8555: Automatic certificate management environment (ACME), Mar 2019.

[8] S. Kirkpatrick, M. Stahl, and M. Recker. Internet Numbers. IETF RFC 1166, 1990.

[9] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and kc claffy. AS Relationships, Customers Cones, and Validations. In *ACM IMC*, 2013.

[10] V. Giotsas, M. Luckie, B. Huffaker, and k. claffy. Inferring Complex AS Relationships. In *ACM IMC*, 2014.

[11] Dan Goodin. Suspicious event hijacks Amazon traffic for 2 hours, steals cryptocurrency. https://arstechnica.com/information-technology/2018/04/suspicious-event-hijacks-amazon-traffic-for-2-hours-steals-cryptocurrency/, Apr 2018.

[12] Ameet Naik. Anatomy of a BGP Hijack on Amazon's Route 53 DNS Service. https://blog.thousandeyes.com/amazon-route-53-dns-and-bgp-hijack/, Apr 2018.

[13] Sharon Goldberg. The myetherwallet.com hijack and why it's risky to hold cryptocurrency in a webapp. https://medium.com/@goldbe/the-myetherwallet-com-hijack-and-why-its-risky-to-hold-cryptocurrency-in-a-webapp-26113
Apr 2018.

[14] The Hunt for 3ve: Taking down a major ad fraud operation through industry collaboration. Technical report, Nov 2018.

[15] Anirudh Ramachandran and Nick Feamster. Understanding the network-level behavior of spammers. In *ACM SIGCOMM Computer Communication Review*, volume 36, pages 291–302. ACM, 2006.

[16] Jérôme Fleuri. Anatomy of a route leak. `https://www.afpif.org/wp-content/uploads/2019/08/Cloudflare.pdf`, August 2010.

[17] Route leak by the big Russian carrier AS8359, February 2020.

[18] Aftab Siddiqui. Big route leak shows need for routing security. `https://www.manrs.org/2020/07/big-route-leak-shows-need-for-routing-security/`, Jul 2020.

[19] Eric C. Rosen. RFC 827: Exterior Gateway Protocol (EGP), October 1982.

[20] Sue Romano, Mary Stahl, and Mimi Recker. RFC1117: Internet Numbers, Aug 1989.

[21] B. R. Smith and J. J. Garcia-Luna-Aceves. Securing the border gateway routing protocol. In *Global Telecommunications Conference, 1996. GLOBECOM '96. 'Communications: The Key to Global Prosperity*, pages 81–85, Nov 1996.

[22] Stephen Kent, Charles Lynn, and Karen Seo. Secure border gateway protocol (S-BGP). *IEEE Journal on Selected areas in Communications*, 18(4):582–592, 2000.

[23] M. G. Gouda, E. N. Elnozahy, Chin-Tser Huang, and T. M. McGuire. Hop integrity in computer networks. *IEEE/ACM Transactions on Networking*, 10(3):308–319, Jun 2002.

[24] Xiaoliang Zhao, Dan Pei, Lan Wang, D. Massey, A. Mankin, S. F. Wu, and Lixia Zhang. Detection of invalid routing announcement in the Internet. In *Proceedings International Conference on Dependable Systems and Networks*, pages 59–68, 2002.

[25] Russ White. Securing BGP Through Secure Origin BGP - The Internet Protocol Journal - Volume 6, Number 3. *The Internet Protocol Journal*, 6(3), Sep 2003.

[26] Geoffrey Goodell, William Aiello, Timothy Griffin, John Ioannidis, Patrick D. McDaniel, and Aviel D. Rubin. Working around BGP: An Incremental Approach to Improving Security and Accuracy in Interdomain Routing. In *ISOC Symposium on Network and Distributed Systems Security*, volume 23, page 156, 2003.

[27] Yih-Chun Hu, Adrian Perrig, and Marvin Sirbu. SPV: Secure path vector routing for securing BGP. *ACM SIGCOMM Computer Communication Review*, 34(4):179–192, 2004.

[28] Lakshminarayanan Subramanian, Volker Roth, Ion Stoica, Scott Shenker, and Randy H Katz. Listen and Whisper: Security Mechanisms for BGP. In *1st Symposium Networked System Design and Implementation,*, page 14, 2004.

[29] Tao Wan, Evangelos Kranakis, and Paul C. van Oorschot. Pretty Secure BGP, psBGP. In *Proceedings of the 2005 ISOC Symposium on Network and Distributed Systems Security*, San Diego, 2005.

[30] Patrick Reynolds, Oliver Kennedy, Emin Gün Sirer, and Fred B. Schneider. Using External Security Monitors to Secure BGP. 2006.

[31] Jian Qiu and Lixin Gao. Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol. page 12, 2006.

[32] Josh Karlin, Stephanie Forrest, and Jennifer Rexford. Pretty Good BGP: Improving BGP by Cautiously Adopting Routes. In *Proceedings of the 2006 IEEE International Conference on Network Protocols*, pages 290–299, Fess parker's Doubletree, Santa Barbara, Ca, USA, Nov 2006. IEEE.

[33] J. Israr, M. Guennoun, and H. T. Mouftah. Credible BGP – Extensions to BGP for Secure Networking. In *2009 Fourth International Conference on Systems and Networks Communications*, pages 212–216, Sep 2009.

[34] Andy Heffernan. RFC 2385: Protection of BGP Sessions via the TCP MD5 Signature Option, Aug 1998.

[35] D. Meyer, J. Heasley, and V. Gill. RFC 3682: The Generalized TTL Security Mechanism (GTSM), Feb 2004.

[36] Joe Touch, Allison Mankin, and Ronald P. Bonica. RFC 5925: The TCP Authentication Option, Jun 2010.

[37] Matt Lepinski, Richard Barnes, and Stephen Kent. RFC 6480: An Infrastructure to Support Secure Internet Routing, Feb 2012.

[38] Matthew Lepinski and Kotikalapudi Sriram. RFC 8205: BGPsec Protocol Specification, Sep 2017.

[39] Meiyuan Zhao, Sean W. Smith, and David M. Nicol. The performance impact of BGP security. *IEEE network*, 19(6):42–48, 2005.

[40] David M Nicol, Sean W Smith, and Meiyuan Zhao. Evaluation of efficient security for BGP route announcements using parallel simulation. *Simulation Modelling Practice and Theory*, 12(3-4):187–216, Jul 2004.

[41] Kevin Butler, Toni R. Farley, Patrick McDaniel, and Jennifer Rexford. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE*, 98(1):100–122, Jan 2010.

[42] M. O. Nicholes and B. Mukherjee. A survey of security techniques for the border gateway protocol (BGP). *IEEE Communications Surveys Tutorials*, 11(1):52–65, 2009.

[43] Geoff Huston, Mattia Rossi, and Grenville Armitage. Securing BGP — A Literature Survey. *IEEE Communications Surveys & Tutorials*, 13(2):199–222, 2011.

[44] Sharon Goldberg, Michael Schapira, Pete Hummon, and Jennifer Rexford. How secure are secure interdomain routing protocols? *Computer Networks*, 70:260–287, Sep 2014.

[45] Robert Lychev, Michael Schapira, and Sharon Goldberg. Rethinking security for internet routing. *Communications of the ACM*, 59(10):48–57, Sep 2016.

[46] M.S. Siddiqui, D. Montero, R. Serral-Gracià, X. Masip-Bruin, and M. Yannuzzi. A survey on the recent efforts of the Internet Standardization Body for securing inter-domain routing. *Computer Networks*, 80:1–26, Apr 2015.

[47] Pavlos Sermpezis, Vasileios Kotronis, Alberto Dainotti, and Xenofontas Dimitropoulos. A Survey among Network Operators on BGP Prefix Hijacking. *ACM SIGCOMM Computer Communication Review*, 48(1):64–69, Apr 2018.

[48] Jian Qiu and Lixin Gao. Hi-BGP: A Lightweight Hijack-proof Inter-domain Routing Protocol. Technical report, 2006.

[49] Mohit Lad, Dan Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang. PHAS: A Prefix Hijack Alert System. In *Proceedings of the 15th Conference on USENIX Security Symposium - Volume 15*, USENIX-SS'06, Berkeley, CA, USA, 2006. USENIX Association.

[50] Xin Hu and Z. Morley Mao. Accurate Real-time Identification of IP Prefix Hijacking. In *2007 IEEE Symposium on Security and Privacy (SP '07)*, pages 3–17, May 2007.

[51] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSPY: Detecting IP Prefix Hijacking on My Own. *IEEE/ACM Transactions on Networking*, 18(6):1815–1828, Dec 2010.

[52] Tongqing Qiu, Lusheng Ji, Dan Pei, Jia Wang, Jun Xu, and Hitesh Ballani. Locating Prefix Hijackers Using LOCK. In *Proceedings of the 18th Conference on USENIX Security Symposium*, SSYM'09, pages 135–150, Berkeley, CA, USA, 2009. USENIX Association.

[53] Xingang Shi, Yang Xiang, Zhiliang Wang, Xia Yin, and Jianping Wu. Detecting Prefix Hijackings in the Internet with Argus. In *ACM IMC*, 2012.

[54] Johann Schlamp, Ralph Holz, Quentin Jacquemart, Georg Carle, and Ernst W. Biersack. HEAP: Reliable Assessment of BGP Hijacking Attacks. *IEEE Journal on Selected Areas in Communications*, 34(6):1849–1861, Jun 2016.

[55] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. ARTEMIS: Neutralizing BGP Hijacking within a Minute. *arXiv:1801.01085 [cs]*, Jan 2018.

[56] Song Li, Haixin Duan, Zhiliang Wang, and Xing Li. Route leaks identification by detecting routing loops. In Bhavani Thuraisingham, XiaoFeng Wang, and Vinod Yegneswaran, editors, *Security and Privacy in Communication Networks*, pages 313–329, Cham, 2015. Springer International Publishing.

[57] M. S. Siddiqui, D. Montero, M. Yannuzzi, R. Serral-Gracià, X. Masip-Bruin, and W. Ramirez. Route leak detection using real-time analytics on local bgp information. In *2014 IEEE Global Communications Conference*, pages 1942–1948, 2014.

[58] Muhammad Shuaib Siddiqui, Diego Montero, René Serral-Gracià, and Marcelo Yannuzzi. Self-reliant detection of route leaks in inter-domain routing. *Computer Networks*, 9, 06 2015.

[59] S. Su, B. Zhang, L. Ye, H. Zhang, and N. Yee. Towards real-time route leak events detection. In *2015 IEEE International Conference on Communications (ICC)*, pages 7192–7197, 2015.

[60] J. Mauch. BGP Routing Leak Detection System. `http://puck.nether.net/bgp/leakinfo.cgi/`.

[61] Jingwei Liu, Bin Yang, Jinju Liu, Yuliang Lu, and Kailong Zhu. A Method of Route Leak Anomaly Detection Based on Heuristic Rules. pages 662–666. Atlantis Press, Jun 2017. ISSN: 2352-5401.

[62] Cisco BGPStream CrossworkCloud. `https://bgpstream.com/`.

[63] Maria Konte, Roberto Perdisci, and Nick Feamster. ASwatch: An AS Reputation System to Expose Bulletproof Hosting ASes. In *ACM SIGCOMM*, 2015.

[64] Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *Proceedings 2015 Network and Distributed System Security Symposium*, San Diego, CA, 2015. Internet Society.

[65] Pablo Moriano, Raquel Hill, and L. Jean Camp. Using bursty announcements for detecting BGP routing anomalies. *Computer Networks*, 188:107835, Apr 2021.

[66] Cecilia Testart. Reviewing a Historical Internet Vulnerability: Why Isn't BGP More Secure and What Can We Do About it? Washinton, DC, Aug 2018. Social Science Research Network.

[67] Yakov Rekhter and Kirk Lougheed. RFC 1163: A Border Gateway Protocol (BGP-2), Jun 1990.

[68] Yakov Rekhter and Tony Li. RFC 1267: A Border Gateway Protocol 3 (BGP-3), Oct 1991.

[69] Sandra Murphy. RFC 4272: BGP Security Vulnerabilities Analysis, Jan 2006.

[70] Gert Doering, Jerome Durand, and Ivan Pepelnjak. RFC 7454: BGP Operations and Security, Feb 2015.

[71] Lianshu Zheng, Keyur Patel, and Mahesh Jethanandani. RFC 6952: Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide, May 2013.

[72] IETF. Secure Inter-Domain Routing (sidr) - Documents.

[73] Geoff Huston, George Michaelson, and Robert Loomans. RFC 6481: A Profile for Resource Certificate Repository Structure, Feb 2012.

[74] Matt Lepinski, Derrick Kong, and Stephen Kent. RFC 6482: A Profile for Route Origin Authorizations (ROAs), Feb 2012.

[75] Randy Bush. RFC 8207: BGPsec Operational Considerations, Sep 2017.

[76] R. Bush and R. Austein. The Resource Public Key Infrastructure (RPKI) to Router Protocol. RFC 6810 (Proposed Standard), Jan 2013. Updated by RFC 8210.

[77] Tony Bates, David Meyer, Daniel Karrenberg, Marten Terpstra, Elise Gerich, and Cengiz Alaettinoglu. RFC 2280: Routing Policy Specification Language (RPSL), Jan 1998.

[78] Akmal Khan, Hyun-chul Kim, Taekyoung Kwon, and Yanghee Choi. A comparative Study on IP Prefixes and their Origin ASes in BGP and the IRR. *ACM SIGCOMM Computer Communication Review*, 43(3), 2013.

[79] John W. Stewart. *BGP4: Inter-domain Routing in the Internet.* Addison Wesley, 1999.

[80] RIPE NCC. RIPE Network Coordination Centre.

[81] Randall Atkinson and Stephen Kent. RFC 2401: Security Architecture for the Internet Protocol, Nov 1998.

[82] Yih-Chun Hu, David McGrew, Adrian Perrig, Brian Weis, and Dan Wendlandt. (R)Evolutionary Bootstrapping of a Global PKI for Securing BGP. page 6, 2006.

[83] RPKI Deployment Monitor. `https://rpki-monitor.antd.nist.gov/`.

[84] Robert Kisteleki and Brian Haberman. RFC 7909: Securing Routing Policy Specification Language (RPSL) Objects with Resource Public Key Infrastructure (RPKI) Signatures, Jun 2016.

[85] Russ White and Bora Akyol. RFC 5123: Considerations in Validating the Path in BGP, Feb 2008.

[86] Kotikalapudi Sriram. RFC 8374: BGPsec Design Choices and Summary of Supporting Discussions, Apr 2018.

[87] Vinay K. Sriram and Doug Montgomery. Design and analysis of optimization algorithms to minimize cryptographic processing in BGP security protocols. *Computer Communications*, 106:75–85, Jul 2017.

[88] Russ White. BGPsec and Reality, Oct 2017.

[89] Alexander Aximov, Eugene Bogomazov, Randy Bush, Keyur Patel, and Job Snijders. Verification of AS_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization. `https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/`, Aug 2021.

[90] Brenden Kuerbis and Milton Mueller. Internet routing registries, data governance, and security. *Journal of Cyber Policy*, 2(1):64–81, Jan 2017.

[91] Danny Cooper, Ethan Heilman, Kyle Brogle, Leonid Reyzin, and Sharon Goldberg. On the risk of misbehaving RPKI authorities. In *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*, HotNets-XII, pages 1–7, New York, NY, USA, Nov 2013. Association for Computing Machinery.

[92] AT&T/as7018 now drops invalid prefixes from peers. `https://mailman.nanog.org/pipermail/nanog/2019-February/099501.html`.

[93] John Kristoff, Randy Bush, Chris Kanich, George Michaelson, Amreesh Phokeer, Thomas C. Schmidt, and Matthias Wählisch. On Measuring RPKI Relying Parties. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20, pages 484–491, New York, NY, USA, Oct 2020. Association for Computing Machinery.

[94] RIPE NCC. Summons of the RIPE NCC Against the State of the Netherlands. Technical report, Mar 2012.

[95] Limiting the Power of RPKI Authorities. Virtual Event Spain.

[96] Matthias Wählisch, Robert Schmidt, Thomas C. Schmidt, Olaf Maennel, Steve Uhlig, and Gareth Tyson. RiPKI: The Tragic Story of RPKI Deployment in the Web Ecosystem. In *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*, HotNets-XIV, pages 1–7, New York, NY, USA, Nov 2015. Association for Computing Machinery.

[97] Bahaa Al-Musawi, Philip Branch, and Grenville Armitage. BGP Anomaly Detection Techniques: A Survey. *IEEE Communications Surveys & Tutorials*, 19(1):377–396, 2017.

[98] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In *Proceedings of the Internet Measurement Conference on - IMC '19*, pages 420–434, Amsterdam, Netherlands, 2019. ACM Press.

[99] RIPE Network Coordination Centre. YouTube Hijacking: A RIPE NCC RIS case study. `https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study`, Mar 2008.

[100] Pierluigi Paganini. BGP hijacking - Traffic for Google, Apple, Facebook, Microsoft and other tech giants routed through Russia. `https://securityaffairs.co/wordpress/66838/hacking/bgp-hijacking-russia.html`, Dec 2017.

[101] Maria Apostolaki, Aviv Zohar, and Laurent Vanbever. Hijacking Bitcoin: Routing Attacks on Cryptocurrencies. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 375–392, San Jose, CA, USA, May 2017. IEEE.

[102] Randy Bush and Rob Austein. The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1. RFC 8210, IETF, Sep 2017.

[103] Mohit Lad, Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang. Understanding Resiliency of Internet Topology against Prefix Hijack Attacks. In *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*, pages 368–377, Jun 2007.

[104] Craig A. Shue, Andrew J. Kalafut, and Minaxi Gupta. Abnormally Malicious Autonomous Systems and Their Internet Connectivity. *IEEE/ACM Transactions on Networking*, 20(1):220–230, Feb 2012.

[105] NANOG mailing list and archives. `https://www.nanog.org/list/archives`.

[106] Mutually Agreed Norms for Routing Security (MANRS). `https://www.manrs.org/`.

[107] Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely Randomized Trees. *Mach. Learn.*, 63(1):3–42, Apr 2006.

[108] Leo Breiman. Out-Of-Bag Estimation. Technical report, Dec 1996.

[109] Trevor Hastie, Robert Tibshirani, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction.* Springer series in statistics. Springer, New York, NY, 2nd ed edition, 2009.

[110] Yanmin Sun, Andrew K. C. Wong, and Mohamed S. Kamel. CLASSIFICATION OF IMBALANCED DATA: A REVIEW. *International Journal of Pattern Recognition and Artificial Intelligence*, 23(04):687–719, Jun 2009.

[111] Carolin Strobl, Anne-Laure Boulesteix, Achim Zeileis, and Torsten Hothorn. Bias in random forest variable importance measures: Illustrations, sources and a solution. *BMC Bioinformatics*, 8(1):25, Jan 2007.

[112] Baptiste Gregorutti, Bertrand Michel, and Philippe Saint-Pierre. Correlation and variable importance in random forests. *Statistics and Computing*, 27(3):659–678, May 2017. arXiv: 1310.5726.

[113] Andre Altmann, Laura Tolosi, Oliver Sander, and Thomas Lengauer. Permutation importance: a corrected feature importance measure. *Bioinformatics*, 26(10):1340–1347, May 2010.

[114] Terence Parr, Kerem Turgutlu, Christopher Csiszar, and Jeremy Howard. Beware Default Random Forest Importances. `http://explained.ai/decision-tree-viz/index.html`, Mar 2018.

[115] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in python. *J. Mach. Learn. Res.*, 12:2825–2830, Nov 2011.

[116] Jon Mitchell. RFC 6996: Autonomous System (AS) Reservation for Private Use, Jul 2013.

[117] Shinyoung Cho, Romain Fontugne, Kenjiro Cho, Alberto Dainotti, and Phillipa Gill. BGP hijacking classification. In *Network Traffic Measurement and Analysis Conference (TMA)*, 2019.

[118] B. Quoitin, C. Pelsser, L. Swinnen, O. Bonaventure, and S. Uhlig. Interdomain traffic engineering with bgp. *Comm. Mag.*, 41(5):122–128, May 2003.

[119] Mattijs Jonker, Anna Sperotto, Roland van Rijswijk-Deij, Ramin Sadre, and Aiko Pras. Measuring the Adoption of DDoS Protection Services. In *ACM IMC*, 2016.

[120] DROP - Don't Route or Peer lists - The Spamhaus Project. `https://www.spamhaus.org/drop/`.

[121] UCEPROTECT. Blacklist Policy LEVEL 2. `http://www.uceprotect.net/en/index.php?m=3&s=4`.

[122] CAIDA - AS Rank. http://as-rank.caida.org.

[123] Krebs on Security. Notorious 'Hijack Factory' Shunned from Web. https://krebsonsecurity.com/tag/bitcanal/.

[124] Doug Madory. Sprint, Windstream: Latest ISPs to hijack foreign networks | Dyn Blog. https://dyn.com/blog/latest-isps-to-hijack/, Sep 2014.

[125] Doug Madory. The Vast World of Fraudulent Routing | Dyn Blog. https://dyn.com/blog/vast-world-of-fraudulent-routing/, Jan 2015.

[126] Doug Madory. Shutting down the BGP Hijack Factory | Dyn Blog. https://dyn.com/blog/shutting-down-the-bgp-hijack-factory/, Jul 2018.

[127] Thomas King. We Care About Data Quality at IXPs. https://ripe75.ripe.net/presentations/54-20171016-TKJS-RIPE-We_Care_About_Data_Quality_at_IXPs.pdf, Oct 2017.

[128] BitCanal hijack factory, courtesy of Cogent, GTT, and Level3. https://seclists.org/nanog/2018/Jun/370.

[129] Ronald Guilmette. RIPE Forum. https://www.ripe.net/participate/mail/forum/anti-abuse-wg/PDU2ODUzLjE1MzM4NTk2NzhAc2VnZmF1bHQudHJpc3RhdGVsb2dpYy5jb20+, Aug 2018.

[130] American Registry for Internet Numbers ARIN WHOIS. http://whois.arin.net/ui/.

[131] Vasileios Giotsas, Ioana Livadariu, and Petros Gigis. A First Look at the Misuse and Abuse of the IPv4 Transfer Market. In Anna Sperotto, Alberto Dainotti, and Burkhard Stiller, editors, *Passive and Active Measurement*, Lecture Notes in Computer Science, pages 88–103, Cham, 2020. Springer International Publishing.

[132] Yuchen Jin, Colin Scott, Amogh Dhamdhere, Vasileios Giotsas, Arvind Krishnamurthy, and Scott Shenker. Stable and Practical {AS} Relationship Inference with ProbLink. pages 581–598, 2019.

[133] Hitesh Ballani, Paul Francis, and Xinyang Zhang. A study of prefix hijacking and interception in the Internet. In *ACM SIGCOMM Computer Communication Review*, volume 37, pages 265–276. ACM, 2007.

[134] Xingang Shi, Yang Xiang, Zhiliang Wang, Xia Yin, and Jianping Wu. Detecting prefix hijackings in the internet with argus. In *Proceedings of the 2012 ACM conference on Internet measurement conference - IMC '12*, page 15, Boston, Massachusetts, USA, 2012. ACM Press.

[135] Andrei Robachevsky. 14,000 Incidents: A 2017 Routing Security Year in Review. https://bgpmon.net/todays-bgp-leak-in-brazil/, January 2018.

[136] Doug Madory. Learning from recent major BGP routing leaks. https://www.slideshare.net/apnic/learning-from-recent-major-bgp-routing-leaks, February 2018.

[137] Cymru BGP Bogon Refence. https://team-cymru.com/community-services/bogon-reference/.

[138] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *ACM SIGCOMM*, 2012.

[139] C. Labovitz, S. Lekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-Domain Traffic. In *ACM SIGCOMM*, 2010.

[140] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In *PAM*, 2008.

[141] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois. Remote Peering: More Peering without Internet Flattening. In *CoNEXT*, 2014.

[142] Romain Fontugne, Anant Shah, and Emile Aben. AS Hegemony: A Robust Metric for AS Centrality. In *Proceedings of the SIGCOMM Posters and Demos on - SIGCOMM Posters and Demos '17*, pages 48–50, Los Angeles, CA, USA, 2017. ACM Press.

[143] Romain Fontugne, Anant Shah, and Emile Aben. The (Thin) Bridges of AS Connectivity: Measuring Dependency Using AS Hegemony. In Robert Beverly, Georgios Smaragdakis, and Anja Feldmann, editors, *Passive and Active Measurement*, volume 10771, pages 216–227. Springer International Publishing, Cham, 2018.

[144] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding bgp misconfiguration. In *SIGCOMM'02, Pittsburgh, Pennsylvania, USA.*, August 2002.

[145] Rahul Hiran, Niklas Carlsson, and Phillipa Gill. Characterizing Large-Scale Routing Anomalies: A Case Study of the China Telecom Incident. In David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Matthew Roughan, and Rocky Chang, editors, *Passive and Active Measurement*, volume 7799, pages 229–238. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[146] Weekend route leak by AS7552, April 2020.

[147] Serious Times — Serious Leaks, March 2020.

[148] Doug Madory. Large European Routing Leak Sends Traffic Through China Telecom. https://blogs.oracle.com/internetintelligence/large-european-routing-leak-sends-traffic-through-china-telecom, June 2019.

[149] Aftab Siddiqui. Route Leak Causes Major Google Outage. https://www.manrs.org/2018/11/route-leak-causes-major-google-outage/, November 2018.

[150] Indian Route Leak or There and Back Again, December 2017.

[151] Andree Toonk. Today's BGP leak in Brazil. https://bgpmon.net/todays-bgp-leak-in-brazil/, October 2017.

[152] Romain Fontugne, Esteban Bautista, Colin Petrie, Yutaro Nomura, Patrice Abry, Paulo Goncalves, Kensuke Fukuda, and Emile Aben. BGP Zombies: An Analysis of Beacons Stuck Routes. In David Choffnes and Marinho Barcellos, editors, *Passive and Active Measurement*, volume 11419, pages 197–209. Springer International Publishing, Cham, 2019. Series Title: Lecture Notes in Computer Science.

[153] Cisco. Configuring the BGP Maximum-Prefix Feature. `https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/25160-bgp-maximum-prefix.html`.

[154] Curtis Villamizar, Ravi Chandra, and Ramesh Govindan. RFC2439: BGP Route Flap Damping, Nov 1998.

[155] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *ACM IMC*, 2014.

[156] Tom Scholl. Maximum Prefix Tripping: A potential workaround for leaking on the Internet. `https://archive.nanog.org/meetings/nanog38/presentations/scholl-maxpfx.pdf`.

[157] Taejoong Chung, Emile Aben, Tim Bruijnzeels, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M. Maggs, Alan Mislove, Roland van Rijswijk-Deij, John Rula, and Nick Sullivan. RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. In *Proceedings of the Internet Measurement Conference*, IMC '19, pages 406–419, Amsterdam, Netherlands, Oct 2019. Association for Computing Machinery.

[158] AS286 Routing Policy. `https://as286.net/AS286-routing-policy.html`.

[159] Dropping RPKI Invalid Prefixes. `https://blog.teliacarrier.com/2020/02/05/dropping-rpki-invalid-prefixes/`.

[160] RPKI Route Origin Validation - Africa. `https://mailman.nanog.org/pipermail/nanog/2019-April/100445.html`.

[161] Yossi Gilad, Avichai Cohen, Amir Herzberg, Michael Schapira, and Haya Shulman. Are We There Yet? On RPKI's Deployment and Security. In *Proceedings 2017 Network and Distributed System Security Symposium*, San Diego, CA, 2017. Internet Society.

[162] Andreas Reuter, Randy Bush, Italo Cunha, Ethan Katz-Bassett, Thomas C Schmidt, and Matthias Waehlisch. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *ACM SIGCOMM Computer Communication Review*, 48(1):9, 2018.

[163] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. To Filter or Not to Filter: Measuring the Benefits of Registering in the RPKI Today. In Anna Sperotto, Alberto Dainotti, and Burkhard Stiller, editors, *Passive and Active Measurement*, volume 12048, pages 71–87, Oregon, US, 2020. Springer International Publishing. Series Title: Lecture Notes in Computer Science.

[164] Tom Strickx. How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today, Jun 2019.

[165] Matt Lepinski, Derrick Kong, and Stephen Kent. RFC 6482: A Profile for Route Origin Authorizations (ROAs), Feb 2012.

[166] S. Kent, D. Kong, K. Seo, and R. Watro. Certificate Policy (CP) for the Resource Public Key Infrastructure (RPKI). RFC 6484 (Best Current Practice), Feb 2012.

[167] G. Huston, G. Michaelson, and R. Loomans. A Profile for X.509 PKIX Resource Certificates. RFC 6487 (Proposed Standard), Feb 2012. Updated by RFCs 7318, 8209.

[168] A. Newton and G. Huston. Policy Qualifiers in Resource Public Key Infrastructure (RPKI) Certificates. RFC 7318 (Proposed Standard), Jul 2014.

[169] G. Huston, G. Michaelson, C. Martinez, T. Bruijnzeels, A. Newton, and D. Shaw. Resource Public Key Infrastructure (RPKI) Validation Reconsidered. RFC 8360 (Proposed Standard), Apr 2018.

[170] Daniele Iamartino, Cristel Pelsser, and Randy Bush. Measuring BGP Route Origin Registration and Validation. In Jelena Mirkovic and Yong Liu, editors, *Passive and Active Measurement*, volume 8995, pages 28–40. Springer International Publishing, Cham, 2015.

[171] Ben Cartwright-Cox. The year of RPKI on the control plane. `https://blog.benjojo.co.uk/post/the-year-of-rpki-on-the-control-plane`, September 2019.

[172] Cloudflare. Is BGP safe yet? `https://isbgpsafeyet.com`.

[173] Chiara Orsini, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In *Proceedings of the 2016 Internet Measurement Conference*, IMC '16, pages 429–444, Santa Monica, California, USA, Nov 2016. Association for Computing Machinery.

[174] RIPE NCC RPKI Validator. `https://rpki-validator.ripe.net/`.

[175] PeeringDB. `https://www.peeringdb.com`.

[176] Ben Maddison. RIPE Forum - Routing Working Group - RPKI Route Origin Validation - Africa, Apr 2019.

[177] Christopher Yoo and David Wishnick. Lowering Legal Barriers to RPKI Adoption. *Faculty Scholarship at Penn Law*, Jan 2019.

[178] Cisco. IP Routing: BGP Configuration Guide, Cisco IOS XE Release 3S. `https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-3s/irg-xe-3s-book/bgp-origin-as-validation.html`.

[179] Lily H. Newman. Why Google Internet Traffic Rerouted Through China and Russia. *Wired*, Nov 2018.

[180] Carl Fredrik Lagerfeldt and Johan Gustawaaon. Routing Security: RPKI Update Q2/20, May 2020.

[181] Radia Perlman. *Network Layer Protocols With Byzantine Robustness*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1988.