Universität des Saarlandes

From the Edge to the Core: Towards Informed Vantage Point Selection for Internet Measurement Studies

A dissertation submitted towards the degree
Doctor of Natural Sciences (Dr. rer. nat.)
of the Faculty of Mathematics and Computer Science of Saarland University

by

Franziska Lichtblau, M.A.

Saarbrücken, 2021

Day of Colloquium:	March, 1st 2022
Dean of the Faculty:	Prof. Dr. Thomas Schuster
Chair of the Committee:	Prof. Dr. Krishna Gummadi
Reporters:	Prof. Anja Feldmann, Ph. D.
	Prof. Dr. Oliver Hohlfeld
	Prof. Dr. Balakrishnan Chandrasekaran, Ph. D.
Academic Assistant:	Devashish Gosain, Ph. D.

Eidesstattliche Versicherung

Hiermit versichere ich an Eides statt, dass ich die vorliegende Arbeit selbstständig und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt habe. Die aus anderen Quellen oder indirekt übernommenen Daten und Konzepte sind unter Angabe der Quelle gekennzeichnet. Die Arbeit wurde bisher weder im In- noch i m Ausland in gleicher oder ähnlicher Form in einem Verfahren zur Erlangung eines akademischen Grades vorgelegt.

Ort, Datum Franziska Lichtblau

Abstract

Since the early days of the Internet, measurement scientists are trying to keep up with the fast-paced development of the Internet. As the Internet grew organically over time and without build-in measurability, this process requires many workarounds and due diligence. As a result, every measurement study is only as good as the data it relies on. Moreover, data quality is relative to the research question—a data set suitable to analyze one problem may be insufficient for another. This is entirely expected as the Internet is decentralized, i.e., there is no single observation point from which we can assess the complete state of the Internet. Because of that, every measurement study needs specifically selected vantage points, which fit the research question.

In this thesis, we present three different vantage points across the Internet topology—from the edge to the Internet core. We discuss their specific features, suitability for different kinds of research questions, and how to work with the corresponding data. The data sets obtained at the presented vantage points allow us to conduct three different measurement studies and shed light on the following aspects: (a) The prevalence of IP source address spoofing at a large European Internet Exchange Point (IXP), (b) the propagation distance of BGP communities, an optional transitive BGP attribute used for traffic engineering, and (c) the impact of the global COVID-19 pandemic on Internet usage behavior at a large Internet Service Provider (ISP) and three IXPs.

Abstract

Seit den frühen Tagen des Internets versuchen Forscher im Bereich Internet Measurement, mit der rasanten Entwicklung des des Internets Schritt zu halten. Da das Internet im Laufe der Zeit organisch gewachsen ist und nicht mit Blick auf Messbarkeit entwickelt wurde, erfordert dieser Prozess eine Meg Workarounds und Sorgfalt. Jede Measurement Studie ist nur so gut wie die Daten, auf die sie sich stützt. Und Datenqualität ist relativ zur Forschungsfrage - ein Datensatz, der für die Analyse eines Problems geeiget ist, kann für ein anderes unzureichend sein. Dies ist durchaus zu erwarten, da das Internet dezentralisiert ist, d. h. es gibt keinen einzigen Beobachtungspunkt, von dem aus wir den gesamten Zustand des Internets beurteilen können. Aus diesem Grund benötigt jede Measurement Studie gezielt ausgewählte Beobachtungspunkte, die zur Forschungsfrage passen.

In dieser Arbeit stellen wir drei verschiedene Beobachtungspunkte vor, die sich über die gsamte Internet-Topologie erstrecken— vom Rand bis zum Kern des Internets. Wir diskutieren ihre spezifischen Eigenschaften, ihre Eignung für verschiedene Klassen von Forschungsfragen und den Umgang mit den entsprechenden Daten. Die an den vorgestellten Beobachtungspunkten gewonnenen Datensätze ermöglichen uns die Durchführung von drei verschiedenen Measurement Studien und damit die folgenden Aspekte zu beleuchten: (a) Die Prävalenz von IP Source Address Spoofing bei einem großen europäischen Internet Exchange Point (IXP), (b) die Ausbreitungsdistanz von BGP-Communities, ein optionales transitives BGP-Attribut, das Anwendung im Bereich Traffic-Enigneering findet sowie (c) die Auswirkungen der globalen COVID-19-Pandemie auf das Internet-Nutzungsverhalten an einem großen Internet Service Provider (ISP) und drei IXPs.

Acknowledgments

Every Ph.D. journey is different. And mine has been no different from that. Yet, I feel very privileged to have been able to make the journey in the way I have. I was lucky enough to have enjoyed the company of inspiring, creative, and open-minded people along the way—people who gave me chances showed me how to take risks and grow as a human being. So I want to thank all of you.

First of all, I would like to thank my advisor Anja Feldmann who gave me the chance to take this junction. You showed me the open-mindedness I grew up with, and even though my professional background was not a "ten out of ten" fit, you gave me a chance. I have learned a lot from you—from paper writing to efficient IP address counting algorithms. But the most important lesson was to experience how a well-lead research group can work together and achieve amazing things without many of the conventions and rules, people expect in a research environment.

The person I have probably spend the most time with and who deserves a big thank you for all the patience and time he spent listening to me is Florian. We went through a lot together, and people often asked how I could work with you around—it was a pleasure (most of the time). Thanks for sharing professional and personal insights and becoming a dear friend over the years. We have learned how to share constructive criticism, deal with crises, and sing songs in a data center. At TU Berlin, our office always was a spot for people to show up and see what was going on—not always to our benefit. Over time, we did many research projects together and learned about our strengths and weaknesses, and managed to form a fantastic team. Thanks for being weird!

I could write another thesis about all the other important people who supported, laughed, and shouted with me along the way. INET was a great place to work over all the years. Special thanks go to the admin team—Sarah, Struck, Seba, Tom, Rainer, Simon, Sabet, and all the other people who set up infrastructure, pulled cables and fixed broken things. Many of you have become my friends, and I am very grateful for that. A big thanks go to Phils, who was among the first group members to open up to me and who made me understand the group's spirit. I can not thank Benthor enough for being my friend and exchanging thoughts, ideas and concepts with me—I will always have a copy of you living in my head. Thank you, Theresa, for sharing your perspective and experience in making your way and constantly reminding me that there is more than one way to do things. Thanks to Enric and Philipp, who accompanied me in some way or the other until now—may it be paper writing, data analytics, or advice on thesis writing, your help has been invaluable. Thanks to all my friends and colleagues at INET for sharing your experience, feedback, and all the work over time—thank you, Thomas, Matthias, Mirko, Arne, Carlo, Niklas, Aniis, Jawad, Mohamad, Seifeddine, Bala, and, all other members of the INET family who have walked parts of the way with me.

All my other collaborators deserve a big thank you as well. For me, the magic of research lies in sharing ideas and creating something out that together. So, thanks a lot to George, Ingmar, Chris, Randy, Cristel, Lars (who also doubled as human literature resource), Oli, Oliver, and Rob—it has been a blast.

A huge thanks goes to all the folks of the RIPE community who kindly took me in and let me see how working on the Internet could be fun. This community did (and still does) a lot to help me grow personally and professionally. Especially I would like to thank Stefan. You showed me around, introduced me to people and share your experience—I am grateful for such good friends.

Taking this journey in the way I did would not have been possible without the many people who have shaped my life until now. I feel deep gratitude for the family I grew up with. Mum and dad, you taught me a lot: Humanism, self-reflection, love, the enjoyments of good food and wine, critical thinking, but most importantly, you have taught me to make my own way and supported me in each step whether you understood it or not. Your first and foremost goal has always been to make my life happy. In my memory, there is no reference to the word "*no*"; it has always been "*nobecause*". This constant recognition of me as a thinking person who knows herself gave me the strength to make my own way.

I could not have made it through all this without my friends. Friends are those you choose to be around, and I am very happy with all my choices. I will not try to list you all, but I love all of you.

Last but not least, I would like to thank my partners who walked that way with me. Especially Arne, I would like to thank you for showing me that I do not need to be afraid and what I could do with my life if I decided to do it. Will, thank you for surviving many crazy moments with me and putting a lot of real-world perspective into a researchy mind. Alex, you have only joined me at the end of the Ph.D. journey, but you taught me to look at the world outside my bubble and how to appreciate how lucky I am.

Publications

Pre-published Papers

Parts of this thesis are based on the following peer-reviewed papers that have already been published. All my collaborators are among my co-authors.

International Conferences

Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann.
"Detection, classification, and analysis of inter-domain traffic with spoofed source IP addresses".
In: Proceedings of the 2017 Internet Measurement Conference. *IMC 2017*.
ACM, 2017, pp. 86–99.
ISBN: 978-1-4503-5118-8/17/11. DOI: https://doi.org/10.1145/3131365.3131367.

Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush.
"BGP Communities: Even more Worms in the Routing Can".
In: Proceedings of the 2018 Internet Measurement Conference. *IMC 2018*.
ACM, 2018, pp. 279–292.
ISBN: 978-1-4503-5619-0. DOI: https://doi.org/10.1145/3278532.3278557.

Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis.
"The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic".
In: Proceedings of the 2020 Internet Measurement Conference. *IMC 2020*.
ACM, 2020, pp. 1–18.
ISBN: 978-1-4503-8138-3. DOI: https://doi.org/10.1145/3419394.3423658.

Journal Articles

Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis.
"A Year in Lockdown: How the Waves of COVID-19 Impact Internet Traffic".
In: Communications of the ACM Volume 64, Issue 7 July 2021
ACM, 2021
ISSN: 0001-0782. DOI: https://doi.org/10.1145/3465212

Contents

List of Figures

List of Tables

1 Introduction

Internet Measurement as a scientific discipline strives to gain an understanding of the current Internet ecosystem. Even though the Internet is man-made, it developed organically. The Internet has reached a level of complexity that makes it impossible to understand it as a whole. Therefore, we constantly need to observe the Internet as a dynamic subject from multiple perspectives: Understanding how an Internet Service Provider (ISP) operates and what impact that may have on its customers is a different problem than analyzing denial-of-service attacks that can take critical services or infrastructure offline. These two examples come from two different areas, but they are essential aspects of the Internet ecosystem. This chapter introduces the importance of carefully crafted Internet measurement campaigns from multiple perspectives to keep track of the rapid developments in the Internet ecosystem and discusses the accompanying challenges.

The Internet has not been designed on a drawing board; instead, it grew and developed organically. As we know it today, the system underwent an unimaginable number of changes in just a couple of decades. Still, it retains a lot of the early design decisions. Initially, the Internet was built and designed in a university-centric closed-bubble world where most participants knew each other. As a result, e.g., securing Internet-based communication was not the top priority, instead "just making it work" was the foremost goal. This strategy was very successful—one could even argue too successful. Now, we are basically "stuck" with a protocol stack and core principles that have been designed for the demands of an earlier Internet. Society today has completely different requirements for a modern Internet. Indeed, the physical layer evolved to accommodate these demands. The amount of content transported by Internet infrastructure increased by several orders of magnitude. Many services are latency-sensitive, we combine different access technologies, and we rely on the Internet for privacy and security-relevant services in daily life. Satisfying these demands implies changes across all layers. We do not argue that all early design principles are outdated and should be abandoned, but all should be open to discussion in good scientific practice. To facilitate an open debate about a future Internet, we need real-world data to make informed decisions. To actually improve the overall stability and service level, rather than "blindly" adding new features resulting in unexpected service degeneration or security issues, we first need to understand the "state of the art". We can only gain these insights by constantly observing the Internet through carefully crafted measurement studies. By that measurement studies greatly inform future protocol design and infrastructure development. Given the complexity of the

measured subject, there is no one-size-fits-all recipe for designing and conducting measurements. As expected, Internet measurement follows good scientific practice: The first step is the problem definition, followed by data collection to answer the proposed research question and concluded by careful analysis and presentation. To make sure the results indeed answer the question at hand, we need to precisely craft every measurement study to fit the given problem [1].

In 2004 Paxson presented a comprehensive summary on how to conduct sound measurement studies [2]. All of the discussed aspects are still valid: dealing with specific measurement devices, dealing with large data volumes, and common misconceptions. However, even 15 years later, as the Internet measurement community, we have neither found a perfect solution to gain a holistic understanding of the Internet. Nor have we fully adopted all-important insights pointed out by Paxson, especially regarding data reproducibility and making measurement data publicly available (where possible). For example, a topic Paxson only touched slightly is vantage point selection for Internet-wide measurements: To correctly interpret observations in a distributed system, a study needs to take the characteristics of each observation point into account and put the results into the correct perspective. Throughout this thesis, we will focus on the relevance of selecting suitable data sources, i.e., vantage points for a given research problem.

As we already highlighted, the Internet constantly grows and changes. Network architects, operators, protocol designers, security specialists, researchers, and many more work together to fix some of the Internet's initial design "mistakes" and introduce new features. But how do you introduce changes in a complex system, and which role do measurements play? Consider those two classes of problems, (a) scenarios in which changes can be rolled out gradually because elements within the system act independently of each other, and (b) systems in which changes need to be done by all participating entities almost simultaneously because of many interdependencies within the system.

The Internet is a hybrid case, but it is close to the second case. A misconfiguration, infrastructure outage, or malicious interference at an important location can affect large parts of the Internet and even take parts of it offline. In 2015 a local misconfiguration at the Amsterdam Internet Exchange (AMS-IX) took the entire IXP offline for a short time [3] resulting in 2.3 TB of dropped traffic [4]. Measurements using RIPE-Atlas [5] show that about 60% of the tested end-to-end paths remained reachable by quickly rerouting and thus avoiding the outage. Still, 30% of the paths suffered from failed connectivity [6]. Similar effects have manifested themselves during the London Internet Exchange (LINX) outage on March 23, 2021 [7]. Giotsas et al. discuss the impact of BGP peering outages in detail [8]. Recently, in June 2021, many large services were not reachable due to an outage triggered by a customer misconfiguration at the content delivery network (CDN) Fastly [9]. In April 2021, AS55410 accidentally hijacked more than 30k prefixes [10] owned by Google, Microsoft, Akamai, Cloudflare, Fastly, and many other ASes. A BGP hijack refers to the illegitimate announcement of IP prefixes by an Autonomous System (AS), that does hold said prefixes. Some BGP hijacks are an intentional attempt to redirect

traffic to, e.g., render services unavailable. While these incidents are malicious, many hijacks are the result of misconfiguration and thus unintentional [11]. The above examples highlight that any change, whether large or small, on relevant Internet infrastructure potentially carries the risk of breaking things badly. This again underlines the importance of understanding the current state, i.e., conduct measurements, before the introduction of change.

Given the Internets' organic development to add a specific new feature, one may need to touch the entire protocol stack. To simplify adaptions, the Internet community designed their protocols quite flexible and left room for extensions, e.g., optional protocol fields, and added negotiation capabilities into the protocol stack. With that, it is sometimes possible to add new features to individual protocols, which do not rely on global adoption without breaking the entire system. However, this kind of flexibility comes at a price—namely, a diverse "zoo" of protocol flavors, versions, options, and derivatives. This diversity increases the number of potential configurations in the wild and the complexity of the entire system. This complexity is directly reflected in the design of measurement studies—a complex system is usually hard to measure because as many corner cases as possible are accounted for.

That most protocols are not designed with mensurability in mind [12] aggravates the situation. Especially early protocols like the Internet Protocol (IP) [13] did not anticipate the complexity we are facing today. In the past, IP addresses have been a valid source to identify a host. Today, multiple hosts can be behind a Network Address Translation (NAT) gateway with one Internet-facing IPv4 address, even hundreds behind a Carrier-Grade NAT (CGN); or a host may have an IPv4 as well as an IPv6 address. While there are efforts to add mensurability as a feature in modern protocol development, e.g., the introduction of the spin bit in QUIC to enable passive latency measurements [14], it is notoriously hard to add such features to old protocols, which simply do not expose this kind of information by design.

Industry-driven protocol development is another aspect that aggravates understanding the current state of the Internet. Naively, one could assume that the Internet runs on standardized protocols as published by the Internet Engineering Taskforce (IETF) in the form of Request for Comments (RFC). These documents result from extensive discussion and procedures of the Internet community and are the go-to references for developers who want to implement protocols. In reality, many changes are forced into production by big players in the industry, e.g., hardware or software vendors, who skip the IETF standardization process. On the one hand, this leads to a faster deployment of new features. On the other hand, it yields all the downsides of non-standardized protocol development: Sometimes, poorly documented vendor-specific features can lead to unexpected behavior and harder to grasp configurations. More importantly, it increases the number of unexpected configurations of connected systems. Then again, standardized protocol development at the IETF is not a guarantee for well-designed protocols nor speedy deployment. The process can take much time as all interested parties can influence the standardization phase. IETF standardization has lead to very long development phases resulting in complex protocol designs to accommodate all participants. There have been situations where initial

design decisions were already outdated by reality once the standardization process concluded. IPv6 is a prominent example. IPv6 became an IETF standard over 20 years ago, but we are still far away from Internet-wide IPv6 adaption despite the depletion of the IPv4 address pool. Many people argue that the tedious process at the IETF and some over-complex aspects of the resulting protocol are part of the factors hindering IPv6 deployment.

While these problems exist across the entire stack, we, in this thesis, focus on two perspectives, namely the routing layer and traffic level analysis. More precisely, we show how to design and run measurement studies from suitable vantage points in the realm of inter-domain routing. The goal of these studies is to uncover the current state and inner workings of the Internet. Additionally, we combine routing data with Internet traffic data to understand traffic flow dynamics and composition in various settings.

The Border Gateway Protocol (BGP), the de-facto standard in inter-domain routing, is a shining example of protocols, which already existed in the early days of the Internet. While it underwent many changes, the basic operational principles are still the same: individual Internet participants announce the reachability of their specific prefixes. Over time, this distributed approach has proven itself to be very robust and scalable, which is why BGP is still the core inter-domain routing protocol today. However, as many of the "first generation" Internet protocols BGP was initially designed without considerations for security or the global importance of the Internet today. Over time, the Internet community extended BGP and published many Best Current Practice (BCP) guides to address those concerns. One would expect quick adoption of vital improvements to such essential protocols. Strangely, this is not the case [15]. This raises two questions: (a) Why are some changes not widely adopted? (b) What are the side effects of adopted changes? As there is no central authority keeping track of which BGP flavors are out there, we need measurement studies to gather that information. These studies can also help to understand which improvements are beneficial. Again, recall that the Internet is not designed on a blank page but changes dynamically. To know if an "improvement" is, in fact, improving the current state, we first need to understand that state and the state after introducing that improvement.

From a routing perspective, the Internet does not have a global or central source of truth. The Internet is a network of networks - and every network is represented by an AS. Every AS announces to all other Autonomous System (AS)es via which path they are reachable using BGP. This process propagates new announcements in seconds to a few minutes on average [16] but can take up to 30 minutes [17]. In principle, all networks can gain a full view on the routing system, either by Peering or buying transit. However, due to the nature of BGP as a distributed system with local mechanisms to raise the preference of certain routes over others, different ASes have different views. This lack of a single source of truth again highlights the importance of carefully crafted measurement studies. For example, if one wants to understand how routing dynamics change in South America, it is not sufficient to look at the local view of a European ISP.

The routing perspective, how traffic traverses the Internet, provides valuable insights about Internet topology and its participants. We need traffic level data, e.g., traffic samples obtained at an ISP, to understand what kind of content is served over the Internet or how the transport layer evolves. Traffic level measurements give us crucial information like the protocol "mix", which strongly correlates to service usage and can estimate traffic volume and timing effects. If the traffic data contains payload, we could also derive what people are communicating. However, gathering payload data has serious privacy implications. All traffic-level data in this thesis consists of aggregated flow data without any payload information; see Section 2 for details. Like routing data, whether traffic data is meaningful for a given problem depends on the vantage point where it has been collected. An Internet Exchange Point (IXP) is an excellent vantage point for studies about the Internet core, whereas an ISP is more suited for eyeball-centric measurements.

In this thesis, we present three measurement studies based on large data sets at suitable vantage points, designed to shed light on different kinds of questions in Internet development, we: (a) demonstrate the importance of passive measurement studies to assess the current state of the Internet and provide a foundation for more specific active measurements; (b) quantify the prevalence of specific Internet-wide observable attacks and highlight difficulties and obstacles in the global rollout of new best practices, and (c) finally show how intertwined the Internet is with societal developments at large. To this end, we discuss two different large vantage points, namely a large European IXP and a large European Tier-1 ISP. These vantage points allow us to capture an approximation of a centralized view on the core of the Internet through the IXP as well as the more eyeball-focused perspective of the ISP. We highlight for which particular questions which vantage point is suitable. To augment the BGP views gathered at the respective vantage point, we also leverage and compare three route collector projects.

1.1 Contributions

The work presented in this thesis covers the following aspects:

Route collector data comparison and sanitation. We compare the views of four different route collector projects based on the fraction of observed IP address space and the number of discovered AS links. Additionally, we present a list of sanitation steps, optional and essential, to clean data obtained by public route collector projects from real-world data artifacts. These steps range from filtering known research prefixes that may exhibit unrepresentative behavior to removing non-routable address space.

Propagation of BGP Communities. We quantify how far BGP communities propagate in the global routing system. We show that almost 50% of BGP communities travel four hops , and depending on the AS path length, 20%-80% of communities travel more than 50% of the AS path distance. The degree of propagation is relevant

for possible attack vectors, which could become a reality with the growing popularity of BGP community-based services. We discuss two possible scenarios in which BGP communities could be abused.

Prevalence and characteristics of IP address spoofing in the wild. We developed and applied a new approach to detect packets with spoofed IP source addresses at a large European IXP. Furthermore, we compared and evaluated different techniques to generate AS-specific lists of valid address space and minimize false positive inferences. Finally, we presented a first in-depth analysis of Internet players emitting spoofed traffic and gave qualitative characteristics of said traffic exchanged in the inter-domain Internet.

Impact of the COVID-19 pandemic on Internet traffic. We studied the Internet traffic shifts at a large European Tier-1 provider as well as at a large European IXP during the COVID-19 pandemic. At both vantage points, we observed a drastic traffic growth of about 20% in a few days with a strong focus on communication-related services, which our vantage points could handle. Furthermore, we found that during the lockdown periods, the typical end-user usage pattern did not change between weekends and weekdays as many people worked from home and show how Internet usage strongly reflects significant societal trends.

1.2 Pre-published work and collaborations

This thesis is based on pre-published papers and articles in collaboration with many authors. We emphasize that our discipline's scientific work is always a joint effort of multiple minds and usually does not occur in isolation. In the following section, we outline the main contributions of the thesis author to pre-published work and collaborations used in this thesis.

Chapter 2: This chapter lays the groundwork for all the subsequent studies. It introduces important terms and concepts necessary to understand the fundamental problems of Internet measurement and the specific questions asked in this thesis. The compilation of information is the author's work, and sources are referenced.

Chapter 3: In this chapter we present a previously unpublished study in collaboration with Florian Streibelt, Lars Prehn, Cristel Pelsser, Randy Bush, and Anja Feldmann. The author's main contributions are (a) a survey of Internet measurement literature regarding the jointly derived sanitation steps, (b) to run analyses on pre-aggregated data sets resulting in the presented plots and descriptions, and (c) narrative conception and writing.

Chapter 4: We include an extract of a conference paper published at IMC 2018 in collaboration with all authors listed in [18]. The author's main contributions are (a) joint development of BGP community-based attack scenarios, (b) analysis of pre-aggregated BGP data regarding community propagation on the public Internet, (c) plot and text creation.

Chapter 5: Here, we present a paper published at IMC 2017 in collaboration with all the authors listed in [19]. The author's main contributions are: (a) Joint conception of the "full-cone", (b) analysis of IXP production traffic to isolate spoofed traffic, (c) analysis of spoofed traffic, (d) attack isolation, (d) join plot creation, narrative conception and writing,(f) infrastructure maintenance.

Chapter 6: This chapter is based on a study published at IMC 2020 and an article published in CACM in 2021 by a collaboration of authors as listed in [20, 21]. The authors main contributions are: (a) Join conception of measurement goals, (b) transport and application-level analyses of ISP production traffic, (c) comparison of IXP and ISP traffic changes, (d) joint plot creation, narrative conception, and writing.

Chapter 7: This chapter contains the authors concluding remarks, outlook to future work, and relevant related work, which has not been covered in previous chapters.

1.3 Structure of this thesis

The remainder of this thesis is structured as follows: Chapter 2 gives relevant information on vantage point and data selection to answer a specific research question, features, and the historical burden of commonly used routing protocols and infrastructural needs. Chapter 3 extensively discusses different route collector projects and best practices in the sanitation of real-world routing data. A case study focused on leveraging passive measurement of the global routing system to check the validity of well-known assumptions and thereby informing active measurement campaigns is presented in Chapter 4. In Chapter 5 we combine routing layer and traffic level analysis to assess how wide-spread source IP address spoofing on the Internet and how it is used. Additionally, Chapter 6 highlights how to design and evaluate a measurement study based on traffic level data of a large eyeball network in response to a global crisis. Finally, Chapter 7 concludes this thesis by summarizing our findings and outlining future work.

2
Background

This chapter introduces important terms and concepts to understand the design and operational building blocks of the Internet. Concepts, design, and operational practices are crucial to craft and carry out meaningful measurement campaigns.

For many IT specialists, the term “networking“ refers to the operation of their Local Area Network (LAN). In this scenario, the operator controls the entire local network, from hardware device selection to specific weights for certain paths within the network. Even though such networks can become quite complex and confusing, in essence, all the devices and configurations are available to gather a “single“ source of truth for the state of the network. Today, such a centralized system does not scale for many large LANs. It did never scale for the Internet. Today’s Internet was designed to connect all participants who do not necessarily know each other. From a researcher’s perspective, it is very tempting to try to understand the Internet by consulting textbooks and the underlying Request for Comments (RFC)s—everything should be documented and specified there. Due to the “organic” nature of the Internet, this is not sufficient. As such, in the following chapter, we focus on both aspects, namely the underlying design principles of the Internet and current operational practices in the wild.

2.1 Holding it all together: A short introduction to BGP terminology

In the following, we will introduce terms and concepts necessary to reason about the Internet structure. An entity that actively participates in the global Internet is called Autonomous System (AS). An AS is the smallest unit in Internet routing and consists of a collection of IP networks under one administrative control [22]. Practically speaking, this usually refers to a company, association, university, or similar real-world entities. Each AS is uniquely identified by a 16 or 32-bit identifier, the Autonomous System Number (ASN) which is assigned by its local Regional Internet Registry (RIR). Within an AS, the “operator” (or, more precisely, the group of routing policymakers) can deploy policies as they fit to realize their local goals. ASes act as outside facing abstraction entity and provide a consistent way to reach networks within the AS. The Border Gateway Protocol (BGP) [23] is the de-facto standard routing protocol on the Internet to route traffic between ASes, i.e., an exterior routing protocol in contrast to interior routing protocols which are used inside ASes. BGP is

a path-vector routing protocol that communicates reachability information between neighboring ASes. BGP neighbors are two BGP speakers (also called "BGP peers"), i.e., ASes (resp. their border routers), having established a stateful BGP session between each other. In a BGP session between the two peers AS_1 and AS_2, where AS_1 is the origin AS, AS_1 sends its prefixes with a set of parameters to AS_2. AS_2 then in return knows how to reach AS_1s prefixes and can communicate its own best path to any other ASes by appending its ASN to the AS path. Simply spoken, if one iterates this process across all ASes, after some convergence time, every AS would have reachability information for all the other directly or indirectly connected ASes and the associated prefixes.

Similar to a street network, on the Internet, one almost always has multiple paths to a destination. Therefore, one of the core features of BGP is to perform Best Path Selection on every BGP speaker, i.e. router, from its local perspective and send their best path to their BGP neighbors[1]. This calculation is based on more than 10 attributes in descending order of relevance. They vary between router vendors. The most common are the following:

Weight[2] , local preference, network or aggregate, shortest AS path length, Interior Gateway Protocol (IGP) or Exterior Gateway Protocol (EGP), lowest Multi Exit Discriminator (MED), eBGP over iBGP, lowest IGP metric, multiple paths, prefer oldest path, lowest router ID, and prefer lowest neighbor IP address.

Recall that BGP is a path-vector protocol, and operators can implement a wide range of policies based on the attributes listed above. For example, one very widely used traffic engineering technique is "AS-Path prepending". As the BGP best path selection algorithm assigns a high weight to the AS path length, operators can artificially lengthen their path by prepending their ASN multiple times. This way, a BGP speaker can lower the preference of, e.g., an expensive link. When peers receive this route, it will unlikely be the best path. This behavior highlights how a feature that was not initially part of the protocol can be added by leveraging existing mechanisms. From just reading the protocol specification, it would not be clear that and to which end such mechanisms are prevalent in the wild. Again, in this respect, measurement studies can broaden our perspective on BGP usage.

2.2 Internet access: The customer perspective

So, how does an end-user connect to the Internet, and what does "connect" refer to in this case? Today, most of us access the Internet daily to consume content. We call this user profile "end-user" or "eyeball". Figure 2.1 illustrates this scenario. Here, end users (or customers from the perspective of an Internet Service Provider (ISP)

[1] In 2016 the IETF defined a BGP extension called "BGP Add Path"[24] which allows the advertisement of multiple paths for the same prefix. This extension is eg used by route collectors to broaden the visibility of the routing infrastructure. As this is an optional extension we will discuss BGP without ADD-PATH unless explicitly stated otherwise.

[2] Cisco proprietary attribute

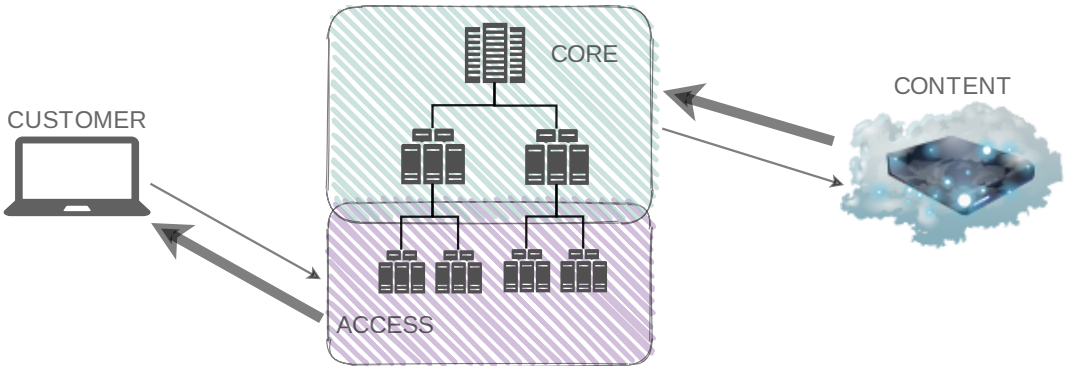

Figure 2.1: Traffic flow in the Internet.

or content provider) connect to the access network of their ISP, traverse the core to access their preferred content. At the borders of an ISPs access network are the handover points to the rest of the Internet and eventually the core. These handover points are called border routers and exchange BGP routes and production traffic with other ASes. An *end user* mainly consumes content and serves only sparely. This behavior is anticipated and reflected in most European ISP contracts where there is a typical ratio of 1 : 10 between upstream traffic and downstream traffic [25]. An end-user is typically a customer of an ISP to get access to the entire Internet. Apart from the typical upstream and downstream ratio, end-users also share a distinct traffic mix on the transport layer. With the success of the World Wide Web (www) and the browser as the most important application today, its supporting protocol, the Hyper Text Transfer Protocol (HTTP), became the dominant protocol in end-user traffic. Nowadays, HTTP is mostly used in combination with some form of transport layer encryption like Secure Socket Layer (SSL) or Transport Layer Security (TLS). With the rise of video streaming, the QUIC protocol also gained much popularity as it is used by, e.g., Google to transport their video traffic. HTTP and QUIC together can make up for up to 90% of end-user traffic [20]. A similar uniformity holds for the ASes a typical end-user contacts where the majority of traffic is served by content delivery network (CDN)s and cloud providers. The end-user manifests a distinct profile that is important to consider when designing and conducting measurements.

In this thesis, we focus on end-user traffic in Chapter 6.

2.3 The Internet from an inter-domain routing perspective

End-users get Internet connectivity via their local ISP. Content providers and all other networks exchange routes for guaranteeing reachability. But how are all those BGP sessions established, and what kind of peering types exist? In the following, we will discuss the Internet as a business model, peering culture as well as describe different relationships between ASes.

2.3.1 The rise of IXPs: Evolution of the Internet business model

As the Internet today reflects complex business and societal relationships, the routing model needs to accommodate them. To understand the different AS relationships, we need to consider the evolution of the Internet graph over time in which economic changes have a direct influence on Internet topology. Specifically, we focus on the rise of Internet Exchange Point (IXP)s as they are a crucial vantage point for the work presented in this thesis. The IP transit market, which underwent similar drastic changes over time, is beyond the scope of this thesis.

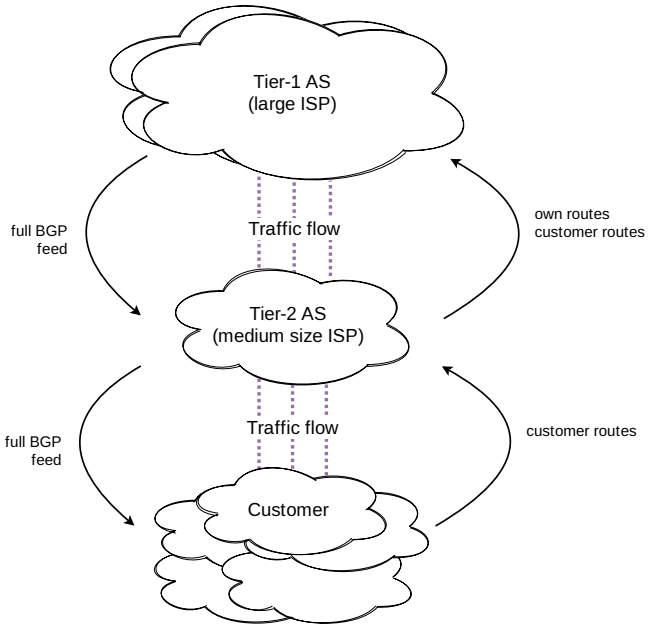

Figure 2.2: Transit relationships from the perspective of a Tier-2 AS.

For a long time, the Internet was organized hierarchically in a Tier model. As Tier-1 ASes one traditionally describes a network that can reach the entire IP address space by its own capacities without any paid peering or transit agreements[3].

Most Tier-1 ASes are large ISPs or carrier who in turn sell transit to Tier-2 ASes via a private interconnect. "Transit" refers to a business agreement as illustrated in Figure 2.2. A Tier-2 AS pays a Tier-1 AS to handle all traffic corresponding to IP address space for which it does not have its own routing information. From a BGP perspective, the Tier-2 AS sends its own and all routes of its customers to the Tier-1 AS and, in turn, receives a BGP full table. Figure 2.3 illustrates the strict Tier model with traditional interconnection fees: A Tier-1 AS sells transit to a Tier-2 AS who then sells to their customers (red arrows). If we depict the Tier hierarchy as a graph, a few Tier-1 networks are the top nodes, with more Tier-2 ASes, e.g.,

[3] As there is no strict definition for a Tier-1 AS, there is no definitive list We refer to the list published on Wikipedia [26]

medium-sized ISPs, on the second level and the majority of networks as Tier-3 or smaller ASes on the lowest level. Typically, money flows from the bottom to the top: Lower tier networks pay a Tier-2 AS for transit, and the Tier-2, in turn, pays a Tier-1. Traditionally, Tier-1 and sometimes Tier-2 ASes exchange traffic among each other for free.

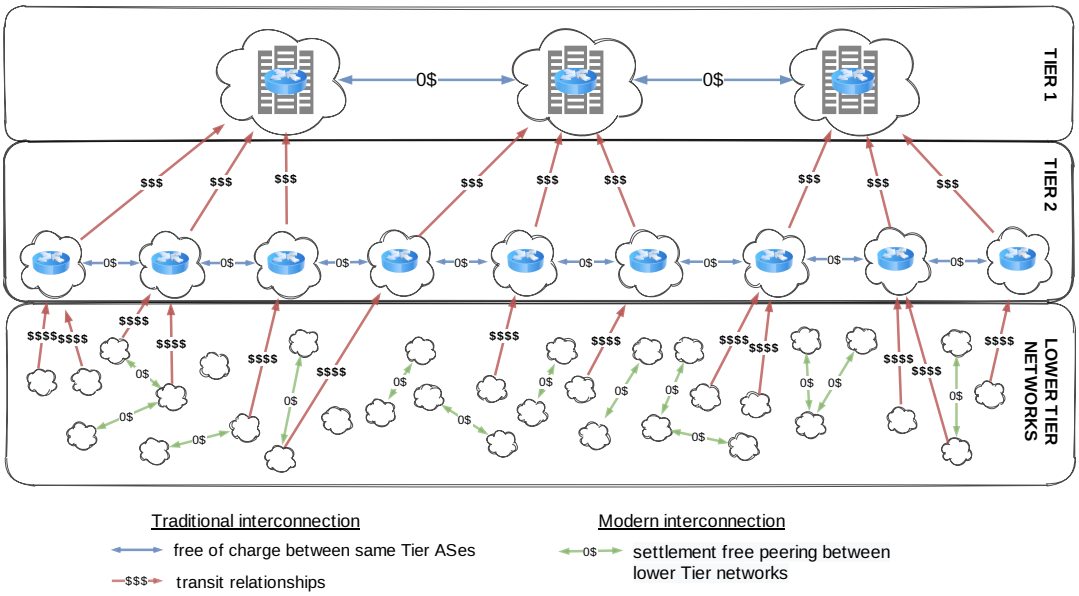

Figure 2.3: Tier model of the Internet with annotated cash flow.

With the growing importance of the Internet, lower Tier networks also started to directly exchange routing information, and as such traffic, without any costs among each other to reduce their transit fees to higher tier ASes as well as double payment[4] as depicted by the green arrows in Figure 2.3. This practice is called "peering". The simplest peering practice is "settlement-free peering" where networks exchange traffic without additional costs, but with negotiated conditions, e.g., traffic ratio or presence at multiple locations. The basic requirement for such traditional peerings is that both networks share a presence, i.e., a BGP speaking router, at the same Point of Presence (PoP). One could imagine that if two networks have routers in neighboring racks in the same data center, it would be simple and cheap for them to peer. In reality, data center providers charge their customers fees for those physical interconnections called "cross connects"[5]. As such, to further reduce interconnection costs, the Internet community came up with the idea of IXPs. While the first IXPs started in the US with data center providers running IXPs, Europe was more open to the idea, and third parties started IXPs on-premise of large data centers.

Fundamentally, an IXP is a layer-2 switching fabric. The IXP operator puts an unmanaged layer-2 switch in a well-connected data center, networks deploy a router

[4]The term "double payment" refers to the situation where, e.g., two Tier-2 ASes both AS_2 and AS_3 pay a Tier-1 AS AS_1 for transit to reach each other.

[5]Prices for cross-connects vary heavily across regions and are subject to the local economic situation

at the same location and connect to the switch. Via this switching fabric, they establish BGP sessions with each other and start peering. Most Internet Exchanges run a route server to accommodate the increasing number of members and efficiently manage route distribution. Members can establish a BGP session to this server and learn all the routes from the BGP of all other members. Route servers drastically reduce the number of individual BGP sessions. The route server also enables new members at the exchange to productively peer from the first day as they receive a large number of routes directly after peering with the route server. Some large ASes do not peer with the route server by policy as the route server does not provide any Service Level Agreement (SLA)s. In this context, SLAs define the level of service two peering partners expect from one another. Typically, they agree on a set of metrics, i.e., traffic ratio or BGP session uptime, and terms what happens if one party does not adhere to the conditions.

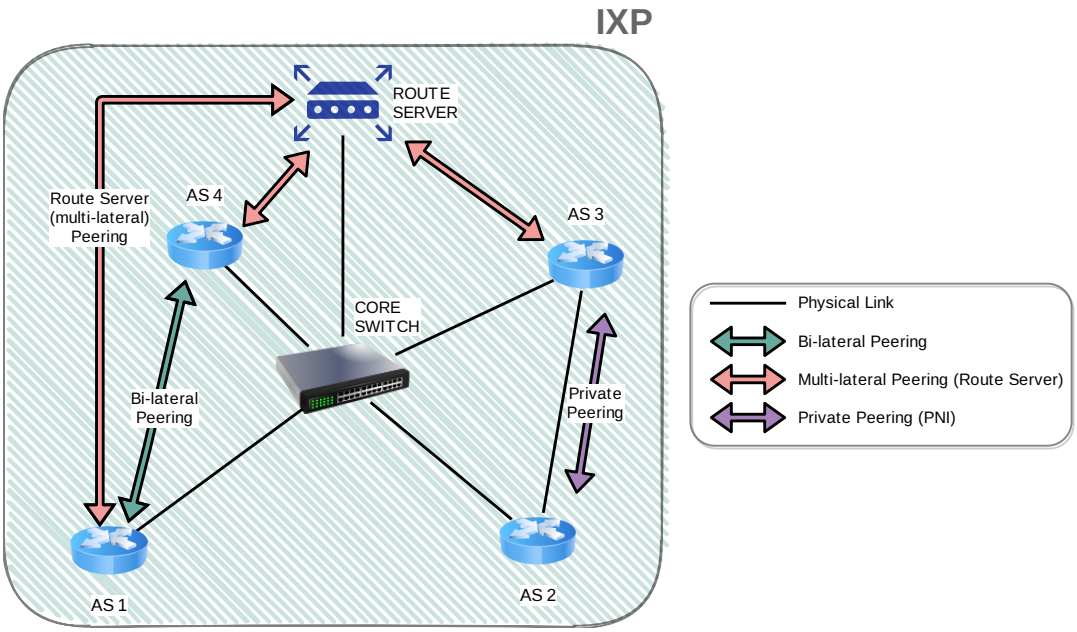

Figure 2.4: Different peering types at an IXP.

While there are no strict definitions, in this thesis, we differentiate between three different types of interconnections on the IXPs infrastructure: (a) bi-lateral (public) peering, (b) multi-lateral peering via the routeserver [27], and (c) Private Network Interconnect (PNI). In Figure 2.4 AS_1 and AS_4 have a bi-lateral peering agreement, i.e., they exchange routes via a BGP session on the public peering fabric of the IXP. AS_1, AS_3 and AS_4 each have a BGP session with the route server, which acts like a route reflector[6]. AS_2 and AS_3 have a private peering via a direct link (PNI and not via the IXP's switching fabric.

[6] iBGP relies on split-horizon for loop prevention which results in a full mesh between all ASes. To reduce the number of BGP connections a route reflector reflects, i.e. redistributes all announcements it receives to all connected BGP speakers and reduces the number of BGP sessions from $n * (n-1)/2$ to n where n is the number of BGP speakers.

Networks usually indicate their willingness to peer at a given PoP, eg as documented in PeeringDB [28]. There are four superclasses of peering policies [29], (a) open: A network with an open peering policy will in principle peer with any other network, (b) selective: In general a network with a selective peering policy will peer with everyone, but adds some conditions (e.g., minimum traffic volume or region), (c) restrictive: This preference indicates the network is not interested in peering except with already existing peers, and (d) no peering: These networks do not peer at all and rely upon transit.

The choice of peering model or IXP is very specific to the individual network. The factors influencing this decision range from technological to political and economical aspects [30]. In the end, it is a cost-benefit calculation weighing the different costs and features, e.g., enhanced network resilience, against each other. The most common factors to consider are transit costs, membership fees at the IXP, equipment cost, the expected amount of traffic in a given peering category, networks present at a given IXP, and colocation fees. This description is, of course, only a high-level description and does not reflect the actual complexity of the situation, which is constantly changing due to new upcoming technologies and business demands, e.g., cloud exchanges [31].

While the first exchanges started as enthusiast projects, the large Internet Exchanges in 2021 are influential business players and move multiple terabits per second of traffic [32, 33]. Moreover, IXPs today offer a wide range of value-added services, like route server, blackholing, or cloud connectivity, to enhance their customers' interconnections. As of 2021 Packet Clearning House (PCH) reports 1015 IXPs worldwide, of which 345 are in Europe[7]. The European Internet Exchange Association (Euro-IX) reports 255 operational IXPs as of 2020 [34] for the European region with steady growth over the last ten years.

Within a few years, this development changed the Internet business model and the Internet graph completely. Until the mid-2000s, the Internet, from a routing perspective, could be described by a strictly hierarchical view as Figure 2.3 shows. With the rise of IXPs and the accompanying peering culture, the Internet graph drastically flattened [35] and as such became even less centralized. At this point, the scientific community needed to update their basic assumptions about traffic flows to the new reality. Again, Internet measurement studies revealed in which way evolution played out.

2.3.2 Beyond the tech: The role of the Internet community

Given the strong economic impact and potential for optimization the peering culture provides, many networks employ dedicated "peering coordinators". Their job is to acquire new peering partners, manage peering agreements and assess which peering partners are of interest-based on the traffic profile of their network. While peerings

[7] The Internet Exchange Directory sometimes lacks accuracy, and for some IXPs are no information available anymore. With that, it is unclear if they are still in operation.

are of critical business relevance, their establishment process is still relatively informal and often conduced in handshake agreements. The Internet community has a long tradition of in-person meetings at various venues on different locality levels to facilitate this culture. Local Network Operators Group (NOG)s meet regularly to exchange experience in running their networks in their specific regions. These local events tighten the community and result in dense peering relationships as well as efficient communication channels. Another typical venue for the settlement of peering agreements are the meetings of the RIRs. They are the organizations that distribute IP address space within their respective service region to the individual ASes: The RIPE NCC for the European and Middle-Eastern region [36], APNIC for the Asia-Pacific region [37], ARIN for the U.S. and Canadian region [38], AFRINIC for the African region [39], and LACNIC for the Latin American region [40]. Their regular meetings provide a venue beyond the local landscape to define policies, exchange experience and gain mutual understanding and collaboration across country borders. A lot of technical background information for this thesis has been gathered at these operator meetings highlighting the importance of the collaboration between research and operational activists. In addition to the non-profit more community-driven venues are also more commercially oriented ones where, eg large IXPs present their specific services to acquire new peers and make their IXP more attractive to potential members [41, 42].

2.4 Vantage points

Given the Internet's design principles we discussed previously, the Internet is inherently decentralized. While this feature, on the one hand, provides a high level of resilience, on the other hand, it makes it impossible to get a single perspective on the complete state of the Internet. Therefore, measurement studies try to approximate this global view by carefully selecting representative data sets or focusing on detailed aspects for which reliable data sources exist that still provide meaningful insights. In both cases, vantage point selection and best practices in data analytics are crucial to interpret the findings.

In the following section, we describe the traffic level vantage points this thesis relies on. First of all, let us consider what makes a good vantage point. To decide that, one needs to know what they are looking for as not every data set can answer all possible questions. Furthermore, every vantage point has its own unique composition of connected networks, position in the Internet graph, hardware configuration, and capturing mechanism. Therefore, the vantage point selection for a particular research question needs to incorporate the following factors: (a) global vs. local study, (b) end-user or business focus, (c) need for traffic data, (d) traffic distribution across networks. For example, a study focusing on service usage by end-users needs a data set from the access perspective to capture end-user traffic. For a detailed discussion of routing information data sets, see Chapter 3.

2.4.1 Flow sampling

All studies presented in this thesis rely on real-world traffic data. This section describes the problem of obtaining this data with a focus on flow sampling techniques. All vantage points we describe in this section handle vast amounts of Internet traffic. We have two principal options available to analyze this traffic: (a) live analysis of all packets as they traverse the vantage point and (b) analysis of captured and stored data. As our studies have a holistic focus that incorporates developments over time, a live analysis does not provide all the information we need. As such, we base our studies on captured and stored traffic information. Given the amount of traffic, e.g., one vantage point transports up to 1Tb/s of peak traffic; it is simply infeasible to dump all the raw packets to disk. Network operators face similar problems while monitoring their networks. Processing this amount of raw traffic is simply impractical. To remedy this issue, the network industry long ago came up with the concept of packet resp. flow sampling. [8].

In general, we call a flow a set of packets with the same 5-tuple, source, and destination port and IP and transport protocol. What denotes a flow is subject to the specific flow collection method and configuration. Still, all packets belonging to a flow need to satisfy common properties based on: packet, transport, or application header fields. All discussed flow collection methods rely on a metering process that exports flow records. As such, the flow collector only exports metadata, e.g., the number of packets and bytes per flow, and not the actual payload. This reduces the amount of stored data and partially addresses potential privacy concerns as no payload is available. To further reduce data size, all flow exporter implement a sampling rate to only consider every 1 out of n packets and then aggregate the exported flows. Trivially, the sampling rate heavily influences the quality of the data set. Consider a small network with very low traffic rates. With a high sampling rate, we may not get a representative data set. When we look at a huge vantage point and consider that network traffic has heavy-tailed properties [43] we can rely on a higher sampling rate without losing representativeness if the intended analysis does not focus on the long tail. For these considerations, a lower sampling rate is advised to capture less frequent patterns. Another issue to consider is flow length. Internet traffic consists of a mix of long-lived and short-lived flows [44]. Short-lived flows are easy to sample as the flow is considered finished by the collector once it stops. Long-lived flows can be order of magnitude larger in terms of flow size [45] than short-lived flows, and if we consider e.g., a large data transfer also have a longer flow duration. In this case, a flow collector has a configurable timeout, after which it considers a flow to be terminated. If it encounters another packet belonging to an already terminated flow, it will create a new one with the same 5-tuple. For an analysis where long-lived flows play a major role, flow stitching, i.e., reconstructing the original flows based on the 5-tuple, is necessary.

[8] In principle, there are two sampling methods, namely, flow sampling and packet sampling. As all vantage points considered in this thesis use flow sampling, we omit packet sampling

Depending on the configuration of the respective flow collectors, different metadata is available, but already with the most common transport header information, we can gain much understanding of a given vantage point, eg number of packets ad bytes per flow, flow duration, transport protocol, application layer protocol, source, and destination IP address, source and destination MAC address, TCP flags, IP version.

2.4.2 Large European Internet Exchange Point

All publications in this thesis use flow data from a collaboration with a large European Internet Exchange Point (LIXP) with more than 900 members and up to 10Tb/s of peak traffic as of July 2021. The data we receive consists of sampled flow data collected at the public peering platform of the LIXP. Recall Figure 2.4 which depicts the three types of interconnection at the IXP. Out of these categories, we can observe only bi-lateral and multi-lateral peering traffic as we do not have access to flows that do not traverse the peering platform, i.e., the core switch in Figure 2.4[9].

All current analyses at the LIXP rely on IP Flow Information Export (IPFIX). IPFIX is a flow export protocol based on NetFlowv9 standardized by the Internet Engineering Taskforce (IETF) [46]. In IPFIX what denotes a flow is subject to configuration, but all packets belonging to a flow need to satisfy common properties based on a template which is sent periodically within the IPFIX stream. In the case of IPFIX, the template can be quite generic and is not per se restricted to network-related data types. As such, one can also gather, e.g., home automation telemetry using IPFIX. As a result, IPFIX flow record generation is very flexible and adaptable to many network topologies. In our case, the IPFIX exporter aggregates packets to flows based on the following criteria: (a) source and destination Media Access Control (MAC) address, (b) source and destination IP address, (c) source and destination transport port, and transport protocol. The source and destination MAC addresses allow it to identify which IXP members are originating and exiting the flow on the IXPs platform. The collector generates a new flow entry if it encounters a packet for which it does not already have an entry with matching flow criteria. If the collector gets a packet for which it already has a flow entry, it increments the number of packets and bytes for the given flow. Upon receiving a packet with `FIN` or `RST` flag set, the exporter considers a flow finished and exports the IPFIX flow to the collector. Alternatively, after 60 seconds have elapsed, the collector considers the flow to be finished. As a result, we can not collect long-living flows directly as the collection process terminates a flow after 60 seconds.

We receive sampled flow data without any payload with a sampling rate of 1 out of 10k collected at the ingress of the e.g., switches of the peering fabric. The traffic is sampled at the e.g., switches rather than at the core as the edge sees more traffic than the core. If two IXP members connect to the same e.g., switch, their traffic will not traverse the core. The exporters then send the IP Flow Information Export (IPFIX)

[9] Note that the depicted core switch is an abstraction for a complex infrastructure of which the exact topology is of no consequence for this explanation.

records via User Datagram Protocol[47] (UDP) to a central collector, which then in turn multiplexes the IPFIX stream to the IXPs monitoring system and to a collector operated by us, which writes it to disk. Due to storage limitations, we can keep one continuous year of flow data. To keep a long-term perspective, we pick the first week of every month and keep the data for a more extended time period.

The sampling is based on ingress traffic as customers may impose blackholing filters on their egress traffic or be the subject of rate-limiting. Especially in the case of blackholing, if we consider only egress traffic, we may miss traffic classes vital to certain studies, e.g., DDoS traffic in the case of security-related research. In general, the less traffic is subject to filtering before sampling, the better for most measurement studies. While it is straightforward to filter out irrelevant traffic, accounting for missing traffic afterward will yield inaccurate results.

We augment the flow data we receive with BGP data from the IXPs routeserver. Every member can establish a BGP session with the routeserver and announce their routes. The route server then performs a shortest prefix match and exports the best paths to all connected members. For all projects, at the IXP we had access to dumps of the route server's routing table as well as `tcpdumps` of the BGP sessions to the routeserver.

The proximity to the Internet core does have a downside. End-user behavior can not be measured reliably at the IXP. Of course, the traffic mix at the IXP contains eyeball traffic, but the IXP is our only vantage point; we can not know if what we observe is representative. We simply do not know how much of the overall amount of end-user traffic traverses the IXP. As such, for all analysis concerning the Internet core, especially business2business relationships, the IXP is an excellent vantage point based on which we can draw solid conclusions. For all research covering networks that are not either directly peering at the IXP or located "near" enough so we can expect their traffic to traverse the IXP, we need additional vantage points.

2.4.3 Large European Internet Service Provider

To do more user-centric studies, like the COVID-19 analysis outlined in Chapter 6, our research group has an ongoing collaboration with a large European Tier-1 ISP that provides service to more than 15 million fixed-line subscribers and operates a transit network. The ISP does not host content delivery servers inside its network, but it has established a large number of peering agreements with all major content delivery and cloud networks at multiple locations. The ISP's view augments the perspective we gain from the IXP. At the ISP, we have access to flow sampled traffic via NetFlow. The ISP operates a Multiprotocol Label Switching (MPLS) network and exports NetFlowv9 traffic statistics at the e.g., of the network on the ingress interfaces of its Label Edge Router (LER)s. As NetFlowv9 is the basis for IPFIX its operational concepts are similar.

3

Route Collectors: A deep dive

Routing-level information is one of the most used data sources in Internet measurement. It allows us to study, eg how routes propagate between different Autonomous System (AS)es in the Internet, which new BGP features are adopted, and how they are used, and how networks use traffic engineering to optimize their traffic flows. In principle, there are multiple ways to obtain routing data. BGP propagates only the best routing path of all those it has received. Therefore, a single measurement point is insufficient to get near a full view of the routing ecosystem. Many networks operate looking glasses that expose their routing tables for debugging purposes, or researchers can collaborate with ISPs or IXPs. Additionally, in 1997 the operations community started to deploy "Route Collectors". Route collectors are dedicated devices that have BGP sessions with routers in multiple ASes, without any exchange of traffic. The route collectors store the received routing information, which is then made publicly available in the form of Routing Information Base (RIB) dumps or files containing BGP updates. Today this infrastructure is the largest source of Internet control plane data and, heavily used by the research community.

In this chapter, we gave an overview of the existing route collector projects and their specific features, differences, and caveats. Additionally, we gave recommendations on how to sanitize real-world routing data for usage in scientific projects.

The following five route collector projects, RIPE/RIS [48], RouteViews [49], Isolario [50], PCH [51], and BGPmon [52], operate route collectors and make their data publicly available. Even though each of those projects operates various, geographically diverse, collectors it is well-known that the Internet topology information derived from this data is incomplete, eg [53–56]. Additional data sources, such as traceroutes and Internet Routing Registry (IRR) data, eg [57–59] can extend the topology. Nevertheless, this does not help uncover root causes for specific routing dynamics, e.g., fully, prefix hijacks, route flap events, path changes, as such events are localized and do not propagate through the whole routing system. Therefore, route collector projects deploy collectors at diverse locations and heavily encourage networks to feed them their full routing views to improve this situation. Real-world routing information has the downside that the data contains a lot of noise due to misconfiguration, e.g., leakage of internal routes, and intentional intervention, e.g., prefix hijacks.

From a measurement perspective route collector projects are essential for BGP routing research as they reveal information about AS-path choices and other BGP attributes,

eg BGP communities. The BGP AS-path attribute has been used to understand the visible AS topology and to infer AS business relationships [60–63] while BGP communities are used to infer infrastructure outages or highlight possible attack vectors [18, 64]. The increase in vantage points raises the question if there are (a) measurable differences between the route collector projects themselves, (b) the number and choice of vantage points, and (c) caveats that the measurement community needs to be aware of.

3.1 Route Collector Projects

Well known route collector projects are: RIPE/RIS [48] and Routeviews [49] which have been operational for over 15 years. Among the upcoming ones are Isolario [50] and Packet Clearning House (PCH) [51]. A route collector project with a focus on live analysis is BGPMon [52].

RIPE/RIS: The RIPE Routing Information System [48] offers full BGP Routing Information Base (RIBs) snapshots, archived every 8 hours, as well as BGP updates from a multitude of BGP peering sessions from 20 active collectors. RIPE operates one multi-hop enabled route collector (rrc0).

RouteViews: Routeviews [49] operated by the University of Oregon includes 23 route collectors as of April 2019. Routeviews stores the routing table dumps every two hours and makes them and the updates publicly available. RV operates five multi-hop enabled route collectors.

PCH: PCH [51] strongly advocates Internet eXchange Points (IXPs) and is a provider of route-servers at major exchange points worldwide. One specialty of the PCH platform is that it maintains route collectors that peer with these route servers at about 180 different IXPs around the Globe (status April 2018). Route servers are typically a value-added service of the IXP that collect routing information in a centralized manner and redistribute it to connected member routers. Thus, PCH offers BGP routing information for many of the IXP members [65].

Isolario: "Isolario offers network operators services based on real-time analysis of inter-domain routing from different points of view in return for full routing tables, following the do ut des principle." All Isolario collectors use multi-hop BGP sessions with the BGP extension `add-path`. It includes 4 collectors with 349 feeders, of which 33 use add-path. Isolario archives the RIBs every two hours and makes them as well as the updates publicly available.[1]

BGPmon: . BGPmon's goal is to help network operators to assess the health of their network via live analysis. It offers only BGP updates from 15 feeders via XML.

[1] Note, that Isolario has announced the discontinuation of its services as of Dec 30th, 2021.

Project	#Feeder	#Feeder AS	#AS	#Prefixes IPv4	#Prefixes IPv6
RIPE	674	335	63575	966K	83K
RV	539	229	63400	1,5M	72K
Isolario	349	108	63798	1.1M	87K
Union	1,562	576	64108	1.8M	101K

Table 3.1: Overview of BGP dataset for May 2018. IPv4 prefixes account for 18.2% of the total prefixes and 25.4% of the observed ASes are IPv6 enabled.

3.1.1 Routing Information: Data selection

While PCH provides its updates in MRT format, they provide their RIB snapshots only as text dumps based on the "sh ip bgp" command. This format does not preserve any information about the actual feeder of a route. We were also not able to infer the feeder based on the BGP path; since PCH's route collectors mainly *peer* at IXPs, it is unclear whether the last AS in the path is a Route Server, the actual feeder, or some other entity along the path between the actual feeder and the route collector. Examples for such other entries include data centers, e.g., WOWRACK [66], that host a colocation facility in which the actual feeder might participate, and that might append their own ASN to the AS path. Thus, we would only be able to utilize PCH's update messages.

Arguably, it is possible to rebuild a RIB snapshot based on BGP update messages. However, it is unclear how many updates, especially across which timeframe (a stable route could have been announced five years ago), need to be assembled to approximate a RIB. Thus, we exclude both BGPmon and PCH to ensure compatibility of the results. In addition, we find accessing the PCH data to be a challenging task since PCH provides routing updates in separate files for every minute and every collector. This results in more than 240k files per day. Thus, all PCH updates for May 2018 reside in 7,324,719 individual files. The project was kind enough to give us access to a particular download server. Still, they asked us to rate-limit our downloads to not overwhelm their service. They suggested shipping disks to transfer the data for more extended periods–a rather "scalable" method.

Thus, we consider the table dumps and all update messages announcing prefixes from RIPE/RIS, Routeviews, and Isolario from 1. to 31. May 2018. We picked this month as it seemed to contain the least number of route collector failures, which could potentially lead to missing or partial routing information. Table 3.1 gives a summary of their key properties.

3.1.1.1 Route Collector Data: Comparing BGP Views

Since the goal of this work is to show the impact of different BGP data sets, we, in this section, identify features that may impact BGP analysis and introduce simple metrics to compare data sets.

By default, a BGP speaker sends its best routes filtered by its egress routing policy to its peers. As such, BGP is an information hiding protocol. With the BGP `add-path` extension [67] a BGP speaker can send multiple paths for the same prefix distinguished by a path identifier. Isolario supports BGP add-path and, since it may increase the number of received routes significantly, we consider this feature.

To compare different BGP views on an IP address space level, we use the IP address space "seen" by a route collector project or individual route collector. Note, we use the entire IP address space as a baseline instead of the routable space which excludes martian and bogon addresses [68].

Even though we agree with the statement by Roughan et al. [69]: "BGP was not designed with AS-level topology discovery/mapping in mind, (...), BGP has to hide information that would otherwise aid topology discovery." we point out that deriving AS links from BGP data which started with the work of Govindan and Reddy [70] is now a common technique: eg [53, 55, 60, 61]. So we use the number of observed AS links as the second metric.

Our baseline is the total number of unique AS links observed across all route collector projects. Since ground truth, i.e., all available AS links reported by each network operator is not feasible to acquire; this is the best baseline we have. While it is possible to discover and confirm additional links by using data sources like traceroutes or the IRR databases we decided against this. Traceroute campaigns (a) have their own artifacts as traceroute path to AS links mapping is problematic due to IP aliasing and uncertainty about which IP is used to reply to a probe, eg [71, 72] and (b) it is doubtful that these campaigns cover the whole AS topology, eg [69]. For details on challenges using the IRR databases, see Section 3.2.

To give more details for some analyses we classified the AS links in terms of link type, i.e., *provider-customer (p2c)* and *peer-peer (p2p)*. To infer these relationships, we use the algorithm proposed by Luckie et al. [60] on our baseline of all observed AS links. Additionally, we classify the ASes according to their level in the Tier hierarchy. As there is no formal authority for either the list of tier 1 networks or the overall classification, we obtained a list of "tier-1" networks from Wikipedia [73]. Even though we, in general, do not consider Wikipedia as a trustworthy source, this list seldom changed in the last 500 revisions (going back till 2012). Every ASN appeared in at least 495 revisions. Thus, the list is (a) continuously monitored by the community, (b) stable, and (c) does not rely on any inference. Moreover, when using the classification algorithm on the baseline AS links, it yields a tier-1 clique that is a subset of Wikipedia's tier-1 list. We then use a simple heuristic [74] to classify the other ASes. We iterate across the unclassified ASes. For each AS, we classify it as tier X+1 if it has an AS-link to an AS from tier X. Afterward, we group all ASes from tiers larger than four into tier 4.

Sanitation	# Prefixes IPv4	# Prefixes IPv6	IP space IPv4	IP space IPv6	AS links IPv4	AS links IPv6
None	1,842,784	101,429	100.0%	100.0%	100% (=461,255)	100% = 167,948
Remove loops	1,842,678	101,416	100.0%	100.0%	- 0.42%	- 0.97%
Remove bogons	1,803,088	101,375	100.0%	12.5%	- 0.02%	- 0.02%
Remove > /8	1,842,775	-	70.2%	-	- 0.00%	-
Remove < /24	938,888	58,844	100.0%	100.0%	- 1.16%	- 0.19%
Remove internal < /24	1,753,598	-	100.0%	100.0%	- 0.00%	-
Recommended	1,802,976	101,362	68.7%	12.5%	- 0.44% (=2036)	- 2.31% (=3896)

Table 3.2: Impact of sanitation steps on # prefixes, IP space visible, and AS links in terms of % removed by each step.

3.2 Route Collector Data Sanitation

Like any other measurement, BGP information is noisy. It may contain both un- as well as intentional artifacts. Thus, a common approach is to sanitize the data and to remove the noise at least partially. By reviewing papers using BGP data published at PAM, IMC, NSDI, and SIGCOMM since 2013, we compiled a list of possible sanitation steps, which we review in this section. We conclude by recommending to use at least the following sanitation steps: no loops, no bogons, nothing > 8, some filters for < 24s and handling of AS-sets. See Table 3.2 for a summary.

Routing loops: While BGP itself does not allow path cycles [75] they can occur naturally during transient states (while BGP converges) [76, 77], or unintended due to human error [78], e.g., while using AS path prepending [79]. Note, some ASes may even announce loops intentionally, e.g., to force certain policy objectives [78]. Using AS_TRANS may also yield valid routing loops. AS_TRANS is an "artificial" AS to help commodity hardware handle 32-bit ASNs. Within all ribs of the 1st Mai 2018, 16:00 UTC, we find that 29K paths contain routing loops of which only one was due to AS_TRANS. Given the small number of "genuine" routing loops and the fact that BGP once converged disallows loops, we suggest, similar to [60, 80], *filtering loops*.

Non-stable routes: Various abnormalities can change routing, eg BGP hijacks, e.g., [81–85], and route leaks, eg [86]. Typically, operators will detect those after "some" time, inform their peers, e.g., via operator mailing lists, which then use various countermeasures to remove such routing abnormalities. Since such abnormalities can skew results, some studies remove (a) all non-stable routing information, e.g., available for less than 2 days, eg [54, 87] or using ad-hoc strategies, eg [88], or non-stable and weird [89], whereby weird refers to other abnormal path features, e.g., extra-long AS paths. While this is likely to filter transient events, it also removes many other routing events, e.g., those related to short outages. Moreover, the approach is unable to handle longer-lasting misconfigurations. Thus, we suggest to *not* use this approach.

Martians: Martian prefixes are prefixes that IANA has allocated for special purposes [90] and include, e.g., private address space [91], loopback addresses [92], and multicast addresses [93]. Martians are static, well documented, and well known. Moreover, they should never be announced publicly. However, over the entire month, we see 35K routes involving Martians. Thus, similar to [80, 94, 95], we suggest *filtering martians*.

Fullbogons: Fullbogon prefixes include Martians as well as prefixes that while allocated from IANA to a Regional Internet Registry (RIR) have not yet been assigned to a Local Internet Registry (LIR) or another organization. Again, none of these prefixes should ever be announced publicly. However, 113K IPv4 routes with fullbogons exist, e.g., due to misconfigurations or cluelessness. Thus, similar to [80, 94, 95], we suggest *filtering fullbogons*.

Yet, different sources for bogon information exist, including the RIRs themselves, which publicly provide status information on their address blocks, as well as Team Cymru [68] which offers a fullbogon reference. The latter augments the RIR information and is less restrictive than the RIR information[2]. Therefore, we suggest using the less restrictive lists, i.e., the one from Team Cymru. Moreover, as the lists are only published once a day, we recommend updating them daily and use the intersection of the two lists from midnight of the same and the next day. This is what we do in this study.

When applying filter lists such as bogons to the BGP data, there are three options. Let us consider a filter list entry: p/24 and that the BGP data includes announcements for both p/24 and the covering prefix p/23. The first option eliminates p/24. The second one eliminates both: p/24 and p/23. The theoretical third one eliminates both but adds the "valid" second /24. Since this tampers with the data, it should not be done. This work uses option 1–filter only exact matches–since we do not want to discard announcements for valid IP address space. The implication is that we may keep prefixes that partially cover bogon addresses. Thus, the observed IP space may not reduce to the whole IP address space minus bogon space, explaining why the % IP space metric, see Table 3.2, is still 100.0% even after filtering bogons.

Too unspecific prefixes: The route collector projects state that any feeder should send their full view of the "Internet" rather than their default route. In reality, this is *not* the case. Traditionally, the largest allocation of prefixes to organizations is for a /8 (former class-A network). Thus, announcements for less specifics than /8s are close to default routes and should be filtered[3]. Still, we see 189 routes involving too unspecific prefixes, of which 80 are for the default route. Thus, similar to [94, 96, 97], we suggest *filtering for less specifics*, i.e., prefixes that are less specific than /8. To understand how many of the observed 313K too specific prefixes with 641K routes are due to blackholing. We leverage previous efforts by Giotsas et al. [94]. However, this only justifies a small fraction: 6K prefixes and 13K routes.

[2] This is in part motivated by the common presumption that RIRs often lack behind in their documentation.

[3] We do not know of any valid /7. Note, an accidental announcement of a /7 by former Worldcom in 2002 choked some routers.

Another reason for too specifics is leaked internal infrastructure address space or leaked more specifics of customer address space [98]. Possible reasons are misconfigured outbound filters by the feeder, which causes it to send its complete internal BGP data. We find 93K such prefixes with 93K routes with AS path length one. Of these, 89K are only revealed to the route collector—we do not see any other announcement for this prefix. Thus, these are likely unintentionally leaked internal subnets. Other reasons for using too specifics are traffic engineering, BGP multi-homing, iBGP policies [99], or perceived increased security. For traffic engineering, BGP multi-homed customer networks often announce a more specific, e.g., a /25, for fine-grained control and the covering prefix to ensure reachability. We obtained ground truth from one ISP "leaking" many more specifics. The policy of this ISP is to reannounce all prefixes of its customers listed in the IRR. The customer, in return, added prefixes more specific than /24 to the IRR databases. Thus, the announcements are due to a policy that a customer (ab)used.

We suggest *not filtering more specifics per se* as, eg [96, 97], but to *filtering leaked internals*. These can be identified by excluding routes for more specifics with AS path length of one for which there is no other route.

Research prefixes: To understand the routing ecosystem better, various routing research projects emerged, which enable researchers to inject routes actively. The most prominent ones are the RIPE BGP beacon project [100] and PEERING [101]. BGP beacons periodically announce and withdraw prefixes to study the route propagation. As such, their dynamics are artificial but do not cause false routing information. Other studies introduce false information on purpose, e.g., in the form of path poisoning [102]. Since PEERING is available, such studies often use this infrastructure [103]. Therefore, we suggest *filtering research prefixes selectively*—only filtering those that can easily be used to inject false routing information rather than all as proposed by, eg [80].

AS sets: AS paths consist of two different elements: AS sequences and AS sets. While AS sequences are strictly ordered, AS sets do not imply any order. Thus, it is usually not possible to use AS sets to infer AS links when looking at AS topology. Only when the AS set consists of a single AS can we replace the AS set with the AS and use the resulting sequence to infer AS links. If the set contains multiple entries, the remaining AS sequences can be used to infer AS links. Within all ribs of the 1st Mai 2018, 16:00 UTC, we see 74K unique paths that contain AS sets. Thus, unlike [60], we *suggest not removing entire paths with AS sets.*

Poisoned AS paths: Path Poisoning [104] implies that an AS adds another ASN (not his own) to its own advertisement. Reasons for path poisoning include (a) preventing a route from being redistributed by a certain provider [102] or (b) appearing "more important" in some widely-used web interface such as `bgp.he.net` or CAIDA's AS Rank [105]. Due to the absence of ground-truth data, detecting poisoned paths is non-trivial.

One possible data source are the IRR databases. Within the IRR, every AS should, in principle, publish who their customers are, with whom they are peering, and what

their inbound and outbound filters are. Unfortunately, different ASes maintain this information at different granularity and accuracy. For a quick check, we extracted a list of ASes that updated their IRR entry on April 7th 2019. Then, we contacted their administrator (as listed in the IRR) to check if the AS uses at least one explicit policy per neighbor and if this list is updated daily (if needed). This resulted in a list of 14 ASes. We then cross-checked the IRR information against the data from the route collectors from April 8th. To our surprise, we find close matches, in terms of neighbor ASes, for only two ASes. Others had many more neighbors than expected—up to 500% for one AS. Moreover, roughly 50% of the expected neighbors according to IRR were not visible in public BGP data. Thus, at this point, we recommend not rely on IRR data.

Previous work suggested removing paths containing three or more "Tier-1" networks [60] due to violating the valley-free assumption, see [106]. Using our list of "Tier-1" networks, we find 1.5M (1.05%) AS paths with three or more "Tier-1"s. However, when accounting for sibling ASes[4] almost none (only 3.3%) still involve more than two tier-1 entities. Thus, we suggest *filtering AS paths with at least 3 Tier-1 ASes only if the ASes belong to three different organizations.*

Unassigned ASNs: Similar to bogus prefixes, there are also not yet allocated or assigned ASNs. In principle, announcements with such ASes should be removed. However, only ARIN distinguishes between assigned and allocated. Using "allocated" as filter flags, we would classify 1.7K (2.7%) ASNs with 1.48M (1%) AS-Paths as bogus—far exceeding our expectations. Close to 50% of these paths involve private ASes and, thus, are likely due to leaked internals, e.g., from using BGP federations. Such private ASes should be removed from the path, or the announcements should be filtered. For the other 50% of the paths with unassigned ASNs, the situation is complicated as RIR documentation can lack behind. Thus, unlike [60], we do *not suggest filtering unassigned ASes* until a more reliable data source is available.

AS hops artifacts: Sometimes, measurement or collection infrastructures insert fictional ASes in the AS path even though they are not observed in the data plane. Del Fiore et al. [107] observe this for the Clemson vantage point as recorded by PEERING [103]. At this point, we *do not yet suggest filtering them* unless the study is vulnerable to this artifact.

3.2.1 Route Collector Data Sanitation: Summary

To encourage reproducibility, every methodology MUST document which sanitation steps are used since missing or wrongly applied sanitation can change any analysis. Yet, this is not always the case, eg [18, 108–111]. Thus, naively, one may assume no sanitation, but there is no certainty.

Next, every sanitation strategy has to match the goal. Trivially, if a study quantifies the number of bogon prefix announcements, it cannot filter them. Similarly,

[4] Siblings are ASes operated by the same entity.

Project	# colls.	IP space IPv4	IP space IPv6	AS links IPv4	AS links IPv6
RIPE	20	68.0%	12.5%	65.4%	69.0%
RV	22	67.7%	12.5%	67.5%	49.1%
Isolario	4	68.4%	12.5%	58.2%	59.0%
RIPE + RV	42	68.0%	12.5%	86.0%	80.2%
RIPE + Iso	24	68.7%	12.5%	82.1%	91.7%
RV + Iso	26	68.5%	12.5%	87.5%	81.5%
Multi	8	68.4%	12.5%	73.9%	65.2%
No-multi	38	68.0%	12.5%	72.2%	73.8%
All	46	68.7%	12.5%	100% =459.219	100% =164.052

Table 3.3: Base metrics: All route collector projects for May 2018 with recommended sanitation.

some studies may need to exclude IXP ASNs [87] or explicitly handle siblings [60]. Projects directly working with the collector feed, e.g., BGPStream [112], should not per-se apply sanitation as this alters the data. Other examples where sanitation may be optional are cases where the data is used to manually validate other measurements [113–116]. Nevertheless, this should always be stated explicitly.

We also find that many (> 20) papers rely on either the pfx2AS mapping [117] or the AS relationships [60, 118] and customer cone data provided by CAIDA. While it is good practice to reuse existing data sets, users should be aware that these are only based on RouteViews and RouteViews plus RIPE/RIS, respectively. Some studies use other route collector subsets, e.g., based on the availability of related resources, such as Atlas probes close to the collectors [119], or after careful analysis, eg [120]. Even others do not justify their selection, eg [87].

For studies aiming at uncovering genuine routing phenomena in the routing ecosystem, we recommend the following sanitation steps: Removal of

1. announcements of bogon & martian prefixes;
2. internally leaked routes, e.g., more specific routes than /24 leaked to the route collectors;
3. prefixes shorter than /8;
4. AS loops;
5. AS sets while keeping path sequences and reconstructing path with sets of size one.

We also suggest removing unassigned ASNs once a validated data source is available.

3.3 Route collector ecosystem

Here, we use the % IP space and % AS links metrics to look for biases after applying sanitation, see Table 3.3. While all route collector projects cover 100% of the % IP space before sanitation, none covers more than 68.4% afterward; all cover roughly the same amount, 67.7% to 68.4%. This corresponds to 85.6% of the routable address space which is roughly the IP space announced by a full BGP feed and is consistent with previous results [96].

Table 3.3 shows that all route collector projects, individually, as well as all pairs, provide only partial views, 65-88%, of % AS links. An obvious conclusion is that for maximum AS link coverage, one should include all projects. There are substantial differences between the route collector projects. RIPE and RouteViews uncover similar percentages of IPv4 links with a comparable number of route collectors. RIPE sees more of the IPv6 AS topology than RouteViews, which may be a result of RIPE having a larger collector deployment in IPv6 heavy regions. Isolario, with only four route collectors, uncovers more than 58% of the visible IPv4 and IPv6 AS links. For possible explanations see below, Subsection 3.3.2. For the rest of the chapter, we focus on IPv4 rather than IPv6 as the overall results are similar and it still is the dominant Internet protocol.

3.3.1 Picking collector subsets

Next, we consider various subsets of route collectors because, if one can get similar results with less data, since this conserves computer, eneg and people resources. Moreover, route collector stability differs greatly according to our experience. While some have consistent data across multiple years, others have significant gaps in their data. Thus, we ask if choosing route collector subsets biases the results, in particular, the visibility of the AS topology. Thus, we pick 5.000 random subsets of different sizes and compute their % AS links metric.

Figure 3.1 shows the results for subsets of 2, 5, 10, and 30 collectors, as empirical cumulative distributions (ECDFs) across the % AS links computed for each randomly chosen subset. We see a huge spread in the ECDFs for each of the collector sets. Indeed, when choosing 10 collectors the spread can be from 40% to 70%, a difference of more than 30% which means that a small number of collectors can uncover a large fraction of the topology. However, other combinations reveal much smaller subsets. The top route collectors are Korriban (Isolario) with 227K AS links, rv3 (RouteViews) with 190K AS links, and Naboo (Isolario) with 184K AS links. Missing or hitting these “important” collectors for a random subset is the reason for the large spread.

Moving from subsets of 2 to 5 route collectors, resp. 5 to 10, 10 to 20, 20 to 30 collectors, the number of uncovered AS links doubles at best. For most subsets, the impact is much smaller and decreases as we move to larger subsets. This is not surprising given that the total number of collectors is 46 and that many AS links are visible to multiple collectors. To confirm this Figure 3.2a shows a histogram of the

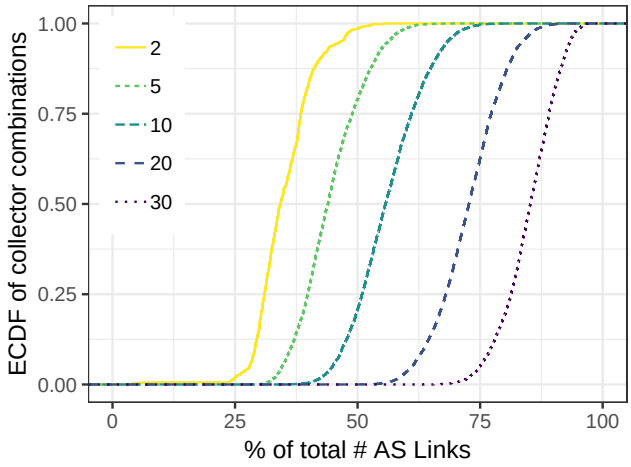

Figure 3.1: ECDF of % AS links: route collectors sets.

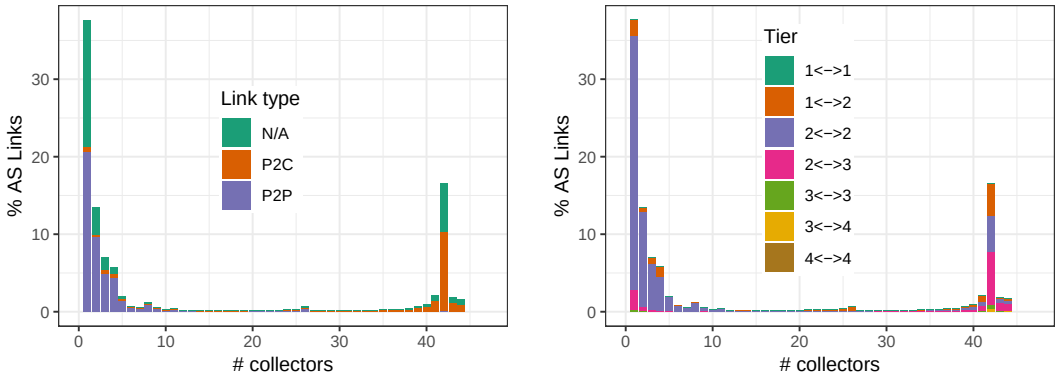

(a) Histogram by Link Type.

(b) Histogram by AS hierarchy.

Figure 3.2: IPv4: Histogram of uncovered AS-links per route collector.

number of different collectors that see an AS link. The distribution is heavily skewed. More than 35% of the AS links are seen by only a *single* collector and 60% are seen by less than five collectors. However, more than 20% of the links are visible to more than 40 collectors.

We colored the histogram by the types of links to distinguish between P2C, P2P, and unclassified links. Popular links (seen by many collectors) are P2C links, while unpopular links are mostly P2P links. If P2P links are essential for an analysis restricting the analysis to a random subset of the route collectors may be problematic. We also see that more than 25% of the links are unclassified because the algorithm by Luckie et al. uses extensive sanitation. Figure 3.2b shows the same histogram but this time colored by the AS hierarchy tier of the links. Popular links are, indeed, either peering links on tier-2 or C2P links going from $1 \rightarrow 2$ or $2 \rightarrow 3$. Most unpopular links are peering links. However, some customer provider links across the AS hierarchy also fall into this category.

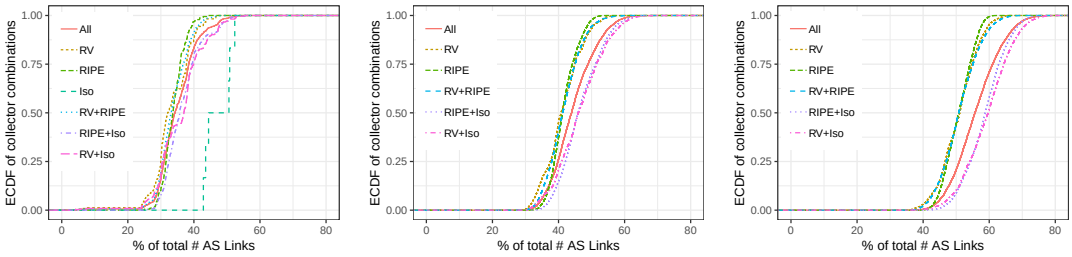

(a) Collector sets of size 2. (b) Collector sets of size 5. (c) Collector sets of size 10.

Figure 3.3: ECDF of % AS links for different route collector sets per route collector project.

3.3.2 Route collector project impact

So far, Figure 3.2 has given us an overview of how and, in part, why visibility changes as the number of route collectors increases. However, it does not address the question of bias by route collector project or geography. In Figure 3.3 we restrict the route collector combinations to a specific project. Figure 3.3a confirms the initial takeaways from Table 3.3. We find that the coverage of RouteViews is larger than that of RIPE. Still, choosing a (un)-favorable combination of two collectors from RouteViews may uncover less/more AS links than RIPE. Yet, any two of Isolario's collectors[5] uncover more than 40% of the AS links the maximum for the other two projects.

Given that the four Isolario route collectors only have 349 feeders compared to 674 feeders for RIPE and 539 feeders for RouteViews, this is somewhat surprising. One explanation is that Isolario supports add-path on all of its route collectors. 33 of Isolario's feeders are using add-path on two of Isolario's collectors, Korriban and Naboo. At the AS-level, 11 ASes have both add-path as well as non-add-path peerings. While add-path is beneficial, feeders send multiple paths for the same prefix, which, e.g., can increase the number of uncovered AS links, we also lose information, namely, which of the path is its best path [67].

To check the impact of add-path, we determined how many different paths we learn via add-path. We find that 40.6% of all Isolario's unique paths are learned only via the add-path extension. This is a considerable percentage. We also find that for feeders with add-path, the % AS links metric is 46.5% compared to 48.9% for all feeders without add-path. This indicates the power of add-path. Still, the overlap among the AS links is 64.0%.

While checking the impact of add-path, we identified an "outlier", namely, NLNOG-RING [121]. This project deploys servers at many different ASes to enable operators to execute various network debugging tools. It also encourages every AS to peer with them. NLNOG-RING then forwards, via BGP add-path, all its BGP data to the route collector Korriban of the Isolario project. The two feeders of this project are responsible for 69.3% of all paths at Korriban (34.5% of Isolario), and by itself,

[5]Since there are only four collectors there are only six possible combinations which explains the steps in the CDF.

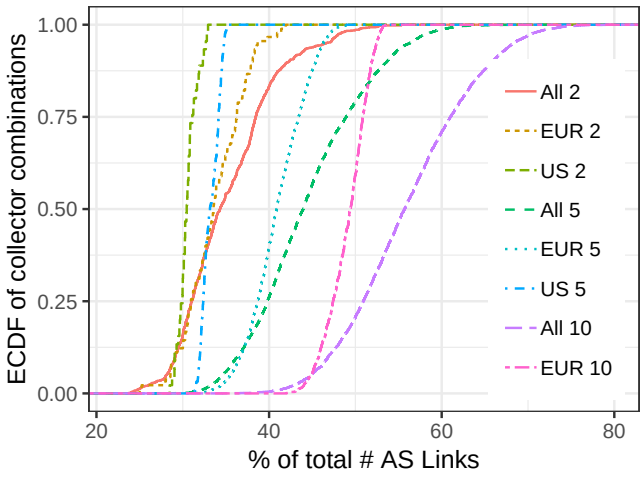

Figure 3.4: ECDF of % AS links for subsets of EU, US, and all route collectors.

the project already has % AS links metric of 44%. To determine if this feeder offers only duplicate information in comparison to all other feeders of Isolario, we check for matching AS-paths after removing the NLNOG-RING-AS as the first hop. The result showed almost no overlap (3.5%). Given this huge number of paths, we next, check if this feeder increases the % AS links metric substantially. This is only partially the case as adding NLNOG increase the metric from 50.2% to 58.2%.

Another explanation for the excellent visibility of the Isolario collectors is that they are BGP multi-hop enabled. Multi-hop simplifies peering and, therefore, these route collectors may have a wider feeder spread. RIPE and RouteViews also operate route collectors that support multi-hop. To further study this effect we summarize all route collectors that support BGP multi-hop under "multi" and those that do not under "no-multi". Interestingly, "multi" with only egt feeders uncovers slightly more of the IP address space as well as the visible AS links, see 3.3.

Moving to larger subsets of route collectors, i.e., from Figure 3.3a to 3.3b we see that more AS links are uncovered but not 2.5 times as many. Interestingly, the difference between RIPE, RouteViews, and RIPE + RouteViews decreases. While the plot does not include Isolario in isolation (it only has 4 collectors with % AS links metric 58%), we see its impact by comparing *all route collectors* vs. *RIPE + RouteViews* and estimate it as less than 7.5%. However, RIPE or RouteViews with Isolario does better. Here, the chances of picking a route collector from Isolario is larger than if we consider all route collectors. This is also, the reason why the curve for "all" is to the left of, e.g., "RV + Iso". Moving to sets of 10 collectors, see Figure 3.3c, these effects are even more pronounced. Isolario again adds diversity and, thus, the line for *all route collectors* is significantly more to the right than the one for the other two projects.

Region	# collectors	IP space IPv4	IP space IPv6	AS links IPv4	AS links IPv6
EU+US	24	67.6%	1.6%	57.6%	51.0%
EU+US+AFR	28	67.6%	12.5%	59.9%	53.8%
EU+US+SA	28	68.0%	1.6%	68.5%	66.6%
EU+US+AUS	26	67.6%	1.6%	58.5%	52.6%
EU+US+AS	28	67.6%	1.6%	58.5%	52.9%

Table 3.4: Impact of none US or European service region.

3.3.3 Geographic diversity

Next, we consider how geographic diversity impacts the % AS links metric. We use the location of each route collector and the knowleg that only a few route collectors support BGP multi-hop (less than 10 in total including all route collectors from Isolario). Thus, most BGP peers are only one IP hop away, which, for most peers implies that they are within The close physical proximity. We group the route collectors into three groups: US for North America, EU for the RIPE NCC service region, and others which we split into Africa, Asia, Australia, and South America. US/EU/Other includes 3/12/4 route collectors of RIPE and 7/2/10 route collectors of RouteViews. This skew is likely a reflection of the regional origin/focus of RIPE, which is EU, and RouteViews which is US.

We see that each region by itself has a limited ability to uncover more than 50% of all visible AS links. Combining EU and US (see Table 3.4) uncovers less than 10% more links (57.6%) than each one, individually. Interestingly, adding route collectors, not in the US or European service region, shows different effects. Adding the 4 African route collectors, the two Australian ones, or the four Asian ones only add a few percent to the % AS links metric. However, the four route collectors in South America add close to 10% to the % AS links metric because South America includes, e.g., multiple route collectors in Sao Paulo with many distinct feeders.

To further explore the impact of geographic diversity when picking only a small subset of the route collectors Figure 3.4 shows the ECDFs of the % AS links metrics for 5000 randomly chosen subsets of 2, 5, and ten collectors from the regions EU and US. We include–for comparison–the ECDFs if we choose from all route collectors. We again see a huge impact of regional focus. Most regional ECDFs (for the same size subsets) are shifted to the left. For subsets of two the US region, in particular, does not do well while EU is slightly and is well outperformed by “all”. The benefit of moving to 5 route collectors is very small for the US adding well less than 5% which is nowhere close to the 10% we see in general. Possible explanations include that different route collectors in the US receive data from the same feeders; thus, the observable link diversity is limited. The European subset does better but still worse than picking from all. This is even more pronounced for subsets of size ten. This demonstrates that geographic diversity does matter.

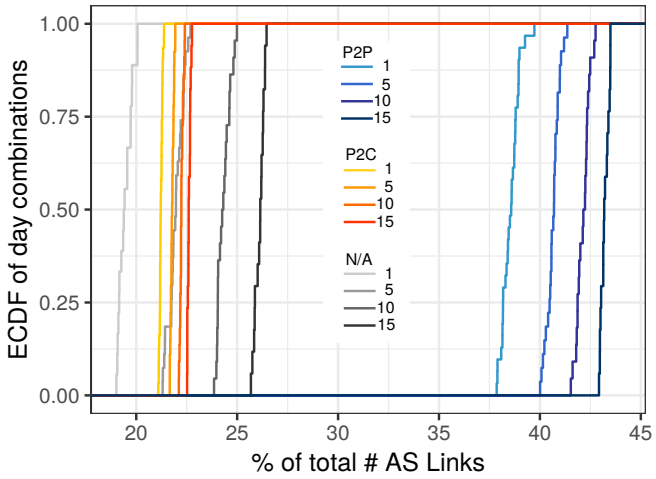

Figure 3.5: Impact of different time window sizes on % AS links by link type (P2P, C2P, N/A).

3.3.4 Time granularity impact

Next, we analyze the effect of time periods by focusing on different time durations, namely, 1, 5, 10, and 15 days within May–using the same month for comparability. We generate different subsets of BGP data using a sliding window approach. For a duration of n days, the start of the window ranges from the 1st ... 31th-n. Overall, we find that the % AS links metric increases from 78% to 82% and 88% to 92% while the variability decreases as we move from 1 to 15 days. The reason is that more links are unveiled due to various routing events. At the same time, the overlap in the input data increases as we move from 1 to 15 days reducing variability.

Figure 3.5 highlights the impact of the measurement duration on the % AS links metric segregated by link type by again plotting the ECDFs of the % AS links metric for different data sets. First, we see that the variability within each period is limited and decreases as we move to longer time durations. However, more interestingly, we do not unveil as many additional C2P as we unveil P2P links as the measurement periods increases. This highlights that to uncover additional P2P links, we need routing events which force path changes, e.g., to a backup link. Alternatively, additional vantage points can help. Moreover, we see that the number of unclassified AS links (N/A) increases substantially from less than 20% for a single day to more than 26% for 15 days. Recall that CAIDA's AS relationship inference uses extensive sanitation which eliminates many paths and may be too conservative, recall Section 3.2. The other reason is an increase in bogus information due to the various possible routing mishaps mentioned earlier.

3.4 Route Collectors: Summary and Recommendations

This chapter presented sanitation strategies for real-world route collector data and multiple comparative analyses between different route collector subsets. In the following, we summarize our findings and recommendations for future research projects and route collector deployment strategies.

Sanitation: First of all, we highlight that it is vital to follow proper data sanitation steps and familiarize oneself with the BGP data. Here, we highlight the following aspects: (a) Documentation can lack behind or may not highlight essential features, e.g., multi-hop collectors at RIPE/RIS, (b) the unexpected impact of the NLRING at Isolario, or (c) the difficulty of even downloading the PCH data.

Choosing collector subsets: Next, various route collectors can uncover different parts of the Internet topology. However, not all route collectors are equal. Thus, if one has limited resources, we recommend picking route collectors that (a) are from a diverse regional set, (b) that use BGP add-path, (c) that support multi-hop, and (d) address the tradeoff between time granularity and route collector diversity. We focus on the BGP views as presented by individual route collector projects and geographic subsets. In general, RIPE/RIS and Routeviews show a similar picture in most settings where a global topology estimation is needed. Nevertheless, they still enrich each other. We recommend research projects with limited computational resources to process routing data to run small-scale versions of their measurements and analytics on a small collector subset to understand general trends. Isolario, for example, yielded almost the same results as both RIPE/RIS and Routeviews, but is heavily biased by the NLRING feed. Sadly, as of 2021 the Isolario project does not exist anymore. Nevertheless, we suggest future work to routinely compute priority lists for different collector combinations to support future measurement studies and reducing unnecessary computational load.

Route collector deployment: We find that geographic diversity, BGP add-path, and multi-hop broaden the visibility, essential for adding new route collectors resp. BGP peers. Indeed, our results show that a few route collectors in a region not yet well covered can increase visibility significantly and impact measurement results. Our recommendation for the deployment of new or updates to existing route collectors is to (a) take advantage of BGP features, such as add-path, and (b) investigate the similarities with existing route collectors add diversity which includes regional geographic diversity as well as the location in the AS topology.

4

Measuring the Routing Layer: BGP Communities

This chapter will discuss how passive measurement studies can inform large-scale active probing campaigns in the domain of Internet routing. We approach this aspect from two perspectives: (a) Internet measurement is a discipline that is strongly intertwined with ongoing developments of the measured subject, i.e., the Internet. One very interesting aspect is the assumptions operators, or researchers make about the Internet. In this chapter, we will visit one such assumption and validate it against real-world data. (b) Many active measurement campaigns have previous conditions which need to be validated before planning and launching the campaign. In this chapter, we will show how to use passively collected routing information to plan and conduct active measurements. To address both, we analyzed the propagation distance of BGP communities across the Internet.

Problems not yet measured on the real-world Internet are often proposed by academics and sometimes disregarded by the operations community because of their focus on the current operation of their infrastructure and not so much on "what-if" scenarios. Additionally, it has to be said that not all problems academics propose are in reality likely to manifest themselves. So, to double-check our drawing board intuitions, we as scientists should try to understand if and under which circumstances potential problems can affect real-world infrastructure. To this end, we studied how BGP communities, an increasingly popular feature, can be abused for malicious purposes in the wild.

In this chapter, we discuss a part of an Internet Measurement Conference 2018 publication [18]. We specifically, highlight the aspect of large scale passive Internet measurements and their importance when it comes to understanding the feasibility of specific attack scenarios[1]. Here we take a look at the communities extension to the Border Gateway Protocol (BGP) and evaluate which potential security-relevant scenarios it can enable and assess which prerequisites need to hold and how widespread they are.

BGP communicates reachability information between neighbors on the Internet. Despite the important role this protocol plays, its design idea is fairly simple: Each participant communicates its reachability to its neighbors, who then further propagate this information. For details see Chapter 2. Back in the 1990s, that was en-

[1]For an in-depth analysis of the proposed scenarios in the wild, refer to [18].

tirely sufficient. As the network evolved, the complexity of connections, policies, and business relationships drove the need for similarly complex and fine-grained routing policies [122–124]. As a result, BGP extensions help to support such policies. This work focuses on one such extension, *BGP communities* [125], and the implications of its real-world implementation and deployment.

BGP communities are an optional transitive BGP attribute used to "tag" advertisements. Operators frequently configure their infrastructure to take different actions depending on community tags. So, communities provide not only a common label for groups of prefixes but also the ability to *signal semantics* between Autonomous System (AS)es and between routers within an AS.

BGP communities are increasingly popular and are used to encode an ever-wider variety of information [126–129]. Within the last year the number of observable communities increased by roughly 20%, see Section 4.4. As we describe in Section 4.1, communities are used to realize routing policies, bias path or peer selection, and steer traffic. ASes also use communities to offer value-added services for customers of Internet Service Provider (ISP)s and members of Internet Exchange Point (IXP)s including tagging of route ingress points and origins [127, 130, 131], selective advertisement [132–135], traffic engineering [126, 136, 137], and Remotely Triggered Blackholing (RTBH), i.e., dropping of traffic to a target destination to mitigate Denial-of-Service-Attacks (DoS) [128, 129, 138, 139]. Some providers even use communities to encode latency information [140, 141].

While BGP communities are a seemingly innocuous feature, they are used to heavily influence traffic flow on the Internet. As such, the question arose whether this feature can be abused maliciously. To this end, we took theoretical attack scenarios and checked whether the Internet provides the right conditions for them to be relevant for real-world Internet operations.[2] We show that the conditions exist to use BGP communities to influence routing in unintended ways. Although the community-based attacks we consider require specific conditions for success, we show that these conditions hold sufficiently widely to warrant operational attention. Importantly, since our extensive measurements show that communities are widely propagated, see Section 4.4, an attacker exploiting the BGP communities of a particular AS does not need to be a directly connected peer.

The attack scenarios rely on weaknesses in the current use and implementation of BGP communities and community-based services. Services enabled by communities are typically relevant only between directly connected ASes, e.g., an AS tagging a backup route with a community to indicate that the remote AS should use a lower local preference. Intuitively, one might expect communities to not propagate through multiple ASes, or beyond their intended destination AS. However, via large-scale analysis of passive BGP datasets, we find that more than 50% of the BGP communities traverse more than four ASes, and we see 10% with a hop count of more than six, see Section 4.4.

[2] For evaluation and active testing of these scenarios in the wild not covered by this thesis, see [142].

4.1 BGP Communities: A Primer

"Communities" are an optional BGP attribute used as a signaling mechanism within and between ASes [125]. While the 32-bit community field[3] can take any value, by convention the first 16 bits represent the AS Number (ASN) of the entity defining the community, while the last 16 bits indicate an action or label. The human-readable community presentation format separates numeric representations of the ASN and label with a colon, e.g., `3130:411`.

There is only a small set of standardized well-known community labels, e.g., NO_EXPORT (`65535:65281`) indicates a route should not leave a BGP confederation, NO_PEER (`65535:65284`) [144] indicates a route should not be propagated via a bilateral peering link, and `65535:666` the standardized blackhole community [139]. These well-known communities cover a tiny subset of all communities in use (Section 4.4) and the complex routing policies network operators realize via BGP communities. Indeed, an AS is free to define (or leave undefined) the semantics of the 2^{16} possible values for its communities. For example, in the previous example, AS 3130 "owns" communities `3130:XXXX` and may define them arbitrarily. It is important to note that there are no explicit mechanisms to enforce this segmentation of the community space, and any AS is free to add, delete, or modify the communities of BGP advertisements that transit its control plane with impunity. Even cryptographic proposals to protect the authenticity and integrity of routing announcements do not cover BGP communities [145–149].

Communities can be added, deleted, or set by an AS on prefix origination, ingress, or egress. Bonaventure et al. were the first to propose a taxonomy of community values [126, 150] and identified two main modes of operation. First, there are AS-internal communities that are set when receiving a route. Second, communities labeled on egress are commonly used to signal or pass information down the path. Such outbound communities carry a broad spectrum of meanings, but most fall into the following categories according to [126]: (a) route selection: adjustment of local_pref and AS path prepending, (b) selective announcement: routes are labeled according to which class of ASes (peer, transit) or even specific ASes they should be announced to, (c) route suppression: same as (b), but states explicitly to whom not to announce a route, (d) blackholing: traffic towards this prefix, mostly /32s (in IPv4) should be dropped, and (e) location: to signal where a route has been learned.

Figure 4.1 illustrates some ways communities are commonly used in practice. Here, AS6 tags incoming routes with the geographic location where the prefix was received, in this case from Los Angeles (LAX) and Frankfurt (FRA). The first part of the community denotes AS6, while the values `201` and `202` are chosen by AS6 to indicate the location. Further, AS3 defines the received community `AS3:103` to prepend its AS three times to path. AS1 can then perform route selection by attaching the

[3] With the advent of 32-bit ASNs, RFC8092 [143] introduces "large" 96-bit communities. This study focuses on traditional 32-bit communities as they already offer many intriguing scenarios. We leave an extended investigation of large, extended, and private communities to future work.

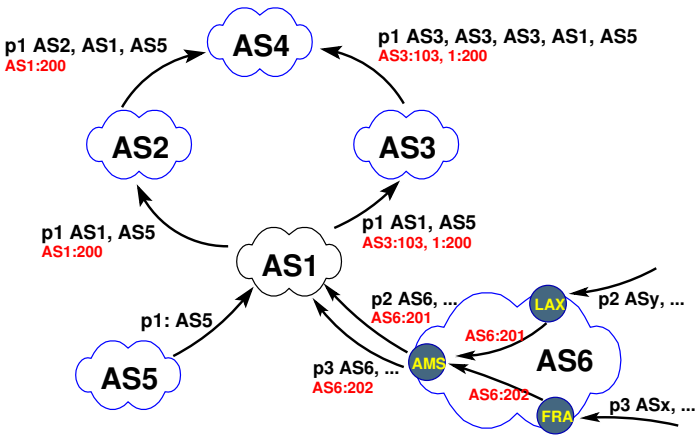

Figure 4.1: Policies with BGP communities: AS1 requests path prepending by tagging `AS3:103` towards AS3, and informs its peers that prefix p1 is a customer prefix, by attaching community `AS1:200`. At the same time AS6 uses communities `AS6:201` and `AS6:202` to signal where a route is learned.

community `AS3:103` to the announcement of $p1$ to AS3. Once AS4 receives both announcements for $p1$, it will prefer the shorter path via AS2.

The level of community support as well as documentation varies considerably among providers. Some networks, especially large ISPs [151, 152] and IXPs [131, 153–155] implement fine-grained semantics using as many as hundreds of communities. Unfortunately, there is no central database of record for providers' communities and associated actions, but rather scattered and incomplete documentation. In reality, this boils down to networks documenting the communities relevant to their peers and customers on their website and/or in Regional Internet Registries (RIR)/Internet Routing Registries (IRR) [156]. We, therefore, lack a definitive understanding of the global definitions and use of communities.

Further complicating the use of communities is that there is no strict policy as to how a network should handle incoming routes tagged with communities. Therefore, there is no consistent behavior in forwarding BGP communities amongst different networks; e.g., some will remove all communities not understood by them, while others will forward everything, and yet others have more complex community propagation policies. We discuss implications of this design choice in Section 4.2, and measure the extent of community propagation in Section 4.4.

4.2 BGP Communities: Can Of Worms

By allowing ASes to extend the semantics of routing updates, BGP communities can significantly simplify policy implementation. As such they are, as we underline in Section 4.4, heavily used in today's routing system. However, as we now show, they also present a can of worms in the sense of "a situation that causes a lot of problems

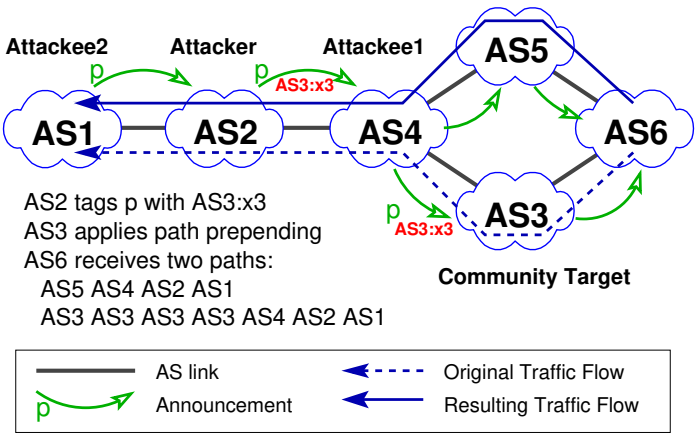

Figure 4.2: BGP communities scenario: AS path prepending.

for you when you start to deal with it" [4]. We, then, discuss why this is too often the case in today's Internet.

4.2.1 Motivating Example Scenario

We use a common community service, AS path prepending, to show the intended use of communities as well as noting the potential for abuse, see Figure 4.2. AS1 announces the prefix p to AS2 and on to AS4, which announces it to AS3 and AS5 and then on to AS6 (see green hollow arrows). Consider traffic from AS6 to p. As the AS paths via AS3 and AS5 have the same length, AS6 may choose to route via AS3 (dotted blue line). AS3 offers AS path prepending via the community `AS3:×n` to prepend n times; where n is typically between 1 and 3. For example, NTT uses `2914:421` for prepending once, `2914:422` for prepending twice, etc. The *intended use* of this service is to enable AS3's peers, e.g., AS4, to do traffic steering.

However, if some AS on the path, e.g., in this case, AS4, does not filter communities, this service can also be (ab)used by other ASes on the announcement path.

Potential abuses include: AS2 or AS1 setting the community `AS3:×3` on the announcement of prefix p; causing AS3 to path prepend three times for the announcement of p to AS6. This changes the traffic flow from AS6 toward AS1 to choose the AS5 (shown via the solid blue line) as opposed to AS3. The motivation for AS2 might be:

Malicious interceptor: If AS5 is a malicious interceptor [157, 158], AS2 is able to steer traffic through it.

Impose additional cost: The link from AS5 to AS4 might be more expensive than the link between AS3 and AS4. AS2 forces AS4's ingress traffic to the "expensive link", that yields high cost for AS4.

[4] Definition of "can of worms" according to Cambridge Advanced Learner's Dictionary & Thesaurus.

Performance improvement: If the service offered by AS6 is popular and the performance via AS5 in terms of bandwidth and/or delay is significantly better, AS2 may improve its service to AS1 by tagging the announcement p with the path prepending community of the provider of its provider, i.e., to steer traffic via AS5 rather than AS3.

Performance impairment: If the performance via AS5 is significantly worse than the performance via AS3, AS2 may slow down an application originating in AS6 that is clogging its network.

Because BGP communities are transitive attributes, the above is fully compliant with the specification. But the actual behavior/use depends on the policies of the involved ASes, in particular, AS3 and AS4.

The above is a teaser example to highlight some potentially unintended consequences of *transitive* BGP community use. In Section 4.3 we show multiple scenarios for traffic steering as well as RTBH (dropping of traffic). When combined with prefix hijacking [159] this raises significant security concerns. Thus, we argue that transitive BGP communities are "a can of worms" for the routing system.

4.2.2 BGP Communities Shortcomings

We believe that BGP communities may be an insufficiently constrained feature for the Internet routing system for the following reasons.

Missing Semantics: Communities are "just tags." This has multiple consequences: (a) Communities do not have a generally agreed upon semantic. Only a few communities and the "expected" community format are standardized via RFCs (Section 4.1). This is analogous to having a program's semantics in the comment statements. (b) Communities are AS specific. Each AS can define their own communities and determine how to publish them, e.g., publicly or only to their peers/customers. (c) The order in which communities are processed by a router is not well-specified and differs by operator configuration as well as by equipment vendor.

No authentication of tagger/community: Any AS on the path can add or modify any of the communities of a routing update. The recipient of a community cannot determine which AS on the path added or modified any of the communities.

Yet, communities are critical for operation since complex routing policies are a reality and unlikely to change. Currently, BGP communities are the most convenient way for signaling information between ASes – an essential component for realizing routing policies. Moreover, an AS may not only mistakenly or maliciously tag a route with a community, it may even free-ride, i.e., hijack a prefix or subprefix[5] by announcing them tagged with a community of their choice.

[5]Hijacking a route corresponds to announcing a prefix for which the AS is not responsible for.

Given the above, one has to ask what this implies for the Internet routing system. First, each AS should define its policy in regard to remote community use and/or install appropriate filters and community parsers. Second, policy implementation should account for ill-specified and misused communities. Misuse of communities can either happen due to malicious intent or by mistake, eg due to fat or thin fingers. Indeed, when considering the above shortcomings, together with the scenarios highlighted in Section 4.3, we urge the community to rethink whether communities are the right mechanism and, if so, how to ameliorate the above shortcomings.

4.3 Unhappy Scenarios

In this section, we outline different scenarios where unintended results can be triggered by the usage of BGP communities, including RTBH and traffic steering

4.3.1 Remotely Triggered Blackholing

High-volume DDoS network attacks can heavily degrade network performance even to the point of making services unavailable [160, 161]. Especially e.g., networks often suffer as they cannot handle such high traffic levels. One mitigation option is *blackholing*, i.e., *dropping* all traffic going to a destination under attack, ideally, as close to the source as possible. As result, the victim IP address or the entire prefix becomes intentionally *unreachable*. Many networks provide their customers with the ability to automatically blackhole traffic using BGP communities as a signaling mechanism, so-called "remotely triggered blackholing" (RTBH) [139]. Networks issue blackholing requests by sending BGP announcements to their direct BGP neighbors for specific destination prefixes with the blackholing community of the respective network. The neighbor, upon receiving such an announcement discards, at its ingress, traffic whose destination address is in the blackholed prefix. In principle, this service should only be used in case of attack and by networks which actually have authority for the blackholed prefix or IP address.

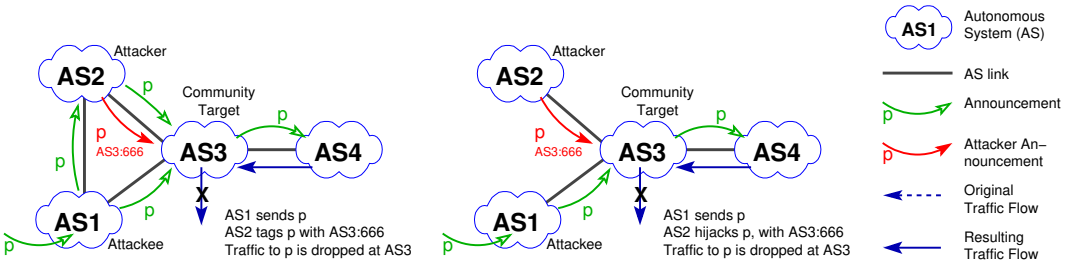

(a) Without prefix hijacking. (b) With prefix hijacking.

Figure 4.3: BGP communities scenario: Remotely triggered blackholing.

However, consider the example shown in Figure 4.3a. Here, $AS1$ announces prefix p to both $AS2$ and $AS3$. $AS3$ offers blackholing service and is the community target in this scenario. If $AS2$, the attacker, adds the blackhole tagged for $AS3$ to its announcement for p to $AS3$, traffic to p may be blackholed at $AS3$ even though the AS path of the tagged route is longer. The reason is often preferred treatment of the blackhole community before best path selection, see, e.g., the suggested configuration in [138]. Alternatively, $AS2$, the attacker, may announce a more specific of p which again has higher priority than the direct announcement from $AS1$, the attackee. Note, if $AS4$ also offers blackholing services via communities the same attack can be launched with $AS4$ as community target as long as $AS3$ propagates communities.

The above example requires the attacker, $AS2$, to be on a path from $AS1$ to $AS3$. However, even if this is not the case $AS2$ may be able to hijack prefix p, especially if $AS2$ and $AS3$ are peering, since strict prefix validation is often not in place, see Figure 4.3b. Indeed, [162] reports 5,295 routing attacks (route leaks and hijacks) alone in 2017 which arguably should not be so frequent if proper filtering would be in place.

Even when prefix validation is in place, it may be possible to hijack prefixes, by tagging them with a blackhole community, depending on the order in which announcements are processed by a router's filters. For example, there are configurations, eg [163], where instead of discarding the announcement (due to hijacking) the router might process the hijacked announcement if tagged with the blackhole community as the community raises the routes precedence.

If AS2 has the ability to hijack prefix p of the attackee (AS1), it can announce p with a short AS path tagged with the blackhole community of AS3. This causes AS3, the community target, to drop all traffic to p. Again, a similar scenario is possible with AS4 as community target if AS3 propagates AS4s blackholing communities. Note, such an attack may be more or less interesting than simply hijacking. First, it may be effective only because of the community tag (validation done after blackholing). Second, whereas hijacking may only partially disrupt traffic (to the poisoned ASes), the hijacking plus blackholing attack disrupts all traffic to the victim.

4.3.2 Traffic Steering

Traffic engineering is one of the essential tasks of a network operator. The generally preferred choice for an AS is selective announcement of prefixes. Sometimes, this is not desired or not sufficient. A common alternative is for remote ASes to provide AS path prepending, Local Preference tuning, Multiple Exit Discriminator (MED) tuning, or partial route announcements, e.g., in specific regions such as Europe only, US only, Asia only. Many ASes accept signals for these tunings via BGP communities; and many ASes are offering these traffic steering services to their customers.

Recall the example from Figure 4.2 in Section 4.2. It highlights that it is possible to intentionally or unintentionally steer traffic over a link that should not be used according to the AS's policy. Indeed, if the involved ASes are susceptible to prefix

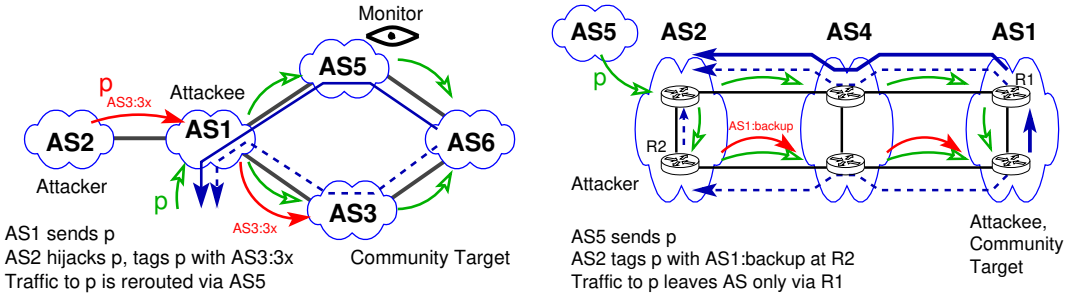

(a) AS path prepending/with hijacking.

(b) Local pref/Without hijacking.

Figure 4.4: BGP communities scenario: Traffic steering.

hijacking this can be further misused as shown in Figure 4.4. An attacker AS may hijack prefix p (which AS1 receives from a peer) and tag it with the prepending community, thus, rerouting traffic via AS5[6]. This can cause trouble for AS1, either due to the unintentional heavy use of the link between AS1 and AS5, e.g., a paid peering link, or if AS2 and AS5 collaborate and AS5 has a malicious traffic tap to inspect all traffic to p.

The next example, see Figure 4.4b, shows how AS2, the attacker, can use the local pref communities of AS1 to force AS1 to route all its traffic to AS2 over a single link via AS4 to AS1. While this may look undesirable at first, this can be highly beneficial for AS2, e.g., if R2 is in Hong Kong and the origin of p is in the US. AS2 in effect forces AS1 to pay for expensive intercontinental transport. In this case, the local pref community can be used to declare the undesired path (from the view of AS2) a backup path. We leave the decision on whether this is an attack or a smart way of reducing cost to the informed reader.

4.4 BGP Communities Propagation

The previous section highlights that communities add even more worms to the routing can. To check their realizability we review the above scenarios and identify the following necessary condition: The above weaknesses of RTBH actions can, in principle, be used if *communities are propagated beyond a single AS* and if the *community service is known.*[7] In the this section we check, if the necessary conditions exist in the wild, i.e., communities are commonly propagated beyond their direct neighbors.

According to RFC1997 [125] BGP communities are an optional transitive attribute. Yet, after discussions with many network operators we concluded, that their expected use is often between two AS neighbors. In this section, we tackle this apparent

[6] Even though AS2 announces a route for p it may not receive much if any traffic if the best route on most routers remains the one to the origin AS for p. If AS2 receives any traffic for p it can loop it back to AS1.

[7] In [142] we also identify an confirm sufficient conidtions for these scenarios to work in the wild and check them in a lab setting.

contradiction as the propagation of BGP communities beyond their intended ASes is a strong prerequisite for the scenarios outlined in Section 4.3.

First, we measure how common BGP communities use is. Then, we show how often communities are propagated beyond a single hop, i.e., are transitive, or if they even on the AS path. Finally, we check for indications that ASes actively strip communities.

4.4.1 Routing Datasets

We rely on a multitude of vantage points within the Internet routing system.

Source	BGP msgs (in Billions)	IPv4 prefixes	IPv6 prefixes	Collectors	IP peers	AS peers	Communities	ASes	Origin	Transit	Stub
RIS	4.80	823,619	76,783	13	275	268	53,208	62,210	61,806	15,016	47,194
RV	9.12	874,054	65,812	15	357	206	57,344	62,424	62,020	9,418	50,991
IS	23.48	830,527	63,584	4	154	97	50,128	62,153	61,754	11,067	51,086
PCH	1.57	802,637	64,136	162	4,640	1,924	40,719	62,033	61,620	10,914	51,119
Total	38.98	967,499	84,953	194	5,158	2,133	63,797	62,681	62,253	15,578	47,103

Table 4.1: Overview of BGP dataset (April 2018). IPv4 prefixes contributed 92% to the total number of prefixes while IPv6 contributes 8%. Therefore, we focus on IPv4 for all other statistics.

BGP routing tables and updates: We rely on the widely-used public datasets of the route collectors from (i) RIPE NCC Routing Information Service (RIS) [48], (ii) University of Oregon Route Views (RV) [49], (iii) Isolario project (IS) [50], and (iv) Packet Clearing House (PCH) [164]. Each of these platforms consists of multiple routers which collect BGP updates from many BGP peers. Some BGP peers send full routing tables, others partial views, and even others only their customer routes. We use the data for the month of April 2018. We remove AS path prepending to not bias the AS path. For an overview see Table 4.1. One specialty of the PCH platform is that it maintains route collectors that peer with the route servers at about 180 different IXPs around the Globe (ca. April 2018) [165]. Route servers are typically a value-added service of the IXP that collect routing information in a centralized manner and redistribute it to connected member routers. Thus, PCH offers BGP routing information for most of the IXP members [132].

Looking Glasses: We use looking glasses of certain ASes, when available, to confirm (i) community availability and propagation, (ii) route changes, as well as (iii) reachability of prefixes.

Active Measurements: We use the RIPE Atlas platform [166] to ping and traceroute to multiple targets during and after routing experiments. RIPE Atlas is an open distributed Internet measurement platform with roughly 10K active measurement nodes. When studying traffic shifting and/or dropping attacks, we use traceroutes along the expected and the altered path to ensure the effect of the routing attack on the data plane.

4.4.2 BGP Communities Use: A first look

As a first step, we measure how wide-spread community use is. Overall, our results validate previous observations [126–129] that it has increased significantly over the last five years, see Figure 4.5. Indeed, today more than 5K ASes offer community-based services[8] and we observe more than 63K different communities in our dataset from April 2018. This is an increase of 18% over 2017.

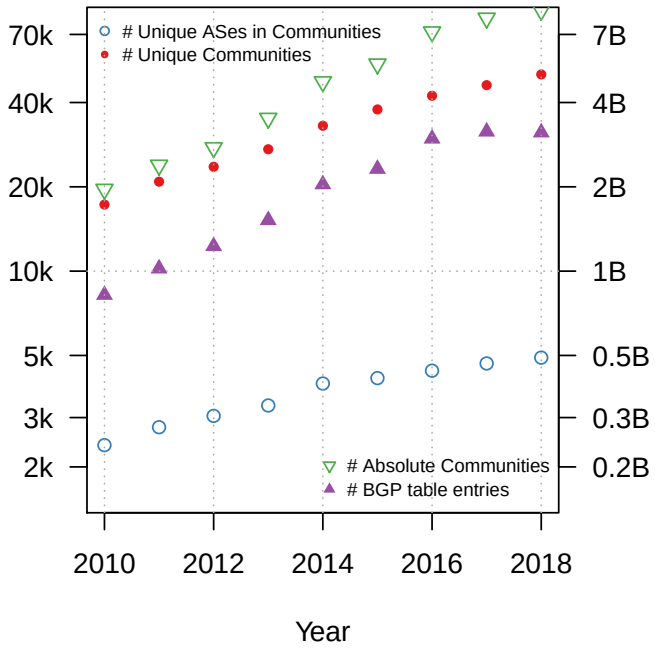

Figure 4.5: BGP communities use over time.

Overall, we find that more than 75% of all BGP announcements at the more than 190 BGP collectors have at least one community set. This means that we can indeed use these collectors to study community use and propagation. Interestingly, some collectors observe more communities than others. Figure 4.6a shows for each BGP collector the fraction of their updates which have at least one community set (in increasing order for each of the four platforms). A large number of our observation points allow us to study community propagation.

We also measure the number of distinct ASes for which we see communities at each BGP collector, see Table 4.2. We see more than 60K unique communities from more than 5.6K ASes which are not directly peering with the respective BGP collector. This suggests that communities are propagated beyond direct BGP neighbors; or one would only see communities associated with direct BGP-peers of the collector.

Next, we measure the number of communities per BGP announcement, see Figure 4.6b. Recall, 75% carry at least one BGP community. Moreover, 51% have more

[8]This statistic is computed under the assumption that communities follow the format convention, namely, `AS:value`.

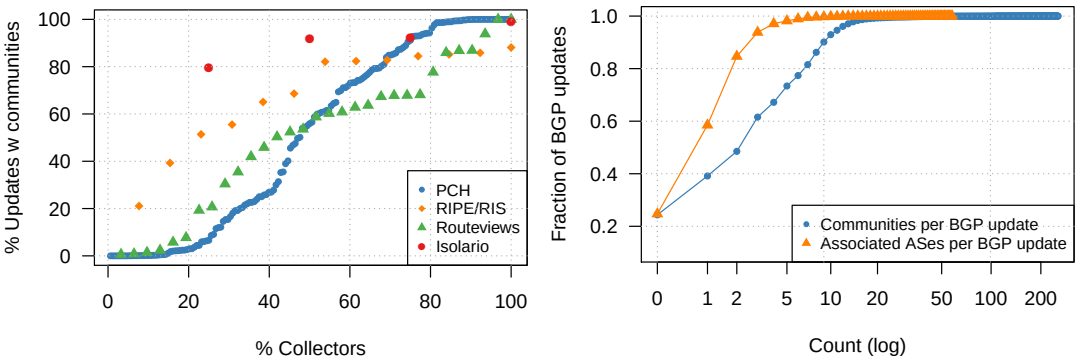

(a) ECDF: Updates with BGP communities by collector per platform.

(b) ECDF: Communities per BGP update and per AS.

Figure 4.6: BGP communities use as observed in the route collector ecosystem.

than two communities set and 0.06% have more than 50 communities set (blue dots). These communities are often (41%) associated with more than a single AS (orange triangle). This is yet another signal that communities are indeed transitive.

4.4.3 BGP Communities Propagation Properties

Next, we measure how far communities propagate. We rely on the format convention, i.e.: `AS:value`. Consider a BGP update for prefix p originated at AS1 and observed at AS5 with AS path `AS5 AS4 AS3 AS2 AS1`. Assume that the update is tagged with `AS1:X` and `AS3:Y`. We assume that AS1 tagged the route with community `AS1:X` since it is the origin AS. The second community is tagged with AS3 and can be either a community received by AS3 from AS2 on ingress or set by AS3 on egress towards AS4. To estimate how far communities propagate we conservatively assume that the route is tagged with the community `AS3:Y` by AS3 rather than by AS2.

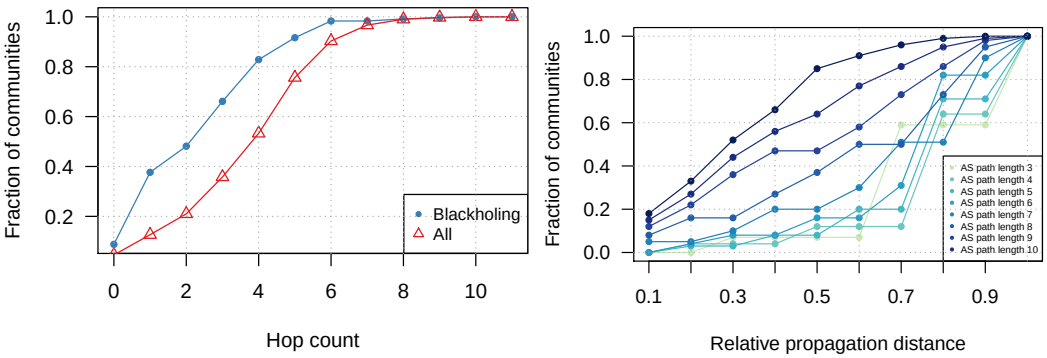

(a) Propagation distance ECDF: path communities.

(b) Propagation distance ECDF: by AS path length.

Figure 4.7: BGP communities propagation properties.

Source	Total # of ASes	w/o collector peer	on-path	off-path	off-path w/o private
RIS	4,931	4,925	3,647	1,826	1,480
RV	5,383	5,375	3,510	1,668	1,279
IS	4,728	4,723	3,513	1,757	1,420
PCH	4,170	4,118	3,002	1,585	1,259
Total	5,659	5,630	3,958	2,154	1,721

Table 4.2: Summary of ASes with observed BGP communities.

However, there is a significant number (21K) of communities of the form `ASX:Y` where ASX is not on the AS path. We call these communities "off-path" and the others "on-path". The former can occur, e.g., at an IXP where the IXP's AS provides the service signaled by the community but, by convention, IXPs are not on the AS path. Other reasons involve widespread tagging (community bundling) to simplify configuration, see, e.g., as reported by Giotsas et al. [128]. Overall, see Table 4.2, we find that 4K ASes are encoded in the on-path communities and 2K in off-path communities. Among the off-path communities there are roughly 400 private ASes [167]. Private ASes are per se off-path as they are not routed. They are often used by networks with large AS numbers which do not fit into the 32-bit community format. Thus, we focus on the communities with public AS numbers.

For on-path communities Figure 4.7a shows an ECDF of the number of AS hops that each community is relayed along the AS path. The red triangles represent the all BGP communities we observed. We find that a significant number of communities are propagated multiple hops. Almost 50% of the communities travel more than four hops (the mean hop length of all announcements [168]). The maximum hop distance we observed is 11 which, given the highly connected AS graph, is rather large.

To check if specific classes of communities are more likely to be propagated we consider blackholing communities as a case study. Hereby, we identify blackholing communities either by the value 666 as defined in RFC7999 [139] or based on the list of verified and inferred blackholing communities from previous work [128]. The resulting ECDF is shown by the purple squares in Figure 4.7a. The difference between the two ECDFs clearly shows that blackholing communities do not travel (on average) as far as other BGP communities. Around 50% of the blackholing communities travel only up to two AS hops, about 80% travel up to four. This is a clear indication that blackholing communities are treated differently by network operators. On the other hand, we still observe some blackholing communities with large hop counts – up to 11.

To check to which extent the above observations are biased by the AS-path length, Figure 4.7b shows the ECDF of the number AS egs that each community is relayed on for different AS path lengths. Hereby, we do not consider communities of the monitor AS but do include the e.g., to the monitor. The color gradient corresponds to the respective AS path length—light green for path length of three up to dark blue with a path length of 10 ASes. This plot highlights that a significant number

of the communities travel more than 50% of the AS-path distance. However, as the path length increases the fraction of the communities that travel longer distances decreases somewhat. The reason for is that each AS on the path can add communities. Therefore, the expected number of communities that can only travel some portion of the AS path is higher. Thus, the plot highlights that communities are propagated significant distances in the Internet independent of the AS path length.

Using the same data we measure how many ASes propagate communities, i.e., are transitive for at least one BGP community of another AS. We do not include the ASes that directly peer with the collector[9]. Thus, for AS2 to be considered transitive we require at least one BGP update for a prefix p tagged with a community `AS1:X` on a path `AS3 AS2 AS1`. We find that there are $2.2K$ transit ASes[10] that relay communities relative to a total of $15.5K$ transit ASes in our dataset.

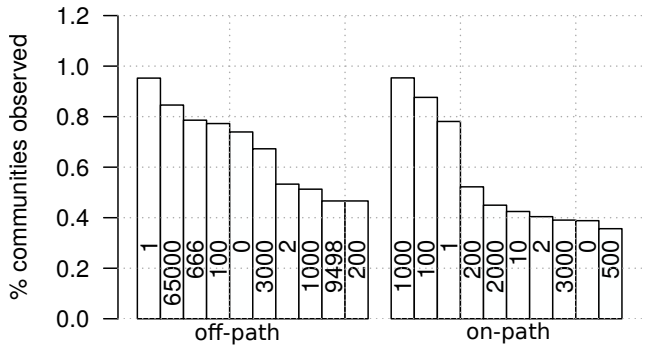

Figure 4.8: Top-10 values for off- and on-path BGP communities.

Next, we explore popular values involved in the observed communities and how these differ for off- vs. on-path communities. Overall, we find that the tails are extremely long—a consequence of the non-standardization of communities. Figure 4.8 shows a histogram of the top-10 most popular values for both off- and on-path communities. Each bar is annotated with the corresponding community values. Note, that their individual contribution is rather small and that they differ significantly. Among the most popular off-path communities is 666 which is used for blackholing. For on-path 666 is not among the top-10 community values. Rather, it is far down in the tail. One explanation is that it is often not observable for on-path since the respective AS should have acted upon receiving the blackhole community. For off-path we see more announcements with blackholing as they are often applied on all peering sessions rather than only selectively [128]. The other values look like convenient values, e.g., for local pref with 100, 200, and 1000.

[9]The configuration for these peerings is often collector specific and may differ from the "regular" policy of the AS.

[10]We consider an AS a transit AS if there is at least one AS path in which it is neither the origin nor the collector.

4.4.4 BGP Communities Filters

So far we focused on how common communities are and if they are forwarded. We have yet to measure if ASes only selectively forward communities or if they actively filter them. As there is no best practice on how to handle communities, networks may filter out all, none, or just specific ones. Measuring this is not straightforward as the only indication of filtering (resp. selective forwarding) is the lack of community propagation as seen in the BGP data. Further compounding the measurement difficulty are that (a) any AS on the path may remove a community, and (b) an AS may receive a "better route" (in the sense of BGP best path selection) not tagged with the community.

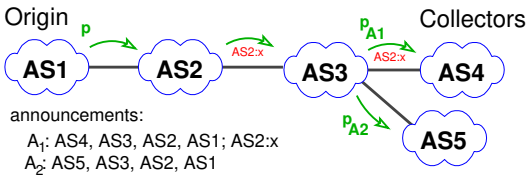

(a) AS egs community change indication counts. After annotating all AS-egs with the observed communities for prefix p from all announcements, we find egs with and without communities, e.g., at segments of the AS-paths. This is repeated for all prefixes.

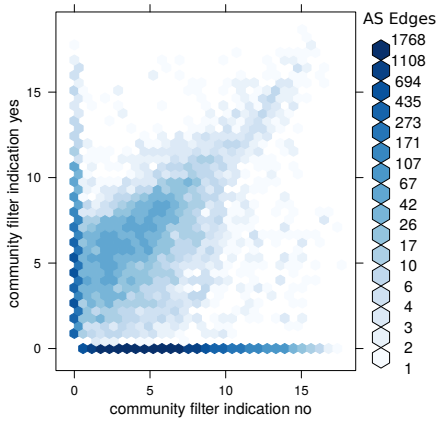

(b) AS neighbor indication counts: Forwarding vs. filtering.

Figure 4.9: BGP community forwarding behavior.

We nevertheless try to identify BGP neighbors where communities are not propagated by collecting indication counts for each directed AS pair. We iterate through all prefixes and, for each prefix p, we consider all updates at the same time and look for ASes where a community that has already been forwarded is propagated to one peer but not to another. The latter is an indication of filtering or selective forwarding. The former is an indication of forwarding, i.e., *no* filtering.

To make this more concrete consider the example shown in Figure 4.9a. We find two announcements A_1, A_2 for prefix p originating in AS1 in the bgpdumps of BGP collectors in AS4 and AS5. Announcement A_1 contains AS-Path `AS1, AS2, AS3, AS4` and carries a community `AS2:X`, while A_2 has AS-Path `AS1, AS2, AS3, AS5` and carries no communities at all. For this analysis we assume the community was not added earlier than AS2. Thus we increase the *community-added* indication on the e.g., (AS2, AS3).

Here, A_1 serves as an indication that AS3 transitively forwards the community from AS2 onwards. Therefore, we increase the *community-forwarded* indication count for the AS pair (AS3, AS4). A_2 allows us to increase the *community-filtered* indication count of AS pair (AS3, AS5). We know from A_1 that for this prefix the community

`AS2:X` is forwarded to AS3 and that AS3 forwards it to some other peers; but we do not see it on the eg (AS3, AS5).

We find signs of transitive forwarding of communities for 4% of the almost 400,000 AS egs and for filtering for roughly 10%. These numbers increase to 6% resp. 15% if we consider AS egs with at least 100 AS paths. We acknowleg that the results of the above heuristic are biased by the BGP collectors which give us different degrees of visibility of the AS egs, as well as by the number of paths observed within the observation period. However, since we consider a full month of BGP updates from four different collector platforms we have reasonable coverage of the AS graph.

Figure 4.9b shows a scatter plot on log-log axis (to the base of 10)[11] of the filtering vs. non-filtering indicators per AS eg. We only include AS edges with at least 100 BGP paths and where we can find either an indication for or against filtering. The count values per AS e.g., range from 0 to 98 million (thus, the values on the x- and y-axis) which comes from the number of different communities and paths that are used in the filter indication computation. The color of the hex-bins correspond to the number of AS egs (darker color indicates more AS egs).

For some AS egs, we find indications that they strip all communities. Those are the ones on the bottom. For others, on the left hand side, we see no indication of filtering, i.e., they forward all communities without touching them. Naturally, we have also many AS egs in the middle of the plot, where we have mixed indications: some communities are forwarded and some are filtered.

The explanation for this mixed picture lies in the absence of best practices regarding BGP communities. After inquiring within the operator community, we found that nearly everyone has a different view on this—some remove all communities, some do not tamper with them at all, while others act upon and remove communities directed at them and leave the rest in place. On the other hand, there are operational reasons to only forward some communities to some BGP neighbors, e.g., different handling of customers and peers.

One natural question in this context is if the relationship type of an AS eg has any influence on filtering. To check this we use the CAIDA AS relationship dataset [169] to distinguish between customer-provider, provider-customer, and peering egs. However, we find that this classification is too coarse-grained to allow for a conclusive picture regarding handling of communities. Thus, we plan in future work to correlate filtering/non-filtering of communities with the role of an AS in the Internet topology.

4.5 BGP Communities: Summary

In this chapter, we specifically focused on investigating BGP communities—a seemingly benign BGP feature. Almost every network operator relies upon communities

[11] The plot uses a logarithmic x- and y-axis. To include zero values we plot the logarithms of the (values + 1).

for a low-overhead simple communication channel between ASes. We showed how a widespread assumption, i.e., BGP communities only travel a few hops, can turn out to be faulty and potentially open an attack vector. Nevertheless, which BGP communities get forwarded or stripped from an announcement is up to the network operator. There is no best practice for community filtering. Establishing global filtering practices would be complex as whether to filter a community or not is highly network-specific.

As such, we recommend to consider the following aspects for the usage and further development of BGP communities:

Need for BGP communities authentication: Given the widespread propagation of BGP communities, there is a strong need for the authentication of the right to attach a community to an announcement or modify one in transit. As of 2021, everyone can attach and remove communities on each announcement traversing their network as BGP communities are by design not covered by Resource Public Key Infrastructure (RPKI). It is solely upon the receiving side to decide whether a received community is legit or not. Unfortunately, there are no known means to do this. Moreover, the adoption of authentication in Internet protocols is a slow process, despite the critical role that the Internet plays in today's economy and society. As cryptographically ensured authenticity of BGP communities is not a realistic goal, we strongly suggest documenting and monitoring the BGP community ecosystem.

Need for proper documentation: Similar to what is ongoing in the IP address space world, see bogon and assigned prefixes lists in Section 5.2.1, the operational community should publish and update well tested best current practices and configuration patterns for community generation, propagation, and action semantics. Proper documentation would ensure everyone knows the effect a community can have on their network and avoid unintended side effects. Operators then could automate their community handling more easily without relying on individual lists of community semantics per peer.

Monitoring the hygiene of BGP communities use: Abuse of communities might be discouraged by monitoring from the points of view of global BGP collectors such as RIPE/RIS and Route Views, analogous to what is being done for BGP hijacks today. This strategy comes with all the problems of BGP monitoring: there is no global BGP view and route collectors only see the announcement they receive. The latter inferences on what happens on the path between the origin and the collector are very difficult. In addition, the lack of structural semantics of BGP communities leaves a lot of room for misinterpretations. Of course, well-known communities can and should be monitored. Yet, this only covers a small fraction of the available community space. Of course, monitoring BGP community behavior is not an active defense; but the attribution of abuse might strongly discourage abuse.

5

Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses

This chapter features a study published at the Internet Measurement Conference 2017 [19]. In this study, we combined both routing and traffic level information to assess the prevalence of IP source address spoofing, the forgery of IP source addresses on the public Internet. We developed a methodology using routing data from RIPE/RIS [48] and Routeviews [49] to derive a list of valid IPv4 prefixes per individual routed Autonomous System (AS). Equipped with this information, we were able to identify spoofed traffic at a large European Internet Exchange Point (IXP). By that, we could provide a lower bound for the amount of spoofed traffic at a vantage point in the Internet's core. We then provide a thorough analysis of the main features of spoofed traffic: (a) originators, (b) spoofed IP source address structure, (c) transport protocol mix, and (c) attack patterns within spoofed traffic.

5.1 The Unsolved Spoofing Problem

The lack of packet-level authenticity of the IP protocol allows for forgery of source IP addresses. This is leveraged by a multitude of attacks that have a vast impact on today's Internet. Despite many ongoing efforts within the research and operations communities to combat IP spoofing, the problem has remained unsolved for more than 30 years [170]. In this section, we provide necessary background on IP address spoofing. We first introduce two common types of attacks involving spoofed source addresses and discuss network filtering practices. We then provide up-to-date perspectives on spoofing and filtering, derived from a network operator survey we conducted. We conclude with a discussion of related work.

We next introduce the two most prominent types of denial-of-service attacks that are enabled by spoofed traffic. We then discuss filtering options that operators have to prevent such attacks.

Flooding Attacks: The attacker overwhelms the victim with packets, either to exhaust the victim's bandwidth resources, or to disrupt the victim's operating system. Here, source IP address forging allows to conceal the true origin(s) of the sender(s) and can cause massive depletion of the victims' operating system resources, e.g., by

flooding with TCP SYN packets from a multitude of source IP addresses, exhausting the state of the victim's TCP stack to the point of disrupting all its network communication [171]. More importantly, randomly spoofing source addresses from a large address range typically makes it impractical, if not impossible, for the victim to filter the offending traffic based on address information alone.

Amplification Attacks: Here, the attackers send crafted packets carrying the source IP address of the intended victim to servers (*amplifiers*) that run a service susceptible to amplification (e.g., NTP or DNS [172]). The servers, in turn, send replies to the victim's IP address that can be orders of magnitude larger than the original requests. This leads to the victim being flooded with a vast amount of unsolicited traffic, potentially disrupting its operation. The ability to forge *specific* IP addresses is essential for this type of attack.

Network Filtering: The decentralized nature of the Internet makes the spoofing problem difficult to address, since there are few topological locations where packet-level sender authenticity can be verified in a straight-forward manner. While it is virtually possible to filter traffic at any given router in the network, the most commonly deployed strategy to prevent spoofing relies on traffic filtering at the AS boundary. In practice, this is achieved by deploying ACLs (Access Control Lists) that only allow traffic with source IP addresses covered by specified prefixes to enter the network. ACLs can be whitelists (i.e., specify a list of allowed prefixes) or blacklists (i.e., specify a list of forbidden prefixes). We synonymously refer to these ACLs as filter lists. Filtering at the AS boundary can be implemented at the *ingress* or the *egress*.

Traffic is most commonly filtered at the ingress, referring to the border router where traffic from other networks (peers) enters the network. Here, the border router maintains a continuously updated list of all prefixes for which it is allowed to accept traffic on a certain interface, from a certain peer. Traffic with IP addresses that are not covered by these prefixes will be dropped before entering the network. Leveraging this strategy to eliminate IP spoofing is documented in detail by Best Current Practices (BCP) documents 38[173] and 84[174]. It is also possible to deploy filtering at the egress where traffic leaves the network, here the same concepts apply as for ingress filtering.

Both strategies rely on prefix lists that must be generated and constantly maintained. In the case of negative filters which mostly refer to a small set of static prefixes (e.g., private address space [91]) the task is trivial since such filters can be statically configured. For fine-grained filtering of valid and routed prefixes that belong to the network and its peers, however, a comprehensive overview of the peering topology as well as constant maintenance are necessary. As of today, no reliable general mechanism for automatically creating these kinds of filter lists exist.

The Internet Protocol (IP) provides a unified and simple abstraction for communication over the Internet. It identifies hosts by their IP addresses, allowing for data exchanges across heterogeneous networks. While the simplicity of the Internet Protocol has proven immensely powerful it comes with inherent limitations, such as the lack of packet-level authenticity. Routers perform only a lookup for the destination

address of incoming packets, the authenticity of source IP addresses of packets is not validated on the path between sender and receiver.

The resulting ability to forge the source IP address of a packet (i.e., *spoofing*) enables a series of cybersecurity threats, ranging from the impersonation of remote hosts to massive denial-of-service Attacks, causing major disruptions of Internet services [172]. In response, the IETF developed best practices for ingress traffic filtering at autonomous system (AS) borders [173]. The spoofing problem also received considerable attention from the research community with systems and architectures that have the potential to either limit or prevent spoofing in the Internet (e.g., [175, 176]). However, these mitigation approaches have not succeeded in eliminating spoofing in production environments: Attacks involving spoofed source IP addresses remain widespread [177, 178].

The measurement community has been very successful in detecting the *ability* to spoof in individual networks using active measurements, i.e., by explicitly crafting packets with spoofed source addresses and measuring the receipt or non-receipt of such packets [179, 180]. While active measurements to assess "spoofability" are indispensable resources to track the deployment of ingress filtering in the Internet, they yield no insight into *if* and *how* the ability to spoof packets is exploited in the Internet. As of today, we still lack a detailed understanding of how to detect spoofed traffic "in the wild". Consequently, little is known about the quantitative and qualitative properties of spoofed traffic, nor about the types of networks that source spoofed traffic into the Internet. The absence of well-tested techniques to detect such traffic as well as detailed measurements documenting the dominant characteristics of spoofed traffic are a major obstacle both for networks operators and for designers of operational systems, who have to rely on best guesses on how to identify such traffic and protect their systems against it.

This chapter, presents a first-of-its-kind study in 2017 that focuses on passive detection and analysis of spoofed traffic as observed in the Internet. To accomplish this, we first developed and evaluated tools that enabled us to detect spoofed traffic in network traces. We then applied our detection method to classify the traffic exchanged between some 700 networks that peer at a major European IXP. Our method, combined with our vantage point, allowed us to provide unprecedented insights into traffic and network characteristics inherent to spoofing in the 2017 Internet. For an update and discussion about spoofing and source address validation in 2021 refer to 7. Our main contributions can be summarized as follows:

(i) We developed a new approach to passively detect packets with spoofed IP addresses in inter-domain traffic. Our approach identifies and leverages sets of valid IP address ranges for individual ASes, derived from transitive AS relationships in BGP data. It allows us to filter out spoofed traffic both with unrouted as well as routed source addresses. We compared and evaluated different techniques to generate AS-specific lists of valid address space and minimize false positive inferences.

(ii) We applied our detection method to classify the traffic exchanged between some 700 networks peering at a major European IXP and provide detailed statistics regarding which networks deploy what kind of address filtering in practice. We then quantified the extent to which individual networks contribute to the different types of spoofed traffic at our vantage point, taking their individual business types and overall traffic shares into account.

(iii) We presented a first in-depth analysis of the qualitative characteristics of spoofed traffic exchanged in the inter-domain Internet. We studied traffic characteristics involving both time-of-the-day effects, spoofed applications, as well as the structure of source and destination addresses. Combining our observations, we identified and studied dominant attack patterns.

Our tools and findings have a number of implications for the networking and research community. Our evaluation of BGP-based spoofing detection yields important considerations and pitfalls for network operators that plan to deploy filtering based on BGP data. Our empirical analysis of the deployment of different filtering techniques as well as spoofing contribution by individual networks can assist network operators when deciding with which networks to peer and under which conditions. Our study of the characteristics of spoofed traffic provides hard-to-get insights that are imperative resources for designing and deploying effective anti-spoofing mechanisms and approaches. We note, however, that our approach is only applicable to inter-domain traffic and, hence, only partially illuminates Internet-wide spoofing. In particular, our approach can not detect "same subnet spoofing", i.e., cases where the spoofed IP addresses belong to the as-legitimate-identified address space of the network sending the traffic. In this work, we consider IPv4 traffic exclusively, as native IPv6 traffic still ranges below 3% at our vantage point.

The remainder of this Section is structured as follows: In Subsection 5.1 we introduce spoofing and provide up-to-date practical insights on spoofing and the resulting challenges from a survey we conducted among network operators. In Subsection 5.2 we introduce our techniques to infer valid address ranges for individual networks and to detect spoofed traffic. We apply and evaluate our methodology in Subsection 5.3, and study network-specific spoofing contributions in Section 5.4. We assess characteristics of spoofed traffic in Subsection 5.5 and highlight attack patterns in Subsection 5.5.2.

5.2 Spoofing Identification: Methodology

In this section, we describe our methodology to passively detect spoofed packets in inter-domain traffic. In contrast to active measurements using deliberately crafted packets, our method does not rely on any explicit information about a given packet beyond its source IP address to detect spoofing. Our approach classifies source IP addresses of packets as either *legitimate* or *illegitimate*. However, not all traffic with illegitimate source IP addresses is necessarily a case of spoofing. We argue for the following distinction among packets with illegitimate source addresses:

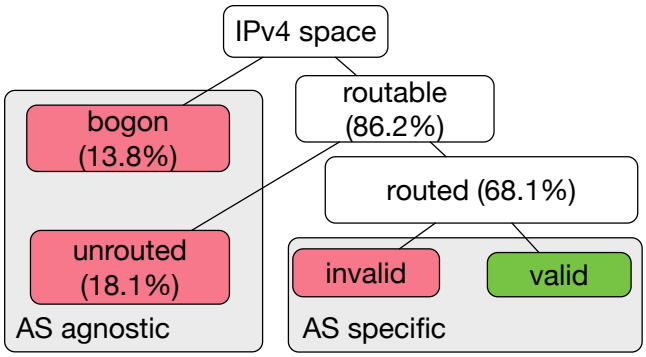

Figure 5.1: Categories of IPv4 Addresses relevant for passive spoofing classification.

Packets with Stray source IP addresses: These are packets with source IP addresses that are the genuine addresses of some interface of the host sending the packet. Yet, packets with such source addresses should either not be forwarded in the inter-domain Internet at all, or not be sourced by a particular AS. The former includes *bogon* IP addresses, e.g., RFC1918. The latter includes traffic with valid routable source IP addresses that we observe on inter-domain links that should not carry them (e.g., routers sending out TTL exceeded over their default route). Typically, stray source IP addresses are the result of misconfiguration without malicious intent.

Packets with Spoofed source IP addresses: In this case the source address is unrelated to any of the genuine IPs of the host sending the packet. Such packets are typically crafted with the intent of misrepresenting the source IP address in a packet to either conceal the identity of the sender or to impersonate another host.

Our goal is to identify traffic with spoofed source IP addresses and to distinguish it from traffic with stray source IP addresses. First, however, we study the categories of IP source addresses that are relevant for our detection method.

5.2.1 Address Space Considerations

To bootstrap our classification approach, we first partition the IPv4 address space into four categories, shown in Figure 5.1: Address space that should not be routed in the inter-domain Internet at all, i.e., reserved ranges, which we refer to as "bogon", and address ranges that are routable, yet we do not find them announced in the global routing table, which we refer to as "unrouted". These source ranges are AS agnostic in the sense that no network should source traffic from these ranges into the inter-domain Internet. The other category includes the IP address space routed in the inter-domain Internet. Here, we distinguish between "invalid" and "valid" address space on a per AS basis.

Bogon Source Addresses: The bogon space captures the address space that is not intended to be used in the public Internet. Bogon source ranges are defined in, e.g.,

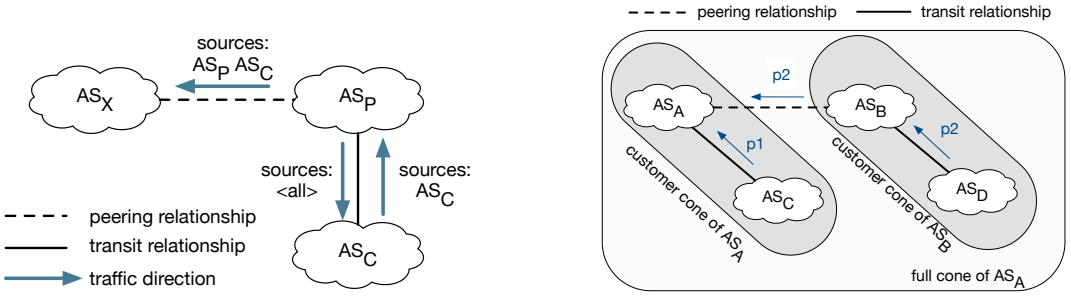

(a) Traffic patterns and source addresses in peering and transit AS relationships.

(b) Full Cone: AS_A can forward AS_B's traffic to AS_C, contrary to customer cone.

Figure 5.2: Inference of valid address space per AS.

RFC1918 [91], RFC5738 [181], and RFC6598 [182]. They include private address ranges as well as multicast and future use.

Unrouted Source Addresses: These addresses *are* part of the routable space, but are not covered by a BGP announcement in the global routing table. We later use extensive BGP datasets to compile a list of currently routed IPv4 prefixes and we consider every address that is not covered by any prefix as unrouted.

Invalid Source Addresses: Naturally, packets with a given IP source address should only be originated from the AS that also announces the prefix covering that address. In accordance with best current practices, they should furthermore only be forwarded by ASes that are in an upstream or peering relation to the announcing AS[1] [174]. We use this observation as a criterion to identify *valid* source ASes for traffic with given source IP addresses. The complexity of determining whether an AS is a valid source for a given IP address depends on its distance from the origin AS in terms of BGP AS distance. In the simplest case, if an AS is a *stub AS*, i.e., not providing transit to any other AS, it is only a valid source for its own prefixes. For *transit ASes*, i.e., networks that forward traffic on behalf of other networks, the situation becomes more complex.

In Figure 5.2a, two networks egge in a *transit* relationship, customer AS_C pays provider AS_P to *(a)* forward traffic it receives from the customer to the rest of the Internet, but also *(b)* to forward traffic from the rest of the Internet towards the customers' routes. Thus, an AS that provides transit typically either offers a full BGP table to its customers or a default route, and is thereby allowed to source the whole routed address space on the links to its customers. If two ASes egge in a *peering* relationship, e.g., AS_P and AS_X in Figure 5.2a, they should only exchange traffic between each other, in particular traffic originating in their own network or in one of their customers. Thus, for the link between AS_P and AS_X, valid sources for AS_P are source IP addresses from AS_P and AS_C. Hence, address range validity for an AS

[1] Some mechanisms take advantage of the fact that this is not strictly realized in the current Internet, e.g., Mobile IP with triangle routing. However, Mobile IP acknowledges this problem and proposes direct routing as an alternative [183].

depends both on its position in the AS-level topology, as well as on the link on which we monitor traffic. Based on the above discussion we will consider any traffic with an invalid source address that is forwarded by an AS to be potentially spoofed.

5.2.2 Inferring Valid IP Space per AS

Our above discussion underlines the need to consider inter-domain information to identify valid source address ranges. We next introduce three approaches for inferring valid IP address space on a per AS basis, ranging from conservative to liberal in terms of the amount of valid IP space per AS.

Naive Approach: As a baseline approach we consider ASes as valid sources for traffic from a given prefix, if we observe the AS on the path of a route announcement of the respective prefix.[2] This information is contained in the BGP AS-path (i.e., the list of all ASes that the announcement has traversed). The naive approach does not account for asymmetric routing or selective announcements; i.e., cases in which an AS does not announce all of its prefixes to all neighbors, but still sends traffic from any of these prefixes to any of them. In such cases, the naive approach tags packets of the partially announced or propagated source prefixes as invalid.

CAIDA Customer Cone: Luckie et al. [184] suggested to use the *CAIDA Customer Cone* [60] for identifying spoofed traffic. The customer cone of an AS is the set of ASes that an AS can reach using provider-customer links. Thus, if *AS A* is the origin of a prefix, then all ASes that include *AS A* in their customer cones may source traffic with source IPs from this prefix. This approach focuses on customer-provider relationships. As such, it intentionally does not take equitable peering relationships into account.

Full Cone: The previous two approaches have the potential to misclassify traffic as invalid, either due to asymmetric routing or due to traffic carried over peering links which the customer cone (intentionally) does not cover. Since we strive to minimize false positive classifications, we develop the *Full Cone*, where we *intentionally sacrifice specificity* compared to the other approaches, by not distinguishing between peering/sibling, customer-provider and provider-customer links. Rather, whenever we see two neighboring ASes on an AS path, we presume a directed link between the two, where the left AS is considered upstream of the right AS. On the resulting directed AS graph (that may indeed contain loops) we calculate for each AS the *transitive closure* containing all its children. Thus, if *AS A* is the origin of a prefix then all ASes that include *AS A* in their transitive closure may source traffic with source IP addresses from this prefix. The Full Cone is the least specific of our approaches, but has the advantage of accounting for peering relationships as well as atypical traffic patterns. Figure 5.2b highlights the potential benefits of this approach when it comes to minimizing false positive classifications: Here, AS_A and AS_B peer with each other. AS_C is a customer of AS_A and AS_D is a customer of AS_B. As such AS_D/AS_C is

[2] This reasoning is also in line with "reverse path forwarding", requiring the reverse route to have been learned from a peer before allowing traffic to be forwarded to it, see BCP84 [174].

in the CAIDA Customer Cone of AS_B/AS_A, respectively. However, these disjoint Customer Cones do not capture the peering relationship. As a consequence, traffic with source IPs in prefix p_2 by AS_D would not be considered valid at AS_A.

Multi-AS Organizations: Our three approaches rely on the existence of visible BGP links. In the case of organizations that use multiple AS numbers, *Multi-AS organizations* [185], peering links between their individual ASes are not necessarily exposed in the global routing table. For the goal of this work, identifying intentionally spoofed traffic, we allow for bidirectional traffic exchange between ASes belonging to the same organization. To identify ASes belonging to the same organization, we rely on CAIDA's *AS to Organization* [186] dataset. This dataset links ASes to organizations based on the available WHOIS information (e.g., email and physical address, name, and contact information). We extract sets of ASes belonging to the same organization, and add a full mesh of links between all ASes within each set. The joint cones and IP address space of each organization is now shared with each constituent AS belonging to the same set, regardless of whether this relationship is reflected in BGP or not. This way, traffic forwarded on behalf of an AS of the same organization is not considered invalid.

5.2.3 Routing Datasets

To determine bogon, unrouted, as well as valid address space for each AS, we rely on the following datasets:

Bogon Lists: We use a list of bogon prefixes as provided by Team Cymru [187], which are widely used by operators for egress filtering. The resulting bogon list contains 14 non-overlapping prefixes corresponding to 218K /24 equivalents.

BGP Datasets: To determine the routable address space as well as to construct the network-specific list of valid address space we rely *(i)* on publicly available BGP datasets as well as *(ii)* on vantage point-specific BGP data. Our measurement period spans 4 weeks from February 5th, 2017 to March 6th, 2017. In particular, we use BGP data from all route collectors from RIPE RIS [188] and RouteViews [49] that have data available for our measurement period (18 out of 21 collectors for RIPE, 16 collectors for RouteViews). RIPE and RouteViews offer snapshots (every 8 hours for RIPE and every 2 hours for RouteViews) of the collector's routing table, as well as all BGP updates that the collector receives from its peers. Note that ASes commonly announce changing sets of prefixes with varying aggregation levels at multiple locations to different networks. To acquire an as-complete-as-possible picture of routed prefixes and of the AS graph, we consider *all* table dumps and update messages within our time period. We disregard announcements for prefixes more specific than /24 and less specific than /8. The latter usually indicates misconfiguration and neither is commonly routed [189]. In total, our announcements cover 11.65M routed /24 equivalents. We extend our BGP datasets with vantage point-specific BGP data from the route server [132] of a major IXP, which will be our vantage point to study spoofed traffic (Section 5.3).

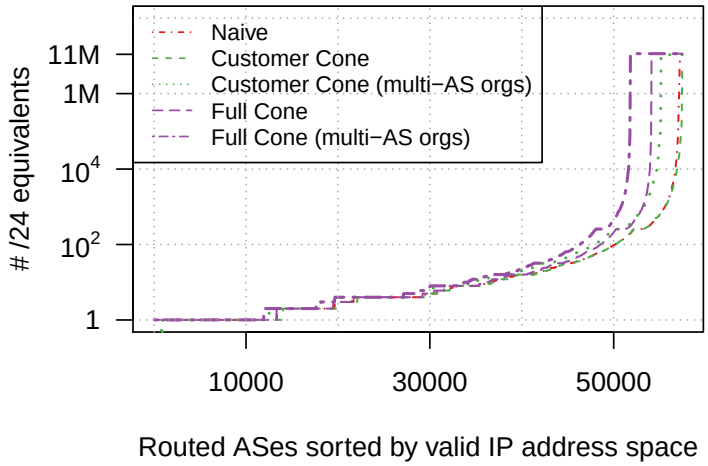

Figure 5.3: Routed ASes sorted by the size of their valid address space based on *Customer Cone* and *Full Cone* inference methods, either with or without considering multi-AS organizations. *Naive* is included as baseline.

5.2.4 Comparison of the Three Approaches

Figure 5.3 shows for each of our approaches the size of the valid IP address space (in /24 equivalents) for each routed AS. Here, we sort the ASes in increasing order according to the size of their valid address space.[3] Furthermore, we show our cone methods both with and without adjustments for multi-AS organizations. We find that, unsurprisingly, all approaches agree on about 12K of the smallest stub ASes. For the remaining ASes, the adjustments for multi-AS organizations consistently cover more address space than the plain, non-adjusted approaches. For the latter, the covered address space only significantly diverges for top 14K ASes. The Full Cone, as expected, yields larger valid address spaces, since it takes any transitive AS relationship into account. Here, we see that for the top 14K ASes the size of their valid address space grows considerably and an upwards of 5K ASes are a valid source for the entire routed address space, roughly 11M /24s. In addition, we also confirmed that the address spaces per AS for the Naive approach as well as for the Customer Cone are fully contained within the Full Cone. Combined with the consistently higher coverage when including adjustments for multi-AS organizations, the Full Cone is the preferred candidate in our endeavor to identify spoofed traffic with an emphasis on minimizing false positive detections.

5.3 Spoofing Detection in Practice

We next apply our method to classify the traffic exchanged between some 700 networks at a major European IXP. While this vantage point provides us with a unique

[3] Note that this figure shows the distribution of valid ranges per AS for each approach individually and does hence not allow for comparison of individual ASes.

	Bogon		Unrouted		Invalid FULL		Invalid NAIVE		Invalid CC	
members	525	(72.0%)	378	(52.0%)	393	(54.06%)	611	(84.04%)	602	(82.81%)
bytes	31.63T	(0.003%)	38.29T	(0.004%)	92.65T	(0.0099%)	10.08P	(1.1%)	1.72P	(0.19%)
packets	304.82G	(0.02%)	217.59G	(0.02%)	387.23G	(0.03%)	17.20T	(1.29%)	4.05T	(0.3%)

Table 5.1: Contributions to each class for our spoofing inference approaches (Traffic scaled to account for sampling).

opportunity to study spoofing at scale we point out that o ur approach is not limited to IXPs: It is applicable for any vantage point that captures inter-domain traffic.

5.3.1 Vantage Point and Traffic Dataset

We use four weeks of continuous traffic traces captured in February 2017 at a major European IXP. IXPs provide a layer-2 switching infrastructure to participating networks, called *members* in the following. Members connect with their border routers to the switching fabric, establish BGP sessions with other members[4] and exchange traffic with each other. At the time of this measurement, the IXP had 727 members that exchanged about 230PB traffic on a weekly basis with peak traffic rates exceeding 5 Tb/s. Our traces consist of IP Flow Information Export (IPFIX) flow summaries which are collected using a random 1 out of 10K sampling of all packets crossing the IXP's switching fabric. The available flow information includes the IP and transport layer headers, as well as flow summaries with packet and byte counts. Thus, at this vantage point, we capture the inter-domain traffic right at the border between ASes.

5.3.2 Classification Pipeline

Our passive spoofing detection mechanism classifies each flow based on its source IP address into either Bogon, Unrouted, Invalid, or valid, see Figure 5.4. Hereby, Bogon and Unrouted refer to the AS agnostic address ranges and Invalid to the AS specific address ranges, recall Section 5.2.1. valid contains all other flows and is not further considered. Our classification is strictly sequential, see Figure 5.4. Once we match a source IP address into a class we stop. Thus, all classes are mutually exclusive. We use the following flow features: source IP address, associated origin AS, and via which IXP member the flow entered the IXP. First, we match the source IP address against the bogon list. Next, we match the source IP address against the routed address space. The following step takes the member AS into account. If we find that the member is not a valid source for this source IP address, we classify the flow as Invalid. To determine this, we check if the IP address is part of the legitimate address space of the member AS, according to each of our three approaches, see

[4]This IXP also provides a route server to its members. Members can opt to establish a single BGP session w ith the route server to immediately egge in *multilateral* peering with a large number of other members [132]. In this study we use BGP snapshots from the IXPs route server in addition to publicly available BGP data.

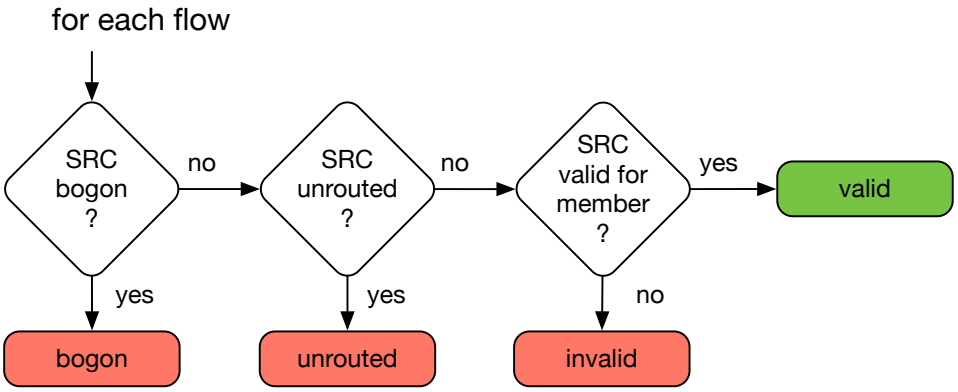

Figure 5.4: Applying our methodology to dissect traffic.

Section 5.2.2. This results in three different sets of invalid traffic, namely, INVALID NAIVE, INVALID CC, and INVALID FULL.

5.3.2.1 Classification Results

The results of applying the above methodology to four weeks of traffic are summarized in Table 5.1. Here, we show both absolute and relative traffic contribution for each class as well as the number of members that contribute traffic to each class.

BOGON and UNROUTED: We observe that more than 72% of the member networks send packets with bogon source IP addresses and 52% send packets with unrouted source addresses. A striking observation, suggesting that the **majority of members do not, or not consistently, filter their outbound traffic**. When taking the relative contribution in terms of packets and bytes into account, however, we see that the share is comparably low, with UNROUTED and BOGON traffic accounting for about 0.02% of the overall traffic. Nevertheless, these traffic contributions sum up to tens of TBs over the course of four weeks. Comparing the contribution of BOGON to UNROUTED, we see that BOGON has more contributing members while UNROUTED has higher traffic volumes though less packets. One possible explanation for packets in BOGON are devices behind misconfigured network address translation devices (NATs). Packets in UNROUTED, on the other hand, are more likely to be caused by intentional source IP address forgery. Apparently, NAT misconfigurations are more common (when seen on a per-network granularity) when compared to source IP address forgery. We point out, however, that UNROUTED traffic contributes more in terms of absolute bytes, suggesting that while fewer networks emit such traffic, they typically emit larger quantities, compared to BOGON.

INVALID: The three right columns of Table 5.1 show the number of members and respective traffic volume classified as INVALID for our three approaches (recall Section 5.2). Here, we observe significant differences across the three approaches. The conservative INVALID FULL, naturally, classifies the smallest portion of traffic as INVALID. Still, more than half of the members contribute traffic to this class. INVALID

NAIVE and INVALID CC identify a significantly larger share of traffic, well exceeding 1% and 0.1% respectively of the total traffic, and including about 80% of members. These observations are in line with the different cone sizes as explored in Section 5.2.4, i.e., the Naive approach and the INVALID CC approach allow less valid address space per AS and hence classify more traffic as INVALID. We observe that the number of members that contribute to INVALID NAIVE and INVALID CC well exceed the number of members that contribute to our classes BOGON and UNROUTED, which are less prone to false positives as they are AS agnostic.

Impact of Multi-AS Organizations: The results shown in Table 5.1 allow bidirectional traffic flow across multi-AS organizations, irrespective of the existence or inferred business type of BGP peerings (recall Section 5.2.2). Allowing such traffic has a different impact on our individual approaches. Allowing inter-organization traffic reduces invalid traffic in INVALID FULL by some 15%, but by almost 85% in the case of INVALID CC. The vast reduction in the case of INVALID CC is due to few heavy traffic-carrying members and closer inspection shows that these members indeed have visible AS links and are thus contained in the INVALID FULL cone, which does not differentiate between different business relationships. INVALID CC only allows customer-provider relationships and hence intentionally discards these relationships. Our results suggest that the customer cone is a promising approach, when refined to take complex AS relationships such as multi-AS organizations into account.

We, in this work, strive to minimize *false positive* classifications. We hence proceed with our analysis with the Full Cone approach, i.e., from now on we will only study INVALID FULL traffic and refer to it as INVALID.

5.3.3 Hunting False Positives

Even our most conservative approach, INVALID FULL, includes false positives, i.e., traffic from source addresses that a member can legitimately source, yet we classify it as INVALID, caused by missing AS relationships. Missing AS relationships can be caused by *(a)* the inherently limited coverage of the AS graph in the available BGP data [63] and *(b)* inter-AS connectivity that is not exposed in the global routing table (e.g., tunnels). To identify traffic that we possibly misclassify as INVALID due to missing AS relationships, we focus on those ASes for which INVALID accounts for a significant share of their overall traffic. Figure 5.5 shows a CCDF of the fraction of the BOGON, UNROUTED and INVALID traffic share of the overall traffic for each member. We note that the largest contribution of any member to BOGON is about 10% and to UNROUTED about 9%. For INVALID, however, we find some few members who contribute close to 100% of their entire traffic to INVALID.

To asses whether we misattribute traffic of these members to be INVALID, we take a closer look at the top 40 member ASes as shown in the CCDF. For these, we generate per-member statistics containing the origin ASes of the source and destination IP addresses in question. Next, we check the databases of the Routing Internet Registries (WHOIS) for missing AS relationships between the member AS sending the traffic,

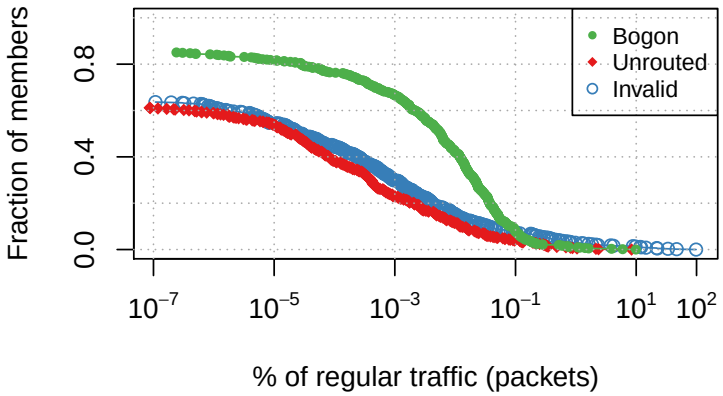

Figure 5.5: Fraction of BOGON, UNROUTED and INVALID of total traffic per IXP member AS.

the origin AS of the source and destination IP addresses, and the member AS receiving the traffic. In particular, we study import and export ACLs that some ASes publish to indicate routing policies. We also leverage information from looking glasses located inside some of these ASes.

Missing AS Links: We identified 15 missing links using WHOIS records, i.e., by matching company names or contact points that are not covered by the AS-to-organization dataset we use, or by matching import/export ACLs for direct peerings. In addition, we find one additional AS relationship based on looking glass information. We identify one instance where two closely related organizations that operate shared network infrastructure exchange internal traffic between parts of their networks via the IXP. Additionally, we encounter several instances where WHOIS data shows that one AS is (or was) an upstream provider, but we do not see evidence in the BGP data at the time we captured the traffic. We currently do not investigate archived BGP data and consider this as future work together with incorporating automated parsing and evaluation of the import and export ACLs to enrich the available BGP data collected.

Uncommon Setups: We also found instances of uncommon routing setups that are not BCP38 compliant. In two cases a customer with multiple upstream providers uses provider-assigned address space from one provider to send traffic via the other provider to the Internet. Analysis of the WHOIS entries reveals that, while the ISP only announces a single covering prefix, an entry in the WHOIS database exists for both customer prefixes naming the customers. In another case we find a cloud-based startup that uses uncommon traffic engineering by tunneling traffic originating at a large cloud provider via their own infrastructure to the IXP.

In future work, we plan to further assess the underlying operational practices that lead to such situations. In this work, we accept such traffic as valid, since we strive to provide an analysis of intentionally spoofed traffic.

After handling all of the above cases, and adding the corresponding IP address ranges to the valid address space of the respective IXP members, we reduce the traffic in INVALID by 59.9% of bytes resp. 40% of packets.

Cross-Check with Active Measurements: Since 2005, the CAIDA Spoofer Project [190] collects active measurement data about the "spoofability" within ASes in a crowd-sourced fashion. In a nutshell, a Spoofer software probe crafts packets with source IP addresses from various ranges and sends them to a measurement server. If the measurement server successfully receives some or all of the intentionally spoofed packets, then spoofing is possible in the AS hosting the probe. These measurements have recently been made publicly available, allowing us to cross-check active inferences of spoofability with our findings.

We leverage the available Spoofer dataset [190], containing results from measurements executed within the last year. In total, we find relevant data for 97 overlapping ASes (i.e., 8% of all IXP members under consideration).[5] Of those 97 ASes, we detected spoofed traffic (INVALID or UNROUTED) for 74%. Spoofer detected spoofability in 30% of the 97 networks. Intersecting our positive detections, we find that Spoofer data agrees with our observations for some 28% of the networks for which we see spoofed traffic. Our passive approach, on the other hand, detects spoofed traffic from 69% of the networks that were tagged as spoofable by Spoofer measurements.

The quantitative differences in our measurements reflect both our different vantage points and the essential difference between the ability to spoof and actual spoofing, as carried out and visible in passive traces. Recall that for a Spoofer probe to reach the target, it has to cross multiple AS boundaries, and is thus subject to filtering, potentially by several ASes on the path. Thus, active measurements provide a lower bound on spoofability in certain networks. Contrarily, Spoofer identified several ASes as spoofable, for which we do not see any spoofed traffic. Reasons here include that there are either no hosts in these networks that do actively perform spoofing, that our inference methodology is too conservative to capture those cases, or recent changes in filtering practices (recall that we compare 4 weeks of passive measurements against one year of crowd-sourced data).

Summary: This concludes our evaluation of the three different approaches to detect spoofed traffic. We chose the most conservative estimation of valid IP address space per AS, the Full Cone. During development of this approach, we encountered various limitations that go along with BGP datasets. As such, in order to get a more fine-grained estimation of the valid IP space per AS, further study of the spatial and temporal characteristics of public BGP data is needed. Our findings also highlight that leveraging external datasets to account for additional AS relationships (e.g., multi-AS organizations) is crucial in order to minimize false positive detections. We acknowledge that our resulting Full Cone considers as many as 5K ASes as legitimate sources for all of the Internet's 11M routed /24 prefixes, which likely results in significant portions of spoofed traffic that remain undetected by our approach.

[5] We only consider ASes in which the Spoofer project conducted direct measurements, i.e., the probes were not located behind a NAT.

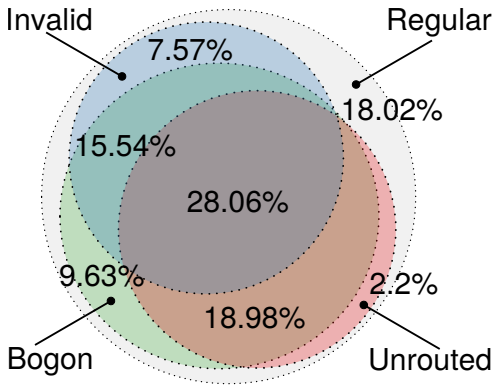

Figure 5.6: Percentage of members contributing traffic to the three classes: Bogon, Invalid, and Unrouted.

However, our conservative approach results in a very distilled traffic dataset that is indeed mostly composed of actual spoofed traffic. Recall that the choice between the three approaches does not affect Bogon and Unrouted. We conduct all further analyses in the remainder of this chapter based on the results this approach yields.

5.4 Network Perspective

In this section, we study which networks send what kind of illegitimate traffic. We first study filtering consistency for individual networks. Then, we take the business types of individual networks into account and contrast them with their individual traffic contribution. Finally, we identify (and remove for later analysis) some members that contribute illegitimate traffic that is not the result of intentional spoofing, but stray traffic originated from routers.

5.4.1 Filtering and Traffic Contribution

Filtering Consistency: The Venn diagram in Figure 5.6 shows what percentage of members at this IXP contribute traffic to our three classes, as well as intersections in contribution to different classes. We next use this to deduce lower bounds on which filtering strategy individual members apply. If we do not observe a member emitting flows falling in one of our categories, we assume this member filters the respective type of traffic. We are aware that this is a soft criterion, eg an AS may simply not emit flows with spoofed source IP addresses traversing its network during our study period. However, we argue that it is still a reasonable approximation to provide tight lower bounds, given the length of the observation period (4 weeks).

In total, we find that only some 18% of members are "clean" in the sense that they do not send any traffic classified as either Bogon, Unrouted, or Invalid. On the other

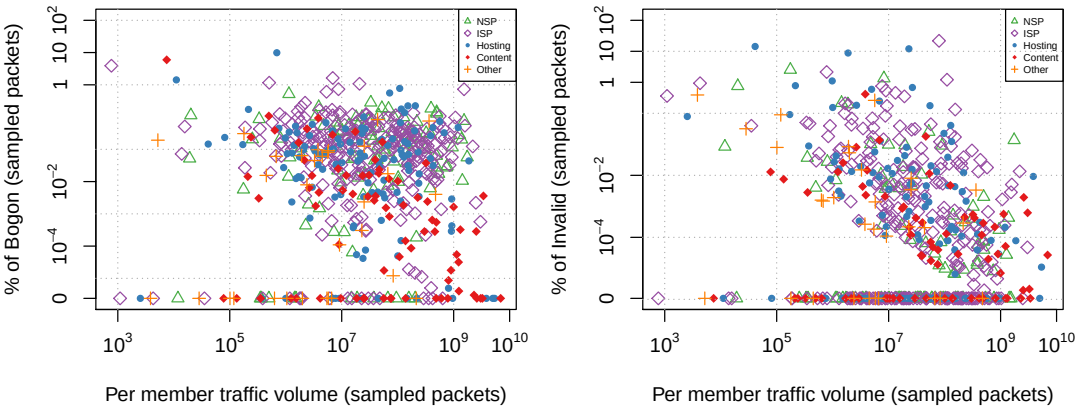

(a) IXP members by network type: Bogon. (b) IXP members by network type: Invalid.

Figure 5.7: Network-wide view of Spoofing: Business Types and Traffic / Filtering.

end of the spectrum, we find around 28% of members contributing traffic to all of our classes. Thus, these networks do not deploy proper filtering. Another interesting case are some 9% of networks that contribute only BOGON traffic. We presume that these networks deploy filtering against spoofing, but lack filtering for bogon ranges. Of the members contributing UNROUTED traffic, the vast majority, 96%, also contribute INVALID or BOGON traffic, highlighting that packets with unrouted source addresses are a good indicator for spoofing detection on a per-network level. Only some 7% of the members contribute only INVALID traffic exclusively, but do not fall in either of the other classes. Here, we can presume that they have best effort filters deployed in the sense that they use appropriate semi-static filters. Still the fact that we see traffic in INVALID suggests that they do not follow BCP38 and BCP84.

Business Types: To understand if the business types of networks directly relate to filtering setup and contribution to illegitimate traffic, Figure 5.7 consists of two scatterplots that show per member the total traffic contribution (x-axis), as well as the share of BOGON respectively INVALID of their individual traffic (y-axis). Note that the general observations for UNROUTED are similar and since only less than 3% of members contribute UNROUTED traffic exclusively, we show only the contributions for BOGON and UNROUTED.

We use different plotting symbols to highlight the different business type of the member ASes, which we derive from PeeringDB [28].[6] Intuitively, members contributing more overall traffic, but a tiny share of BOGON or INVALID are located in the bottom right corner, while members with large fractions of BOGON/ INVALID traffic, but low overall traffic volume are in the upper left corner. Generally, we find that most networks with significant overall traffic shares show a comparably low fraction of illegitimate traffic, according to our classes. Indeed, most large content providers do not contribute any traffic to BOGON and only few to UNROUTED. This is reasonable, since most content providers have full control over their network and almost no end-user machines.

[6] We classified ASes without PeeringDB entries manually.

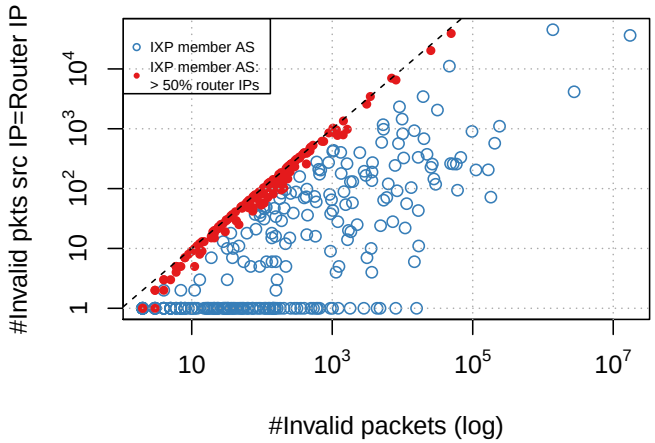

Figure 5.8: Router IP addresses among invalid packets per IXP member.

In terms of members that have significant shares (> 1%) of BOGON, UNROUTED, and INVALID traffic, we predominantly find hosting companies (highlighted with blue dots), end-user ISPs and, to a lesser extent, transit providers. These network types have in common that they typically provide connectivity (and possibly hardware), but have little control over how the provided resources are *used* by individuals (e.g., virtual machines within a hoster). In the absence of proper filtering, spoofing is more likely to be carried out from hosts within such networks, compared to eg large content providers. ISPs provide service for end users, which may indeed have incentives to originate spoofed traffic and they also may as well suffer from misconfigurations, which can lead to leaked traffic, e.g., from CPE NAT devices.

5.4.2 Spoofing vs. Stray

So far, we focused on the contribution of illegitimate traffic from individual networks, as well as their filtering strategies. We want to recall, however, that not all illegitimate traffic is in fact the result of spoofing, but can also be the result of uncommon routing policies or misconfigurations (i.e., stray traffic, recall Section 5.2). While we are not able to comprehensively identify and remove stray traffic from our analysis (e.g., BOGON traffic as result of misconfiguration), we found a prominent case of stray traffic that contributes to INVALID: Traffic from router IP addresses. Recall that routers have multiple interfaces each with its own IP address. A router that sends out a packet (i.e., an ICMP packet) chooses one of these IPs, often arbitrarily [191]. Since the prefixes and corresponding IP addresses for transit links between ASes are not necessarily routed at all or captured by our cone methodology, such packets contribute to INVALID. Using the CAIDA Ark traceroute dataset[192], we extracted router IP addresses from some 500M available traceroutes conducted in February 2017, and tag the corresponding traffic originated from router IP addresses in INVALID.

The scatterplot in Figure 5.8 shows for each member INVALID packets vs. the number of packets with a router source IP address. We find that many members are on, or close to, the diagonal, indicating that most of their INVALID traffic comes from

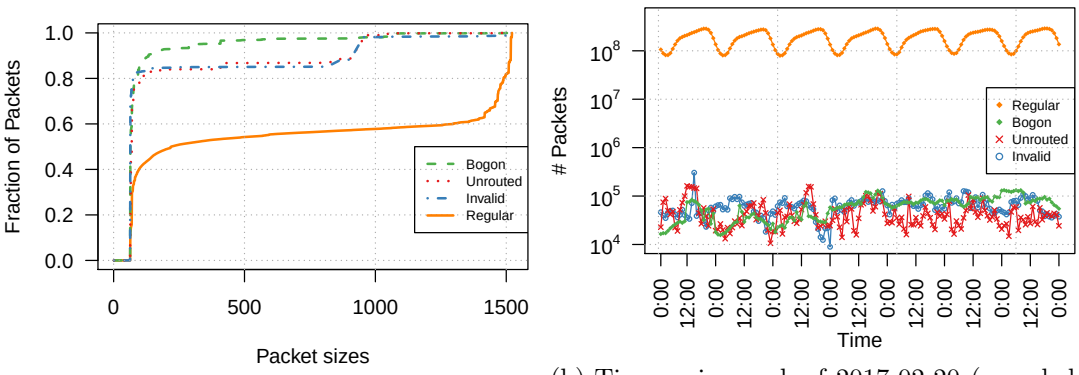

(a) CDF: packet sizes.

(b) Time series week of 2017-02-20 (sampled pkts).

Figure 5.9: Traffic characteristics.

router IPs. While the overall contribution of router IP addresses to our INVALID class is small (less than 1%), we find it highly unlikely that a member whose INVALID packets are dominated by router IP addresses is otherwise a heavy carrier for spoofed traffic. We hence omit members whose INVALID packets consist of 50% or more packets with router IP addresses from our following analysis. This reduces the percentage of members contributing INVALID traffic from 57.68% to 39.59% . Note that this significantly reduces the number of considered members, but not the amount of INVALID traffic.

When looking at the transport layer protocol breakdown of traffic from router IP addresses, we find that about 83% of the packets are ICMP, while UDP and TCP make up for only 14.4% resp. 2.3%. The high percentage of ICMP suggests that a large fraction of this traffic is indeed stray traffic (e.g., *ping* replies from routers). We point out, however, that not all traffic from router IP addresses is necessarily stray traffic: Analysis of the UDP flows shows that 76.3% are destined towards NTP servers with only a small number of source IP addresses, which could indicate attempted reflection attacks on these particular routers (we study amplification attacks in Section 5.5.2). We acknowledge that we might discard some spoofed traffic by not considering members whose INVALID traffic consists primarily of packets with router source IP addresses.

5.5 Traffic Perspective

In this section, we study quantitative and qualitative characteristics of BOGON, UNROUTED, and INVALID traffic. To put our findings into perspective, we contrast characteristics of spoofed traffic with regular traffic exchanged at our IXP.

Timeseries and Packet Sizes Figure 5.9a shows a CDF of packet size distributions for the different traffic classes. While regular traffic shows a typical bimodal distribution, i.e., large data-carrying packets and small ACK packets [193], spoofed traffic consists almost exclusively of small packets. In fact, more than 80% of packets in

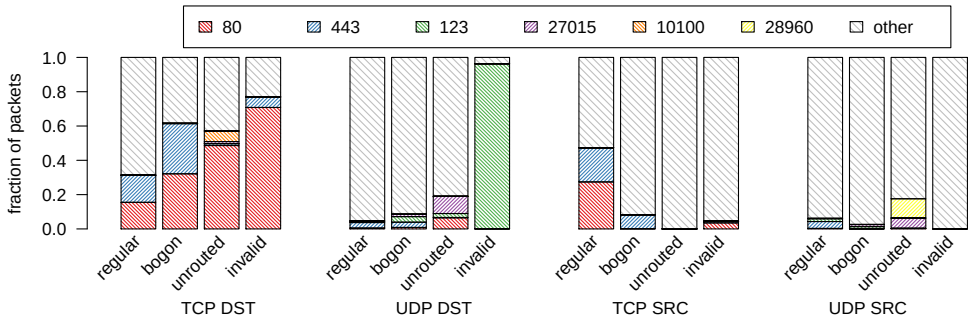

Figure 5.10: Traffic mix for regular, BOGON, UNROUTED and INVALID traffic.

all three classes have a size less than 60 bytes. In the case of TCP, this indicates that these packets do not carry actual data, but are mere connection attempts. This strongly suggests that spoofing is less often used for volume-based attacks, but rather for SYN flooding and amplification attacks, whose return traffic (if any) is regular traffic.

With regards to time-of-day patterns (Figure 5.9b), our three classes of traffic again vastly deviate from regular traffic. While regular traffic shows a typical day pattern, UNROUTED and INVALID traffic show a very unsteady pattern, including significant spikes. This is a first indication that this traffic is mainly caused by attacks, and not part of regular user interaction. BOGON traffic, on the other hand, shows similar irregularities, but a slight time-of-day pattern (pronounced, especially during the first three days). This suggests that BOGON does not exclusively consist of attack traffic, but also contains some stray traffic, i.e, that is likely related to unsuccessful TCP connection attempts from devices in misconfigured NAT environments, triggered by regular user behavior.

Figure 5.10 shows a port-based application classification of packets of our three classes, contrasted with regular traffic exchanged at the IXP. Here, we partition our port-based classification according to *(i)* direction, i.e., SRC and DST port numbers, and *(ii)* the respective transport protocol, i.e., TCP vs. UDP. We only show the six most popular port numbers, and aggregate the remaining port numbers into "other". We note that port numbers in "other" are mostly randomly distributed, suggesting ephemeral port numbers.

In the case of regular HTTP(S) traffic, we expect to see *both directions*: traffic from clients to servers, as well as traffic from servers to clients. Hence, packets from clients to servers carry 80/443 in their DST field, and reply packets from servers to clients carry 80/443 in their SRC field, and an ephemeral port number in the DST field. This interaction is well-reflected when comparing TCP DST and TCP SRC statistics for regular traffic. In the case of spoofed traffic, however, the situation is different: Here, we expect to see only *one direction*, i.e., the spoofed packet to its respective destination. Replies from the server (if any) will not fall into our spoofing categories, since they naturally carry a valid SRC IP address, the servers' address. This observation is well-reflected in our port statistics for spoofed traffic:

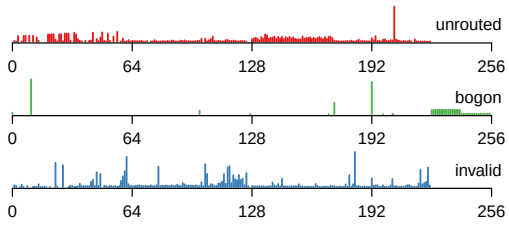

(a) Source Addresses

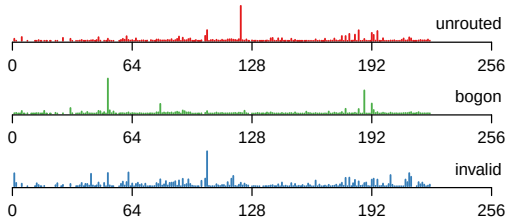

(b) Destination Addresses.

(c) Traffic distribution across address space

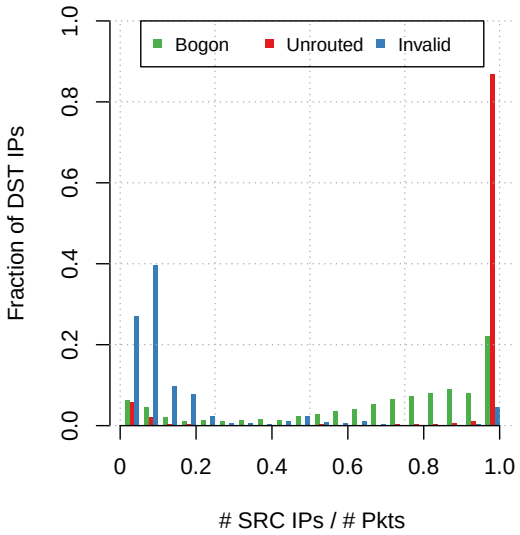

(d) Ratio of #spoofed SRC IPs vs. #packets to DST (DSTs with > 50 sampled pkts).

Figure 5.11: Attack indicators.

The majority of BOGON, INVALID, and UNROUTED packets carry HTTP(S) as DST address, hinting towards flooding attacks destined to HTTP(S) servers. Indeed, we find a corresponding attack pattern, which we study in Section 5.5.2.

The case of UDP traffic is even more intriguing: In the regular case, we mostly find randomly distributed SRC and DST port numbers.[7] Stunningly, we see that the DST port numbers in the case of INVALID traffic are far from randomly distributed: More than 90% of all INVALID UDP packets carry port number 123 as DST and are, hence, destined to NTP servers. Recall that NTP is prone to amplification attacks. We study the related attack patterns in detail in Section 5.5.2. Interestingly, we also notice that while UNROUTED UDP traffic carries mostly random DST port numbers, port 27015 (Steam, online gaming) stands out. A recent study [178] identified this port as commonly attacked.

5.5.1 Address Structure

We next study spatial characteristics of the source and destination IP addresses. Figure 5.11 shows the distribution of packets for each class across the IPv4 address space. Here, we partition the address space in 256 /8 bins and show, for each /8, the number of sampled packets. We observe pronounced differences between our three classes of traffic.

For UNROUTED packets, we find that their source addresses are mostly randomly distributed across the entire address space. The higher density of source addresses

[7] BitTorrent is the dominant UDP-based protocol seen at this vantage point and primarily uses random port numbers [194].

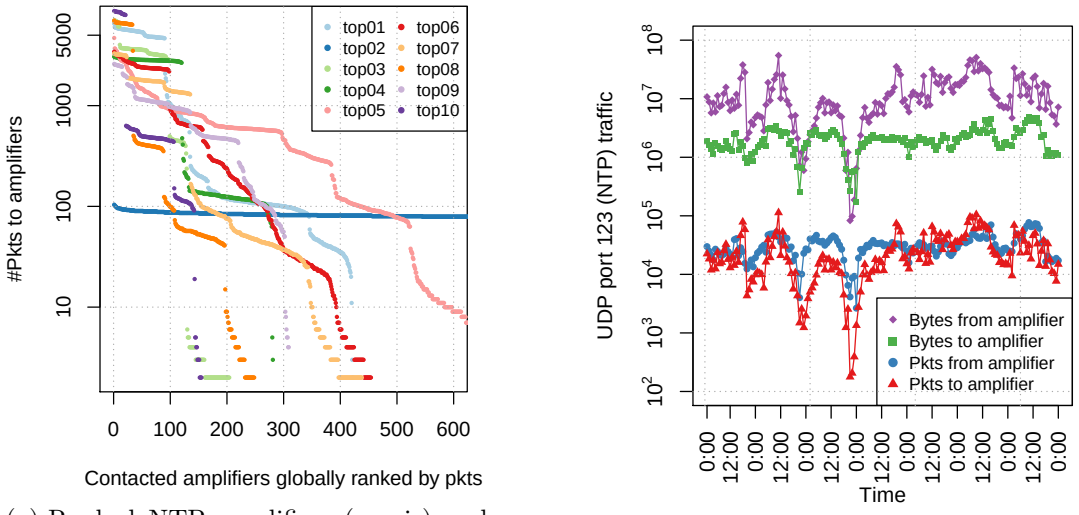

(a) Ranked NTP amplifiers (x-axis) and corresp. pkts (y-axis) for top10 victims.

(b) Time series: NTP attack (INVALID) and response (regular) for week 3.

Figure 5.12: Attack patterns: Selectively vs. uniformly spoofed source IPs.

in some ranges (e.g., from 1/8 to 50/8 and 128/8 - 160/8) is caused by the fact that those address ranges simply have larger amounts of unrouted addresses than others [195]. Address uniformity is a common assumption for spoofed traffic [196]. However, we note that this is not always the case, since we observe one pronounced spike at around 200/8. Destination addresses of UNROUTED packets, however, show strong concentrations on particular address blocks (and, in fact, single addresses, as we show in Section 5.5.2).

For BOGON packets, we see that the source ranges are inherently concentrated on a small subset of the address space (after all, there are only few bogon ranges). The majority falls in private address ranges (with spikes at 10/8 and 192/8). Additionally, we see multicast and, to a lesser extent, "Future Use" (right end). Hence, source addresses are not uniformly distributed across bogon ranges. This suggests that BOGON contains both shares of traffic from likely misconfigured devices (strong concentration in RFC1918 ranges), as well as traffic related to randomly spoofed source addresses (rather uniform distribution in multicast/future use ranges). Indeed, we find that the spikes in destination addresses at 192/8 and 80/8 mostly receive traffic from random IP addresses in the multicast/future use space, suggesting attacks with random BOGON source addresses.

INVALID source addresses differ significantly from the other classes. The distribution shows several peaks, indicating that some specific source addresses are spoofed much more often than others. This is a typical signature of amplification attacks and underlines that address uniformity can not be unanimously accepted as a criterion to identify spoofed traffic [172]. We find large peaks at 183/8 and 61/8. In one of the cases most of the traffic share is due to spoofed addresses routed by a large hosting company which is known to be often targeted by DDoS attacks. Closer inspection of

this traffic shows that it is indeed destined to NTP servers, suggesting amplification attacks. Destination addresses again show large peaks on particular destinations. Here, we expect to see both victims of spoofing with random addresses (potentially in INVALID), as well as targeted amplifiers, e.g., NTP servers. For a detailed analysis see Section 5.5.2.

Summary: Our observations regarding traffic packet sizes, time-of-day-effects, application mix and address structure highlight large differences between regular and spoofed traffic. The characteristics we observe are well in line with different attack patterns carried out with spoofed source addresses. This indicates that our approach is effective in isolating spoofed traffic.

5.5.2 Attack Patterns

The dominant characteristics of traffic with spoofed source addresses suggest the presence of different attacks. We now take a closer look at two common attack patterns, namely amplification and flooding (recall Section 5.1). Recall that flooding attacks are often carried out using a wide range of source IP addresses (random spoofing), while amplification attacks require selective spoofing of source IP addresses of victims.

Selective vs. Random Spoofing To study selective vs. random spoofing events, we first isolate the set of destination IP addresses for which we sampled more than 50 UNROUTED, BOGON, or INVALID packets (8.4K, 19.7K and 9.7K, respectively).[8] Then, we calculate for each destination the ratio of source IP addresses vs. received packets. Figure 5.11d shows a breakdown of this ratio for the destinations, partitioned by the class of traffic. Destinations falling in the leftmost bin received traffic only from very few (or one single) IP addresses. Consequently, IP addresses in the rightmost bin received every single packet from a different source address.

Here, we observe a striking difference when comparing the three classes: Close to 90% of destinations of UNROUTED traffic receive every single packet from a unique source IP address. This highlights that the vast majority of packets with UNROUTED source addresses are due to random spoofing attacking a single destination. In fact, the top 5 destinations receive an upward of 2.3 billion packets over the course of four weeks (sampling extrapolated) from random source addresses. Interestingly, we also see that a significant share of destinations in BOGON addresses show high degrees of source address uniformity, yet with a lesser extent and with some 2.51% that receive significant traffic only from one single IP address. INVALID is the most intriguing case: Here, we see that some comparably small fraction of destination IP addresses receive uniformly spoofed addresses (rightmost bin), but the majority of target addresses receive INVALID traffic from a small set of source addresses (see spikes in the left area of the plot). This is the signature of amplification attacks, where attackers specifically craft packets with spoofed source addresses of their victims, and send packets towards amplifiers.

[8] Note, 50 sampled packets extrapolate to some 500K packets exchanged via the IXP.

NTP Amplification: Recall that INVALID traffic is typically *selectively* spoofed and that the vast majority of INVALID UDP packets is directed to NTP servers. We also found that a single member at the IXP is responsible for 91.94% of all INVALID NTP traffic and the top 5 members together emit more than 97.86% of INVALID NTP. During our observation period, we see NTP trigger traffic from 7,925 individual IP addresses sent by 44 members towards 24,328 possible amplifiers. We compare the list of our 24,328 destinations against a list of some 1.3M NTP servers derived from ZMap scans [197] executed February and March 2017 and find an overlap of 3,865 addresses. Comparing with ZMap scans from December 2016 and January 2017 we find less than 1.8K and 2K hits.

To gain a better understanding of the underlying strategy of some of the largest amplification attacks, we plot in Figure 5.12a for the top 10 victims (i.e., source addresses of trigger traffic) the number of amplifiers ranked by packets (x-axis) and the number of trigger packets sent to each amplifier (y-axis). Here, we observe different attack patterns: Some amplification attacks involve only a handful of amplifiers (90) receiving the bulk of trigger traffic. Other strategies involve using a large number of amplifiers and distributing trigger traffic uniformly across them (as in the case of top-2, 13,377 amplifiers contacted). To assess the effect of amplification, we isolate those IP pairs, for which we are able to see both the trigger traffic to the amplifier, as well as the amplifiers' response packets to the victim. Figure 5.12b shows a timeseries of packets and bytes sent towards amplifiers (trigger traffic), as well as the responses. Here, we see that amplification indeed works: While the number of packets in both directions is similar (and tightly correlated), the number of bytes returned by the amplifiers exceeds the trigger traffic by an order of magnitude. An interesting observation of how amplification attacks manifest at our vantage point.

Summary: Our analysis of attack patterns allows us to illuminate both how attackers carry out flooding and amplification attacks, as well as how these attacks manifest in inter-domain traffic. We see evidence of both random spoofing attacks as well as sophisticated amplification attacks, where attackers rely on different strategies to select amplifiers. In the case of amplification attacks, our vantage point allows us to not only study attack strategies, but to also partially expose their eventual effect on victims, i.e., the resulting traffic from amplifiers.

5.6 Summary

In this chapter, we presented a first approach for passive detection of spoofed traffic. Our method enables us to detect if individual networks allow spoofing, isolate spoofed traffic, and study its properties. We apply and evaluate our approach in practice, studying spoofed traffic exchanged between some 700 networks peering at a major European IXP. Due to its central position in the Internet graph, the IXP gives us a broad perspective with a very diverse member composition. We find that the majority of connected networks do not filter consistently and allow traffic with spoofed source IP addresses to be injected into the Internet. Our analysis of the properties of spoofed

traffic "in the wild" yields hard-to-get insights into both the dominant characteristics of this type of traffic as well as into detailed patterns of attacks carried out with such traffic. Future work includes better recognition of stray traffic and refining the construction of AS-specific prefix lists to achieve tighter bounds when estimating the valid IP space per network. This entails a thorough study of the size and completeness of the BGP-derived address spaces per AS, as well as improving methods to derive additional AS relationships from external data.

6

A year in lockdown: COVID-19 and the Internet

In this chapter, we focus on the insights that traffic-level data can provide. We utilize four data sets, namely eyeball traffic of a large European Internet Service Provider (ISP) and traffic captured at the private peering fabric of one large European, one medium-sized Southern European, and one medium-sized US Internet Exchange Point (IXP) to understand the changes in Internet usage behavior during the COVID-19 pandemic.

The worldwide pandemic caused by the Corona Virus 2019 (COVID-19) is a once-in-a-generation global phenomenon that changed billions of people's lives and destabilized the interconnected world economy. What started as a local health emergency in Asia at the end of 2019 turned into a global event at the beginning of 2020 when the first cases appeared on other continents. By March 2020, the World Health Organization (WHO) declared COVID-19 as a pandemic, causing many governments around the globe to impose strict lockdowns of economic and social activities to reduce the spread of COVID-19. These measures changed the habits of a large fraction of the global population, who from then on depended more than ever on residential Internet connectivity for work, education, social interaction, and entertainment. As of the writing of this thesis, April 2021, the pandemic still affects the entire world.

With the Internet as a crucial entity to keep societies running as well as possible, many people, including leading politicians, were concerned whether the Internet could withstand this extra load. During the pandemic, the importance even made politicians aware of the crucial role the Internet plays in our lives. In March 2020, a commissioner of the European Union urged major streaming providers to reduce their video resolution to the standard definition from March 19, 2020 onward and at least Netflix followed [198, 199]. According to mainstream media, they upgraded their services back to high definition or 4K around May 12, 2020 [200]. This is an exemplary case for the importance of continuous measurement studies because nobody knew if the Internet could withstand the new pandemic-induced demands. The lessons we have learned over the years on how people use the Internet and how usage patterns and demands typically change over time now gives us the perfect base to understand how the pandemic impacts Internet usage behavior.

The profile of a typical residential user—in terms of bandwidth usage and traffic destinations—is one of the most critical parameters that network operators use to

drive their network operations and inform investments [201–203]. Changes in Internet user behavior are common, but they normally occur gradually and over long periods of time. Notable examples of such changes are the increase in demand for peer-to-peer applications that happened in the early 2000s; the increase of traffic served by content delivery networks – including an increase in streaming – that took place in the 2010s; and, more recently, the elevated demand for mobile applications. In all of these cases, the telecommunications industry and network operator community reacted by increasing the investment in network infrastructure. However, the changes in Internet user behavior during the pandemic has been unique because the shifts took place within weeks, leaving hardly any time to react. This raised questions of whether user behavior changes yield to changes in Internet traffic and, more importantly, concerns if the Internet can sustain this additional load. Unexpectedly, the Internet held up to this unforeseen demand [204] with no reports of large-scale outages or failures in more developed countries. This unique phenomenon allows us to observe changes that may be expected within months or years in a matter of days.

In two publications, namely [205] and [21], we investigated the impact of the COVID-19 pandemic on the Internet traffic by analyzing more than two years of Internet traffic data, including the first year of the pandemic. More specifically, we characterize the overall traffic shifts and the changes in demand for particular applications that became very popular in a short amount of time. During the process, we try to understand if there is a "new normal" in Internet traffic and to see how the Internet reacted in these unprecedented times.

This chapter is organized into two parts. First, we summarize our initial observations for the spring 2020 wave (February 2020 to June 2020)[1] highlighting the first drastic changes we see in correlation with the lockdown periods across Europe. In the second part, we show additional observations during the fall 2020 wave of the pandemic (September 2020 to February 2021) and try to understand which changes are persistent.

To that end, we collected and analyzed network traffic data from a large Internet Service Provider (ISP) in Europe, three Internet Exchange Points (IXPs) in Europe and the US, as well as a mobile operator[2].

Our two main observations across one year of pandemic can be summarized as follows:

COVID-19-induced traffic growth. We observe a significant traffic evolution in 2020 across all our vantage points. Figure 6.1 gives an overview of the timeline of the pandemic and highlights the week of the initial lockdowns in the countries where the ISP and IXPs in Central and Southern Europe, and the mobile operator are located. Although the exact lockdown dates differed across Europe, these dates are very close to each other in mid of March 2020.

[1] We use "spring" and "fall" from the viewpoint of the Northern hemisphere, where our vantage points are located. Exchange both terms for the Southern hemisphere.

[2] The educational network, which we analyze in both publications, is beyond the scope of this thesis

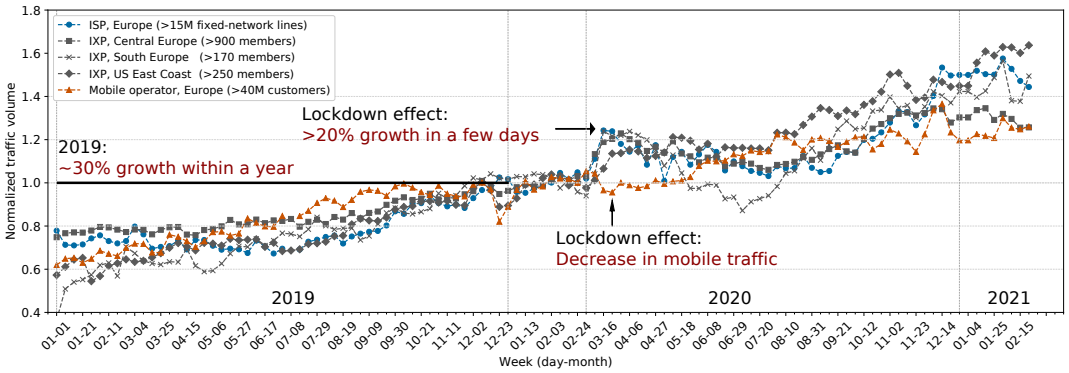

Figure 6.1: Traffic changes during the COVID-19 pandemic's spring and fall waves at our Internet vantage points.

The COVID-19 outbreak reached Europe in late January (week 4) and first lockdowns were imposed in mid March (starting on week 11). Thus, we normalize weekly traffic volumes by the median traffic volume of the first ten weeks of 2020 (pre-lockdown period). We can clearly identify drastic changes in the data collected at multiple and diverse vantage points (see Section 6.1 for details). The first observation is that the ISP (blue squares) and IXPs in Central and Southern Europe (gray lines) show a more than 20% traffic increase within a week after the official announcement of the lockdown. This could be perceived as a "moderate" surge in traffic. However, in Internet-reality, this is a substantial increase in only a short period of time. To put it into perspective, the figure shows that the annual increase of 2019 was around 30%, which is similar to the annual increase in previous years. This means that the expected traffic increase in one year happened only within a couple of weeks in March 2020 following the spring 2020 lockdown. This observation alone justifies a closer investigation of the new dynamic of Internet traffic patterns as it highlights (a) that our day to day life and the Internet are so intertwined that a change in one may be reflected in the other and (b) the engineering challenges in keeping the Internet running.

During the fall 2020 lockdown we see, that demand changed provoke changes in traffic volume, causing a traffic surge of 15-20% for the ISP/IXPs in our study. In summer 2020, after the first reopening of the economy and educational facilities in this region, an increase of about 20% at one IXP, but only 6% at the Tier-1 ISP, are still visible. The fall 2020 wave also had an impact, with the annual traffic increase in 2020 being about 10% higher than in previous years. Our findings are aligned with the insights offered by mobility reports published by Google [206] and the increased digital demand as reported by Akamai [207, 208], Comcast [209], Google [210], Nokia Deepfield [211], and TeleGeography [212].

When we look at the mobile operator (orange triangles), we observe an anti-pattern. While the traffic at all other vantage points increased after the lockdown in March 2020, the mobile operator experiences a drastic decrease in traffic. The same pattern repeats itself at the end of 2020 during the next lockdown periods. This observation

reflects the reduced mobility across Europe and aligns with reports by the COVID-19 Mobility Project [213].

Drastic shift in usage patterns. In light of the global COVID-19 pandemic a total growth of traffic is expected. More relevant for the operations of networks is how exactly usage patterns are shifting, e.g., , during the day or on different days of a week. In Section 6.2 we discuss shifts in the diurnal pattern for the spring 2020 wave. Especially during the first lockdown in March 2020, we see that high traffic levels already occur earlier in the day compared to February 2020, see Figure 6.2. This correlates well with people working from home: Instead of leaving for school and the office, they start their workday at home relying on their local Internet connection. In Section 6.5 we augment these findings with an analysis how usage patterns on workdays vs. weekdays change. The Internet's regular workday traffic patterns are significantly different from weekend patterns [214–216]. On workdays, traffic peaks are concentrated in the evenings as shown in Figure 6.8 for the IXP in Central Europe and the ISP-CE. We find, that up to mid-March, most weekend days are classified as weekend-like days and most workdays as workday-like days. The only exception is the holiday period at the beginning of the year in Figure 6.8b. This pattern changes drastically once the confinement measures are implemented: Almost all days are classified as weekend-like. This change persists across all periods where contacts were drastically limited. Both vantage points show, that the shift towards a weekend-like pattern becomes less dominant as countermeasures were relaxed in September 2020 and manifests again in December 2020 when new restrictions were in place.

These observations raise the question of the extent and cause for this significant traffic growth and shift in patterns. Given that many people are staying at home for all purposes, it is likely, that working from home, remote education, performing online social activities, or consuming entertainment content, are the main contributors to the observed growth. In the following sections, we approach these questions from different angles: (a) transport port and application mix, (b) changes in hypergiant traffic, (c) port capacity upgrades at the IXP, and (d) a closer look at shifts in the diurnal pattern.

6.1 COVID-19: Datasets

This section describes the network traffic datasets that we used for our analysis discussed in this chapter. We utilize vantage points at the core of the Internet (Internet eXchange Points) and the backbone and peering points of a major ISP, all which we describe below.

ISP-CE: Network flows from a large Central European ISP that provides service to more than 15 million fixed-line subscribers and also operates a transit network (Tier-1). The ISP does not host content delivery servers inside its network, but it has established a large number of peering agreements with all major content delivery and cloud networks at multiple locations. This ISP uses NetFlow [217] at all

	ISP-CE	IXP-CE	IXP-SE	IXP-US
base	Feb 20–26 '20	Feb 20–26 '20	Feb 20–26 '20	Feb 20–26 '20
March	Mar 19–25 '20	Mar 19–25 '20	Mar 12–18 '20	Mar 19–25 '20
April	Apr 09–15 '20	Apr 23–29 '20	Apr 23–29 '20	Apr 23–29 '20
June	Jun 18–24 '20	Jun 18–24 '20	Jun 18–24 '20	Jun 18–24 '20
September	Sept 10–16 '20	Sept 10–16 '20	-	-
November	Nov 12–18 '20	Nov 12–18 '20	-	-
December	Dec 10–16 '20	Dec 10–16 '20	-	-
January	Jan 14–20 '21	Jan 14–20 '21	-	-

Table 6.1: Summary of the dates used in weekly analyses in the spring wave (March '20—Jun '20) and the fall wave (Sept '20—Jan '21). Dates in Southern Europe vary due to different courses of the pandemic. For the fall wave we only show data from the ISP-CE and IXP-CE.

border routers to support its internal operations. We rely on two different sets of NetFlow records for this study. First, we use NetFlow data collected at ISP's Border Network Gateways [218] to understand the impact of changing demands of the ISPs' subscribers. Second, we use NetFlow records collected at the ISP's border routers to gain a better understanding of how companies running their own ASNs are affected by these changes.

IXPs: Network flows from the public peering platform of three major Internet Exchange Points (IXPs). The first one has more than 900 members, is located in Central Europe (IXP-CE), and has peak traffic of more than 8 Tbps. The IXP-CE is located in the same country as the ISP-CE. The second one has more than 170 members, is located in Southern Europe (IXP-SE), and has peak traffic of roughly 500 Gbps. It covers the region of the EDU network. The third one has 250 members, is located at the US East Coast (IXP-US) and has peak traffic of more than 600 Gbps. At the IXPs, we use IPFIX data [46].

We augment our analysis with NetFlow records from a large mobile operator that operates in Central Europe, with more than 40 million customers.

Normalization: Since all data sources exhibit vastly differing traffic characteristics and volumes, we normalize the data to make it easier to compare. For plots where we show selected weeks only, we normalize the traffic by the minimum traffic volume. For plots spanning a larger timeframe, we normalize the traffic by the median traffic volume of the first ten weeks of 2020, depending on the availability of data.

Time frame: We use two methods to reflect the developments since the beginning of the COVID pandemic: (a) for general trends over time we use continuous data from *Jan 1, 2020—Jun 24, 2020*, (b) to highlight detailed developments we compare 7-day periods as shown in Table 6.1 from before, during, after and well after the lockdown in 2020.[3]

[3]Due to data availability, the ISP-CE is using Apr 09–15 which covers the Easter holiday period.

6.1.1 Ethical Considerations

Both NetFlow and IP Flow Information Export (IPFIX) data provide only flow summaries based on the packet header and do not reveal any payload information. To preserve users privacy, all data analyses are done on servers located at the premises of the ISP and IXPs. The output of the analyses are the aggregated statistics as presented in this chapter. The data at the ISP and IXPs is collected as a part of their routine network analysis.

6.2 A glance at the first wave

In this section we focus on the spring wave 2020 of the COVID-19 pandemic. When the first lockdown was put into place in March, nobody was really aware of the upcoming developments. As such, we especially focus on the transition from February 2020 to mid-March 2020. Namely, we analyze traffic data from March 2020 to June 2020. First we take a look at the overall traffic shifts before, during and after the first lockdown period. Based on this observations, we then present our more detailed analyses on hypergiant ASes, ASes relevant for remote working, and shifts in link utilization at the IXP-CE.

6.2.1 Macroscopic Analysis

Figure 6.2 plots the aggregated normalized traffic volume in bytes at the granularity of one hour for the ISP-CE, IXP-CE, IXP-US, and IXP-SE in four selected weeks in spring 2020 (see Table 6.1). For the ISP-CE, Figure 6.2a shows the time series using normalized one-hour bins. For the IXPs, Figure 6.2b reports the hourly average for workdays and weekends.

First of all, we see that the overall traffic after the lockdown increased by more than 20% for the ISP-CE and 30%/12%/2% for the IXP-SE/IXP-CE/IXP-US, respectively. Once the lockdown measures were relaxed, the growth started declining for the ISP-CE but persisted for the IXP-CE and the IXP-SE. These differences are most likely attributed to the fact that the ISP-CE traffic pattern is dominated by end-user and small enterprise traffic—recall, we are not analyzing any transit traffic—while the IXP-CE has a wider customer base. Traffic persistently increased for the IXP-US where the lockdown was put into place later.

At the ISP-CE a new pattern emerges. The time series shows a shift from workday to weekend traffic patterns starting with the lockdown in mid-March. In the past user behavior at the ISP showed distinct differences between workdays (Monday-Friday) and weekends (Saturday-Sunday). This is a natural consequence of the classic western work week: Between Monday and Friday people leave their homes for work

As partial lockdowns and travel restrictions were still in place, the introduced bias may be very small.

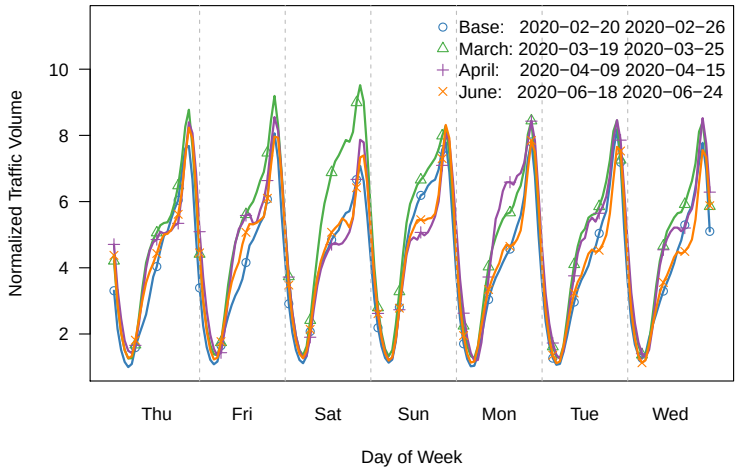

(a) L-ISP (Central Europe).

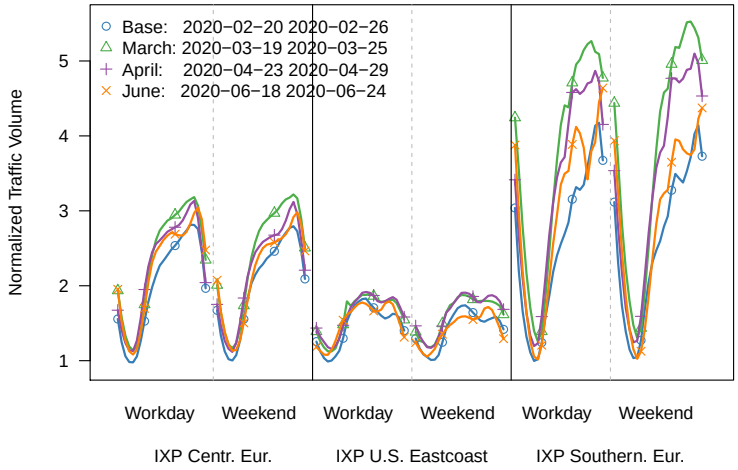

(b) IXPs (Central Europe/US Eastcoast/Southern Europe).

Figure 6.2: Time series of normalized aggregated traffic volume per hour for ISP-CE and three IXPs for four selected weeks: before, just after, after, and well after lockdown (base/March/April/June).

and school, while on the weekends they spend more time at home. In accordance with that observation, traffic increases much earlier in the day with a small dip at lunchtime. However after lunch hours, traffic grows to roughly the same volume during the evening time, spiking late in the evening. This change persists throughout the lockdown. Once this was relaxed, the pattern became less pronounced and the shift to a weekend like pattern became less dominant. Additionally, it is important to note 1) the Easter vacations in the April week, and 2) the seasonal effects in the weekend of the June week (an increase of outdoor activities).

For all IXPs, see Figure 6.2b, not only do we see an increase in peak traffic but also in the minimum traffic levels. This correlates with link capacity upgrades of many IXP members leading to overall increases of 3% at IXP-CE, 12% at IXP in Southern Eu-

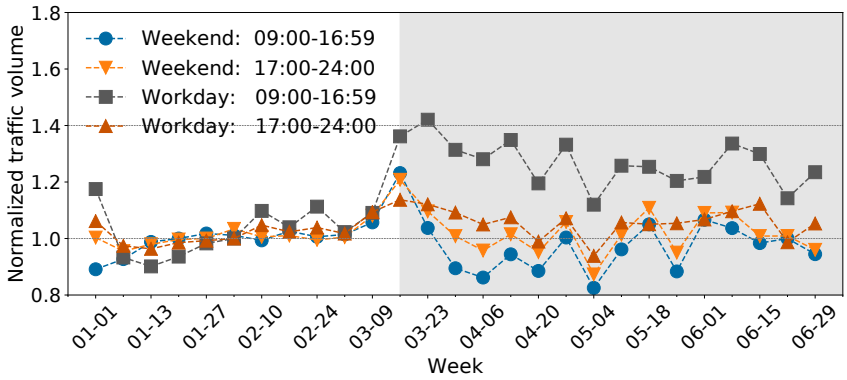

(a) Hypergiants

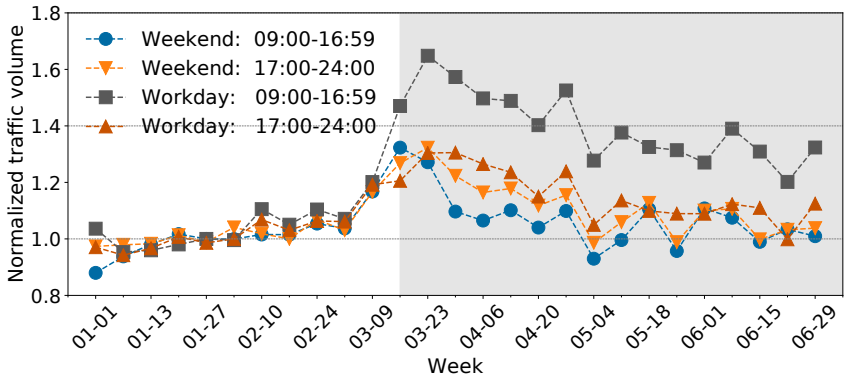

(b) Other ASes

Figure 6.3: ISP-CE: Normalized daily traffic growth for hypergiants vs. other ASes across time.

rope and 20% at IXP at the US East Coast. In addition, we see the increase in traffic during daytime, which is very pronounced at the IXP-CE. However, the differences between weekends and workdays are not as apparent as at the ISP. Interestingly, as lockdown measures were mandated, the daytime traffic again decreases but stays well above the pre-lockdown level. In contrast, traffic at the IXP-US barely changes in March and increases only in April, otherwise showing similar effects as the other IXPs. The delayed increase in volume is likely due to the later lockdown in the US. Overall, the effects of the time of day at this IXP are less pronounced compared to the two others because it (a) serves customers from many different time zones, and (b) members are diverse and include eyeball as well as content/service providers. In contrast, the IXP-SE interconnects more regional networks, and as such the traffic patterns are closer to the ones of the IXP-CE.

6.2.2 Hypergiants

To understand the composition of residential traffic, we investigate who is responsible for the traffic increase at the ISP-CE. The first step is to look at the top 15

hypergiants [219–221]. Hypergiants are networks with high outbound traffic ratios that deliver content to approximately millions of users in the locations at which we have vantage points. The 15 hypergiants we consider in this study are responsible for about 75% of the traffic delivered to the end-users of the ISP-CE, which is consistent with recent reports in the literature [201, 222, 223]. We note that the fraction of hypergiant traffic vs. traffic from other ASes does *not* change drastically for the ISP-CE as well as all IXPs.

Given that the overall traffic has increased, we next report the relative increase of the two AS groups compared to the median traffic volume during the pre-lockdown period, see Figure 6.3. In detail, we focus on different times of day and days within the week. We find that the relative traffic increase is *significantly* larger for *other ASes* than for hypergiants.

Both sets of time series are more or less on top of each other until the lockdown. This observation also holds for data from 2019 (not shown). However, after the lockdown, the time series for the other ASes present higher deviations from the reference value than those of the hypergiants. The most visually striking difference occurs during working hours of work-days: Hypergiants experience a 40% increase whereas the remaining ASes grow by more than 60%. While this difference is significantly reduced around mid-May, the relative increase for both sets of ASes is still substantial. In fact, except for the working hours during work-days, the traffic surge seems to normalize around mid-May, especially for *other ASes*. Notice the fluctuations during weekends mornings starting around the end of April—they can be also observed in 2019 (not shown).

A plausible explanation for the increase of daily traffic volumes in this vantage point are family members being forced to continue their professional and educational activities from home. ainly video streaming—explains the increase in traffic volume associated with hypergiants, many of which offer such services. The increase in traffic by the other ASes has more facets and it requires a more thorough analysis that incorporates traffic classification methods. Before doing that, the next subsections investigate the impact that these ASes have on parts of the infrastructure of some of our vantage points.

6.2.3 Link Utilization Shifts

We analyze to which extent the observed changes are reflected in our link utilization dataset to assess how many networks suffer changes in their traffic characteristics. For this, we look at changes in relative link utilization between the base week in February and the selected week in March. We choose IXP-CE as reference vantage point as it houses the greatest variety of connected ASes, thus allowing a more complete and meaningful analysis. Our dataset reflects link capacity upgrades as well as customers switching to PNIs. We plot the minimum, average and maximum link utilization for all members at IXP-CE in Figure 6.4.

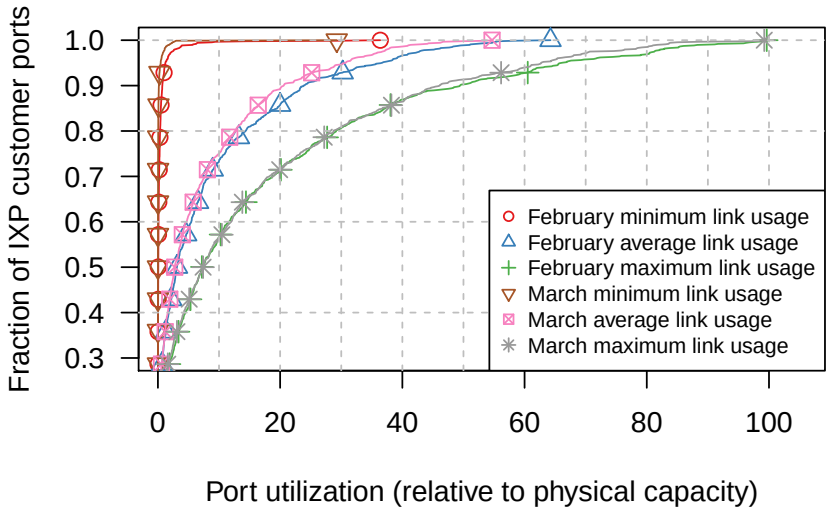

Figure 6.4: IXP-CE: ECDF of link utilization before and during the spring lockdown.

Figure 6.4 shows a slight shift to the left during lockdown. This denotes a tendency towards decreased link usage across many IXP members which could be caused by link capacity upgrades or members switching to PNIs in response to increased traffic demand [211]. It is important to note that increased link usage of a network can be concealed by another network upgrading its port. However, the main takeaway is that many of the non-hypergiant ASes show changes in their link usage due to the lockdown-induced shifts in Internet usage. To gain a better understanding of this phenomenon, we reconsider the non-hypergiant ASes and their role in the Internet for further analysis.

6.2.4 Remote-work Relevant ASes

Having observed that the relative increase in traffic during working hours is more pronounced for non-hypergiants ASes, we study temporal patterns to identify which ASes are relevant for remote work, e.g., large companies with their own AS or ASes offering cloud-based products to be used by their employees. To this end, we use the ISP-CE dataset, including its transit traffic, to compute the received and transmitted traffic per ASN.[4] In addition, we compute the traffic that each one of them sends and receives to/from manually selected eyeball ASes, i.e., the large broadband providers in the region. Using this data, we define three distinct groups of ASes: those whose traffic ratio of workday/weekend traffic is dominated by workdays, those who are balanced, and those in which weekend traffic patterns dominate.

We focus on the first group, as we expect companies and enterprise subscribers deploying remote working solutions for their employees to fall into this class. We cross-check their AS numbers with the WHOIS database. We find that a small number of

[4]We are aware of limitations of this vantage point, e.g., companies may have additional upstream providers.

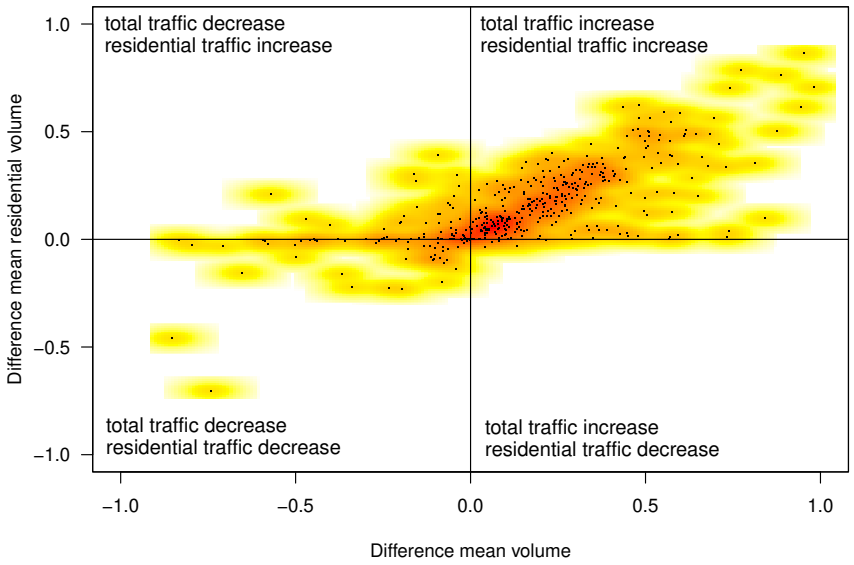

Figure 6.5: ISP-CE: Heatmap of traffic shift vs. residential traffic shift (Feb. vs. Mar.).

content-heavy ASes also fall in this category. In Figure 6.5 we show the difference in normalized traffic volumes between a base week in February and one in March after the lockdown began (x-axis) vs. the normalized difference in traffic from/to eyeball ASes. We observe that some ASes experience major traffic shifts, but with almost no residential traffic (scattered along the x-axis, and close to 0 in the y-axis). However, for a majority of the ASes, there is a correlation between the increase in traffic involving eyeball networks and the total increase. This suggests that most of the traffic change is due to eyeball networks. Interestingly, some ASes suffer a decrease in total traffic, yet residential traffic grows (top-left quadrant). These are likely companies that either offer online services that became less popular and relevant during the lockdown or that do not generate traffic to the Internet "internally". When looking at the other AS groups (not shown), the correlation still exists but is weaker.

These observations help us to put the implications of the lockdown measures in perspective: Some ASes need to provision a significant amount of extra capacity to support new traffic demands in an unforeseen fashion. In the following sections, we will explore which specific traffic categories have experienced most dramatic changes.

6.3 Transport-Layer Analysis

Based on the overall traffic pattern shifts identified in Section 6.2.1, in this section we focus on differences in raw transport port-protocol distributions.

We analyze the shift in application traffic due to the lockdown at two vantage points, the ISP-CE and the IXP in Central Europe (IXP-CE). At both networks, we aggregate traffic volume statistics from four weeks described in Table 6.1. For each hour of the day, we keep separate traffic volume statistics and then compare these to the

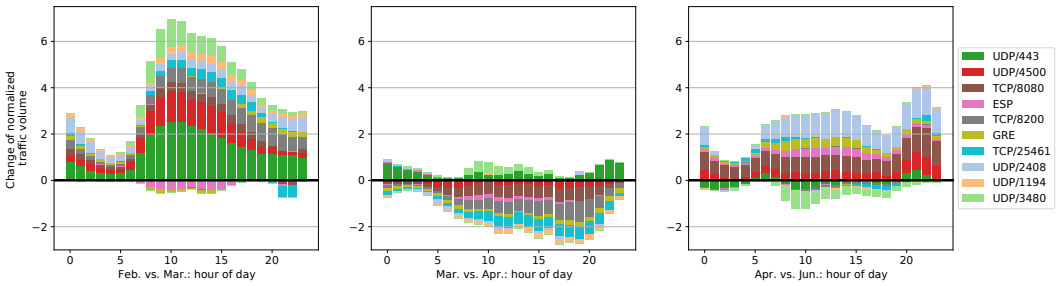

(a) IXP-CE traffic difference by top application ports: normalized aggregated traffic volume difference per hour comparing the workdays of February, March, April, and June. We omit TCP/80 and TCP/443 traffic for readability purposes.

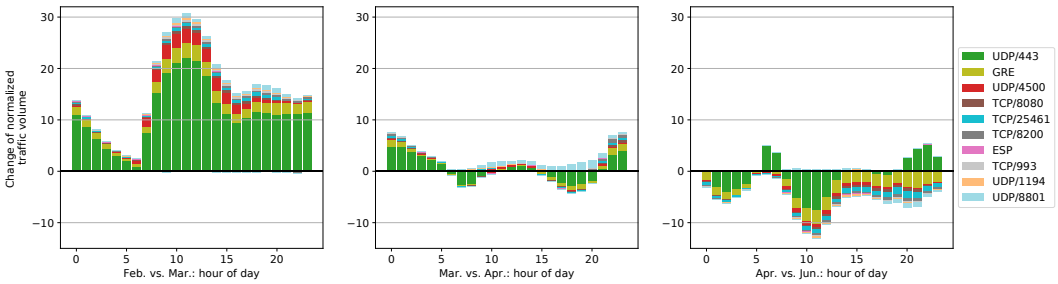

(b) ISP-CE traffic difference by top application ports: normalized aggregated traffic volume difference per hour comparing the workdays in February, March, April, and June. We omit TCP/80 and TCP/443 traffic for readability purposes.

Figure 6.6: Top application ports at the ISP-CE and IXP in Central Europe during the spring wave.

respective day and hour of the previous month, which allows us to identify diurnal patterns, and more importantly, changes therein.

We plot the top transport ports for each vantage point. As the two most common ports TCP/443 and TCP/80 make up 80% and 60% of traffic at the ISP-CE and IXP-CE, respectively, any small changes in their traffic volume would dominate the plot. Therefore, we omit those from Figures 6.6a and 6.6b for clarity purposes. [5]

We instead focus on the top 3–12 ports. Figure 6.6a depicts changes in traffic volume per transport-layer port for the IXP-CE, and Figure 6.6b for the ISP-CE. Note that we aggregate the hours of day of all working days of a week into a single subplot.

While both networks share similar top ports, their distribution, and the changes in these distributions over time, are very different. This reflects the different types of customers present at these vantage points. Recall, that the ISP-CE dataset consists of subscriber traffic, which is largely composed of end-users and small enterprises, while the IXP-CE one has a very diverse set of members across the entire Internet economy exchanging traffic over its platform. In general, we see a very strong increase at the IXP-CE as well as at the ISP-CE when comparing the changes in March (leftmost

[5] We also consider alternative HTTP port TCP/8080, rendered in the figures, but we do not observe any significant change in its usage.

subplots), compared to the more gradual changes in the following months (middle and rightmost subplots).

Next, we analyze in-depth specific ports to more accurately attribute overall changes in diurnal patterns:

QUIC: Running on port UDP/443, QUIC is mainly used for streaming purposes by e.g., Google and Akamai [224]. QUIC traffic increases 30%–80% at the ISP-CE and about 50% at the IXP-CE when comparing traffic volumes in March with the base week of February. Once the lockdown starts, we see the largest increase at the ISP-CE in the morning hours. Moreover, at the IXP-CE the increase is more gradually distributed over the day. This likely reflects the behavior of entire families staying at home. In the months of April and June the traffic volumes of QUIC remain relatively stable, with some hours gaining traffic while other losing some.

NAT traversal / IPsec / OpenVPN: Port UDP/4500 is registered at IANA for IPsec NAT traversal and is commonly used by VPN solutions, Port UDP/1194 is OpenVPN's default port. As more people are working from home and using VPNs to access their company or university network, we see an increase of both ports during working hours at the two vantage points in March. In the following weeks in April and June the traffic volumes for UDP/4500 stay above the traffic volume of the February base week, whereas OpenVPN's volume recedes. Interestingly, GRE and ESP, which transport the actual IPsec VPN content, decrease at the IXP-CE in March during the lockdown, while GRE traffic sees a slight increase at the ISP-CE. To summarize, more people are using VPNs from their homes resulting in an increased need of NAT traversal, but VPN connections between companies which are the primary source of GRE and ESP traffic decrease over time. [6].

TV streaming: On port TCP/8200 at the IXP-CE we see, similar to QUIC, how changes in user behavior affect the traffic profile. This port is used by an online streaming service for Russian TV channels. In March, we notice traffic volumes increasing throughout the day, shifting away from an evening centric traffic profile. We mainly observe this at the IXP-CE as it serves a broader and more international customer base. Additionally, the strong increase in March is not persisting over the following months.

Cloudflare: Port UDP/2408 is used by the CDN Cloudflare for their load balancer service [225]. We verify that the traffic indeed originates from Cloudflare prefixes. During our observation period, we see an increase in Cloudflare load balancer traffic at the IXP-CE in March and in June.

Video conferencing: The video communication tool Skype and the online collaboration service Microsoft Teams both use port UDP/3480, most likely for STUN purposes [226, 227]. We confirm this by verifying that the addresses reside in prefixes owned by Microsoft. Additionally, we find a small number of non-Microsoft addresses in our data. During the lockdown in March, we see a large increase in UDP/3480 traffic at the IXP-CE, especially during working hours on workdays. At the ISP-CE

[6] For an in-depth analysis of VPN traffic shifts, we refer to [21, 205]

it does not show up among the top 12 transport layer ports. Zoom, another video conferencing solution, uses UDP/8801 for its on-premise connector which companies can deploy to route all meeting traffic through it [228]. At the ISP-CE this traffic increases by an order of magnitude from February to April. Since Zoom only became popular in Europe due to the lockdown, this drastic increase reflects the adoption of a new application by companies deploying connectors in their local network. These changes once again underline the fact that people working from home do change the Internet's traffic profile. Zoom traffic decreases again in June, which might also be related to the vacation period resulting in fewer online office meetings.

Email: At the ISP-CE, especially during working hours, we find a 60% increase in TCP/993, which is used by IMAP over TLS to retrieve emails. While the overall amount of traffic is small compared to, e.g., QUIC, it is nevertheless an additional indicator for people conducting their usual office communication from their homes.

Unknown port: We could not map TCP/25461 to any known protocol or service. The addresses using this port mostly reside in prefixes owned by hosting companies.

To summarize, we find significant changes in the traffic profile for some popular transport-layer ports at both vantage points. This highlights the impact of drastic human behavior changes on traffic distribution during these weeks. We see an increase in work-related as well as entertainment-related traffic, reflecting the lockdown where people had to work and educate from home. This rationale is supported by the significant shift in workday patterns, especially at the ISP-CE from February to March when the lockdown began. As more people stay at home, the traffic levels which are dominated by residential customers increase steeply in the morning, compared to the steady growth observed over the whole day in February.

6.4 Application classes

Building on the analysis of the raw ports presented in the previous section, we now provide a more in-depth analysis of traffic shifts for different *application classes*. This is especially relevant for traffic using protocols such as HTTP(S), where a single transport-layer port number hides many different applications and use cases.

To investigate application layer traffic shifts, we apply a traffic classification based on a combination of transport port and traffic source/sink criteria. In total, we define more than 50 combinations of transport port and AS criteria based on scientific-related work [220, 229], product and service documentations [226, 227, 230, 231], and public databases [28, 232].

We aggregate the filtered data into 8 meaningful application classes representing applications consumed by end-users on a daily basis (see Table 6.2): *Web conferencing and telephony (Web conf)* covers all major conferencing and telephony providers, *Collaborative working* captures online collaboration applications, *Email* quantifies email communication, *Video on Demand (VoD)* covers major video streaming services,

Application class	# of filters	# of distinct ASNs	# of distinct transp. ports	Notes
Web conferencing and telephony (Web conf)	7	1	6	Conferencing audio/video ports, AS-based for pure conferencing provider (TCP:444, UDP:3478-3481, UDP:8200, UDP:5005, UDP:1089, UDP:10000)
Video on Demand (VoD)	5	5	-	Large to medium VoD provider ASes
Gaming	8	5	57	Transport ports of popular games , AS-based for large gaming providers (e.g. TCP:1716, TCP:4001, TCP:3074, ...), includes cloud gaming services
Social media	4	4	1	Social networks including their respective CDNs (HTTPs+respective AS)
email	1	-	10	Typical mail transport ports (TCP:25, TCP:587, TCP:109, TCP:110, TCP:143, TCP:220, TCP:645, TCP:585, TCP:993, TCP:995)
Educational	9	9	-	ASes of universities close to respective vantage points
Collaborative working	8	2	9	Collaborative editing, file sharing, versioning, VPN, remote administration (e.g. TCP:1194, UDP:1194, UDP:1197, UDP:1198, ...)
Content Delivery Network (CDN)	8	8	-	Dominant CDN providers (excluding social network CDNs) by AS

Table 6.2: Overview of filters for the application classification. Filters are based on transport ports or ASes , either in combination or separately. Used to classify data in Figure 6.7.

Gaming captures traffic from major gaming providers (cloud and multiplayer), *Social media* captures traffic of the most relevant social networks, *Educational* focuses on traffic from educational networks, and *Content Delivery Networks (CDN)* classifies content delivery traffic. Note that social networks, e.g., Facebook, also offer video telephony and content delivery services for their own products, which may be captured by this class but not by the more specific other classes.

We perform the application classification for the different IXP vantage points (IXP-SE, IXP-CE, IXP-US) and for the ISP-CE.[7] To clearly present the large amount of information, we transform the data as follows.

Week-wise comparison: We focus our analysis on four weeks, a *base week* well before the lockdown, to which we compare three weeks representing the different

[7] In case of the ISP-CE we analyzed upstream as well as downstream traffic. As the differences between the weeks manifest in both directions in a very similar fashion we only show the downstream direction.

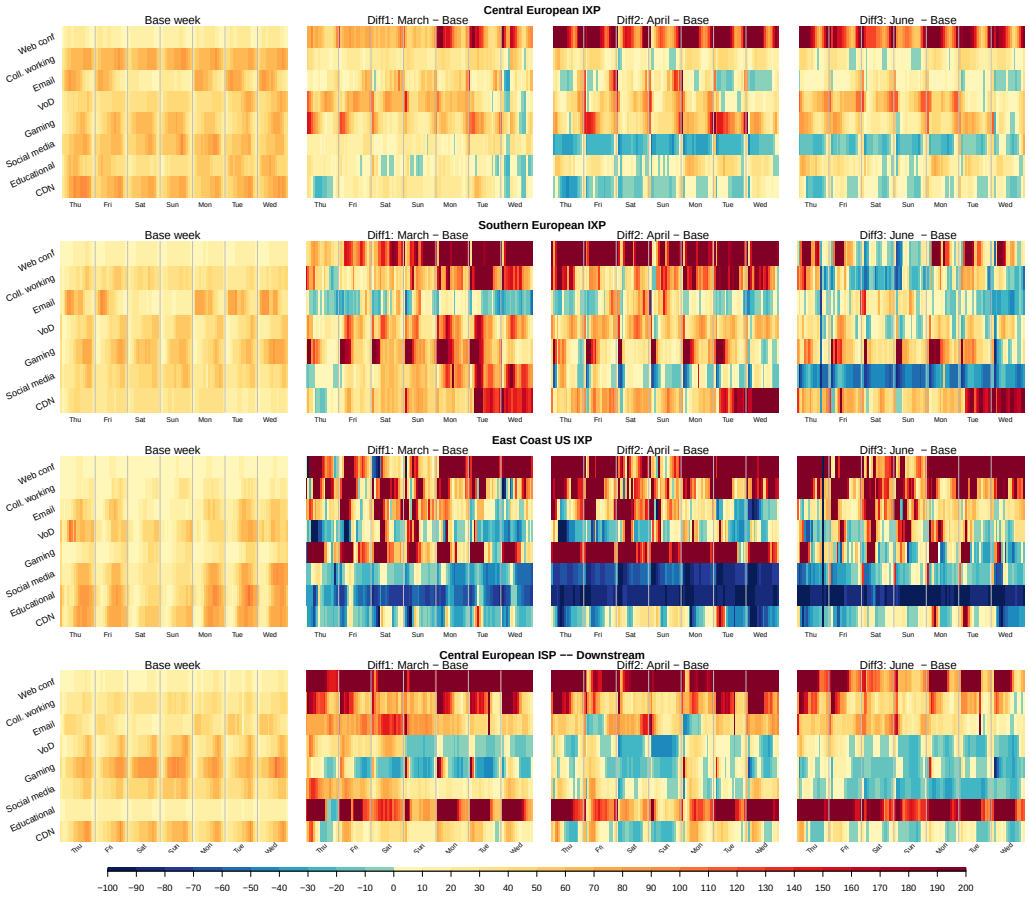

Figure 6.7: Heatmaps of application class volume for three different IXP locations and the ISP-CE.

stages of the COVID-19 measures as they were imposed throughout Europe—see Table 6.1 in Section 6.1.

Normalization and filtering: After normalization as outlined in Section 6.1, we remove the early morning hours (2–7 am). The *total* volume of the vantage points hits its daily minimum during these hours, but does not change much during the lockdown. Removing these hours allows us to visualize more details of traffic shifts during the day in order to compare application classes of different traffic volumes as well as the relative growth between the *base week* and the other weeks.

Difference to base week: We visualize each week as the difference of the respective week and the *base week*. This enables quick visual identification of increased/decreased application class usage compared to pre-COVID times. We remove any growth above 200% and any decrease below 100%.

The condensed timelines of the different application classes are shown in Figure 6.7 for all four vantage points. We highlight our main observations next:

Communication-related applications: At all vantage points, *Web conferencing* applications show a dramatic increase of more than 200% during business hours, and at the ISP-CE, IXP-SE, and the IXP-US also on the weekends. In this category the ISP-CE experiences the largest growth in *March* right after the lockdown across all hours of the day. In *June* this trend is less pronounced, which corresponds with people slowly going back to their offices. *Collaborative working* mainly increases at the IXP-SE and the IXP-US, at the ISP-CE we see a vast increase on Thursday and Friday morning which persists until *June*—this might be due to coordination between work partners before the weekend. While in a lockdown situation one might expect a lot of additional *Email* communication, we see a different trend. At the IXP-CE and the IXP-SE *Email* actually declines during the lockdown and in *June* remains on a lower level than before the lockdown. Instead, *Email* rises at the ISP-CE it, but not as high as other traffic classes as *Web conferencing*. One possible explanation could be that many companies start connecting their remote employees via Virtual Private Networks (VPNs) and users connect to the mail systems via the VPN [8]. For the IXP-US the trend is less pronounced, and we see phases of usage increase and decrease over time.

Entertainment related applications: *VoD* streaming application usage shows high growth rates at the European IXPs of up to 100%. Interestingly, ISP-CE only sees a slight growth of about 10% during the lockdown, while in *June* – well after the lockdown – the traffic volume drops back to the February level. Recall that the major streaming companies reduced their streaming resolution in Europe by mid-March [198] for 30 days. In the case of the ISP-CE that covers the *March* as well as the *April* week.[9] In the US, the trend is the other way around. Notably, this may be a biased measurement, as at the IXP-US the measurement of the *VoD* class is based on only three ASes, one of which is very large. Consequently, the decrease may reflect a traffic engineering decision of the large AS, e.g., establishing a private network interconnect instead of peering. The strong growth of *gaming* applications is more coherent across all three IXP vantage points, especially during the day. While the ISP-CE shows a significant increase during morning hours, it generally leans towards declining. Note, that this effect is mainly caused by unusually high traffic levels in this category in February. Gaming applications, typically used in the evening or at weekends, are now used at any time. The trend starts to flatten in June—this may in relation with people going on vacation or spending more time outside. Moreover, we see an increase at the IXPs for *Social media* application traffic during the *March* week, while the effect quickly diminishes in *April*. In *March* the ISP experiences a 70% growth, which slows down in *April* but not as drastic as at the IXPs. The effects in this class correlate with the gradual de-escalation of the lockdown restrictions in Europe: as people are allowed to leave their homes freely again and resume social live, this traffic decreases. In June, social media usage has returned to figures slightly below the level of March across all vantage points.

[8] For an in-depth discussion of VPN traffic we refer to the full publications [21, 205]

[9] The necessary measurements to quantify the impact of the resolution change by the VoD providers are beyond the scope of this work.

Other applications: *Educational* networks and applications behave completely different at all vantage points. At the IXP-CE, their traffic remains relatively stable —as would be expected given students attending classes from home—, but at the ISP-CE, instead, it drastically increases by up to 200%. This growth could be attributed to some European educational networks providing video conferencing solutions, which are now being used by customers of the ISP-CE. Due to the lack of connected educational networks at the IXP-US, we omit this category at this vantage point [10]. Likewise, *CDN* traffic increases in Europe, but does not grow much—even decreasing at times—in the US. Similar to *VoD*, there is a skewed distribution of CDNs present at the vantage point. Thus, a rerouting decision of a large player may explain the moderate loss of CDN traffic at the IXP-US.

To summarize, the use of communication-related applications increase during working hours, especially in *Web conferencing*. Entertainment related applications such as *gaming* and *VoD* are also consumed at any time of the day, as the become more demanded during the lockdown. *Social media* shows a strong initial increase which flattens over time. These observations complement and strengthen those made in Section 6.3. Together, they demonstrate the massive impact that the drastic change in human behavior caused by the COVID-19 pandemic had on application usage.

6.5 The new normal? Revisiting COVID-19 in fall 2020

After the COVID-19 spring wave, many countries in central Europe could return to a relatively normal summer. People could visit the outdoor areas of restaurants, travel with only a few restrictions, schools reopened, and many people returned to their offices as local incidences were relatively low. However, at the end of the summer month, the pandemic returned full scale with an exponential rise of infections. As a result, most governments again imposed more or less strict lockdowns. This section analyzes the second wave of the COVID-19 pandemic in light of our previous observations. Especially, we focus on new patterns which already emerged in the spring wave and try to highlight those, which are here to stay. We call this trend *the new normal* as many changes, which have been triggered through the pandemic may outlive it. Now that companies had to establish efficient ways to organize distributed teams, some may stick to remote work to enlarge their potential employees' pool. Similar trends may arise for academic or industry conferences and meetings. As a society, we now can see which gatherings that traditionally happen in person also work well with remote attendance. Especially with the global climate crisis in mind, this could be one way to make society more climate-friendly and sustainable. All of these considerations only work with the support of the Internet.

To understand which new patterns from the spring wave remain, we again characterize overall traffic shifts and changes in demand for particular applications that became very popular in a short amount of time. During the process, we try to understand

[10] For an in-depth study of the traffic shifts in a large educational network we refer to the full publications [21, 205]

if there is a "new normal" in Internet traffic and to see how the Internet reacted in these unprecedented times. We then extend our previous study for the fall[11] 2020 wave (September 2020 to February 2021).

6.5.1 New persistent patterns emeg

Beyond the macroscopic observations, our analysis sheds light on the shifts in Internet usage patterns relevant to network operation and management. The Internet's regular workday traffic patterns are significantly different from weekend patterns [214]. On

[11] We use "spring" and "fall" from the Northern hemisphere's viewpoint, where our vantage points are located. Exchange both terms for the Southern hemisphere.

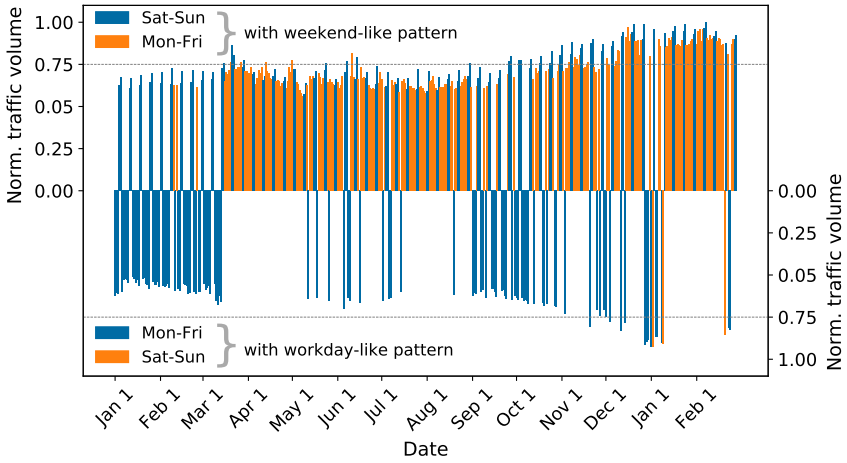

(a) ISP-CE: Weekend-like (top) vs. workday-like (bottom).

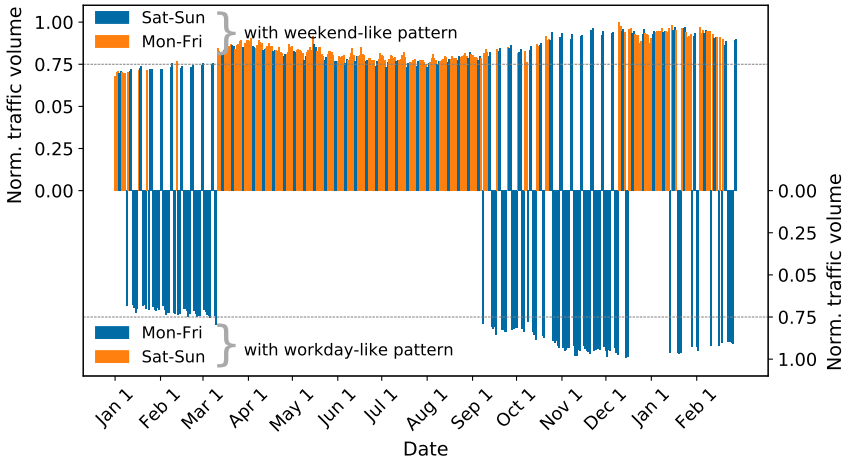

(b) IXP-CE: Weekend-like (top) vs. Workday-like (bottom).

Figure 6.8: Drastic shifts in Internet usage patterns during the COVID-19 pandemic. Classification of weekends- and workdays-like pattern.

workdays, traffic peaks are concentrated in the evenings, typically between 18:00 and midnight, also called "peak hours". However, during the weekend, the activity is more distributed in the non-peak hours as more people are at home and using the Internet.

With the pandemic lockdown in March, this workday traffic pattern shifts towards a continuous weekend-like pattern. More specifically, we call a traffic pattern a workday pattern if the traffic spikes in the evening hours and a weekend pattern if its primary activity gains significant momentum from approximately 9:00 to 10:00 am. Figures 6.8a and 6.8b show the normalized traffic for days classified as weekend-like on the top and for workday-like on the bottom. If the classification is in line with the actual day (workday or weekend), the bars are colored in blue; otherwise, they are colored orange. We find that up to mid-March, most days are classified correctly. The only exception is the holiday period at the beginning of the year in Figure 6.8b. This pattern changes drastically once the lockdown measures are implemented. Indeed, almost all days are classified as weekend-like. This change persists in Figure 6.8b until the end of August due to the vacation period, which is consistent with the behavior observed in 2019 (not shown).

In contrast, Figure 6.8a shows that the shift towards a weekend-like pattern becomes less dominant as countermeasures were relaxed in mid-May, but in August, the pattern resembles again the weekend-pattern due to the vacation period.

From August to December 2020, the patterns both at the ISP and the IXP is back to the usual weekday and weekend pattern. When the first lockdowns of the fall COVID-19 wave are imposed in December 2020, this pattern is disrupted, more noticeably at the IXP. In the first two months of 2021, there is a mixed pattern for both the ISP and the IXP. We conclude that we still observe a transient behavior in 2021, and it is unclear whether the changes of daily usage patterns are here to stay.

6.5.2 Effect on the Traffic Asymmetry

As we discussed in the previous sections, residential traffic surged both during the spring and fall COVID-19 waves. In this section we take a closer look on the directionality of the traffic and comment on new patterns in upstream and downstream traffic. Recall that residential traffic is asymmetric in nature, i.e., downstream traffic is typically many times higher than the upstream one. This is to be expected as users send less traffic than they receive when using applications like video streaming and browsing. In Figure 6.9 (top) we show the aggregated upstream traffic from October 2019 to end of February 2021. There is a slight increase in upstream traffic after the first lockdown in mid March 2020. This trend manifests itself in the following months: The minimum, but more noticeably the maximum upstream level increase across the rest of the observation period.

As a result of the general elevated traffic levels, the downstream traffic also increases during this period. To assess if there is a change in the established ratio between upstream and downstream traffic, in Figure 6.9 (bottom) we plot the ratio of upstream

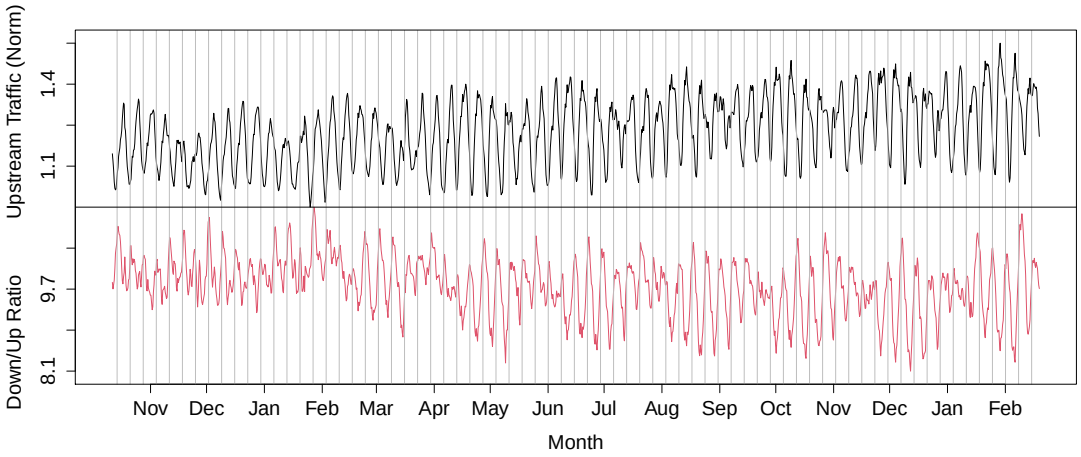

Figure 6.9: ISP aggregated over 8 hours traffic during October 2019 - February 2021: upstream traffic normalized growth (top), and downstream vs. upstream traffic ratio (bottom).

vs. downstream traffic. Before the COVID-19 pandemic, typical values of this ratio were around 9.8 with some noticeable variation. After the initial lockdown in the beginning of March 2020 and until the end of February 2021, this pattern changes. Indeed, the ratio of upstream vs. downstream traffic decreases significantly, with typical values around 9 and very high variation. During the weekdays, this ratio is as low as 8.1. This shows that the relative increase of the upstream traffic is up to 18% higher than that of the increase of the downstream traffic. An independent study that analyzes traffic data from US ISPs reported even higher upstream vs. downstream traffic ratios during the pandemic [233]. We attribute this to the increase in remote working and teleconferencing applications that utilize user upstream bandwidth much more than other popular user applications, e.g., video streaming and browsing. This is an important observation as ISPs in general allocate way less upstream than downstream capacity to end users. If we see a persistent change in demand from end users to push more traffic towards the Internet, ISPs may need to adapt their handling of subscriber lines. This is a notable result, because last mile capacity is notoriously expensive for ISPs and hard to replace with new technology.

6.5.3 Application usage revisited

We now turn our attention towards the traffic shifts for different application classes that were expected to be affected by the COVID-19 pandemic, namely, Web conferencing applications, Video-on-Demand streaming, online gaming, and traffic that originates from university service networks.

In Figure 6.10, we visualize two weeks in the Spring and Fall waves, namely, the second week in March 2020, June 2020, December 2020, and January 2021, as the difference of the respective week. We compare them to a base week before the initial lockdowns began, i.e., February 20–26, 2020. As the traffic classes we are considering

show growth way beyond the expected natural increase over one year we do not factor out that increase. Each column represents one hour of a day. This approach enables quick visual identification of increased/decreased application class usage compared to pre-COVID-19 times. We focus on the observations gathered at the ISP and the IXP in Central Europe (IXP-CE) vantage points.

Web conferencing: Web conferencing applications have seen a dramatic surge during the lockdown periods. In this category the ISP and IXP-CE experience a large traffic growth in March – right after the first lockdown began – spanning across all hours of the day, especially during weekdays. This trend accelerates in June and culminates in December and January with an increase exceeding 300% compared to the base week at both vantage points. Notably, in December and January, the extreme growth also persists at weekends. This indicates that not only work life has moved online but private social activities did as well.

Video-on-Demand: video streaming applications' usage shows high growth both in the Spring and the Fall wave. Interestingly, the ISP only sees a moderate growth during the lockdown in the first half of March followed by a reduction of volume in the second half of March below the pre-COVID-19 reference time frame. We attribute this to major streaming companies reducing their streaming resolution in Europe by mid-March for 30 days [198]. In the case of the IXP a similar but not that much pronounced trend can be observed in March. However, there is a significant increase of the traffic related to Video-on-Demand in June, December and January, that exceeds 200% (IXP) and 100% (ISP) for some days, especially on weekends indicating that more people stayed at home during leisure time instead of going outside.

Gaming: The strong growth of gaming applications is more coherent at the IXP vantage point, especially during the day. While the ISP shows a significant increase during morning hours, it generally leans towards declining in the Spring wave. Note, that this effect is mainly caused by unusually high traffic levels in this category during our baseline week in February 2020. As the initial download of a game nowadays supersedes the amount of data transferred while playing these high levels may relate to new releases or updates of popular games. Gaming applications, typically used in the evening or at weekends, are now used at any time. The trend starts to flatten in June—this may in relation with people going on vacation or spending more time outside. The ISP sees an increase up to 300% in gaming related traffic during the fall wave across all weekdays, but with emphasis to the first half of the day. A similar pattern unfolds at the IXP, but with smaller increases. One explanation for the strong increase at both vantage points in the morning hours is that schools were closed during the fall wave.

University networks: Traffic that originates from such networks behaves similar at both vantage points with the ISP showing a more pronounced trend. Both vantage points see a high increase in traffic especially during the fall wave with a growth of 100% and more. This growth could be attributed to some European educational networks providing video conferencing solutions, which are now being used by customers of the ISP/IXP. In December 2020 and January 2021 most academic collaboration

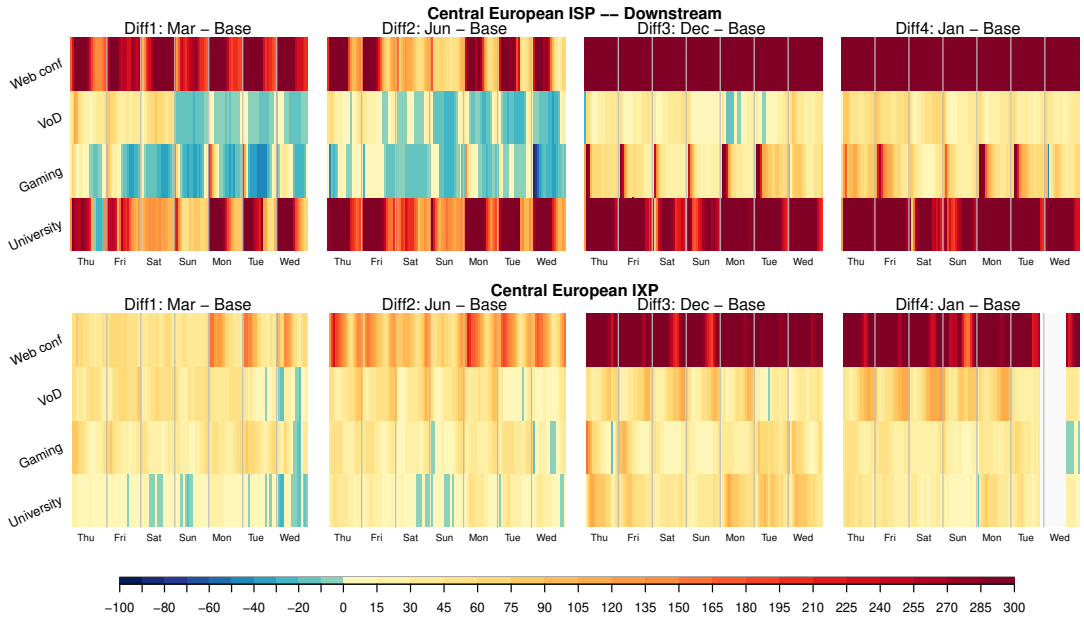

Figure 6.10: ISP (top), IXP-CE (bottom) heatmaps of application classes' traffic at the ISP and IXPs during COVID-19 pandemic: spring and fall waves. Each subplot shows the change in the aggregated traffic volume per hour for the respective class compared to the base week in February 2020. White areas mark missing data.

and teaching activities moved to an online setting. This is in line with the smaller surge of activity at weekends.

6.6 Internet operation during the pandemic: a success story.

The COVID-19 pandemic "underscored humanity's growing reliance on digital networks for business continuity, employment, education, commerce, banking, healthcare, and a whole host of other essential services" [234]. At the beginning of the pandemic and the first lockdown measures to control its spread, sudden changes in user demand for online services raised concerns for network operators, e.g., to keep networks running smoothly especially for life-critical organizations such as hospitals [235]. In fact, the pandemic increased the demand for applications supporting remote teaching and working to guarantee social distancing which manifests itself in our analysis across all vantage points. We observe this trend over our entire analysis period—almost one year. The Internet handled this increased load thanks to its original design concept to find efficient routes [204], the flexibility and elasticity that cloud services offer, and the increasing connectivity of cloud providers [219, 222, 236–238]. Our results confirm that most of the applications with the highest absolute and relative increases are cloud-based.

In our work, we report a traffic increase of more than 20% a week after the lockdown began. This is in line with reports of ISPs and CDNs [207–209, 211] as well as IXPs [239]. Typically, ISPs and CDNs are prepared for a traffic increase of 30% in a single year period [222, 240, 241]. While networks perform yearly plannings, the pandemic has created substantial shifts within only a few days. As a result, ISPs either needed to benefit from over-provisioned capacity—e.g., to handle unexpected traffic spikes such as attacks or flash-crowd events—or add capacity very quickly. The latter was possible due to the adoption of best practices on designing, operating, and provisioning networks which contributed to the smooth transition to the new normal. Due to the advances in network automation and deployment, e.g., automated configuration management and robots installing cross-connects at IXPs without human involvement, it was possible to cope with the increased demand. For example, DE-CIX Dubai managed to quickly enable new ports within a week for Microsoft, which was selected as the country's remote teaching solution for high schools [242].

Our study reveals the importance of covering different lenses to gain a complete picture of these phenomena. Additionally, our observations highlight the importance of approaching traffic engineering with a focus that looks beyond Hypergiant traffic and popular traffic classes to consider "essential" applications for remote working. In fact, our study demonstrates that over-provisioning, proactive network management, and automation are key to provide resilient networks that can sustain drastic and unexpected shifts in demand such as those experienced during the COVID-19 pandemic. Yet, as the pandemic is still ongoing, it is critical to continue studying the traffic activity to understand usage shifts during these unprecedented times.

7
Conclusion

This thesis sheds light on the usage of different data sets collected at a diverse group of vantage points for the design of Internet measurement studies. We combined four studies—ranging from a meta-analysis of commonly used data sets to direct observation of user-behavior dynamics—to highlight the importance of aligning a research question with the available data sets. In this chapter, we summarize our main findings and outline perspectives for future research.

Public BGP data: Route collector projects provide publicly available BGP data sets with good coverage of the global routing table. These data sets are a suitable base for many Internet-wide routing studies and are easy to obtain. We highlight that for many purposes, already a small subset of the entire route collector ecosystem is sufficient to gain valuable insights. We emphasize that combining ten collectors of RouteViews and RIPE/RIS only yields 5-10% more AS links than a combination of five collectors. The data set size of these collectors is in the same order of magnitude. As such, choosing a smaller subset for initial analyses can drastically reduce computation time and provide faster results. If computational resources are not a limiting factor relying on the Isolario Project data will uncover many additional AS links as Isolario supports BGP ADD-PATH. Yet, their collector dumps are an order of magnitude larger compared to RIPE/RIS or RouteViews. Sadly, on Feb 9th, 2021, the Isolario project sent out a termination notice [243] and will permanently discontinue its services on Dec 30th, 2021. Given the rich data set and the support of BGP ADD-PATH, we strongly encourage other BGP practitioners to launch and support similar services. In addition, we strongly suggest that projects with valuable and unique data sets like PCH [164] make their data more readily available to facilitate and inform further routing studies. On a more general note, we recommend a periodic in-depth review of the worldwide route collector ecosystem to systematically identify where the placement of new collectors or the acquisition of additional BGP feeders is beneficial. Moreover, these analyses can yield up to date lists of the most representative collectors to help early studies or research projects with limited resources to save time and eneg (see Section 3.4). Based on the findings in this thesis, we suggest future research efforts to update our list of recommended sanitation steps for real-world BGP routing data to account for new developments across the routing ecosystem.

While our analysis covered the routing ecosystem, we also extend this advice to traffic-level data. Every vantage point has its own features. Those need to be col-

lected and reviewed together with the local egneers. We, for example, encountered the following characteristics throughout different research projects which we handled with respect to thesgiven research question: (a) differentiate between transit versus the customer traffic at an Internet Service Provider (ISP), (b) filtering MAC addresses local to an Internet Exchange Point (IXP)'s infrastructure, or (c) identify characteristics of broken flow records specific to the collection infrastructure. In addition to the identification of vantage-point-specific properties, we recommend a review of existing literature on traffic level measurements to derive common sanitation strategies.

BGP Communities: Based on publicly available BGP routing information, we demonstrated how a widely used protocol feature could yield unexpected results. BGP communities are a popular, benign tool for traffic engineering, and many network operators rely upon them to realize policies with their peers. From our discussions within the operations community, we learned that most operators assume that BGP communities only travel a relatively small distance through the Internet. Through a carefully crafted passive measurement study, we discovered that almost 50% of the observed communities travel more than more hops and that a significant fraction travels for 50% of the AS-path. (depending on the actual AS-path length, see Section 4.4). With that, we demonstrated the existence of preconditions for BGP community-based attacks as outlined in Section 4.3. Whether these scenarios are a scalable real-world threat or not, it is vital for researchers and operators alike to base their work on the correct underlying assumptions to avoid unintended mishaps.

For future research projects, we highlight the importance of contributing to the Internet's overall health by regular validation of "well-known" assumptions, like the propagation distance of BGP communities. The foremost responsibility of network operators is ensuring the stability, performance, and security of their individual networks. On the one hand, these aspects also contribute to the stability and performance of the Internet as a whole. On the other hand, some issues pose Internet-wide security and stability risks beyond the scope of an individual network. In the latter case, well-designed measurement studies in collaboration with operators can inform the Internet community where new best operational practices can be helpful.

IP source address spoofing at a major IXP: In a first of its kind study, we quantified the prevalence of IP source address spoofing at a large European IXP. We conducted this first large-scale study at an IXP due to its location in the Internet graph and the number of connected members to gain as much understanding of the global spoofing phenomenon as possible. Nevertheless, our method is not limited to an IXP as a vantage point. In principle, every network on the inter-domain Internet can apply it to filter its incoming traffic or detect spoofing. For now, our methodology provides a very conservative overestimation of the valid IP address space per AS. We intentionally sacrificed the specificity of a closer estimation to reduce misclassifications in Invalid.

Following up on our approach, Müller et al. conducted a similar study at two Brazilian IXPs [244] and compared it to our results. Their methodology relies on Autonomous System (AS) relationships and customer cones. As such, it takes customer-provider relationships into account where one AS may be a legitimate source for the entire

IP address space. By leveraging a more restrictive classification, the authors identify an order of magnitude less traffic as spoofed than our approach. With that, their approach drastically reduces the possible amount of false positives and gives a more conservative approximation of the actual amount of spoofed traffic. Any classification problem always faces these kinds of trade-offs. Similar to the careful design of measurement studies, we face the question of the purpose of the classification. An automated spoofing detection system should focus on minimizing false positives and not accidentally discarding legitimate production traffic. If the classification is, for example, used in a scientific study to understand the composition of spoofed traffic, it is legitimate to accept a higher level of false positives.

The next step in an analysis pipeline is to search for indications of misclassification and remove legitimate traffic manually. It is very challenging to automate this process. Manual verification relies on cross-checking different data sources that are not always up to date. Sometimes it is even necessary to directly contact operators and ask them about their routing policies to understand some of the effects manifesting in the data sets. In the end, the process of manual verification is neither perfect nor complete and relies on human understanding and case-by-case plausibility check. For any future studies, we strongly recommend documenting the manual verification process as detailed as possible[1].

Irrespectively of the chosen classification problem and methodology, we encourage future work to validate the correctness of spoofing identification results. One approach may be to trigger active measurements to addresses identified as spoofed for a given IXP member and to then check whether we receive a legitimate answer via the same member. This verification process could, in principle, identify false positives. Yet, not receiving a response would not give any insights as there are many reasons for the return packet to be filtered.

Impact of the COVID-19 pandemic on Internet usage: Using the combined perspectives of a large European ISP, a large European IXP and two medium-sized IXPs in southern Europe and the US, we showed how the COVID-19 pandemic changed global Internet usage behavior over one year.

Despite the disruption due to COVID-19, life continued thanks to the increased digitization and resilience of our society, with the Internet playing a critical support role for businesses, education, entertainment, purchases, and social interactions. We analyzed Internet flow data from multiple vantage points in several developed countries. Together, they allow us to gain a good understanding of the impact that the COVID-19 waves and the lockdown measures caused on Internet traffic. One year after the first lockdown measures, the aggregated traffic volume increased by around 40%, well above the typical expected annual growth.

We found that ISP subscribers as well as IXP members adapted their Internet access patterns to the "new normal" of the pandemic where most people were working and socializing from home. Workday traffic patterns have rapidly changed and the relative

[1]Depending on the non-disclosure agreements with the respective vantage points, it may not be possible to name specific AS relationships.

difference to weekend patterns has almost disappeared during lockdowns. Sociologists already identified the Internet the substrate for a new media epoch and establish a new sociology for the modern world [245]. In line with these new developments, our work highlights how societal changes reflect themselves in Internet usage patterns. Beyond societal aspects, we also confirmed that the Internet core could indeed withstand the increased pandemic-induced demands. The new needs we observed are in line with the change in the way of life for most people: (a) the increased popularity of video conferencing services, (b) elevated usage of Virtual Private Network (VPN) and collaborative working solutions (c) increased activity by academic networks to support online teaching, and (d) higher consumption of digital entertainment content, e.g., video on demand services, some of which are transported via QUIC or online gaming. To satisfy these demands, the vantage points we considered have successfully established procedures to quickly deploy additional capacity. With the evidence collected across a diverse set of vantage points provided through our work, we conclude that the Internet—from the perspective of our vantage points—did its job and coped well with unseen and rapid traffic shifts. Related work, however, reported performance degradation in less developed regions [246]. As such, we emphasize that to gain a complete perspective on the impact of the COVID-19 pandemic on the Internet, it is vital to leverage representative vantage points from a diverse set of regions, economic backgrounds, and maturity levels. The unseen traffic shifts we observe due to the implementation of confinement measures underline the importance of the Internet's distributed nature to react amicably to such events and enhance society's resilience.

Summary: The four studies discussed in this thesis show how different data sets can be employed to answer different classes of research questions. We started with a meta-analysis of commonly used data sets, namely BGP route collector data, to highlight artifacts in these data sets and share our experience working with that data source. Then, we used this data set to validate a widespread assumption regarding the propagation distance of BGP communities. Furthermore, we combined the same routing information sources with passively obtained traffic-level measurements to understand IP source address spoofing—a mechanism used to leverage high-volume Distributed Denial of Service (DDoS) attacks. Lastly, we demonstrate how we used traffic-level data alone to understand the impact of the COVID-19 pandemic on the Internet. To that end, we combined the perspectives of one large eyeball-centric vantage point, namely an ISP with vantage points located in the core of the Internet, namely three IXPs.

Bibliography

[1] Vern Paxson and Sally Floyd. "Why we don't know how to simulate the Internet". In: *Proceedings of the 29th conference on Winter simulation*. 1997, pp. 1037–1044 (cit. on p. 2).

[2] Vern Paxson. "Strategies for sound internet measurement". In: *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*. 2004, pp. 263–271 (cit. on p. 2).

[3] *Outage at Amsterdam internet hub affects much of Netherlands*. 2015. URL: https://nltimes.nl/2015/05/13/outage-amsterdam-internet-hub-affects-much-netherlands (visited on 06/09/2021) (cit. on p. 2).

[4] *AMS-IX global outage (Status)*. 2015. URL: https://noc.nforce.com/notifications/item/250 (visited on 06/17/2021) (cit. on p. 2).

[5] RIPE NCC. *RIPE Atlas*. https://atlas.ripe.net (cit. on p. 2).

[6] Emile Aben. *Does the Internet Route Around Damage? A Case Study Using RIPE Atlas*. 2015. URL: https://labs.ripe.net/author/emileaben/does-the-internet-route-around-damage-a-case-study-using-ripe-atlas/ (visited on 06/17/2021) (cit. on p. 2).

[7] Emile Aben. *Does The Internet Route Around Damage? - Edition 2021*. 2021. URL: https://labs.ripe.net/author/emileaben/does-the-internet-route-around-damage-edition-2021/ (visited on 06/17/2021) (cit. on p. 2).

[8] Vasileios Giotsas et al. "Detecting peering infrastructure outages in the wild". In: *Proceedings of the conference of the ACM special interest group on data communication*. 2017, pp. 446–459 (cit. on p. 2).

[9] *Summary of June 8 outage*. 2021. URL: https://www.fastly.com/blog/summary-of-june-8-outage (visited on 06/09/2021) (cit. on p. 2).

[10] Aftab Siddiqui. *A major BGP route leak by AS55410*. 2021. URL: https://blog.apnic.net/2021/04/26/a-major-bgp-route-leak-by-as55410/ (visited on 06/22/2021) (cit. on p. 2).

[11] Shinyoung Cho et al. "BGP hijacking classification". In: *2019 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE. 2019, pp. 25–32 (cit. on p. 3).

[12] Mark Allman, Robert Beverly, and Brian Trammell. "Principles for measurability in protocol design". In: *ACM SIGCOMM Computer Communication Review* 47.2 (2017), pp. 2–12 (cit. on p. 3).

[13] *Internet Protocol*. RFC 791. Sept. 1981. DOI: 10.17487/RFC0791. URL: https://rfc-editor.org/rfc/rfc791.txt (cit. on p. 3).

[14] Jana Iyengar and Martin Thomson. *QUIC: A UDP-Based Multiplexed and Secure Transport*. RFC 9000. May 2021. DOI: 10.17487/RFC9000. URL: https://rfc-editor.org/rfc/rfc9000.txt (cit. on p. 3).

[15] Matthew Luckie et al. “Network hygiene, incentives, and regulation: deployment of source address validation in the Internet”. In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. 2019, pp. 465–480 (cit. on p. 4).

[16] Siyuan Jia et al. “Tracking the deployment of IPv6: Topology, routing and performance”. In: *Computer Networks* 165 (2019), p. 106947 (cit. on p. 4).

[17] Renata Teixeira, Steve Uhlig, and Christophe Diot. “BGP Route Propagation Between Neighboring Domains”. In: vol. 4427. Apr. 2007, pp. 11–21. ISBN: 978-3-540-71616-7. DOI: 10.1007/978-3-540-71617-4_2 (cit. on p. 4).

[18] Florian Streibelt et al. “BGP Communities: Even more Worms in the Routing Can”. In: *ACM IMC*. ACM. 2018, pp. 279–292 (cit. on pp. 6, 22, 28, 37).

[19] Franziska Lichtblau, Florian Streibelt, Thorben Krüger, Philipp Richter, and Anja Feldmann. “Detection, classification, and analysis of inter-domain traffic with spoofed source IP addresses”. In: *Proceedings of the 2017 Internet Measurement Conference*. 2017, pp. 86–99 (cit. on pp. 7, 55).

[20] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, et al. “The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic”. In: *Proceedings of the ACM Internet Measurement Conference*. 2020, pp. 1–18 (cit. on pp. 7, 11).

[21] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. “A Year in Lockdown: How the Waves of COVID-19 Impact Internet Traffic”. In: *Commun. ACM* 64.7 (June 2021), pp. 101–108. ISSN: 0001-0782. DOI: 10.1145/3465212. URL: https://doi.org/10.1145/3465212 (cit. on pp. 7, 80, 91, 95, 96).

[22] J. Hawkinson and T. Bates. *Guidelines for creation, selection, and registration of an Autonomous System (AS)*. RFC 1930. IETF, Mar. 1996. URL: http://tools.ietf.org/rfc/rfc1930.txt (cit. on p. 9).

[23] Y. Rekhter, T. Li, and S. Hares. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. IETF, Jan. 2006. URL: http://tools.ietf.org/rfc/rfc4271.txt (cit. on p. 9).

[24] D. Walton et al. *Advertisement of Multiple Paths in BGP*. RFC 7911. IETF, July 2016. URL: http://tools.ietf.org/rfc/rfc7911.txt (cit. on p. 10).

[25] Martino Trevisan et al. “Five years at the edge: Watching internet from the isp network”. In: *IEEE/ACM Transactions on Networking* 28.2 (2020), pp. 561–574 (cit. on p. 11).

[26] *List of Tier-1 networks*. URL: https://en.wikipedia.org/wiki/Tier_1_network#List_of_Tier_1_networks (visited on 05/23/2021) (cit. on p. 12).

[27] Philipp Richter et al. “Peering at peerings: On the role of IXP route servers”. In: *Proceedings of the 2014 Conference on Internet Measurement Conference*. 2014, pp. 31–44 (cit. on p. 14).

[28] PeeringDB. *PeeringDB*. https://www.peeringdb.com. 2020 (cit. on pp. 15, 70, 92).

[29] William B Norton. *The Internet peering playbook: connecting to the core of the Internet*. DrPeering Press, 2014 (cit. on p. 15).

[30] Hyunseok Chang, Sugih Jamin, and Walter Willinger. “To peer or not to peer: Modeling the evolution of the Internet’s AS-level topology”. In: *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*. IEEE. 2006, pp. 1–12 (cit. on p. 15).

[31] Reza Motamedi et al. “On mapping the interconnections in Today’s Internet”. In: *IEEE/ACM Transactions on Networking* 27.5 (2019), pp. 2056–2070 (cit. on p. 15).

[32] DE-CIX. *DE-CIX Frankfurt statistics*. https://www.de-cix.net/en/locations/germany/frankfurt/statistics. Last accessed: 27th June, 2021. Archived version available at:https://web.archive.org/web/20210620110006/https://www.de-cix.net/en/locations/germany/frankfurt/statistics. 2021 (cit. on p. 15).

[33] AMS-IX. *Total Traffic Statistics*. https://stats.ams-ix.net/index.html. Last accessed: 27th June, 2021. Archived version available at:https://web.archive.org/web/20210627072325/https://stats.ams-ix.net/index.html. 2021 (cit. on p. 15).

[34] *Euro-IX Report 2020*. URL: https://www.euro-ix.net/media/filer_public/93/77/937743c9-1996-4586-8daf-f5fcd2640be3/ixp_report_2020_.pdf (visited on 05/24/2021) (cit. on p. 15).

[35] Nikolaos Chatzis et al. “There is more to IXPs than meets the eye”. In: *ACM SIGCOMM Computer Communication Review* 43.5 (2013), pp. 19–28 (cit. on p. 15).

[36] RIPE NCC. *The RIPE NCC*. https://www.ripe.net/. (Visited on 07/05/2021) (cit. on p. 16).

[37] APNIC. *APNIC*. https://www.apnic.net/. (Visited on 07/05/2021) (cit. on p. 16).

[38] ARIN. *ARIN*. https://www.arin.net/. (Visited on 07/05/2021) (cit. on p. 16).

[39] AFRINIC. *AFRINIC*. https://www.afrinic.net/. (Visited on 07/05/2021) (cit. on p. 16).

[40] LACNIC. *LACNIC*. https://www.lacnic.net/. (Visited on 07/05/2021) (cit. on p. 16).

[41] *European Peering Forum*. URL: https://www.peering-forum.eu/ (cit. on p. 16).

[42] *Global Peering Forum*. URL: https://www.globalpeeringforum.org/ (cit. on p. 16).

[43] Will E Leland et al. “On the self-similar nature of Ethernet traffic (extended version)”. In: *IEEE/ACM Transactions on networking* 2.1 (1994), pp. 1–15 (cit. on p. 17).

[44] Liang Guo and Ibrahim Matta. “The war between mice and elephants”. In: *Proceedings Ninth International Conference on Network Protocols. ICNP 2001.* IEEE. 2001, pp. 180–188 (cit. on p. 17).

[45] Vibhaalakshmi Sivaraman et al. “Heavy-hitter detection entirely in the data plane”. In: *Proceedings of the Symposium on SDN Research*. 2017, pp. 164–176 (cit. on p. 17).

[46] B. Claise, B. Trammell, and P. Aitken. *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information*. RFC 7011. IETF, Sept. 2013. URL: `http://tools.ietf.org/rfc/rfc7011.txt` (cit. on pp. 18, 83).

[47] J. Postel. *User Datagram Protocol*. RFC 768. IETF, Aug. 1980. URL: `http://tools.ietf.org/rfc/rfc0768.txt` (cit. on p. 19).

[48] *RIPE Routing Information Service*. `http://www.ripe.net/ris/` (cit. on pp. 21, 22, 46, 55).

[49] *Routeviews Project – University of Oregon*. `http://www.routeviews.org/` (cit. on pp. 21, 22, 46, 55, 62).

[50] *Isolario Project*. `https://isolario.it/` (cit. on pp. 21, 22, 46).

[51] Packet Clearing House. *Resources: Raw Routing Data*. `https://www.pch.net/resources/Raw_Routing_Data/` (cit. on pp. 21, 22).

[52] BGPmon Project. *BGPmon*. `https://bgpmon.net/` (cit. on pp. 21, 22).

[53] Yihua He et al. “Lord of the links: a framework for discovering missing links in the internet Topology”. In: *IEEE/ACM Transactions On Networking* 17.2 (2009), pp. 391–404 (cit. on pp. 21, 24).

[54] Ricardo Oliveira et al. “The (in) completeness of the observed internet AS-level structure”. In: *IEEE/ACM Transactions on Networking* 18.1 (2010), pp. 109–122 (cit. on pp. 21, 25).

[55] Matthew Roughan, Simon Jonathan Tuke, and Olaf Maennel. “Bigfoot, sasquatch, the yeti and other missing links: what we don’t know about the as graph”. In: *ACM SIGCOMM*. ACM. 2008, pp. 325–330 (cit. on pp. 21, 24).

[56] Bernhard Ager et al. “Anatomy of a large European IXP”. In: *ACM SIGCOMM*. ACM, 2012, pp. 163–174 (cit. on p. 21).

[57] Vasileios Giotsas, Shi Zhou, Matthew Luckie, et al. “Inferring multilateral peering”. In: *CoNEXT*. ACM. 2013, pp. 247–258 (cit. on p. 21).

[58] Kai Chen et al. “Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes from P2P Users”. In: *IEEE Trans. on Computers* 63.4 (2014), pp. 1021–1036 (cit. on p. 21).

[59] H. Chang et al. “Towards capturing representative AS-level Internet topologies”. In: *Computer Networks* 44.6 (2004), pp. 737–755 (cit. on p. 21).

[60] Matthew Luckie et al. "AS Relationships, Customer Cones, and Validation". In: *ACM IMC*. Barcelona, Spain: ACM, 2013, pp. 243–256. ISBN: 978-1-4503-1953-9. DOI: 10.1145/2504730.2504735. URL: http://doi.acm.org/10.1145/2504730.2504735 (cit. on pp. 22, 24, 25, 27–29, 61).

[61] Yuchen Jin et al. "Stable and Practical AS Relationship Inference with ProbLink". In: *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. Boston, MA: USENIX Association, 2019, pp. 581–598. ISBN: 978-1-931971-49-2. URL: https://www.usenix.org/conference/nsdi19/presentation/jin (cit. on pp. 22, 24).

[62] Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, et al. "Inferring complex as relationships". In: *ACM IMC*. ACM. 2014, pp. 23–30 (cit. on p. 22).

[63] V. Giotsas and S. Zhou. "Improving the discovery of IXP peering links through passive BGP measurements". In: *IEEE INFOCOM*. IEEE. 2013 (cit. on pp. 22, 66).

[64] Vasileios Giotsas et al. "Detecting Peering Infrastructure Outages in the Wild". In: *ACM SIGCOMM*. SIGCOMM '17. Los Angeles, CA, USA: ACM, 2017, pp. 446–459. ISBN: 978-1-4503-4653-5. DOI: 10.1145/3098822.3098855. URL: http://doi.acm.org/10.1145/3098822.3098855 (cit. on p. 22).

[65] Philipp Richter et al. "Peering at Peerings: On the Role of IXP Route Servers". In: *ACM IMC*. 2014, pp. 31–44. DOI: 10.1145/2663716.2663757. URL: https://doi.org/10.1145/2663716.2663757 (cit. on p. 22).

[66] WOWRACK. *Data Centers built to withstand the elements*. https://www.wowrack.com/about-wowrack/data-center. 2019 (cit. on p. 23).

[67] D. Walton et al. *Advertisement of Multiple Paths in BGP*. rfc7911. July 2016. URL: http://tools.ietf.org/rfc/rfc7911.txt (cit. on pp. 24, 32).

[68] Team Cymru. *BOGON ROUTE SERVER PROJECT*. http://www.team-cymru.com/bogon-reference.html (cit. on pp. 24, 26).

[69] Matthew Roughan et al. "10 lessons from 10 years of measuring and modeling the internet's autonomous systems". In: *IEEE Journal on Selected Areas in Communications (JSAC)* 29.9 (2011), pp. 1810–1821 (cit. on p. 24).

[70] Ramesh Govindan and Anoop Reddy. "An analysis of Internet inter-domain topology and route stability". In: *IEEE INFOCOM*. Vol. 2. IEEE. 1997, pp. 850–857 (cit. on p. 24).

[71] Walter Willinger, David Alderson, and John C Doyle. "Mathematics and the internet: A source of enormous confusion and great potential". In: *Notices of the American Mathematical Society* 56.5 (2009), pp. 586–599 (cit. on p. 24).

[72] Yu Zhang et al. "A framework to quantify the pitfalls of using traceroute in AS-level topology measurement". In: *IEEE Journal on Selected Areas in Communications* 29.9 (2011), pp. 1822–1836 (cit. on p. 24).

[73] Wikipedia. *Tier 1 network*. https://en.wikipedia.org/wiki/Tier_1_network (cit. on p. 24).

[74] Vinay Aggarwal, Anja Feldmann, and Christian Scheideler. "Can ISPs and P2P users cooperate for improved performance?" In: *ACM SIGCOMM Computer Communication Review (CCR)* 37.3 (2007), pp. 29–40 (cit. on p. 24).

[75] Yakov Rekhter and Tony Li. *A Border Gateway Protocol 4 (BGP-4)*. RFC 1654. RFC Editor, July 1995, pp. 1–56. URL: https://www.rfc-editor.org/rfc/rfc1654.txt (cit. on p. 25).

[76] Urs Hengartner et al. "Detection and analysis of routing loops in packet traces". In: *ACM IMC*. ACM. 2002, pp. 107–112 (cit. on p. 25).

[77] Dan Pei et al. "A study of BGP path vector route looping behavior". In: *Int. Conf. on Distriubted Computing Systems*. IEEE. 2004, pp. 720–729 (cit. on p. 25).

[78] Ratul Mahajan, David Wetherall, and Tom Anderson. "Understanding BGP misconfiguration". In: *ACM SIGCOMM Computer Communication Review (CCR)*. Vol. 32. ACM. 2002, pp. 3–16 (cit. on p. 25).

[79] Enrico Gregori et al. "BGP and inter-AS economic relationships". In: *International Conference on Research in Networking*. Springer. 2011, pp. 54–67 (cit. on p. 25).

[80] Thomas Krenc and Anja Feldmann. "BGP Prefix Delegations: A Deep Dive". In: *ACM IMC*. Santa Monica, California, USA: ACM, 2016, pp. 469–475. ISBN: 978-1-4503-4526-2. DOI: 10.1145/2987443.2987458. URL: http://doi.acm.org/10.1145/2987443.2987458 (cit. on pp. 25–27).

[81] Hitesh Ballani, Paul Francis, and Xinyang Zhang. "A study of prefix hijacking and interception in the Internet". In: *ACM SIGCOMM Computer Communication Review (CCR)* 37.4 (2007), pp. 265–276 (cit. on p. 25).

[82] Changxi Zheng et al. "A light-weight distributed scheme for detecting IP prefix hijacks in real-time". In: *ACM SIGCOMM Computer Communication Review (CCR)* 37.4 (2007), pp. 277–288 (cit. on p. 25).

[83] Zheng Zhang et al. "iSPY: Detecting IP prefix hijacking on my own". In: *IEEE/ACM Transactions on Networking* 18.6 (2010), pp. 1815–1828 (cit. on p. 25).

[84] Johann Schlamp et al. "Investigating the Nature of Routing Anomalies: Closing in on Subprefix Hijacking Attacks". In: *Network Traffic Measurement and Analysis Conference (TMA)*. 2015, pp. 173–187. DOI: 10.1007/978-3-319-17172-2_12. URL: https://doi.org/10.1007/978-3-319-17172-2%5C_12 (cit. on p. 25).

[85] Jian Qiu et al. "Detecting bogus BGP route information: Going beyond prefix hijacking". In: *International Conference on Security and Privacy in Communications Networks*. IEEE. 2007, pp. 381–390 (cit. on p. 25).

[86] K. Sriram et al. *Problem Definition and Classification of BGP Route Leaks*. RFC 7908. IETF, June 2016. URL: http://tools.ietf.org/rfc/rfc7908.txt (cit. on p. 25).

[87] Riad Mazloum et al. "Violation of Interdomain Routing Assumptions". In: *International Conference on Passive and Active Network Measurement (PAM)*. Los Angeles (CA), Mar. 2014 (cit. on pp. 25, 29).

[88] Giovanni Comarela, Evimaria Terzi, and Mark Crovella. "Detecting Unusually-Routed ASes: Methods and Applications". In: *ACM IMC*. 2016, pp. 445–459. URL: http://dl.acm.org/citation.cfm?id=2987478 (cit. on p. 25).

[89] Enrico Gregori et al. "Improving the reliability of inter-AS economic inferences through a hygiene phase on BGP data". In: *Computer Networks* 62 (2014), pp. 197–207 (cit. on p. 25).

[90] J. Durand, I. Pepelnjak, and G. Doering. *BGP Operations and Security*. RFC 7454. IETF, Feb. 2015. URL: http://tools.ietf.org/rfc/rfc7454.txt (cit. on p. 26).

[91] Y. Rekhter et al. *Address Allocation for Private Internets*. RFC 1918. IETF, Feb. 1996. URL: http://tools.ietf.org/rfc/rfc1918.txt (cit. on pp. 26, 56, 60).

[92] M. Cotton et al. *Special-Purpose IP Address Registries*. RFC 6890. IETF, Apr. 2013. URL: http://tools.ietf.org/rfc/rfc6890.txt (cit. on p. 26).

[93] M. Cotton, L. Vegoda, and D. Meyer. *IANA Guidelines for IPv4 Multicast Address Assignments*. RFC 5771. IETF, Mar. 2010. URL: http://tools.ietf.org/rfc/rfc5771.txt (cit. on p. 26).

[94] Vasileios Giotsas et al. "Inferring BGP Blackholing Activity in the Internet". In: *ACM IMC*. London, UK, Nov. 2017 (cit. on p. 26).

[95] George Nomikos and Xenofontas Dimitropoulos. "traIXroute: Detecting IXPs in traceroute paths". In: *International Conference on Passive and Active Network Measurement (PAM)*. Springer. 2016, pp. 346–358 (cit. on p. 26).

[96] Franziska Lichtblau et al. "Detection, classification, and analysis of inter-domain traffic with spoofed source IP addresses". In: *ACM IMC*. 2017, pp. 86–99. DOI: 10.1145/3131365.3131367. URL: https://doi.org/10.1145/3131365.3131367 (cit. on pp. 26, 27, 30).

[97] Matthew J. Luckie et al. "bdrmap: Inference of Borders Between IP Networks". In: *ACM IMC*. 2016, pp. 381–396. URL: http://dl.acm.org/citation.cfm?id=2987467 (cit. on pp. 26, 27).

[98] Luca Cittadini et al. "Evolution of internet address space deaggregation: myths and reality". In: *IEEE Journal on Selected Areas in Communications (JSAC)* 28.8 (2010), pp. 1238–1249 (cit. on p. 27).

[99] Stefano Vissicchio, Luca Cittadini, and Giuseppe Di Battista. "On iBGP routing policies". In: *IEEE/ACM Transactions on Networking* 23.1 (2015), pp. 227–240 (cit. on p. 27).

[100] Z Morley Mao et al. "BGP beacons". In: *ACM IMC*. ACM. 2003, pp. 1–14 (cit. on p. 27).

[101] Brandon Schlinker et al. "Peering: An as for us". In: *ACM Workshop on Hot Topics in Networks (HotNets)*. ACM. 2014, p. 18 (cit. on p. 27).

[102] Ethan Katz-Bassett et al. “Machiavellian routing: improving internet availability with bgp poisoning”. In: *ACM Workshop on Hot Topics in Networks (HotNets)*. ACM. 2011, p. 11 (cit. on p. 27).

[103] *PEERING: The BGP Testbed*. https://peering.usc.edu (cit. on pp. 27, 28).

[104] Randy Bush et al. “Internet optometry: assessing the broken glasses in internet reachability”. In: *ACM IMC*. ACM. 2009, pp. 242–253 (cit. on p. 27).

[105] *The CAIDA AS AS-Rank Dataset*. http://as-rank.caida.org/ (cit. on p. 27).

[106] Jianhong Xia and Lixin Gao. “On the evaluation of AS relationship inferences [Internet reachability/traffic flow applications]”. In: *IEEE Global Telecommunications Conference (GLOBECOM)*. Vol. 3. IEEE. 2004, pp. 1373–1377 (cit. on p. 28).

[107] Julian Martin Del Fiore et al. “Filtering the Noise to Reveal Inter-Domain Lies”. In: *Network Traffic Measurement and Analysis Conference (TMA)*. June 2019 (cit. on p. 28).

[108] Yves Vanaubel et al. “MPLS under the microscope: Revealing actual transit path diversity”. In: *ACM IMC*. ACM. 2015, pp. 49–62 (cit. on p. 28).

[109] G. Moura et al. “Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event”. In: *ACM IMC*. 2016 (cit. on p. 28).

[110] Mattijs Jonker et al. “A First Joint Look at DoS Attacks and BGP Blackholing in the Wild”. In: *ACM IMC*. 2018, pp. 457–463. URL: https://dl.acm.org/citation.cfm?id=3278571 (cit. on p. 28).

[111] Alexander Marder et al. “Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale”. In: *ACM IMC*. 2018, pp. 56–69. URL: https://dl.acm.org/citation.cfm?id=3278538 (cit. on p. 28).

[112] Chiara Orsini et al. “BGPStream: a software framework for live and historical BGP data analysis”. In: *ACM IMC*. ACM. 2016, pp. 429–444 (cit. on p. 29).

[113] Philipp Richter et al. “Advancing the Art of Internet Edge Outage Detection”. In: *ACM IMC*. 2018, pp. 350–363. URL: https://dl.acm.org/citation.cfm?id=3278563 (cit. on p. 29).

[114] Rahul Hiran, Niklas Carlsson, and Phillipa Gill. “Characterizing large-scale routing anomalies: A case study of the china telecom incident”. In: *International Conference on Passive and Active Network Measurement (PAM)*. Springer. 2013, pp. 229–238 (cit. on p. 29).

[115] Robert Beverly et al. “Measuring and characterizing IPv6 router availability”. In: *International Conference on Passive and Active Network Measurement (PAM)*. Springer. 2015, pp. 123–135 (cit. on p. 29).

[116] Arpit Gupta et al. “An industrial-scale software defined internet exchange point”. In: *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. 2016, pp. 1–14 (cit. on p. 29).

[117] CAIDA. *RouteViews Prefix to AS mappings*. http://www.caida.org/data/routing/routeviews-prefix2as.xml (cit. on p. 29).

[118] *The CAIDA AS Relationships Dataset*. http://www.caida.org/data/active/as-relationships/ (cit. on p. 29).

[119] Massimo Rimondini, Claudio Squarcella, and Giuseppe Di Battista. "Towards an automated investigation of the impact of bgp routing changes on network delay variations". In: *International Conference on Passive and Active Network Measurement (PAM)*. Springer. 2014, pp. 193–203 (cit. on p. 29).

[120] Romain Fontugne, Anant Shah, and Emile Aben. "The (thin) bridges of AS connectivity: measuring dependency using AS hegemony". In: *International Conference on Passive and Active Network Measurement (PAM)*. Springer. 2018, pp. 216–227 (cit. on p. 29).

[121] *The NLNOG RING project*. https://ring.nlnog.net/ (cit. on p. 32).

[122] G. Huston. "Peering and Settlements: Part I". In: *The Internet Protocol Journal* 2.1 () (cit. on p. 38).

[123] G. Huston. "Peering and Settlements: Part II". In: *The Internet Protocol Journal* 2.2 () (cit. on p. 38).

[124] P. Faratin et al. "The Growing Complexity of Internet Interconnection". In: *Communications and Strategies* (2008) (cit. on p. 38).

[125] R. Chandra, P. Traina, and T. Li. *BGP Communities Attribute*. RFC 1997. IETF, Aug. 1996. URL: http://tools.ietf.org/rfc/rfc1997.txt (cit. on pp. 38, 39, 45).

[126] B. Donnet and O. Bonaventure. "On BGP Communities". In: *ACM CCR* 38.2 (Mar. 2008), pp. 55–59 (cit. on pp. 38, 39, 47).

[127] Vasileios Giotsas et al. "Detecting Peering Infrastructure Outages in the Wild". In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. SIGCOMM '17. Los Angeles, CA, USA: Association for Computing Machinery, 2017, pp. 446–459. ISBN: 9781450346535. DOI: 10.1145/3098822.3098855. URL: https://doi.org/10.1145/3098822.3098855 (cit. on pp. 38, 47).

[128] V. Giotsas et al. "Inferring BGP Blackholing Activity in the Internet". In: *ACM IMC*. 2017 (cit. on pp. 38, 47, 49, 50).

[129] C. Dietzel, A. Feldmann, and T. King. "Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild". In: *PAM*. 2016 (cit. on pp. 38, 47).

[130] V. Giotsas et al. "Mapping Peering Interconnections at the Facility Level". In: *CoNEXT*. 2015 (cit. on p. 38).

[131] *DE-CIX: Informational BGP Communities: Origin tagging*. https://www.de-cix.net/de/resources/informational-bgp-communities (cit. on pp. 38, 40).

[132] P. Richter et al. "Peering at Peerings: On the Role of IXP Route Servers". In: *ACM IMC*. 2014 (cit. on pp. 38, 46, 62, 64).

[133] V. Giotsas et al. "Inferring Multilateral Peering". In: *CoNEXT*. 2013 (cit. on p. 38).

[134] N. Chatzis et al. “There is More to IXPs than Meets the Eye”. In: *ACM CCR* 43.5 (2013) (cit. on p. 38).

[135] V. Giotsas et al. “Inferring Complex AS Relationships”. In: *ACM IMC*. 2014 (cit. on p. 38).

[136] P. Smith. *BGP Techniques for Internet Service Providers*. NANOG 50. 2010 (cit. on p. 38).

[137] M. J. Levy. *Using BGP communities to control your transit providers*. APRICOT 2013 (cit. on p. 38).

[138] CISCO. *Remotely Triggered Black Hole Filtering - Destination Based and Source Based*. Cisco White Paper, `http://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf`. 2005 (cit. on pp. 38, 44).

[139] T. King et al. *BLACKHOLE Community*. RFC 7999. IETF, Oct. 2016. URL: `http://tools.ietf.org/rfc/rfc7999.txt` (cit. on pp. 38, 39, 43, 49).

[140] *ECIX’s New Route Server RTT Communities*. https://www.ecix.net/about-us/news/ecixs-new-route-server-rtt-communities (cit. on p. 38).

[141] *Euro-IX: Large BGP Communities*. https://www.euro-ix.net/en/forixps/large-bgp-communities/ (cit. on p. 38).

[142] Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush. “Bgp communities: Even more worms in the routing can”. In: *Proceedings of the Internet Measurement Conference 2018*. 2018, pp. 279–292 (cit. on pp. 38, 45).

[143] J. Heitz et al. *BGP Large Communities Attribute*. RFC 8092. IETF, Feb. 2017. URL: `http://tools.ietf.org/rfc/rfc8092.txt` (cit. on p. 39).

[144] G. Huston. *NOPEER Community for Border Gateway Protocol (BGP) Route Scope Control*. RFC 3765. IETF, Apr. 2004. URL: `http://tools.ietf.org/rfc/rfc3765.txt` (cit. on p. 39).

[145] R. Bush and R. Austein. *The Resource Public Key Infrastructure (RPKI) to Router Protocol*. RFC 6810. IETF, Jan. 2013. URL: `http://tools.ietf.org/rfc/rfc6810.txt` (cit. on p. 39).

[146] E. Heilman et al. “From the Consent of the Routed: Improving the Transparency of the RPKI”. In: *ACM SIGCOMM*. 2014 (cit. on p. 39).

[147] G. Goodell et al. “Working Around BGP: An Incremental Approach to Improving Security and Accuracy of Interdomain Routing”. In: *NDSS*. 2003 (cit. on p. 39).

[148] Y. C. Hu, A. Perrig, and M. Sirbu. “SPV: Secure Path Vector Routing for Securing BGP”. In: *SIGCOMM*. 2004 (cit. on p. 39).

[149] O. Nordstrom and C. Dovrolis. “Beware of BGP attacks”. In: *ACM CCR* 34.2 (2004) (cit. on p. 39).

[150] O. Bonaventure and B. Quoitin. *Common utilizations of the BGP community attribute*. IETF draft, work in progress, draft-bq-bgp-communities-00.txt, June 2003. (Cit. on p. 39).

[151] *NTT routing policies*. https://www.us.ntt.net/support/policy/routing.cfm (cit. on p. 40).

[152] *KPN BGP communities*. https://as286.net/AS286-communities.html (cit. on p. 40).

[153] *DE-CIX Frankfurt Route Server Guide*. https://www.de-cix.net/en/locations/germany/frankfurt/rc guide (cit. on p. 40).

[154] *AMS-IX Deployment guides*. https://ams-ix.net/technical/specifications-descriptions/ams-ix-route-servers/deployment-guides (cit. on p. 40).

[155] *Seattle Internet Exchange route servers*. https://www.seattleix.net/route-servers (cit. on p. 40).

[156] *IRR - Internet Routing Registry*. `http://www.irr.net` (cit. on p. 40).

[157] Y. Sun et al. "RAPTOR: Routing Attacks on Privacy in Tor". In: *NSDI*. 2015 (cit. on p. 41).

[158] M. Apostolaki, A. Zohar, and L. Vanbever. "Hijacking Bitcoin: Routing Attacks on Cryptocurrencies". In: *IEEE Symposium on Security and Privacy*. 2017 (cit. on p. 41).

[159] S. Goldberg. "Why is It Taking So Long to Secure Internet Routing?" In: *Comm. of the ACM* 57.10 (2014) (cit. on p. 42).

[160] M. Antonakakis et al. "Understanding the Mirai Botnet". In: *USENIX Security Symposium*. 2017 (cit. on p. 43).

[161] D. Gillman et al. "Protecting Websites from Attack with Secure Delivery Networks". In: *IEEE Computer Magazine* 48.4 (2015) (cit. on p. 43).

[162] A. Robachevsky. *14,000 Incidents: A 2017 Routing Security Year in Review*. Internet Society, `https://www.internetsociety.org/blog/2018/01/14000-incidents-2017-routing-security-year-review/`. Jan. 2018 (cit. on p. 44).

[163] J. Soricelli and W. Custavus. *NANOG Tutorial: Option for Blackhole and Discard Routing*. `https://www.nanog.org/meetings/nanog32/presentations/soricelli.pdf`. Oct. 2004 (cit. on p. 44).

[164] *Packet Clearing House routing archive*. `https://www.pch.net/resources/data.php` (cit. on pp. 46, 103).

[165] *Packet Clearing House Peering*. `https://www.pch.net/about/peering` (cit. on p. 46).

[166] *RIPE Atlas*. `https://atlas.ripe.net/` (cit. on p. 46).

[167] J. Mitchell. *Autonomous System (AS) Reservation for Private Use*. RFC 6996. IETF, July 2013. URL: `http://tools.ietf.org/rfc/rfc6996.txt` (cit. on p. 49).

[168] B. Quoitin, S. Uhlig, and O. Bonaventure. "Using Redistribution Communities for Interdomain Traffic Engineering". In: *QofIS'02/ICQT'02*. 2002 (cit. on p. 49).

[169] *CAIDA - AS Rank*. `http://as-rank.caida.org` (cit. on p. 52).

[170] R. T. Morris. *A Weakness in the 4.2BSD Unix TCP/IP Software*. 1985 (cit. on p. 55).

[171] W. Eddy. *TCP SYN Flooding Attacks and Common Mitigations*. RFC 4987. IETF, Aug. 2007. URL: http://tools.ietf.org/rfc/rfc4987.txt (cit. on p. 56).

[172] C. Rossow. "Amplification hell: Revisiting network protocols for DDoS abuse". In: *NDSS*. 2014 (cit. on pp. 56, 57, 75).

[173] P. Ferguson and D. Senie. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. RFC 2827 (Best Current Practice 38). Updated by RFC 3704. Internet Engineering Task Force, May 2000. URL: http://www.ietf.org/rfc/rfc2827.txt (cit. on pp. 56, 57).

[174] F. Baker and P. Savola. *Ingress Filtering for Multihomed Networks*. RFC 3704 (Best Current Practice). Internet Engineering Task Force, Mar. 2004. URL: http://www.ietf.org/rfc/rfc3704.txt (cit. on pp. 56, 60, 61).

[175] D. G. Andersen et al. "Accountable internet protocol (aip)". In: *ACM SIGCOMM*. 2008 (cit. on p. 57).

[176] B. Liu, J. Bi, and Y. Zhu. "A deployable approach for inter-AS anti-spoofing". In: *IEEE ICNP*. 2011 (cit. on p. 57).

[177] R. Miao et al. "The Dark Menace: Characterizing Network-based Attacks in the Cloud". In: *ACM IMC*. ACM. 2015 (cit. on p. 57).

[178] J. Czyz et al. "Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks". In: *ACM IMC*. 2014. DOI: 10.1145/2663716.2663717 (cit. on pp. 57, 74).

[179] R. Beverly and S. Bauer. "The Spoofer project: Inferring the extent of source address filtering on the Internet". In: *USENIX SRUTI*. 2005 (cit. on p. 57).

[180] R. Beverly et al. "Understanding the Efficacy of Deployed Internet Source Address Validation Filtering". In: *ACM IMC*. 2009 (cit. on p. 57).

[181] M. Cotton and L. Vegoda. *Special Use IPv4 Addresses*. RFC 5735. IETF, Jan. 2010. URL: http://tools.ietf.org/rfc/rfc5735.txt (cit. on p. 60).

[182] J. Weil et al. *IANA-Reserved IPv4 Prefix for Shared Address Space*. RFC 6598. IETF, Apr. 2012. URL: http://tools.ietf.org/rfc/rfc6598.txt (cit. on p. 60).

[183] C. Perkins. *IP Mobility Support for IPv4*. RFC 3344. IETF, Aug. 2002. URL: http://tools.ietf.org/rfc/rfc3344.txt (cit. on p. 60).

[184] M. Luckie et al. *Software Systems for Surveying Spoofing Susceptibility*. 2016. URL: https://www.caida.org/publications/presentations/2016/software_systems_surveying_spoofing_nanog/software_systems_surveying_spoofing_nanog.pdf (cit. on p. 61).

[185] X. Cai et al. "Towards an AS-to-Organization Map". In: *ACM IMC*. 2010 (cit. on p. 62).

[186] B. Huffaker et al. *CAIDA inferred AS to organization mapping dataset*. URL: http://www.caida.org/research/topology/as2org/ (cit. on p. 62).

[187] Cymru Team. "IP to ASN mapping". In: *http://www.team-cymru.org/IP-ASN-mapping.html* () (cit. on p. 62).

[188] RIPE NCC. *RIPE Routing Information Service (RIS)*. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris (cit. on p. 62).

[189] J. Durand, I. Pepelnjak, and G. Doering. *BGP Operations and Security*. RFC 7454 (Best Current Practice). Internet Engineering Task Force, Feb. 2015. URL: http://www.ietf.org/rfc/rfc7454.txt (cit. on p. 62).

[190] CAIDA. *Spoofer Project*. https://www.caida.org/projects/spoofer/. URL: https://www.caida.org/projects/spoofer/ (cit. on p. 68).

[191] B. Augustin et al. "Avoiding traceroute anomalies with Paris traceroute". In: *ACM SIGCOMM*. 2006 (cit. on p. 71).

[192] *Archipelago (Ark) Measurement Infrastructure*. http://www.caida.org/projects/ark/ (cit. on p. 71).

[193] R. Sinha, C. Papadopoulos, and J. Heidemann. *Internet Packet Size Distributions: Some Observations*. Tech. rep. ISI-TR-2007-643. johnh: pafile: USC/Information Sciences Institute, May 2007. URL: http://www.isi.edu/%5C%7ejohnh/PAPERS/Sinha07a.html (cit. on p. 72).

[194] P. Richter et al. "Distilling the Internet's Application Mix from Packet-Sampled Traffic". In: *PAM*. 2015 (cit. on p. 74).

[195] P. Richter et al. "A Primer on IPv4 Scarcity". In: *ACM Computer Communication Review* 45.2 (2015) (cit. on p. 75).

[196] D. Moore et al. "Inferring internet denial-of-service activity". In: *ACM Transactions on Computer Systems (TOCS)* 24.2 (2006), pp. 115–139 (cit. on p. 75).

[197] *Rapid7 Labs, Project Sonar UDP Scans*. https://scans.io/study/sonar.udp. URL: https://scans.io/study/sonar.udp (cit. on p. 77).

[198] Netflix. *Reducing Netflix traffic where it's needed while maintaining the member experience*. https://media.netflix.com/en/company-blog/reducing-netflix-traffic-where-its-needed. 2020 (cit. on pp. 79, 95, 100).

[199] European Commission. *Commission and European regulators calls on streaming services, operators and users to prevent network congestion*. https://ec.europa.eu/digital-single-market/en/news/commission-and-european-regulators-calls-streaming-services-\operators-and-users-prevent-network. 2020 (cit. on p. 79).

[200] Forbes. *Netflix Starts To Lift Its Coronavirus Streaming Restrictions*. https://www.forbes.com/sites/johnarcher/2020/05/12/netflix-starts-to-lift-its-coronavirus-streaming-restrictions/#7bcba5bf4738. 2020 (cit. on p. 79).

[201] M. Trevisan et al. "Five Years at the Edge: Watching Internet from the ISP Network". In: *CoNEXT*. 2018 (cit. on pp. 80, 87).

[202] G. Maier et al. "On Dominant Characteristics of Residential Broadband Internet Traffic". In: *ACM IMC*. 2009 (cit. on p. 80).

[203] J. L. Garcia-Dorado et al. "Characterization of ISP Traffic: Trends, User Habits, and Access Technology Impact". In: *IEEE Transactions on Network and Service Management* 9 (2 2012) (cit. on p. 80).

[204] C. Timberg. *Your Internet is working. Thank these Cold War-era pioneers who designed it to handle almost anything*. The Washington Post, April 6, 2020 https://www.washingtonpost.com/technology/2020/04/06/your-internet-is-working-thank-these-cold-war-era-pioneers-who-designed-it-handle-almost-anything/. 2020 (cit. on pp. 80, 101).

[205] A. Feldmann et al. "The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic". In: *ACM IMC*. Oct. 2020 (cit. on pp. 80, 91, 95, 96).

[206] Google. *COVID-19 Community Mobility Report*. https://www.google.com/covid19/mobility/. 2020 (cit. on p. 81).

[207] M. McKeay. *Parts of a whole: Effect of COVID-19 on US Internet Traffic*. https://blogs.akamai.com/sitr/2020/04/parts-of-a-whole-effect-of-covid-19-on-us-internet-traffic.html. 2020 (cit. on pp. 81, 102).

[208] M. McKeay. *The Building Wave of Internet Traffic*. https://blogs.akamai.com/sitr/2020/04/the-building-wave-of-internet-traffic.html. 2020 (cit. on pp. 81, 102).

[209] Comcast. *COVID-19 Network Update*. https://corporate.comcast.com/covid-19/network. 2020 (cit. on pp. 81, 102).

[210] Google. *Keeping our network infrastructure strong amid COVID-19*. https://www.blog.google/inside-google/infrastructure/keeping-our-network-infrastructure-strong-amid-covid-19/. 2020 (cit. on p. 81).

[211] C. Labovitz. *Pandemic Impact on Global Internet Traffic*. NANOG 79. 2020 (cit. on pp. 81, 88, 102).

[212] Telegeography. *State of the Network: Updates on Covid-19*. https://www2.telegeography.com/network-impact. 2020 (cit. on p. 81).

[213] Frank Schlosser et al. "COVID-19 lockdown induces disease-mitigating structural changes in mobility networks". In: *Proceedings of the National Academy of Sciences* 117.52 (2020), pp. 32883–32890. ISSN: 0027-8424. DOI: 10.1073/pnas.2012326117. URL: https://www.pnas.org/content/117/52/32883 (cit. on p. 82).

[214] A. Lakhina et al. "Structural Analysis of Network Traffic Flows". In: *ACM SIGMETRICS*. 2004 (cit. on pp. 82, 97).

[215] BLINC: Multilevel Traffic Classification in the Dark. "T. Karagiannis and D. Papagiannaki and M. Faloutsos". In: *ACM SIGCOMM*. 2005 (cit. on p. 82).

[216] M. Z. Shafiq et al. "Characterizing and ModelingInternet Traffic Dynamics of Cellular Devices". In: *ACM SIGMETRICS*. 2011 (cit. on p. 82).

[217] Cisco. *Introduction to Cisco IOS NetFlow - A Technical Overview*. https://www.cisco.com/c/en/us/products/collateral/ios-nx-os-software/ios-netflow/prod_white_paper0900aecd80406232.html. 2012 (cit. on p. 82).

[218] Cisco. https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r4-2/bng/configuration/guide/b_bng_cg42asr9k/b_bng_cg42asr9k_chapter_01.pdf. 2012 (cit. on p. 83).

[219] C. Labovitz et al. "Internet Inter-Domain Traffic". In: *ACM SIGCOMM*. 2010 (cit. on pp. 87, 101).

[220] T. Böttger, F. Cuadrado, and S. Uhlig. "Looking for Hypergiants in PeeringDB". In: *ACM CCR* 48.3 (2018) (cit. on pp. 87, 92).

[221] T. Böttger et al. "A Hypergiant's View of the Internet". In: *ACM SIGCOMM CCR* 47.1 (2017) (cit. on p. 87).

[222] C. Labovitz. *Internet Traffic 2009-2019*. APRICOT 2019. 2019 (cit. on pp. 87, 101, 102).

[223] E. Pujol et al. "Steering Hyper-Giants' Traffic at Scale". In: *Proceedings of ACM CoNEXT 2019*. 2019 (cit. on p. 87).

[224] J. Rüth et al. "A First Look at QUIC in the Wild". In: *PAM*. 2018 (cit. on p. 91).

[225] Cloudflare. *General best practices for load balancing at your origin with Cloudflare*. https://support.cloudflare.com/hc/en-us/articles/212794707-\General-best-practices-for-load-balancing-at-your-origin-with-Cloudflare. 2020 (cit. on p. 91).

[226] Microsoft. *Which ports need to be open to use Skype on desktop?* https://support.skype.com/en/faq/FA148/which-ports-need-to-be-open-to-use-skype-on-desktop. 2020 (cit. on pp. 91, 92).

[227] Microsoft. *Prepare your organization's network for Microsoft Teams*. https://docs.microsoft.com/en-us/microsoftteams/prepare-network. 2020 (cit. on pp. 91, 92).

[228] Zoom Video Communications. *Configure Meeting Connector Controller Port Forwarding*. https://support.zoom.us/hc/en-us/articles/204898919-Configure-Meeting-Connector-Controller-Port-Forwarding. 2020 (cit. on p. 92).

[229] Arunan Sivanathan, Hassan Habibi Gharakheili, and Vijay Sivaraman. "Can We Classify an IoT Device using TCP Port Scan?" In: Dec. 2018, pp. 1–4. DOI: 10.1109/ICIAFS.2018.8913346 (cit. on p. 92).

[230] Riot Games. *League of Legends: Troubleshooting Connection Issues*. https://support-leagueoflegends.riotgames.com/hc/en-us/articles/201752664-Troubleshooting-Connection-Issues. 2020 (cit. on p. 92).

[231] Cisco. *Network Requirements for Webex Teams Services*. https://help.webex.com/en-us/WBX000028782/Network-Requirements-for-Webex-Teams-Services. 2020 (cit. on p. 92).

[232] RIPE NCC. *RIPE Database Query*. https://apps.db.ripe.net/db-web-ui/query. 2020 (cit. on p. 92).

[233] F. Bronzino et al. "Mapping the Digital Divide: Before, During, and After COVID-19". In: *TPRC 48*. 2021 (cit. on p. 99).

[234] ITU. *Press Release: New 'State of Broadband' report warns of stark inequalities laid bare by COVID-19 crisis*. https://www.itu.int/en/mediacentre/Pages/PR20-2020-broadband-commission.aspx. 2020 (cit. on p. 101).

[235] J. Snijders. *Internet Network Operations during Pandemics*. Mar. 2020. URL: https://www.youtube.com/watch?v=tFeVlzBxICc (cit. on p. 101).

[236] Y. Chiu et al. "Are We One Hop Away from a Better Internet?" In: *SIGCOMM HotNets*. 2015 (cit. on p. 101).

[237] B. Schlinker et al. "Engineering Egress with Edge Fabric: Steering Oceans of Content to the World". In: *ACM SIGCOMM*. 2017 (cit. on p. 101).

[238] K-K. Yap et al. "Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering". In: *ACM SIGCOMM*. 2017 (cit. on p. 101).

[239] B. Sanghani. *COVID-19 & IXPs*. RIPE 80, https://ripe80.ripe.net/wp-content/uploads/presentations/27-ripe80-covid-ixp-1.pdf. May 2020 (cit. on p. 102).

[240] T. Leighton. *Can the Internet keep up with the surge in demand?* https://blogs.akamai.com/2020/04/can-the-internet-keep-up-with-the-surge-in-demand.html. 2020 (cit. on p. 102).

[241] Cisco. *Cisco Annual Internet Report*. https://www.cisco.com/c/en/us/solutions/executive-perspectives/annual-internet-report/index.html. 2020 (cit. on p. 102).

[242] DE-CIX. *DE-CIX Virtual Get-together - Focus Middle East & Asia 22 Apr 2020*. Apr. 2020. URL: https://www.youtube.com/watch?v=DfPt10aopns (cit. on p. 102).

[243] *Isolario Termination Notice*. https://www.isolario.it/web_content/php/site_content/join_us.php. Last accessed: 12th July, 2021. Archived version available at:https://web.archive.org/web/20210712080537/https://www.isolario.it/web_content/php/site_content/join_us.php (cit. on p. 103).

[244] Lucas Müller et al. "Challenges in inferring spoofed traffic at IXPs". In: *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*. 2019, pp. 96–109 (cit. on p. 104).

[245] Sabine Maasen and Jan-Hendrik Passoth. *Soziologie des Digitalen-Digitale Soziologie?: Soziale Welt-Sonderband 23*. Nomos Verlag, 2020, p. 20 (cit. on p. 106).

[246] T. Boettger, G. Ibrahim, and B. Vallis. "How the Internet reacted to Covid-19 — A perspective from Facebook's Edge Network". In: *ACM IMC*. 2020 (cit. on p. 106).