SAARLAND UNIVERSITY

Faculty of Mathematics and Computer Science
Department of Computer Science
Dissertation

Improving Reactive Capabilities of Internet Peering Infrastructure in Stressful Situations

Dissertation zur Erlangung des Grades des
Doktors der Naturwissenschaften (Dr. rer. nat.)

der Fakultät für Mathematik und Informatik
der Universität des Saarlandes

vorgelegt von
Daniel Wagner

Saarbrücken, 2023

Date of the colloquium:	June 3rd, 2024
Dean:	Professor Dr. Roland Speicher
Chairman of the examination board:	Professor Dr. Jens Dittrich
Reporter:	Professor Anja Feldmann, PhD Professor Georgios Smaragdakis, PhD Professor Paul Barford, PhD Dr. Christoph Dietzel
Scientific Assistant:	Dr. Tiago Heinrich

Notes on style:
As most of the work presented in this dissertation was done in collaboration with other researchers, the scientific plural "we" is used.

Abstract

The Internet has revolutionized communication, entertainment, and access to information since its inception. All of this depends on a resilient Internet infrastructure. Resilience is threatened both by analog and digital stress situations. Analog stress situations include, e.g., a global pandemic, and digital stress situations include, e.g., attacks in the Internet.

In this thesis, we investigate how the Internet can handle both analog and digital stress situations. We use a diverse set of Internet vantage points to measure the stress situations and propose conceptual, reactive, and proactive solutions to improve the overall resilience of the Internet.

In the early 2020, the COVID-19 pandemic led to an analog stress situation. People were forced to stay at home and use their residential Internet connection in their daily lives. As a consequence, we observe significant traffic surges within a few days. We find that for some networks, this increase corresponds to the traffic growth expected during a whole year. To maintain resilience, the Internet needs to enable fast increase of capacity. This requires flexibility on many levels. Yet, existing network hardware often lacks flexibility and automation.

To add more flexibility to the Internet infrastructure, we leverage programmability within the network and propose a P4-enabled Internet Exchange Point. Programmability adds flexibility as well as automation and, as such, can lead to reduction in deployment time and eases updates to the network. This, in turn, helps to reduce reaction time under stress situations.

Internet attacks, a digital stress situation, also threaten the resilience of the Internet. Typically, mitigation techniques are deployed close to the victim of the attack. Unfortunately, this means that intermediate networks still have to carry the attack traffic close to the target. To be able to drop traffic at intermediate points, we propose to utilize Internet Exchange Points. To do this effectively, we propose that multiple Internet exchange points can collaborate.

To be able to drop traffic, we have to detect attacks as early as possible. To further improve our attack detection capabilities, we propose to use "Meta-Telescopes". A Meta-Telescope uses address blocks on the Internet that are likely to be unused. Such address blocks can be inferred by utilizing Internet Exchange Points as vantage points.

Zusammenfassung

Das Internet hat seit seiner Gründung Kommunikation, Unterhaltung und Zugang zu Informationen revolutioniert. All dies braucht eine verlässliche Internet-Infrastruktur. Diese Verlässlichkeit wird sowohl von analogen als auch digitalen Stresssituationen bedroht. Zu den analogen Stresssituationen zählen beispielsweise eine Pandemie und zu den digitalen Stresssituationen beispielsweise Angriffe im Internet.

Wir nutzen unterschiedliche Internet-Standpunkte, um zu beobachten, wie das Internet auf sowohl analoge als auch digitale Stresssituationen reagiert. Wir nutzen diese Beobachtungen, um konzeptionelle, reaktive und proaktive Lösungen zu entwerfen, die die Verlässlichkeit des Internets verbessern.

Die COVID-19 Pandemie führte Anfang 2020 zu einer analogen Stresssituation. "Lockdowns" führten unter anderem dazu, dass Menschen ihren heimischen Internetzugang für ihre täglichen Routinen nutzten. Konsequenterweise wurde in wenigen Tagen ein deutlicher Anstieg von Datenverkehr im Internet beobachtet. Für einige Netzwerke entsprach dieser Anstieg dem, was deren Netzwerkbetreiber über ein ganzes Jahr verteilt zu sehen erwarteten. Um in einer solchen Situation die Verlässlichkeit des Internets weiterhin zu gewährleisten, muss das Internet imstande sein, in kurzen Zeitfenstern neue Kapazitäten zur Verfügung zu stellen. Dazu wird viel Flexibilität benötigt, an der es oft mangelt.

Um mehr Flexibilität in die Internet-Infrastruktur zu bringen, nutzen wir Flexibilisierung via Software. Mit Software können wir sowohl die Flexibilität als auch die Automatisierung erhöhen. Das kann unter anderem die Einführung von neuen Netzwerkkonfigurationen vereinfachen und beschleunigen, die in Stresssituationen dringen benötigt werden.

Attacken im Internet, eine digitale Stresssituation, bedrohen gleichermaßen die Verlässlichkeit des Internets. Typischerweise wird Attacken-Verkehr durch Schutzmaßnahmen nah am angegriffenen Netzwerk eingeschränkt. Bedauerlicherweise bedeutet das, dass Netzwerke entlang des Angriffspfades den Attacken-Datenstrom weiterleiten müssen. Um Pakete der Attacke bereits früher entlang des Pfades verwerfen zu können, schlagen wir vor, dass Internet Exchange Points zusammenarbeiten sollen.

Um Attacken möglichst früh erkennen zu können, müssen wir die Erkennung der Attacke verbessern. Dazu schlagen wir das Konzept eines Meta-Telescopes vor. Ein Meta-Teleskop ist ein Adressblock im Internet, der womöglich ungenutzt ist. Solche Blöcke können mittels Datenanalyse an Internet Exchange Points erkannt werden.

Acknowledgments

It is neither the submission of the thesis nor is it the degree awarded by the university that fills my heart with true joy. It is the experience of meeting unbelievably talented, well-meaning, and ambitious people who are willing to accompany and support me along the path to become the person I am today. I am convinced that my success depends to a large extent on how people around me motivate, shape, and teach me. Hence, my success is a direct reflection of their success. I wish to thank each and every one of you involved. Without you, none of this would be possible.

First of all, I would like to thank Anja Feldmann for supervising me on my journey. I am deeply grateful for the chance to listen to your words, learn from them, and grow into a person I never thought I'd be one day. The effort you invested in supporting me is beyond imagination. Your ability to impart such vast amounts of knowledge while maintaining a sense of humor and fostering positive motivation in me at the same time is far more than impressive. Further, I am very grateful for the opportunity to join your research group INET at the Max-Planck Institute for Informatics and become a part of the family. The whole experience is more than a personal enrichment; it is a fun ride!

I want to express my gratitude to Georgios Smaragdakis. It feels like you do not distinguish between your doctoral students and other fellow researchers in your field like me. I am very impressed by your knowledge, commitment, character, and goodwill. The number of ideas that you develop reflects your amazing ambition — and you know how to sell them.

Thank you, Christoph Dietzel. You trusted me from the beginning and gave me the chance to get to know Anja Feldmann and her research group. Thank you for leading the research department at DE-CIX that makes such outstanding collaborations possible. I want to thank Matthias Wichtlhuber and Daniel Kopp for forming an amazing research team and supporting me by any means. You make my time at work a delight.

I want to thank my collaborators for all their great work and ideas. You all contributed to this: Oliver H., Alberto, Franziska, Oliver G., Michalis, Sahil, Harm, Jeremias, Narseo, Enric, Ingmar, and Juan.

I want to thank my family and especially my parents, Sabine Parschau and Michael Wagner for the creation of my life and supporting me throughout every day of it. I want to thank Mateusz Jablonski for convincing me that computer science is the subject I want to study instead of chemistry. I thank Nadine Flach for spending your time with me and making it more than a lovely pleasure. Lastly, I want to thank all my friends and the music I enjoy for giving me the balance and confidence I need to keep going.

List of Publications

Parts of this dissertation are based on pre-published work. These works are co-authored with other researchers as listed below.

International Conference Publications

A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, **D. Wagner**, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. 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: doi.org/10.1145/3419394.3423658
Results are presented in Chapter 3.

D. Wagner, D. Kopp, M. Wichtlhuber, C. Dietzel, O. Hohlfeld, G. Smaragdakis, A. Feldmann. "United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale." In: Proceedings of the 2021 Conference on Computer and Communications Security. CCS 2021. ACM SIGSAC, 2021, pp. 970–987. ISBN: 978-1-4503-8454-4. DOI: doi.org/10.1145/3460120.3485385
Results are presented in Chapter 5.

D. Wagner, S. Ranadive, H. Griffioen, M. Kallitsis, A. Dainotti, G. Smaragdakis, A. Feldmann. "How to Operate a Meta-Telescope in your Spare Time." In: Proceedings of the 2023 Internet Measurement Conference. IMC 2023. ACM, 2023, pp. 328–343 ISBN: 979-8-4007-0382-9. DOI: doi.org/10.1145/3618257.3624831
Results are presented in Chapter 6.

Peer-reviewed Journal Publications

A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, **D. Wagner**, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. 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, pp. 101–108. DOI: doi.org/10.1145/3465212
Results are presented in Chapter 3.

Workshops

D. Wagner, M. Wichtlhuber, C. Dietzel, J. Blendin, A. Feldmann. "P4IX: A Concept for P4 Programmable Data Planes at IXPs." In: Proceegns of the 2021 SIGCOMM Workshop on Future Internet Routing & Adressing. FIRA 2022. ACM SIGCOMM, 2022, pp. 72–78. ISBN: 9-781-4503-9328-7. DOI: doi.org/10.1145/3527974.3545725
Results are presented in Chapter 4.

Contents

Chapter 1
Introduction

The Internet has evolved into the largest digital information network ever built by mankind. It is nowadays capable of connecting more than 5 billion people [40, 227, 188, 32, 125]. It supports their daily routines by enabling communication and information exchange to any digital system at an unprecedented scale. The whole society, e-commerce, education, research, leisure activities, video streaming, gaming, finance, healthcare, industry, and sports benefit from the vast and diverse information accessible via the Internet.

Still, the Internet was developed in the past century based on assumptions, that may no longer hold. While the technical foundations of the Internet are continuously reshaped, refined, and improved, some old protocols are still in use. The longer some of these protocols are in operation, the more their importance increases, and they become harder to replace. A well-known example is the Border Gateway Protocol (BGP) which first standardized in 1989 [149], reworked in 1990 [150], 1991 [151], 1995 [203], 2006 [204], 2012 [78], 2017 [158], and in 2021 [195]. BGP enables networks in the Internet to exchange reachability information and is sometimes called "the most important part of the Internet" [90, 105, 16]. Even though, it has fundamental flaws which can threaten the operation of the Internet [218, 237]. Despite endeavors to replace BGP [253], it is still kept alive and additional mechanisms are continuously added to resolve its flaws [10, 244].

BGP is not the only protocol on which the Internet relies. Another example is the Transmission Control Protocol (TCP) [82]. TCP was first standardized in 1981 [122] and is still essential for reliable data transmission while avoiding congestion situations between communicating hosts in the Internet. TCP, just like BGP, has design flaws that can be used to launch attacks in the Internet [127]. For example, attackers perform a large amount of incomplete connection establishments to exhaust the attacked host's resources that is maintaining a state to await the complete establishment of connections [216]. Having the goal of denying the targeted host to service, it falls in the category of Denial of Service (DoS) attacks. If multiple hosts are orchestrated to attack the same target, we have an Distributed Denial of Service (DDoS) attack. This type of attack can be devastating and is becoming increasingly prominent [179, 9, 31, 245]. Depending on the type of provided service, this can affect the target financially [224], be part of warfare [220, 232], or even threaten human lives directly [120].

Besides routing protocols, such as BGP, and transport protocols, such as TCP, there are many application protocols that have vulnerabilities. One example is "memcached", an application that is intended to improve content delivery latency for database-driven websites by means of caching frequently requested contents [69]. Several reports exist on using "memcached" to launch DDoS attacks [2, 222, 173, 48]. Exploiting these protocols can lead to stress situations in the Internet that have a digital cause.

Stress situations in the Internet can as well have analog cause. Among these are war scenarios, natural catastrophes, or pandemics. They can lead to physical damage to the infrastructure of the Internet [49] or a pandemic forces mankind to drastically change its Internet usage to a pattern that no operator has configured their network for [143, 161]. To be well prepared for such scenarios a deeper understanding of the underlying infrastructure and its operation is necessary. This poses multiple challenges. On the one hand, the Internet was initially not designed to be neither deployed at the scale it nowadays is, nor was it designed to be extensively measured as a whole to derive necessary insights to get a deeper understanding of the same. On the other hand, the Internet changes and evolves every day, making it hard to grasp the "big picture" that needs to be investigated for a fundamental understanding. To obtain this, we must have a look at the Internet at large.

The goal of this thesis is to understand both stress factors from the analog and digital world and how they impact the Internet and derive, realize, and evaluate possible solutions that can improve the resilience of the Internet.

To address the goal overall, we are using a measurement-driven approach using data from multiple globally distributed Vantage Points (VPs) of the Internet over multiple years. We study both, analog and digital unforeseeable events, hereby we use the COVID-19 pandemic as a representative of a stress situation from the analog world. As a representative of the stress situations from the digital world, we are using malicious activities exploiting flaws in the design of Internet protocols, namely DDoS attacks.

As mentioned above, to first gain insights into an unforeseeable stress situation in the Internet from the analog world, we focus on the implications of the COVID-19 pandemic on Internet traffic. Thus, we phrase our first research question as follows: **(1) How does the core of the Internet react to unprecedented traffic volume?** We find that the Internet was able to handle the traffic shifts and surges well. On the one hand, this was due to network operators planning with headroom for the annual traffic growth and the traffic surges did not necessarily occur during the busy hours. On the other hand, Networks that did not have enough spare resources quickly provisioned additional capacity.

The process of provisioning additional resources can be slow due to missing flexibility in both the networking equipment and its operation. In principle, flexibility is the key enabler for timely reaction to unforeseeable events. Hence, we next focus on how to add flexibility to the network. This leads us to the second research question: **(2) Can novel programming paradigms improve network configuration flexibility?** We propose a concept for a Programming Protocol-independent Packet Processors (P4) enabled Internet Exchange Point (IXP) to enrich the core of the Internet with flexibility and automation. Both help to reduce the reaction time to unforeseeable stress events in the Internet. This, in turn, improves the resilience of the Internet overall.

Regarding resilience improvements in the light of stress situations from the digital world, we turn our attention to attacks in the Internet. Traditional DDoS attack mitigation techniques are typically deployed close to the victim of attacks. This allows the attack traffic to travel through the Internet and, thereby, stress the networks along the path.

Hence, we propose a distributed DoS attack mitigation and frame our next research question as follows: **(3) Can the collaborative exchange of traffic characteristics benefit the distributed mitigation of DDoS attacks?** Using a collaborative information exchange DDoS attack traffic mitigation can be improved and take place at an earlier point in the Internet. This allows intermediate networks to be relieved from forwarding malicious traffic.

Attackers perform Internet scans to locate vulnerable hosts that can be exploited for launching their attacks. Such insights in malicious activity can help to deploy proactive DDoS attack mitigation mechanisms. So-called "Internet Telescopes", i.e., an inactive address space, often capture such scanning activities. Insights about recent scans can help to deploy attack mitigation techniques even before the actual attack is launched. Hence, we phrase our last research questions as follows: **(4) Can we distill traffic anomaly trends by focusing on inactive destinations?** Since sophisticated and devastating attacks have a highly distributed nature, both in terms of network type and geographical origin, scanning activity has to be captured in multiple parts of the Internet. We, therefore, build an Internet telescope that captures scanning activity of multiple prefixes, in multiple countries, and multiple networks without being in control of the subnets.

In the remainder of this chapter, we shed light on the precise contributions made in this thesis in § 1.1. They address the presented research questions. We then give an overview of the thesis' structure in § 1.2.

1.1 Contributions

This section gives an overview of the contributions made by this thesis.

With regard to the first research question, i.e., how the Internet reacted to unprecedented traffic volumes, we make the following contributions:

- We observe that traffic changes follow demand changes. We report a surge of 15-20% during the spring 2020 lockdown. The fall 2020 wave also had an impact, with the annual traffic increase in 2020 being higher than in a typical year.
- We find that traffic increases mostly take place during non-traditional peak hours. Daily traffic patterns are moving to weekend-like patterns, especially during the spring 2020 lockdown.
- We observe that online entertainment demands are related to hypergiant traffic surge.
- We observe that the port utilization at IXPs increases which is mostly explained by a higher traffic demand from residential users.
- We find that traffic related to remote working applications, such as Virtual Private Network (VPN) connectivity applications and video-conferencing applications, surge by more than 200%.

Regarding the second research question, i.e., whether novel programming languages can help to improve the flexibility in the core of the Internet, we make the following contributions:

- We provide a characterization of the IXP market landscape and implications for programmable, i.e., P4-based IXPs (P4IX).
- We outline technical and operational P4IX requirements and a P4IX pipeline concept to enable high flexibility and low time-to-market.
- We provide a critical discussion of operational, technical, and organizational P4IX challenges.

Regarding the third research question, i.e., whether collaborative exchange of traffic characteristics can benefit a distributed mitigation of DDoS attacks, we make the following contributions:

- We detect and analyze more than 120 thousand amplifications DDoS attacks using a collaboration of 11 IXPs.
- We show that more than 80% of the observable attacks are not mitigated because local detection mechanisms aren't triggered.
- We show the critical role that IXPs, located in the core of the Internet, can play in dropping DDoS attack traffic earlier in the network.

With regard to the last research question that focuses on distilling traffic anomalies from inactive destinations, our main contributions are:

- We develop and evaluate a methodology to identify meta-telescope prefixes, i.e., globally advertised but unused address space to cumulatively contribute to a distributed meta-telescope.
- Meta-telescope prefixes can be identified according to various requirements regarding geographical footprint, network location, and address block size.
- We show that in a single day, more than 350K /24 IPv4 prefixes can be identified as meta-telescope prefixes.
- We comment on our experience in detecting and operating a meta-telescope in the wild.
- We provide a discussion on how a meta-telescope can be used to shed light on scanning and other network activity across the Internet to know proactively about possible DDoS attacks.

1.2 Overview and Structure

The aforementioned topics are presented in this thesis according to the following structure:

Chapter 2 provides background information related to the topics of this thesis. Among these are the foundations of the Internet ecosystem, the COVID-19 pandemic, Software Defined Networking (SDN), DDoS attacks, and lastly Internet telescopes.

To investigate how the Internet reacts to a stress situation from the analog world, Chapter 3 provides an in-depth analysis of the Internet resilience during the COVID-19 pandemic. The results are published in [91] and [92].

Chapter 4 presents a concept to enable data plane programmability at IXPs. This achieves the necessary flexibility to react faster to stress events like the COVID-19 pandemic. The results are published in [243].

We then turn our attention to stress situations from the digital world, i.e., DDoS attacks, in Chapter 5. We combine data from 11 IXPs in different continents to evaluate the effectiveness of a collaborated DDoS mitigation in the core of the Internet. The results are published in [241].

In Chapter 6, we identify unused Internet Protocol (IP) address spaces in the public Internet using 14 IXPs to build a large distributed Internet telescope. The results are published in [242].

Lastly, Chapter 7 summarizes the contents of this thesis, concludes the answers to the research questions, and positions them in a greater context. Finally, some thoughts on future work are presented.

Chapter 2
Background

This chapter provides background information for this thesis. We start off with the Internet architecture in Section 2.1. After discussing Autonomous System (AS) relationships in the Internet, we position IXPs in the Internet eco-system, and explain their role for peering. Next, we present the concept of SDN in Section 2.3. We then turn our attention to events that stress the Internet. This includes the COVID-19 pandemic (Section 2.2) as well as malicious attacks in the Internet (Section 2.4). Finally, we provide an introduction to Internet telescopes in Section 2.5.

2.1 Internet Eco-System

The Internet we know today is the product of about 60 years of continuous research and development (as of June 19, 2024). It evolved into the largest data network in the world. The Internet is a network of networks. Every network is referred to as an "AS" that has a globally unique Autonomous System Number (ASN) for identification. An AS is operated by an entity, e.g., a company or an individual. Every host on the Internet has an IP address so that packets can be sent to it via the network of networks. ASes bundle IP addresses into prefixes and announce them to the other networks via the Border Gateway Protocol (BGP).

2.1.1 Internet Addressing

Just like BGP, IP has gone through multiple stages of evolution and standardization. Today, the fourth version of IP, i.e., the Internet Protocol Version 4 (IPv4) [121], and the sixth version, i.e., Internet Protocol Version 6 (IPv6) [71] are in use. Both versions use different types of IP addresses. IPv4 uses 32 bit addresses, whereas IPv6 uses 128 bit addresses. By lengthening the addresses, the available address space is increased. Groups of IP addresses are summarized using IP prefixes. An (IP) prefix consists of two parts. The first part is the first address of the address space. The second part starts with a slash symbol (/), followed by a number N indicating the "size of the prefix". The size of

a prefix corresponds to the number of addresses inside the address space of the prefix. This number can be calculated using the following equation:

$$prefix\ size = 2^{address\ bits - N}$$

For example, the IPv4 prefix `172.16.16.0/24` contains 256 IPv4 addresses, i.e., `172.16.16.0` through `172.16.16.255`. The IPv6 prefix `fe80::/110` contains 262,144 addresses, i.e. `fe80::` through `fe80::3:ffff`.

2.1.2 Internet Routing

Information about address space reachability is exchanged from AS A to AS B by AS A announcing a prefix to AS B. A prefix announcement contains, among others, the prefix itself, the ASN of the announcing AS, and the AS path. The connectivity between two ASes is established using "Border Routers", i.e., BGP speaking network devices. ASes operate and physically connect their border routers and announce their prefixes to other neighboring ASes. ASes not only announce their own address space to their neighbors but also prefixes that they have learned from other neighbors. By announcing a prefix, the AS adds its own ASN in the AS path of the announcement. Whether or not an AS announces a prefix to its neighbor is determined by routing policies. The routing policy depends on, e.g., the financial arrangement between the two ASes, their AS relationships, and considerations of traffic engineering. Border routers store the prefixes in their Routing Information Base (RIB) which they use to determine the next router to forward the packet towards its destination.

2.1.3 Autonomous System Relationships

Recall the Internet is a network of networks. This means that ASes have to interconnect to exchange traffic. When exchanging traffic, network operators have different interests in mind. This leads to different AS-relationships. Over time, two major AS-relationship classes have emerged, i.e., "Peering" and "Transit". In addition, we discuss two more different types of AS-relationships. This list is not exhaustive and there are many variants of these possible AS-relationships:

- **Peering.** An interconnection strategy is referred to as "peering", if two ASes have a mutual interest in exchanging traffic. Typically, they do not charge their peering neighbor for any traffic that they send or receive. The only costs that are incurred are related to the physical cost of the port at the border router, cabling required for physically reaching the neighboring AS, and the cost of maintaining the BGP session. ASes announce prefixes to peering neighbors according to their mutual interest. Thus, peering relations are typically not transitive, meaning that AS B will not announce prefixes learned form AS A to AS C. To reach full connectivity to the Internet one would have to peer with every other AS. As of March 2021, there are 99,868 ASes in the Internet [169].
- **Transit.** This AS-relationship helps to overcome the aforementioned shortcoming of peering, i.e., to reach global reachability. Transit requires two ASes to enter a customer-to-provider relationship, where one AS provides Internet connectivity (provider AS) to the other (customer AS). The customer AS is charged by the provider AS based on the traffic exchanged via the provider AS, typically by

the 95th percentile of traffic volume [75]. The provider AS is often also referred to as "upstream" AS. Usually, the international Tier-1 networks sell upstream to their "downstream" ASes, enabling them to reach the global Internet with a single interconnection. The Tier-1 networks themselves form a full-clique by peering with each other to offer connectivity to the entire Internet.
Depending on the requirements of an AS, it might be the cheapest solution to exchange as much traffic as possible via peering links. Thus, networks strive to have a large number of peers.

- **Sending Party Network Pays.** Another AS-relationship is called Sending Party Network Pays (SPNP) [239]. This relationship is de facto not in use as of 2023, but discussed as part of the "Network cost distribution debate" in the European Parliament [87]. Here, the sending network is charged by the receiving network based on how much data is sent. The rationale for this is that sending data imposes load on other networks, which should in turn be compensated. This concept leads to high benefits for Internet Service Providers (ISPs) as they get paid from large Content Delivery Networks (CDNs) that send their content through the ISPs' network to the end-users. However, this pricing model was strongly criticized by various entities. The Body of European Regulators for Electronic Communications (BEREC) stated, that the "traffic is requested and thus 'caused' by ISPs' customers" [85]. Netnod claimed that "the commission potentially tries to fix an imaginary issue" [186] and lastly the European Internet Exchange Association (Euro-IX) concluded that SPNP "would be detrimental for the entire Internet ecosystem" [86].

- **Sibling ASes.** Sometimes an organization operates more than a single AS, e.g., when an organization bought another organization. The relationships between such ASes are often neither strictly peering nor strictly transit. Thus, they are referred to as siblings. Sibling ASes can have an arbitrary economic and routing policy.

2.1.4 Network Classes

ASes differ greatly in size and purpose. As such, it is not straightforward to group them into classes. Among the well known classes accepted by the community are the following:

- **CDNs**. The purpose of these ASes is to distribute content to users. They typically host huge amounts of data. This data is often replicated across multiple geographically distributed Points of Presence (PoPs) and served on behalf of the original content creator. Content creators often use the service of the CDNs as their focus is on creating content and not on the technical challenges of distributing content, i.e., serving millions of consumers with low latency. CDNs deliver both (1) static content, i.e., data that has been created in the past and is not subject to any alterations, e.g., software images and updates, games, and movies, as well as (2) dynamic content, i.e., data that is created in the present and subject to immediate delivery, e.g., a video live stream of a sports event. CDN servers are mostly, if not exclusively, located in data centers, i.e., facilities specialized on housing computing resources and offering excellent connectivity to other networks.
Representatives are: Akamai [187] and Fastly [88].

- **Cloud Networks**. The purpose of these ASes is to provide Virtual Machine (VM) resources. Virtual machines are the result of virtualizing compute resources and

are typically operated in data centers. Cloud networks are usually using resources in multiple data centers over the globe, they provide access to VMs. VMs have different hardware and network specifications. Cloud Networks are often used to outsource services of a company. Their motivation is having flexibility, reduction of Capital Expenditures (CAPEX) and Operational Expenditures (OPEX), and that they are too small to operate their own equipment. Cloud networks received great popularity by many companies that either can't afford the costs of the hardware or the operation thereof, lack the experience or time to maintain the machines, or require their service only for a short period so that it is not profitable to purchase an own instance of these machines.
Representatives are: Microsoft Azure [172] and Amazon AWS [6].

- **ISPs**. The purpose of these ASes is to provide Internet connectivity to other ISPs and end users. This class further divides into three sub-categories, depending on their network size, geographical reach, and offered connectivity. The three sub-categories are referred to as "Tiers". Tier-1 networks offer global connectivity and typically operate large networks consisting of multiple thousands of routers around the globe. Tier-2 networks connect to Tier-1 networks for global connectivity and are smaller networks with national reach. Tier-3 networks use Tier-2 networks for global connectivity and typically offer connectivity within a region [55].
Representatives are: Deutsche Telekom [231], Tele2 [229], and SachsenGigaBit [212].

- **Educational Networks**. The purpose of these ASes is to provide Internet connectivity to educational institutes. This can be restricted to a single university, such as, e.g., REDIMadrid [202], or a set of research institutes including universities, such as Deutsches Forschungsnetz (German Research Network) (DFN) [97]. The network characteristics of this class can vary greatly depending on the purpose of the connected institutions.
Representations are: REDIMadrid [202], DFN [97].

- **Enterprise Networks**. The purpose of these ASes is to connect office computers to a central server of the enterprise. Since the cloudification, these enterprise servers are no longer operated in-house, but rather in cloud networks. Hence, enterprise networks often have great connectivity to cloud networks.
Representatives are: Deutsche Börse AG [111], Daimler Trucks [63].

2.1.5 Internet Exchange Points

Peering is an attractive AS relationship. ASes typically favor traffic exchange via peering over transit. But it is difficult to establish many peering sessions. IXPs solve this problem. IXPs provide a Layer-2 interconnection fabric, i.e., a so-called "Peering Local Area Network (LAN)", to which ASes can connect via one or multiple ports. They can then use the peering LAN to establish peering relationships with other networks in the peering LAN.

IXPs vary greatly in terms of the number of physical switches, the number of connected networks (also referred to as a "member" or "customer" of the IXP), the peak of exchanged traffic volume, and the number and type of services provided. To ease the establishment of peering sessions, IXPs often operate Route Servers (RSs). Route servers collect routes from IXP members and can, if desired, forward them to other members. Figure 2.1 shows an IXP consisting of four switches that form the peering LAN with five connected ASes.

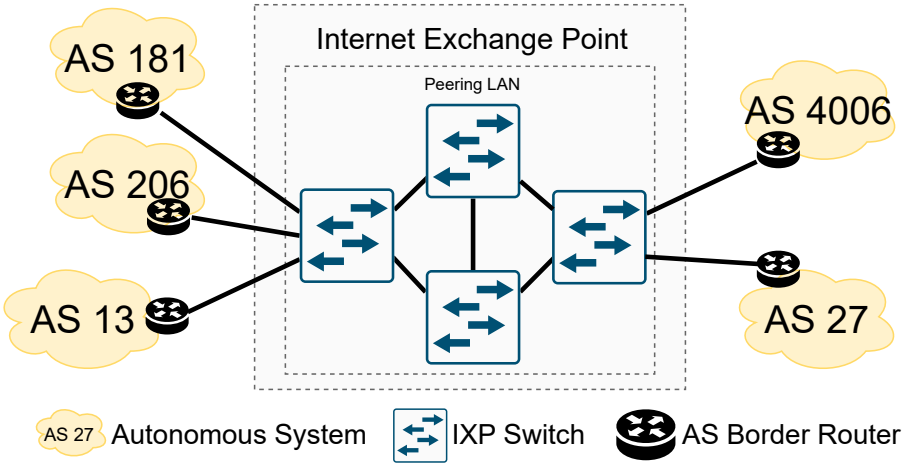

Figure 2.1: An IXP with five connected ASes.

It shows that ASes operate a border router to connect to the peering LAN and establish peerings.

We use IXPs as vantage points in this thesis. In particular, we rely on flow data collected by the IXPs' switches. Flow data is collected by the IXP for operational purposes in compliance with the legal regulations of the countries of operation. It is generated at the switches in the peering LAN by inspecting sampled packet headers. For every inspected packet header, a flow-record is created, unless it exists already. A flow record consists of information about (1) the involved border routers (identified by their source and destination Media Access Control (MAC) address), (2) the source and destination of the packet in the Internet (identified by the source and destination IP address), (3) the transport protocol used (identified by the "protocol" field in IP header), and (4) the communicating applications (identified by the source and destination transport ports). Once a flow record is generated, it stays for at least 15 seconds in the switch before it is exported as flow data. If additional packets of the same flow are sampled, they are aggregated in the flow record. The format of the exported flow data is IPFIX [44].

2.2 COVID-19 Pandemic and Remote Working

In late 2019, a hazardous and highly contagious virus, named Severe Acute Respiratory Syndrome Coronavirus Type 2 (SARS-CoV2), first occurred in China. Despite immediate measures, called "lockdowns", imposed by the Chinese government, the outbreak of the resulting Coronavirus Disease 2019 (COVID-19) could not be stopped to reach the extent of a pandemic [102]. An increasing number of confirmed COVID-19 cases was observed in all parts of the world. This led many governments to impose lockdown measures on their population. This aimed at minimizing the risk for the population in terms of suffering from severe symptoms and to minimize the risk of spreading the virus. Lockdowns occurred in early 2020 all over the globe, which imposed "social distancing" measures to reduce the physical contact between people. This helped to slow down the infection rates and ensured that the health system provided enough capacity to accommodate severe cases. This minimized the number of causalities due to the virus.

These measures included an appeal to all employers to let their employees work remotely from home. Such lockdowns, while well intended, had significant implications on the Internet traffic; people were relying on the Internet in even more daily routines than ever before [230, 160, 109].

The term "remote work" describes office work that is done from a remote location that is not the conventional office. This remote location is in most cases the home of the employee. This can only be possible, if the work does not require physical presence at the worksite and Internet connectivity is given at the remote location. Office work, which mainly consists of interactions with a computer, can be transferred to the employee's home using various solutions. Among these are, e.g., the deployment of VPNs, which allow an employee to securely access the remote office network through his residential connection to the public Internet. Many interactions with colleagues moved to web-conferencing tools that stream the microphone and camera input of the employee through the public Internet to the co-workers. On the one hand, this enabled many people to continue their daily work during lockdowns due to the COVID-19 pandemic. On the other hand, this led to an analog stress situation for the Internet infrastructure. Some can see this as collateral damage of an unforeseeable but highly justified decision.

2.3 Software Defined Networking

Due to the growing demand for more network flexibility, networking hardware has been augmented by the concept of software programmable network devices, creating the field of Software Defined Networking. Previously, hardware vendors focused on developing high-performance networking devices with a fixed and inflexible set of functions. Thus, operators are often forced to rely on workarounds to address problems for which their hardware does not offer a solution. While there is a feedback loop between operators requesting new features and hardware manufacturers realizing new features there is a significant time lag and not all features are realized. This feedback loop consists of multiple sometimes very time consuming steps—for example, RFC4271 [204] took over 10 years of standardization work to become an Internet standard. Even after the publication of a standard, e.g., through an organization like the Internet Engineering Task Force (IETF) [95] that maintains its standards in Requests for Comments (RFCs) [96], hardware vendors may or may not choose to implement the new standard. Either they adapt it in their upcoming firmware, or, if new hardware is required, in their upcoming series of devices. In the latter case, operators need to upgrade their hardware to use the requested features. This often leads to significant expenses, extensive testing, and risk of service disruption.

In SDN, the operator has the freedom to implement new features in software that is then deployed on SDN-enabled switches in the network. SDN also offers to consolidate the control plane at a centralized software defined controller. Traditionally, the control plane packets that enable, e.g., routing, are transmitted using the same network as the data plane packets, i.e., packets that carry actual end-user data. In SDN, the software defined controller manages a network of forwarding devices, i.e., SDN switches. They are equipped with simple hardware, i.e., hardware with less control plane logic, that receives flow rules instantiated by the SDN controller. Packets for which no flow rule matches are relayed to the SDN controller. It processes the packet and instantiates the missing forwarding rule in the SDN switches which can then forward the packet. A SDN network can be implemented using various technologies. The most prominent technology is OpenFlow [162].

One of the latest SDN technologies is P4 [18]. P4 moves beyond the control plane by adding programmability to the data plane. It is a possible solution to reduce the deployment time of new features to the network. Operators can deploy their own protocols and features as fast as a software development team can implement new requirements. Instead of relying on a year-long standardization process with the risk of having to replace hardware, prototypes can be deployed immediately to the existing network hardware. An example for a P4-programmable network device is the Intel Tofino [123].

SDN is a concept that can be used in different networks, e.g., data center networks, home networks, or IXP networks. If the idea of SDN is applied to IXP networks, it is referred to as a Software Defined IXP (SDX) [114]. To improve related work that implements an SDX using OpenFlow [23, 113, 112, 114, 137, 157, 34, 181] we conceptualize a P4-based SDX in Chapter 4.

2.4 Distributed Denial of Service Attacks

Countless servers are online in the public Internet. They host numerous services and applications using various protocols. Almost all of them have vulnerabilities. Some of these vulnerabilities are rooted in the design of their software and can be exploited with malicious intent. Indeed, new exploits appear over time. While some of these exploits can be fixed via security updates, attackers can abuse these exploits until they are fixed. This especially applies to systems that are reachable through the public Internet but are no longer maintained, i.e., supplied with security updates. Some vulnerabilities allow attackers to take over the system, others allow the attackers to amplify traffic. We refer to such systems as "infected" systems.

Attackers can orchestrate infected systems to impair the availability of any targeted online service. If successful, the target's resources are depleted so that its service is denied to benign users. We use the term "Denial of Service (DoS)" attack for this class of attacks. If an attacker leverages multiple hosts simultaneously to attack the same target, we have a "Distributed DoS" attack. This thesis focuses on the sub-class of "amplification reflection DDoS attacks". Here, the attacker exploits a service that relies on a vulnerable protocol, e.g., Network Time Protocol (NTP). A protocol needs two properties to qualify for an amplification reflection DDoS attack: (1) the protocol must generate responses upon request and (2) these responses must contain more bytes than the request. For example, if an attacker locates an NTP server on the web, it only requires a single 236-byte `req_get_monlist` [57] request to generate 10 response packets each of length 4,460 bytes [45]. In this case, the responses generated by the NTP server are about 189 times larger than the initial request. Another vulnerable protocol, i.e., "memcached" allows an attacker to generate responses up to 51,200 times larger than the initial request [48]. By spoofing the source address of the request, the response of the servers can be redirected to any target in the Internet.

Mitigation techniques that are commonly implemented to defend against a DDoS attack typically operate at the targeted network [4, 47]. As all attack traffic accumulates at this point, detection is straightforward and it is possible to mitigate the attack by dropping the traffic. However, this mitigation technique has the drawback that networks in the core of the Internet still have to forward the attack traffic to the target. This effectively leads to stress situations in the intermediate networks. However, parts of the attack traffic traverse different intermediate networks, which makes the detection of an attack in the

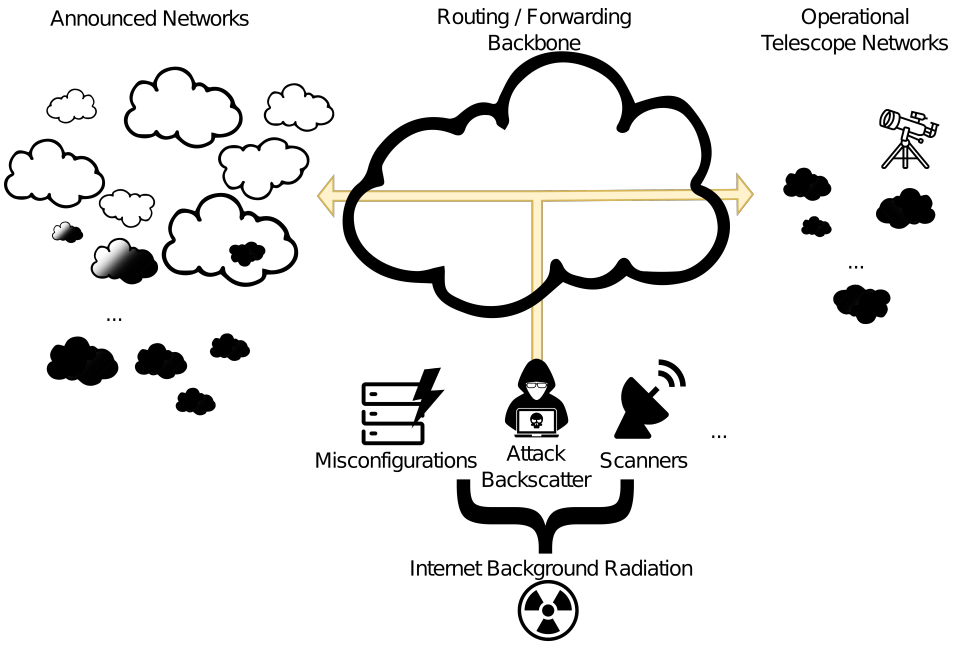

Figure 2.2: Sources and reach of IBR in the Internet.

core of the Internet more difficult. To overcome this, multiple intermediate networks can combine their local visibility to improve the detection of ongoing attacks in the Internet. This helps to drop attack traffic earlier in the Internet and also reduces the collateral attack damage dealt to intermediate networks. In Chapter 5, we study the benefit of such a collaboration.

2.5 Internet Telescopes

Reliable information about upcoming attacks is hard to obtain. But indications can be derived from abnormal or malicious activity in the Internet. So-called "Internet Telescopes" [176] record incoming unsolicited network traffic. A network telescope, or simply "telescope", is an infrastructure that passively monitors traffic reaching Internet address space that is not assigned to any hosts but is advertised to the global routing system (i.e., *dark address space*). This traffic is by definition *unsolicited* and is constituted of an evolving mix of diverse traffic components originating from across the whole Internet [13]. We refer to such unsolicited network traffic as "Internet Backround Radiation (IBR)".

2.5.1 Internet Background Radiation

Over the years, researchers found ways to extract insights into various Internet properties and phenomena from IBR, such as, e.g., identifying misconfigurations [13] and large-scale malicious activities [226, 174, 178, 175, 67], monitoring Internet connectivity [68], inferring the utilization of the IPv4 space [66], etc. Such traffic is fundamentally part of nonproductive traffic and abnormal activity. In Figure 2.2, we sketch how IBR travels through the Internet, denoted by yellow arrows. It shows the three different causes for

IBR and how it spreads through the Internet. Operational telescope networks (right) are exposed to IBR just like other networks in the Internet (left).

There are three main causes for IBR:
Scanning Activity. As explained in Section 2.4, an amplification reflection attack requires systems in the Internet that host an application that relies on an exploitable protocol. For an attacker to find these systems, they can send scanning packets to the whole IP address space and see what the responses are. This practice causes scanning packets to arrive at the Internet telescope's IP address space.
Attack Backscatter. In the event of a DDoS attack, the attacker typically veils the source IP address to remain unidentifiable and instead picks a random IP address for outgoing packets. If the attacked target replies to the incoming, spoofed packets, they might be destined to the Internet telescope's address space.
Misconfigurations. Another reason why packets arrive in an Internet telescope's address space are misconfigurations. Some services are manually configured to contact certain IP addresses for coordination, e.g., Domain Name System (DNS), or synchronization, e.g., NTP. If the configuration file contains typos, the Internet telescope's address space might be contacted, although to no avail.

The main requirement to receive IBR is to announce an address space that does not host any services to the public Internet. Some research groups dedicate large chunks of address space to record incoming IBR [27, 166]. Ideally, this address space has never hosted any service before. As address space is scarce, leaving it unused may not be what network operators desire, as such we need alternative ways to study IBR. In Chapter 6, we propose to use IXPs to build an Internet meta-telescope.

2.5.2 Visualization of IP Address Space

For our analysis, we have to find a way to visualize IPv4 address space. This is challenging, as there are over 4 billion IPv4 addresses in the IPv4 address space. A possible method for visualization is using the Hilbert Curve, i.e., a continuous fractal space-filling curve [117]. In its 12th order, the Hilbert Curve takes 16,777,216 turns which corresponds to the number of /24 subnets in the IPv4 address space. Rendered as a 4096 by 4096 pixel grid, each of which corresponds to a /24 subnet, it can be enriched with annotations about the owning party of a block of subnets. In Figure 2.3, we provide the empty Hilbert Curve of the IPv4 address space.

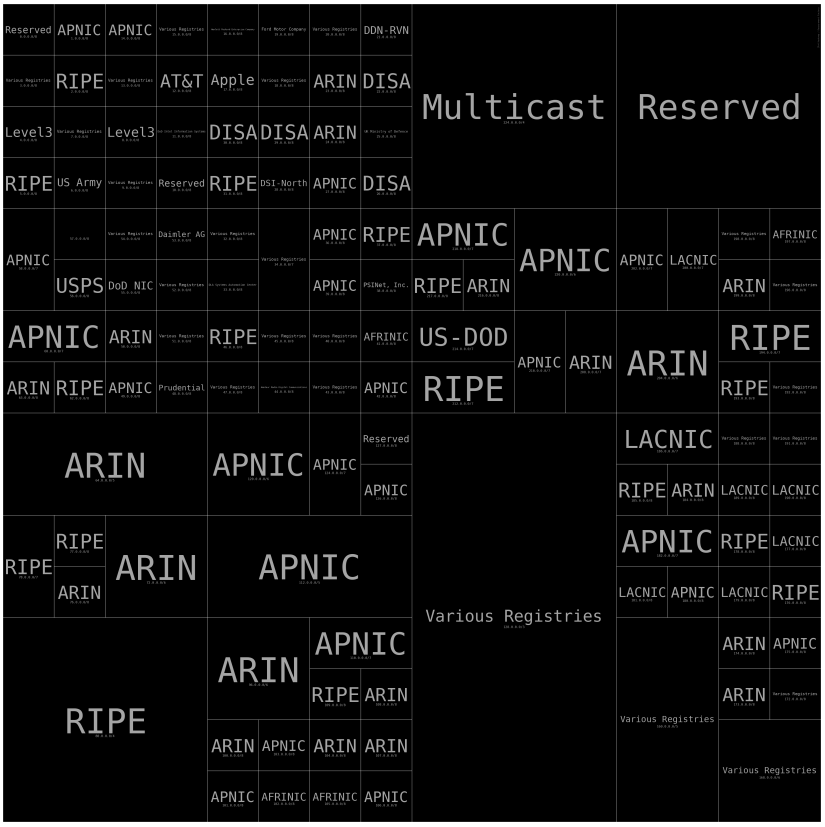

Figure 2.3: Hilbert Curve of the whole IPv4 address space.

Chapter 3
The Internet Under Stress

The COVID-19 pandemic is most likely a once in a generation global phenomenon that drastically changed the habits of millions of Internet users around the globe. As a result of the government mandated lockdowns, a large fraction of the population had to depend on their residential Internet connectivity for work, education, social activities, and entertainment. Unexpectedly, the Internet held up to this unforeseen demand with no reports of large scale outages or failures in more developed countries [233]. This unique phenomenon allows us to observe changes that may be expected within months or years in a matter of days.

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 [234, 155, 100]. In the last twenty years, user profiles have changed significantly. We observed user profile shifts from peer-to-peer applications in the early 2000s [191, 76, 247], to content delivery and streaming applications in 2010s [145, 77, 140, 196, 142], and more recently to mobile applications [248, 119]. Although changes in user profiles are a moving target, they typically have time scales of years. Thus, staying up to date, e.g., via measurements, was feasible.

In this chapter, we study the effect that government-mandated lockdowns had on the Internet by analyzing network data from a major Central European ISP (ISP-CE), three IXPs located in Central Europe, Southern Europe, the US East Coast, and a Spanish educational network (EDU). This enables us to *holistically* study the effects of the COVID-19 pandemic both from the network edge (ISP-CE/EDU) and the Internet core (IXPs).

Overall, we find that:

- Changes in traffic volume follow demand changes, causing a traffic surge of 15-20% during the spring 2020 lockdown for the ISP/IXPs in our study. In summer 2020, after the reopening of the economy, an increase of about 20% at one IXP. The fall 2020 wave also had an impact, with the annual traffic increase in 2020 being higher than in a typical year.

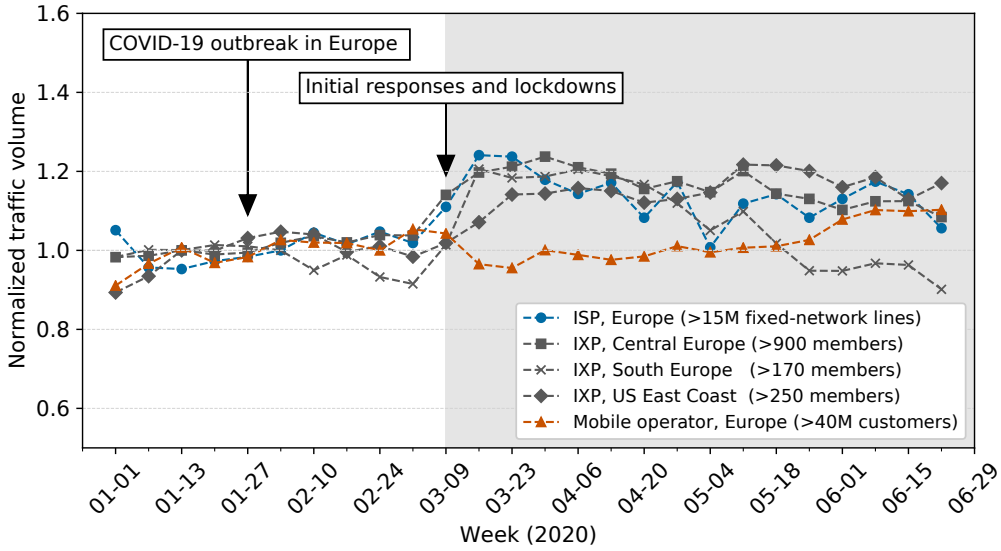

Figure 3.1: Traffic changes during 2020 at multiple vantage points. Daily traffic is averaged per week and normalized by the median traffic volume of the first up to ten weeks.

- The observed traffic increase mostly takes place during non-traditional peak hours. Daily traffic patterns are moving to weekend-like patterns, especially during the spring 2020 lockdown.
- Online entertainment demands account for hypergiant traffic surge. Yet, the need for remote working increases the relative traffic share of many "essential" applications like VPN and conferencing tools by more than 200%. At the same time, the traffic share for other traffic classes decreases substantially, e.g., traffic related to education, social media, and—for some periods—CDNs.
- At the IXP-level, we observe that port utilization increases. This phenomenon is mostly explained by a higher traffic demand from residential users.

In this chapter, we start by presenting (1) the results of an initial study that was done shortly after the lockdowns were imposed(Section 3.1) and (2) a long-term study that was done after the lockdowns were in place for more than a year (Section 3.2). The ethical considetarions for both studies can be found in Section 3.3, followed by related work in Section 3.4, and a discussion section (Section 3.5), and a summary in Section 3.6).

3.1 Implications of the COVID-19 Pandemic on Internet Traffic

Shortly after the confinement measures were imposed by many governments due to the COVID-19 pandemic, we performed an initial analysis on the increased Internet traffic demands of residential users. In particular, we measured Internet traffic for remote working, entertainment, commerce, and education that all showed traffic shifts as a result

of the lockdowns. In this study, using data from a diverse set of vantage points (one ISP, three IXPs, and one metropolitan educational network), we examine the effect of these lockdowns on traffic shifts.

COVID-19-induced weekly growth. We observe a significant traffic evolution in 2020 at multiple Internet vantage points in Figure 3.1. 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 3.1.1 for details): Traffic demands for broadband connectivity, as observed at an ISP in Central Europe as well as at a major IXP in Central Europe and an IXP in Southern Europe increased slowly at the beginning of the outbreak and then more rapidly by more than 20% after the lockdowns started. The traffic increase at the IXP at the US East Coast trails the other data sources since the lockdown occurred several weeks later. While we observe this phenomenon at the ISP and IXP vantage points, one difference between them is that the relative traffic increase at the IXP seems to persist longer while traffic demand at the ISP decreases quickly towards May. This correlates with the first partial opening of the economy, including shop reopenings in this region in mid-April and further relaxations including school openings in a second wave in May. Our findings are aligned with the insights offered by mobility reports published by Google [107] and the increased digital demand as reported by Akamai [160, 161], Comcast [51], Google [109], Nokia Deepfield [143], and TeleGeography [230].

Drastic shift in usage patterns. In light of the global COVID-19 pandemic a total growth of traffic is somewhat 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. To this end, we show the daily traffic patterns at two of the above mentioned vantage points in Figure 3.2. The Internet's regular workday traffic patterns are significantly different from weekend patterns [146, 132, 219]. On workdays, traffic peaks are concentrated in the evenings, see Figure 3.2(a). For instance, Wed., February 19 vs. Sat., February 22, 2020: With the pandemic lockdown in March, this workday traffic pattern shifts towards a continuous weekend-like pattern, as can be seen in the daily pattern for Mar. 25, 2020 in Figure 3.2(a). More specifically, we call a traffic pattern a workday pattern if the traffic spikes in the evening hours and a weekend pattern if its main activity gains significant momentum from approximately 9:00 to 10:00 am. For our classification, we use labeled data from late 2019 and use an aggregation level of 6 hours. Then, we apply this classification to all available days in 2020. Figures 3.2(b) and 3.2(c) 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 blue, otherwise they are colored in orange. 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 3.2(c). This pattern changes drastically once the confinement measures are implemented: Almost all days are classified as weekend-like. This change persists in Figure 3.2(c) until the end of June due to the vacation period, which is consistent with the behavior observed in 2019 (not shown). In contrast, Figure 3.2(b) shows that the shift towards a weekend-like pattern becomes less dominant as countermeasures were relaxed in mid-May.

These observations raise the question of the cause for this significant traffic growth and shift in patterns, given that many people are staying at home for all purposes, e.g., working from home, remote education, performing online social activities, or consuming

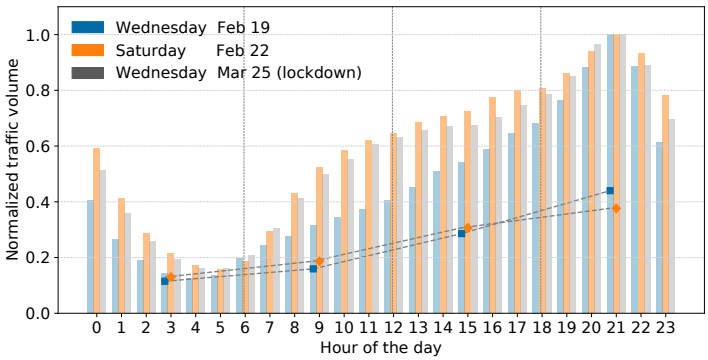

(a) ISP-CE: Hourly traffic increase and workday vs. weekend pattern for February 19 (Wed), February 22 (Sat), March 25(Wed).

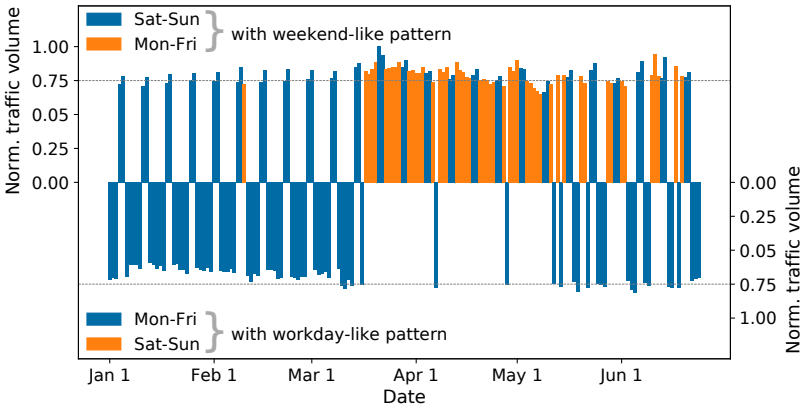

(b) ISP-CE: Workday-like (bottom) vs. weekend-like (top) January 1–June 24.

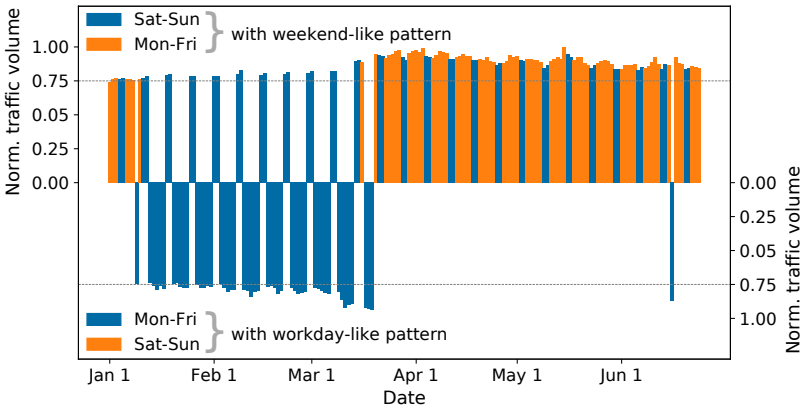

(c) IXP-CE: Workday-like (bottom) vs. weekend-like (top) January 1–June 24.

Figure 3.2: Drastic shift in Internet usage patterns for times of day and weekends/workdays.

entertainment content. The increased *demand in entertainment*, e.g., video streaming or gaming, may imply an increase in hypergiant traffic. This is in accordance with a statement by a commissioner of the European Union which stated that major streaming companies reduced their video resolution to the standard definition from March 19, 2020 onward [185, 84]. According to mainstream media, some started to upgrade their services back to high definition or 4K around May 12, 2020 [94]. Furthermore, *the need for remote working* may imply an increased demand for VPN services, usage of video conference systems, email, and cloud services.

3.1.1 Data Sets

This section describes the network traffic data sets that we used for our analysis. We utilize vantage points at the core of the Internet (IXPs), at the backbone and peering points of a major Internet Service Provider, and at the edge (a metropolitan university network), all which we will 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 [39] at all border routers to support its internal operations. We rely on two different sets of NetFlow records for this chapter. First, we use NetFlow data collected at ISP's Border Network Gateways [38] 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 about 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 a 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 a peak traffic of more than 600 Gbps. At the IXPs we use IPFIX data [44].

EDU: Network flows from the REDImadrid [202] academic network interconnecting 16 independent universities and research centers in the region of Madrid. It serves nearly 290,000 users including students, faculty, researchers, student halls, WiFi networks (including Eduroam), and administrative and support staff. The network operator provided us with anonymized NetFlow data captured at their border routers (captured at all ingress interfaces) during 72 days in the period of Feb 28 to May 8, 2020. The final data set contains 5.2B flows entering or leaving the educational network.

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.

	ISP-CE	IXP-CE	IXP-SE	IXP-US	EDU
base	Feb 20–26	Feb 20–26	Feb 20–26	Feb 20–26	Feb 20–26
March	Mar 19–25	Mar 19–25	Mar 12–18	Mar 19–25	Mar 12–18
April	Apr 09–15	Apr 23–29	Apr 23–29	Apr 23–29	Apr 23–29
June	Jun 18–24	Jun 18–24	Jun 18–24	Jun 18–24	n/a

Table 3.1: Summary of the dates used in weekly analyses. Dates in Southern Europe vary due to different courses of the pandemic.

Time frame: We use two methods to reflect the developments since the beginning of the COVID pandemic: (1) for general trends over time we use continuous data from *Jan 1, 2020—Jun 24, 2020*, (2) to highlight detailed developments we compare 7-day periods as shown in Table 3.1 from before, during, after and well after the lockdown in 2020.[1]

3.1.2 Aggregated Traffic Shifts

To understand traffic changes during the lockdown we first look for overall traffic shifts before, during, and after the strictest lockdown periods. Moreover, we take a look at hypergiant ASes vs. other ASes, shifts in link utilization, and ASes relevant for remote working.

3.1.2.1 Macroscopic Analysis

Figure 3.4 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 (see Table 3.1). For the ISP-CE, Figure 3.4(a) shows the time series using normalized one-hour bins. For the IXPs, Figure 3.4(b) 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.

As previously noted, the ISP-CE time series shows the same workday to weekend traffic patterns shifts starting with the lockdown in mid-March. 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 3.4(b), not only do we see an increase in peak traffic but also in

[1]Due to data availability, the ISP-CE is using Apr 09–15 which covers the Easter holiday period. As partial lockdowns and travel restrictions were still in place, the introduced bias may be very small.

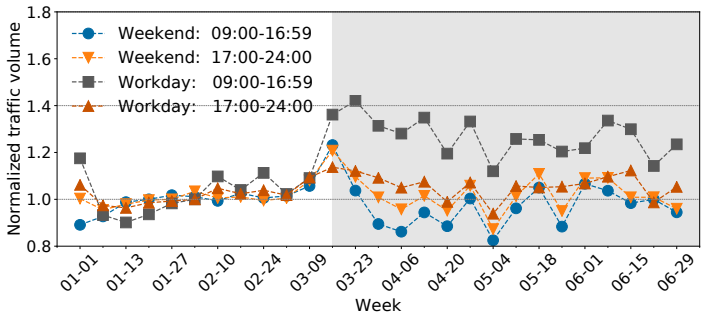

(a) Hypergiants

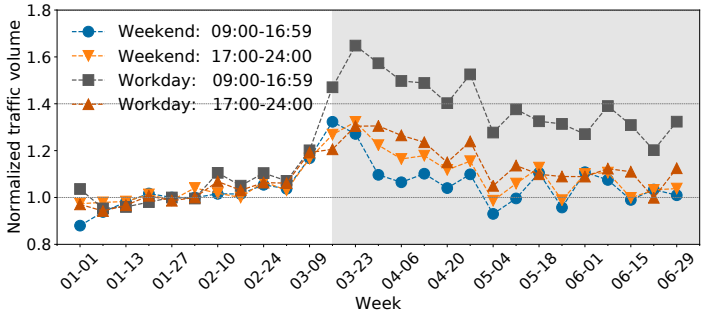

(b) Other ASes

Figure 3.3: ISP-CE: Normalized daily traffic growth for hypergiants vs. other ASes across time.

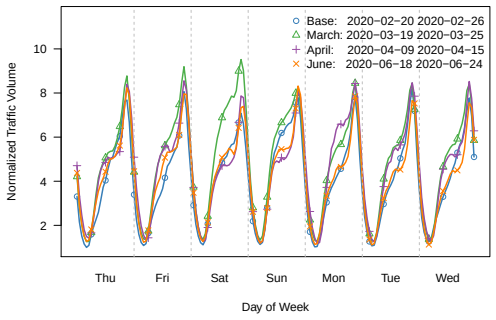

(a) L-ISP (Central Europe).

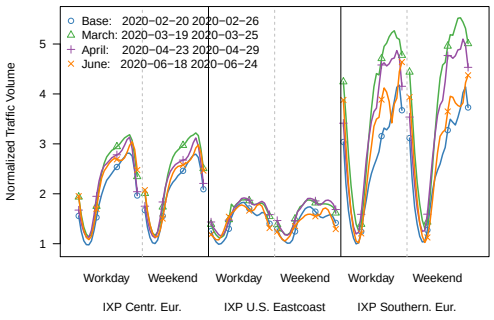

(b) IXPs (Central Europe/US Eastcoast/Southern Europe).

Figure 3.4: 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.

Org. Name	ASN
Apple Inc	714
Amazon.com	16509
Facebook	32934
Google Inc.	15169
Akamai Technologies	20940
Yahoo!	10310
Netflix	2906
Hurricane Electric	6939
OVH	16276
Limelight Networks Global	22822
Microsoft	8075
Twitter, Inc.	13414
Twitch	46489
Cloudflare	13335
Verizon Digital Media Services	15133

Table 3.2: List of Hypergiant ASes Used to classify data in Figure 3.3.

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 Europe 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 (1) serves customers from many different time zones, and (2) 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.

3.1.2.2 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 [145, 21, 20]. Table 3.2 reports the full list of ASes considered for this category. 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 in Central Europe which is consistent with recent reports in the literature [142, 196, 234]. 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 3.3. 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. Yet, the demand for entertainment content—mainly 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.

3.1.2.3 Link Utilization Shifts

We analyze to which extent the observed changes are reflected in our link utilization data set 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 data set 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 3.5.

Figure 3.5 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 [143]. 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.

Figures 3.6 and 3.7 show the relative link utilization at IXP-CE for weeks in April and June, respectively. We also plot the link utilization from the reference week in February for comparison. These plots show, in contrast to Figure 3.5, an increased overall link utilization at IXP-CE.

3.1.3 Application Classes

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 [21, 223], product and service documentations [99, 41, 170, 171], and public databases [184, 193].

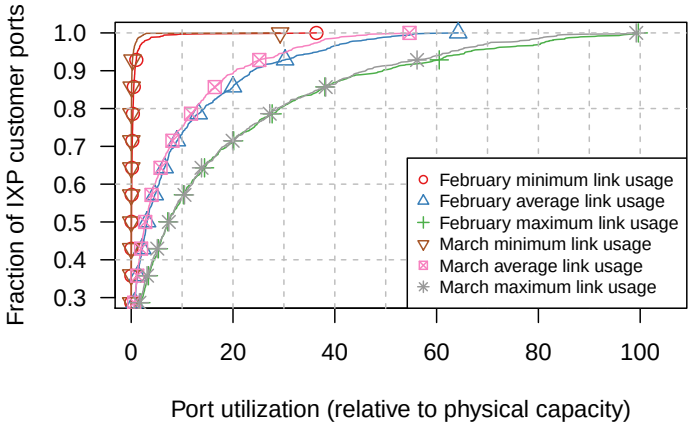

Figure 3.5: IXP-CE: ECDF of link utilization before and during the lockdown.

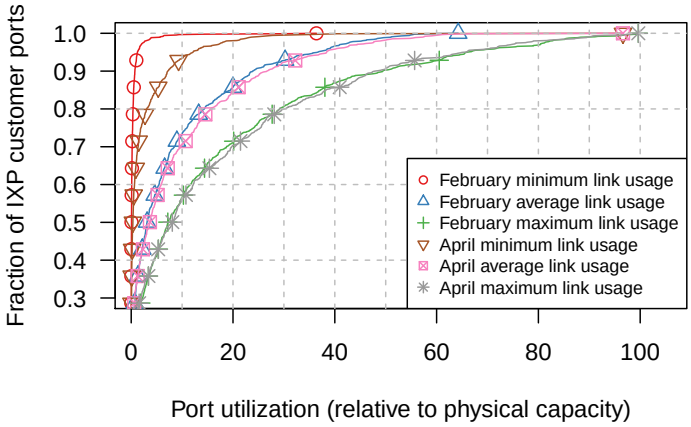

Figure 3.6: ECDF of minimum, average and maximum link utilization at IXP-CE, February week vs. April week.

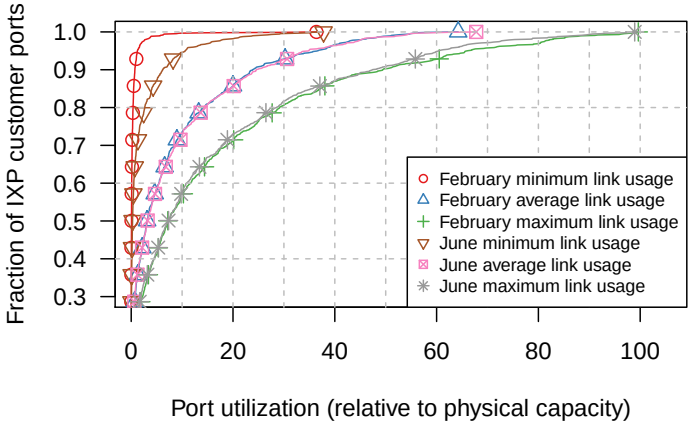

Figure 3.7: ECDF of minimum, average and maximum link utilization at IXP-CE, February week to June week.

To the best of our knowledge, there is no established and comprehensive classification of flow data into traffic classes. Even if such a classification existed, it would be a constantly moving target and highly dependent on the vantage point. These classifications have the largest possible overlap, but may differ between vantage points for one or more of the following reasons.

Local differences. We are investigating vantage points from a total of three countries on two continents. There exist local content providers and ISPs in each country that play a dominant role in their respective home market (e.g., digital offers of local broadcasting networks, national ISPs). Likewise, for IXPs, not every network is present at every IXP, which makes defining a common classification across different IXPs difficult.

Different types of Networks. We investigate different types of networks attracting different traffic mixes. For instance, cloud gaming does not play a major role in academic networks , and Video on Demand is usually not consumed via mobile providers. Consequently, different traffic classes are relevant for different networks leading to a different classification.

Ease of Classification. Not all traffic classes can be classified easily and they are not mutually exclusive. An example is the VPN classification in subsection 3.1.4 requiring the additional use of DNS information. Moreover, the number and size of the data sets used in this work is exceptional, so certain classifications cannot be performed on all data in reasonable time.

Notably, the goal of the classifications defined in this work is not to catch *all* traffic for a certain traffic class, but rather a *representative* subset of traffic allowing to reason about trends during the pandemic. As ISP and IXP networks have a comparable traffic mix, we compiled a joint classification for the ISP/IXP vantage points allowing for a high comparability. The classification is based on combinations of ASes (at IXPs by port, at ISP by IP ranges) and transport protocol ports if characteristic protocols exist. While the transport protocols are disclosed in 3.3, the measured ASes cannot be disclosed due to non-disclosure agreements. In the following we disclose as many details of the classifications used in this work as possible.

We aggregate the filtered data into 8 meaningful application classes representing applications consumed by end-users on a daily basis (See Table 3.3): *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, *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.

Figure 3.8 showcases the *Gaming* class at the IXP-SE vantage point. For this application class, we filter data of five gaming software/services providers and 57 typical gaming transport ports in various combinations (see Table 3.3). We then analyze the changes in usage behavior using two metrics: (1) the number of distinct source IP addresses, as a way to approximate the order of households, and (2) the traffic volume. Figure 3.8 shows clear changes when comparing multiplayer and cloud games before and during the lockdown. From week 10 on, i.e., when the local government imposed a lockdown, the number of unique IPs seen in the trace as well as the delivered volumes rose steeply with substantial gains of the daily minimum, average, and maximum. Notably, during

Application class	# of filters	# of distinct ASNs	# of distinct transp. ports	Notes
Web conferencing & telephony	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)
E-mail	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	8	8	-	Dominant CDN providers (excluding social network CDNs) by AS

Table 3.3: 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 Figures 3.8, 3.9.

the first lockdown week, the accounted volume plunges for two days to the lowest values observed in that time frame. We verified that this is not a measurement artifact. Instead, the drop correlates with an outage of a large gaming provider, which may be related to the sudden increase in users.

We perform the application classification for the different IXP vantage points (IXP-SE, IXP-CE, IXP-US) and for the ISP-CE.[2] 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 stages of the COVID-19 measures as they were imposed throughout Europe—see Table 3.1 in Section 3.1.1.

Normalization and filtering: After normalization as outlined in Section 3.1.1, 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

[2]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.

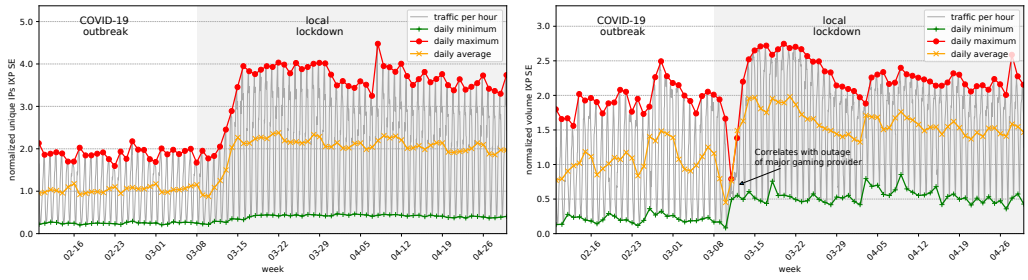

(a) (b)

Figure 3.8: IXP-SE: Application class Gaming before and during lockdown. It shows a steep increase in # IPs and traffic volume.

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 3.9 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. We discuss VPN traffic in Section 3.1.4. 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 [185] for 30 days. In the case of the ISP-CE that covers the *March* as well as the *April* week.[3]

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

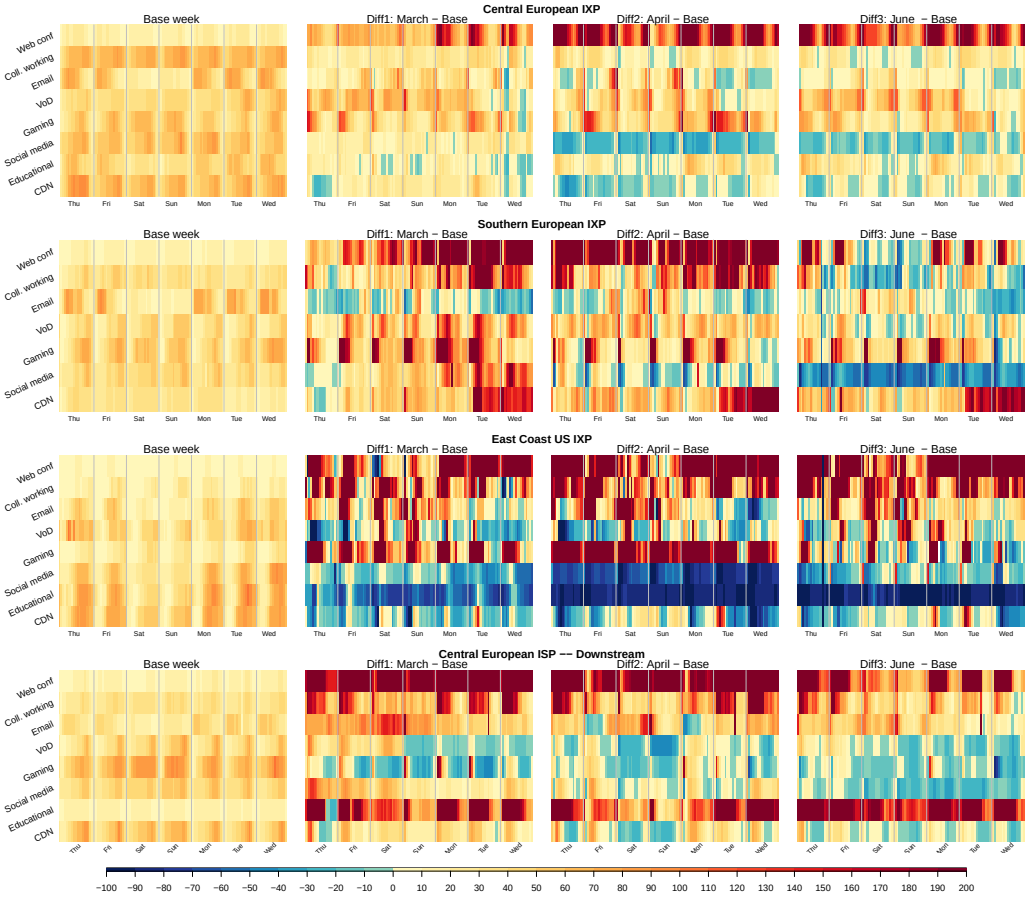

Figure 3.9: Heatmaps of application class volume for three different IXP locations and the ISP-CE.

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.

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. 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 they become more demanded during the lockdown. *Social media* shows a strong initial increase which flattens over time. Together, they demonstrate the massive impact that the drastic change in human behavior caused by the COVID-19 pandemic had on application usage.

3.1.4 VPN Traffic Shift

As a response to the pandemic, many institutions asked their employees to work from home. A typical way to access *internal* company infrastructure from home is by using VPN services. As a result, we expect VPN traffic to increase after the lockdown.

Port-based classification. We apply a twofold approach to identify VPN traffic. First, we classify traffic as VPN traffic if the well-known transport ports and protocols are used exclusively by a VPN service. We only focus on major VPN protocols and identify IPsec (port 500, 4500), OpenVPN (1194), L2TP (1701), and PPTP (1723)—both on TCP and UDP.

Identifying VPN usage on TCP/443. Since there are, however, many VPN services using TCP/443 to tunnel VPN traffic, a pure port-based identification approach cannot distinguish this traffic from HTTPS. To *limit* the potential for misclassification, we employ a second approach using DNS data to identify IPs labeled as `*vpn*` but not as `www.` in the DNS. That is, we identify potential VPN domains by searching for `*vpn*` in any domain label *left* of the public suffix [182] (e.g., `companyvpn3.example.com`) in (1) 2.7B domains from TLS certificates that appeared in CT Logs during 2015–2020 and (2) 1.9B domains from Rapid7 Forward DNS queries of reverse DNS, zonefiles, TLS certificates from the end of March 2020, and (3) 8M domains found in the Cisco Umbrella toplist in 2020. We resolve all matching domains to 3M candidate IP addresses. In order to get a conservative estimate of VPN traffic over TCP/443, we then also resolve the domains from the same public suffix prepended with `www` (e.g., `www.example.com`). If the returned addresses of the `*vpn*` domain and the `www` domain match, we eliminate them from our candidates. This approach *limits* misclassifying Web traffic destined to the `www` domain as VPN traffic to the `*vpn*` domain, if they share the same IP address. After removing shared IP addresses, we end up with 1.7M candidate VPN IP addresses. We classify TCP/443 traffic to these VPN addresses as VPN traffic.

VPN traffic on the rise. In Figure 3.10 we report our findings using the port-based and domain-based VPN traffic identification approach. We use four weeks of flow data from the IXP in Central Europe and aggregate them into workdays and weekends. Interestingly, we see almost no change in port-based VPN traffic before and after the lockdown. When looking at the VPN traffic identified with the domain-based technique,

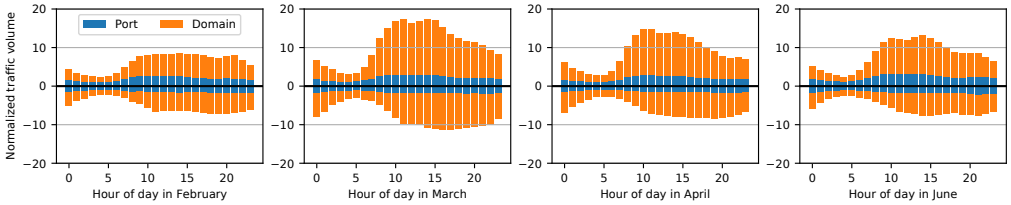

Figure 3.10: VPN traffic at the IXP-CE: normalized aggregated traffic volume per hour at the IXP-CE vantage point for four selected weeks. Aggregated workdays are shown as positive values, aggregated weekends as negative values. VPN servers are identified by ports and `*vpn*` label in the domain name.

we see a significant increase in VPN traffic. During workdays at working hours, VPN traffic increases by more than 200% in March compared to the base week in February. The increase on weekends is not as pronounced as during workdays, further indicating that these traffic shifts occur due to changes in user behavior (i.e., people working from home). When looking at the week in April, we still see a gain in VPN traffic compared to February, although not as large as in March. In June, VPN traffic decreases further compared to previous months, although its traffic volume on workdays remains well above the levels observed for the base week of February. This is likely due to the gradual lifting of lockdown restrictions in Central Europe and the beginning of the summer holiday season, resulting in fewer people working from home in June compared to March.

In conclusion, we see a clear pattern of VPN traffic increase during working hours due to lockdown restrictions. Moreover, as the visible increase of VPN traffic was limited to TCP/443 on `*vpn*` domains, we argue that VPN identification solely on a transport port basis vastly undercounts actual VPN traffic. To mitigate this problem, we propose to identify seemingly HTTPS flows as VPN traffic using domain data. This allows for a more accurate picture of the VPN landscape.

3.2 A Year in Lockdown

After presenting the results from our initial study we next expand our study by analyzing more than two years of Internet traffic data including the first year of the pandemic. We are interested to see if there is a "new normal" in Internet traffic and whether it is here to stay. We again characterize the overall traffic shifts and the changes in demand for particular applications that became very popular in a short amount of time. We extend our previous study for the fall[4] 2020 wave (September 2020 to February 2021). To that end, we collect and analyze network traffic data from the same vantage points. Recall, that this includes a large Internet Service Provider (ISP) in Europe, three Internet Exchange Points (IXPs) in Europe and the US, as well as a mobile operator and a metropolitan academic network in Europe (REDIMadrid).

[4]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.

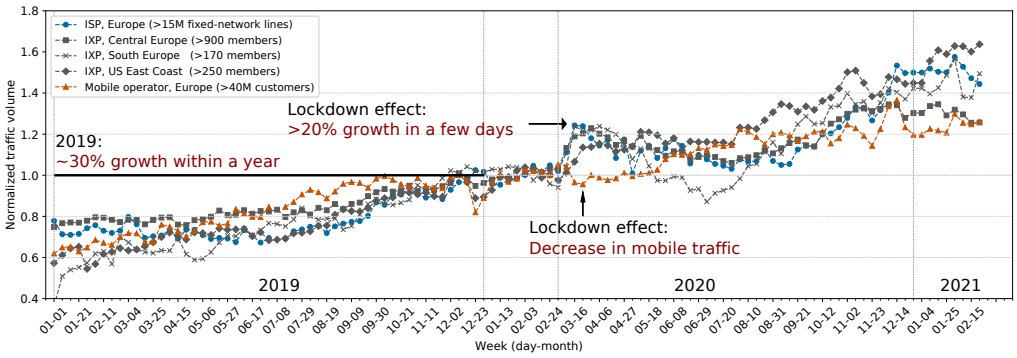

Figure 3.11: Traffic changes during the COVID-19 pandemic's spring and fall waves at our Internet vantage points.

3.2.1 Network Traffic Shifts

To understand traffic changes over a year of the COVID-19 pandemic, we first look for overall changes in well-established traffic patterns before, during, and after the strictest lockdown periods for both the spring and fall 2020 waves. Since all data sources exhibit very different traffic characteristics and volumes, we normalize the data to make it comparable. In Figure 3.11 we show the normalized aggregated traffic of the ISP, the IXPs, and mobile operator vantage points from January 2019 until the end of February 2021. We normalize this by using the traffic of the first week of January 2020 for each corresponding vantage point.

3.2.1.1 Macroscopic Observations

In Figure 3.11, we annotate the week of the initial lockdowns in the countries that host the ISP and IXPs in Central and Southern Europe, and the mobile operator. Although the exact dates of when the lockdown was imposed differs across Europe, these dates are very close to each other and followed the declaration of COVID-19 as a pandemic by the WHO. A first observation is that the ISP and IXPs in Central and Southern Europe 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 traffic 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. Well-provisioned networks, like the ones we measured for our study, could cope with this surge. However, networks that "ran hot" may have faced problems as this increase is significant and takes place in a relatively short period of time. In sharp contrast, the traffic of the European mobile operator decreased as users switched to WiFi and reduced their commute and traveling. This aligns with reports in other studies [154]. It is worth noting, that, during the same period, the traffic at the IXP at the US East Coast surged only by 2% as there was no announcement of strict lockdowns in the US at that time.

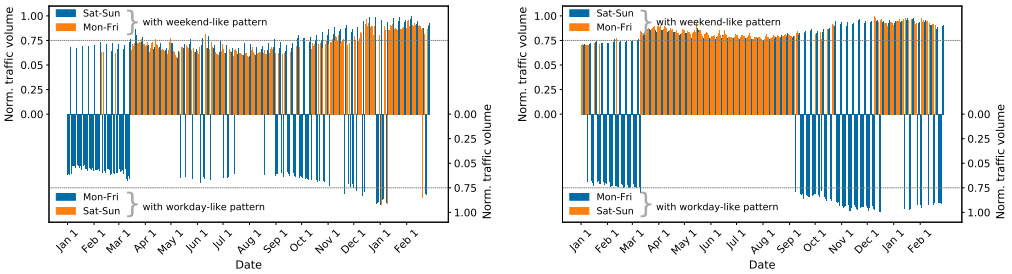

(a) ISP-CE: Weekend-like (top) vs. workday-like (bottom).

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

Figure 3.12: Drastic shifts in Internet usage patterns during the COVID-19 pandemic. Classification of weekends- and workdays-like pattern.

The impact of the spring wave was significant as the traffic levels remain at same elevated level during the lockdown. The traffic of the IXP at the US East Coast increased significantly when lockdown measures took place by the state authorities. However, when the economy opened again after June 2020, we observe a slight decrease of ISP and IXPs traffic as well as an increase of the mobile operator's traffic.

The impact of the fall wave is clearly visible in traffic patterns beginning in September 2020. The traffic at the ISP and all IXPs surged again, while the traffic of the mobile operator declined, except for the holiday period at the end of 2020. Although the lockdowns in the fall of 2020 differ significantly from country to country, and in some cases there were lockdowns with on-off periods, the impact of the fall wave was significant. The 2020 annual increase for the ISP and the IXPs varied between 35%–50%, i.e., higher than the expected annual increase. The mobile operator showed an annual increase of around 20%, which is lower than expected. As the fall wave of COVID-19 continues into 2021, we observe similar trends until February 2021. It is also worth noting that in some countries, the fall wave was a superposition of multiple waves of COVID-19 and its mutations, which were faced with harder lockdown restrictions. Additionally, the severe weather conditions in Southern Europe with historic snow volumes in January and February 2021 may have also played an additional role in keeping people at home and the corresponding increase of Internet traffic at the ISP and IXPs.

The traffic increases we have seen across vantage points can arise unexpectedly and may create a need for capacity increases by network operators. We observed capacity increases in the order of 1,500 Gbps (3%) across many IXP members at the IXP-CE alone. Beyond our data sets, some networks publicly reported that traffic shifts due to the pandemic resulted in partial connectivity issues and required new interconnections [72, 214]. The vantage points in this chapter range from extremely large to moderate sizes with sufficient resources and a lot of experience in network provisioning and resilience. In general, smaller networks with limited resources may not be able to plan with sufficient spare capacities and fast enough reaction times to compensate for such sudden changes in demand. In fact, performance degradation issues have been reported in less developed regions [22], that also highlights the digital divide.

3.2.1.2 Drastic Shifts in Internet Usage Patterns

Beyond the macroscopic observations, our analysis sheds light on the shifts in Internet usage patterns that are also relevant to network operation and management. The Internet's regular workday traffic patterns are significantly different from weekend patterns [146]. On workdays, traffic peaks are concentrated in the evenings, typically between 18:00 and midnight, also referred to as "peak hours". During the weekend, the activity is more distributed also 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 main activity gains significant momentum from approximately 9:00 to 10:00 am. Figures 3.12(a) and 3.12(b) 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 in 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 3.12(b). This pattern changes drastically once the lockdown measures are implemented. Indeed, almost all days are classified as weekend-like. This change persists in Figure 3.12(b) until the end of August due to the vacation period, which is consistent with the behavior observed in 2019 (not shown). In contrast, Figure 3.12(a) 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.

During the period of August to December 2020 the patterns both at the ISP and the IXP are 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.

3.2.2 Application Traffic Shifts

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. We refer to 3.9for technical details on how we classified traffic in any of these categories.

3.2.2.1 Application Classes' Traffic Shift

In Figure 3.13, 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.

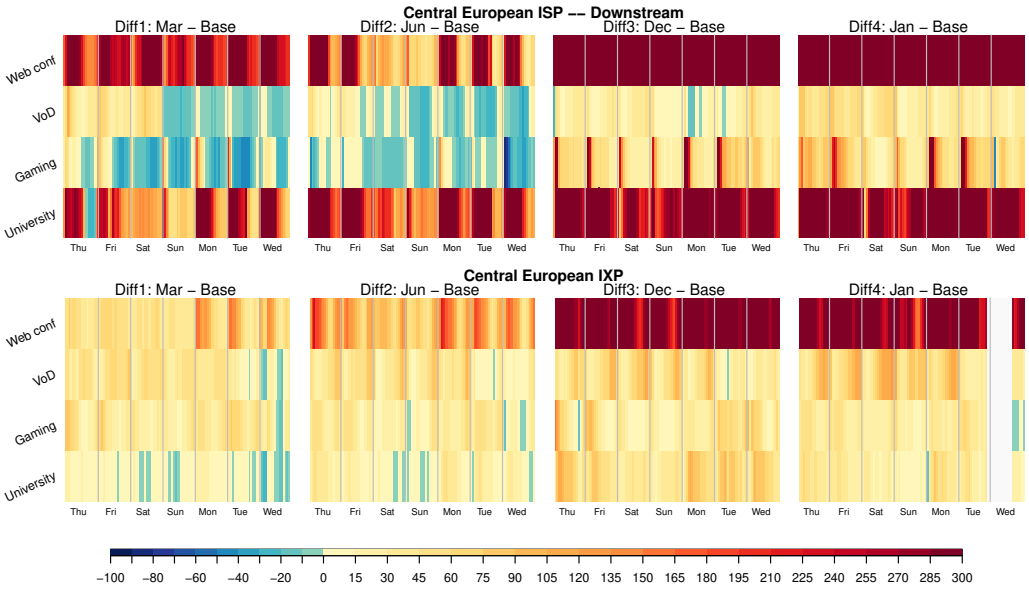

Figure 3.13: ISP (top), IXP-CE (bottom) heatmaps of application classes' traffic at the ISP and IXPs during COVID-19 pandemic: spring and fall waves. Each subplot shows the change in the aggregated traffic volume per hour for the respective class compared to the base week in February 2020. White areas mark missing data.

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 [185]. 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. The initial download of a game nowadays supersedes the amount of data transferred although 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

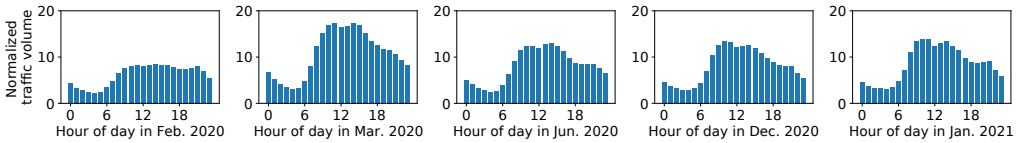

Figure 3.14: VPN traffic evolution during the COVID-19 pandemic.

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

3.2.2.2 VPN Traffic Shift

Working from home leads to a higher demand for Virtual Private Network (VPN) solutions as employees need to access firewall-protected resources hosted in internal company networks. We identify VPN traffic using a novel technique based on transport port data as well as DNS data as in 3.1.4. In Figure 3.14 we show the changes in VPN traffic during the spring and fall COVID-19 waves in 2020. We use five weeks of data from the IXP in Central Europe from February 2020 to January 2021, each in different months, to highlight the differences. February 2020 serves as a baseline, i.e., to show the state of VPN traffic before COVID-19 restrictions were enforced.

In March 2020—after the first lockdown restrictions were authorized in Europe—we notice a large increase in VPN traffic during working hours. This growth partially recedes in June 2020 as lockdown restrictions are relaxed again and employees could return to their workplaces. Nevertheless, VPN traffic volume is still over the February 2020 baseline levels. In the Fall wave, VPN traffic increases again but not as high as in March 2020.

We also investigate the *share* of VPN traffic among the total volume across waves. We find that the VPN traffic share remains stable from February to March 2020. This suggests that overall traffic volumes increase, regardless of the application. However, in June 2020, VPN traffic share increases whereas overall traffic volume decreases. Towards December 2020 and January 2021 we see a slight decrease in the VPN traffic share as the overall traffic gains traction again after the summer holiday.

In summary, we see that VPN traffic increases during working hours since the first lockdown measures were implemented. The increase is higher in the spring wave than in the fall one. This finding aligns with reports indicating that more people in Central Europe were working from home in the Spring wave compared to the fall wave [246].

3.3 Ethical Considerations

Both data sources used in this chapter, i.e., NetFlow and 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, IXPs, and the academic network. IP addresses are hashed to prevent information leaks and raw data being transferred. The output of the analyses are the aggregated statistics as presented in the chapter. The data at the ISP and IXPs is collected as a part of their routine network analysis.

3.4 Related Work

Our study provides a testimonial of the impact of an unprecedented medical crisis in recent human history on the operation of the Internet. Previous studies followed a similar approach to ours, i.e., collect measurements at different vantage points, to understand the impact of other events on the Internet. Partridge et al. collected and analyzed routing and protocol data during and after 9/11 to understand the resilience of the Internet under stress [24]. Their findings showed that, overall, the Internet operation was robust: Although unexpected outages did happen, they only had a local impact. Notice, however, that the penetration and importance of the Internet in our life has significantly increased in the last twenty years, and the global nature of the COVID-19 pandemic crisis makes this case unique. Other studies focus on physical phenomena, e.g., earthquakes [36] or severe weather conditions [192, 81], and power outages [11, 17] to understand the Internet behavior and the change on Internet user activity. Beyond physical phenomena, also human-triggered events such as major update roll-outs can cause substantial traffic shifts [15].

The study of the impact of the COVID-19 pandemic to the performance and traffic of the Internet has attracted significant attention in the form of blogs posts [160, 161, 51, 230, 109] and more recently in presentations at network operator conferences [143]. By the time of our submission, a limited number of research studies have been already published. Favale et al. report and analyze the impact of the remote learning activity by 16k students on the Politecnico di Torino campus network due to the lockdown enforcement [89]. The university utilized an in-house online teaching solution. Thus, although the impact of remote learning on the campus network shares similarities with our analysis of the academic and research network in our study, there are also significant differences. Another study [251] analyzed Wi-Fi network data collected at university campuses in Singapore and the US during the pandemic. Their results show that the activity on campuses decreases, but mobility did not. In our study, we found that the mobility patterns reduced drastically in Europe, most likely due to the stricter measures and complete lockdowns. A study of the access patterns of Wikipedia shows that during the pandemic Web visitors had an increased interest in topics such as health [205]. Parallel to our work, researchers evaluated (1) the impact of the pandemic on traffic of a UK mobile network operator reflecting changes in users' mobility [154], (2) changes in traffic demand at a major social network [22], (3) transactions volumes at an underground market during the pandemic concluding that the observed higher transaction volumes are a market stimulus rather than an effect of the pandemic [240], and (4) the impact of the pandemic on Internet latency in various European countries, finding an increase in the variance of additional latency and packet loss [28].

3.5 Discussion

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" [124]. At the beginning of the pandemic, 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 [225]. In fact, the pandemic increased the demand for applications supporting remote teaching and working to guarantee social distancing as shown in our analysis across all vantage points. The Internet could handle this new load due to the flexibility and elasticity that cloud services offer, and the increasing connectivity of cloud providers [145, 142, 35, 215, 249]. Our results confirm that most of the applications with the highest absolute and relative increases are cloud-based. Moreover, the adoption of best practices on designing, operating, and provisioning networks 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 [70]. In summary, our study demonstrates that over-provisioning, 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.

Taming the traffic increase. In this chapter, we report an increase in traffic in the order of 15-20% within days after the lockdown began. This is in line with reports of ISPs and CDNs [143, 51, 160, 161] as well as IXPs [213]. Typically, ISPs and CDNs are prepared for a traffic increase of 30% in a single year period [142, 147, 40]. While these are yearly plannings, the pandemic 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. We observed port capacity increases in the order of 1,500 Gbps (3%) across many IXP members at the IXP-CE alone (see Section 3.1.2.1). Beyond our data sets, some networks publicly reported that traffic shifts due to the pandemic resulted in partial connectivity issues and required new interconnections [72, 214]. When we turn our attention to traffic peaks, we notice that the increase is even smaller. Traffic engineering focuses on peak traffic increase as this requires more network resources. The effect of the pandemic fills the valleys during the working hours and has a moderate increase in the peak traffic, which can be handled by well-provisioned networks that are prepared for sudden surges of peak traffic by 30% or more, due to attacks, flash-crowds, and link failures that shift traffic to other links. One concern that network operators raised in March brought awareness to network instabilities that might occur due to traffic shifts [225]. While on the one hand we find no evidence that the traffic shifts due to the pandemic impact network operation of our vantage points, individual links experience drastic increases in traffic—way beyond the overall 15-20%. Such increases arise unexpectedly to some network operators and may create a need for port upgrades. On the other hand, the vantage points in this chapter range from extremely large to moderate sizes with sufficient resources and a lot of experience in network provisioning and resilience. In general, smaller networks with limited resources may not be able to plan with sufficient spare capacities and fast enough reaction times to compensate for such sudden changes in demand.

Substantial shift in traffic pattern. From a network operator perspective, coping with the pandemic has required some port capacity upgrades but otherwise does not appear to impact operation. The ability of network operators to quickly add capacity when needed highlights that the Internet infrastructure works well at large, despite some challenges to access data centers imposed by the lockdown. From the perspective of the traffic mix, the pandemic, however, results in substantial changes in traffic, ranging from shifted diurnal pattern to traffic composition. This represents a remarkable shift in Internet traffic that is, based on our observations, handled surprisingly well by the Internet core at large supposedly because many operators are prepared and can react quickly to new demands. While the pandemic represents a rather extreme and exceptional case, one may argue that with the growing intertwining of the Internet and our modern society such events can occur more often. In any case, the COVID-19 pandemic highlights that user behavior can change quickly and network operators need to be prepared for sudden demand changes.

3.6 Summary

In this chapter, we investigated how the Internet responds to a stress situation with analog causes. The COVID-19 pandemic is a—hopefully once in a lifetime—event that drastically changed working and social habits for billions of people. 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. In this chapter, we analyze 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 were enforced, the aggregated traffic volume increased by around 40%, well above the typical expected annual growth. Additionally, workday traffic patterns have rapidly changed and the relative difference to weekend patterns has almost disappeared during lockdowns. Applications for remote working and education, including VPN and video conferencing, experienced traffic increases beyond 200%.

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. Our study demonstrates that over-provisioning, proactive network management, and automation is key to providing resilient networks that can sustain drastic and unexpected shifts in demand such as those experienced during the COVID-19 pandemic.

Chapter 4
Bringing Flexibility to the Core of the Internet

To address analog stress situations in the Internet, we need a large degree of flexibility in the core of the Internet, e.g., to add resources when necessary quickly. However, most of today's infrastructure does not offer such flexibility. This is in particular problematic for IXPs which are a crucial part of today's Internet infrastructure. Large IXPs connect thousands of ASes and facilitate the exchange of more than 10 Tbps of traffic during peak hours. Still, their specific technical requirements (e.g., large Layer-2 domains, complex traffic filtering) are not well addressed by today's networking hardware, as vendors optimize for the ISP market due to revenues that are orders of magnitude higher. IXP operators have to find workarounds that are inflexible and difficult to operate.

Software Defined Internet eXchanges (SDXes) are a promising solution since they enable tailored hardware and software stacks to satisfy the specific IXP requirements. They combine a high degree of automation with the flexibility to implement value-added services and, thus, may reduce IXPs' costs. Since previous work is based on the OpenFlow standard, which was last updated in 2017, we revisit the idea by leveraging the flexibility of P4 networking hardware.

In this chapter, we present the P4IX, a technical concept for a generic P4 packet processing pipeline for IXPs. The P4IX concept is built upon a comprehensive requirements analysis: we characterize the IXP landscape and provide first-hand insights of a large IXP operator (more than 1000 well-distributed ports). Moreover, we use our insights to critically discuss the P4IX from an operational, technical, and organizational perspective.

4.1 Conceptualizing a P4-Enabled Internet Exchange Point

Internet eXchange Points (IXPs) are traffic hubs between Autonomous Systems (ASes) and facilitate the settlement free exchange of traffic over a Layer-2 platform (peering) [33]. The largest IXPs ensure low-latency interconnection for hundreds or even thousand ASes

and exchange 10 Tbit/s or more during peak times[5]. The Software Defined Internet eXchange (SDX) concept, first introduced by Gupta et al. [114], shows how Software Defined Networking (SDN) can benefit Internet Exchange Points (IXPs) [1]. The basic idea of an SDX is to tailor a soft-/hardware stack to the requirements of IXPs by relying on SDN building blocks. Such specific requirements include the need to realize one large *Layer-2* domain with the inherent broadcast problems, e.g., for Address Resolution Protocol (ARP), as well as sophisticated inbound and outbound traffic filtering, while providing the reliability expected from critical infrastructures.

The body of SDX works [23, 113, 112, 114, 137, 157, 34, 181] relies on the OpenFlow paradigm [162]. In 2017, the OpenFlow standard was augmented by the Open Networking Foundation with the increased capabilities and flexibilities of a P4-enabled stack [18, 19]. P4 is a domain-specific language which defines how packets are processed by the data plane, i.e., switches or routers. The language allows the definition of custom packet header parsing and assembly as well as match/action pipelines to perform non-trivial operations on packets in line rate. To the best of our knowledge, there has been no work on a holistic SDX concept that takes advantage of the P4 capabilities[6]. Thus, we revisit the question of how to realize a P4IX. Our motivation is two-fold: (1) we have first-hand experience from operating a very large distributed IXP (more than 1,000 ports across many data centers), which allows us to precisely scope P4IX requirements and (2) we find that OpenFlow's limitations have lead to a number of non-optimal design choices.

Firstly, we review some of the limitations of using OpenFlow: (1) previous solutions enabled multi-hop IXPs by using MAC headers for encoding routing information (VMAC concept) [23, 114, 8]—this implies a loss of compatibility to the existing Layer-2 switching paradigms and complicates debugging. Rather, this should be realized in the data plane and an external controller should only be required when the set of IXP members changes or additional hardware is added or removed. Indeed, relying on an external control for IXPs is complex as they are critical infrastructure and, thus, require a fail-safe complex controller setup (e.g., Martins et al. [157] uses a distributed ONOS controller). (2) Mechanisms implementing fast rerouting on link failure are notoriously hard to solve using OpenFlow only, since they either require large state space on the switches or involve the OpenFlow controller [8, 23, 114] for rerouting, which introduces latency and packet loss. Katta et al. [135, 134], in a data center context, show how to tackle this problem using P4 in the data plane, a solution we incorporate in our concept. Additionally, our approach is transparent and does not require changes in the member networks, in contrast to member operated SDN controllers [181].

We use our experience from operating a large distributed IXP to outline how to take advantage of P4-enabled hardware for a P4IX. We conceptualize a P4IX in the context of realistic technical, operational, and business requirements of IXPs. It provides a solid basis for adding more sophisticated services such as (1) BGP policy enforcement in the data plane [112], (2) data plane integrated DDoS mitigation [62, 183], and (3) other value-added services as outlined in [34]. More precisely, we make the following contribution:

- A characterization of the IXP landscape/market and implications for P4IXes and an overview of operational and technical peculiarities of IXPs.

[5] E.g., AMS-IX (`https://stats.ams-ix.net/index.html`), DE-CIX (`https://de-cix.net/en/locations/frankfurt/statistics`), and LINX (`https://portal.linx.net/okta-login`).

[6] Silva et al. [60, 61] focus coping with elephant flows in IXP networks.

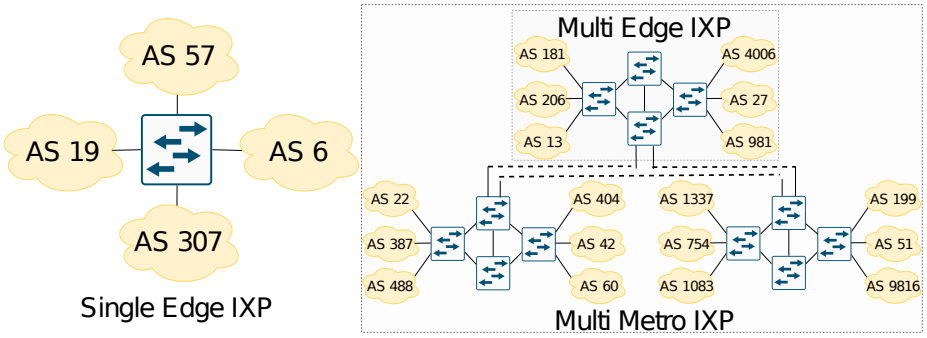

Figure 4.1: Different IXP setups ranging from a single edge IXP to an interconnected multiple metropolitan area IXP. The setup implements a Layer-2 domain even for multi metro IXPs.

- An outline of technical and operational P4IX requirements and a P4IX pipeline concept.
- A critical discussion of operational, technical, and organizational P4IX advantages and disadvantages.

4.2 IXP Landscape Characterization

In this section we touch upon IXP market specific properties and their implications for a P4IX. We start by looking at the IXP landscape using data from PeeringDB[7] —a central database for IXPs and peerings in general. Figure 4.2(a) plots a CDF of the number of ASes that are members at each IXP. As such, the x-axis shows the log-scaled number of connected ASes per IXP and the y-axis the fraction of IXPs with that number of member ASes. Already this simple figure provides a number of relevant insights.

The network hardware market size of IXPs is negligible compared to the one of ISPs and Cloud providers. 40% of all IXPs have less than ten member ASes, 80% have less than 50 member ASes, see Figure 4.2(a). In addition, given that there are 920 IXPs in the PeeringDB that interconnect 42,303 ASes, we get a ratio of 46 ASes per IXP. Even if we assume that each IXP member AS operates only one router–the edge router visible to the IXP–this implies that the hardware footprint of IXPs is at least 46 times smaller compared to ISP/Cloud provider ASes that they interconnect. In practice, this gap is even orders of magnitude larger. Thus, since the IXP networking hardware market segment is small, hardware vendors tend to tailor solutions for ISPs instead of IXPs. However, IXPs' network architectures differ substantially from ISPs'. While ISPs predominantly operate Layer-3 architectures, IXPs operate large Layer-2 architectures.

IXP scalability requirements. On the one hand, the majority of IXPs can at least theoretically be realized on a small hardware platform, i.e., a single 32 port switch, if we ignore redundancy requirements, and a 64 port switch if we keep a redundant port for each member (single edge IXPs, see Figure 4.1). On the other hand, large and very large IXPs have a share of 15% of the overall IXP market. Their setups need to provide connectivity for hundreds to thousands of ports. Moreover, they are often distributed across multiple data centers in a geographical region (multi edge IXP, see Figure 4.1).

[7] `https://www.peeringdb.com`, last accessed 11/24/2021.

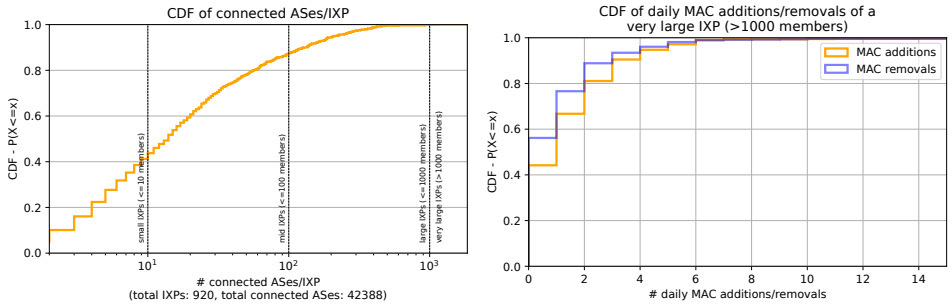

(a) CDF of connected ASes per IXP, data from peeringDB.com, 11/24/2021.

(b) Daily MAC additions/removals in the peering LAN of a very large IXP from 07/2017-11/2021.

Figure 4.2: IXP landscape

Very large IXPs, e.g., LINX, AMS-IX, and DE-CIX, operate separate IXP locations around the globe which are interconnected by a global backbone network (multi metro IXP, see Figure 4.1).

As such P4IX platforms have to be scalable from very small setups to very large ones. Moreover, they need to support low-end hardware as well as high-end hardware. Ideally, a P4IX platform supports scale out: if an IXP requires more ports, additional switches can be added, thus eliminating the need to replace existing hardware.

IXP business challenges. Acquiring additional member ASes is becoming increasingly difficult for IXPs. The pool of potential non-peering ASes is shrinking as peering has become a common practice over the last decade. Moreover, transit traffic prices have been declining for over a decade [8]. This, in turn, applies pressure on peering prices, as both are to some extent competing products for the same traffic. In the foreseeable future, IXPs have the following options to generate growth: (1) geographical expansion/internationalization, i.e., establishing new IXP locations; (2) expanding to new customer segments beyond ISPs, CDNs, and Cloud providers; (3) providing additional services on top of peering to increase members' added value and, thus, the IXP's revenue. The geographical expansion strategy underlines the P4IX scalability requirements. Addressing new customer segments and implementing value-added services require (1) a high degree of automation and (2) the ability to realize proprietary services on top of the P4IX platform.

IXP operational challenges. The large asymmetry in member ASes among IXPs implies a large asymmetry in revenue and, thus, in resources available for research and development of a P4IX architecture. Even very large IXPs are no multi-billion dollar enterprises. They often operate under non-profit membership governance models. It is unlikely that any single IXP can amortize the cost for developing its own P4IX platform. Consequently, an open source community effort is needed. Still, at the same time there is a need for service differentiation. Thus, the base open source P4IX platform has to be able to support custom P4IX services on top to foster competition and innovation.

IXP networks evolve slowly. While the IXP community has introduced a common API for members to adjust configurations at the IXPs[9], some human interactions between

[8] `https://drpeering.net/FAQ/What-are-the-historical-transit-pricing-trends.php`, data up to 2015, last visited 29/11/2021; the industry is not publishing transit/peering prices, so the absolute values have to be taken with a grain of salt. However, the trend is clear.

[9] Standardized IX-API (`https://ix-api.net/`)

the IXP operator and the member ASes are still required. Consequently, IXP networks evolve slowly and constitute a mostly static network environment. To underline this slow rate of change Figure 4.2(b) shows a CDF of the daily MAC address changes at a very large IXP over the last 4 years. On about 50% of days, no MAC addresses were added or removed; on 99% of days fewer than 8 MACs were changed. Since any MAC change results in a routing change, this corresponds to the number of required routing changes due to members joining/leaving the IXP. A P4IX architecture can take advantage of this mostly static setup to simplify operations, by, e.g., replacing difficult-to-administer routing protocols at IXPs with multiple switches with offline optimized static routes.

4.3 Challenges When Realizing Large IXPs Using ISP Hardware

Before discussing how to build a P4IX, we next outline the technical challenges of realizing very large IXPs with traditional ISP-focused hardware.

Large IXPs are victims of feature bloat. Traffic rates at large IXPs easily reach an average of multiple terabits per second. At the same time, IXPs are typically built using a comparably small number of hardware units (if compared to ISPs). Thus, they require compact hardware to reach the required port density at a reasonable energy consumption. These requirements make off-the-shelf hardware, in particular data center switches, inadequate. Operators of large IXPs are thus left to build their platforms using high-end (expensive) service routers tailored for ISP requirements. These routers offer a massive feature set, e.g., they can be used as Broadband Remote Access Servers (BRAS) that terminate a large number of end customer subscriber lines, while offering services such as admission control, traffic shaping, and integration into billing infrastructures. However, IXPs only need a very small subset of these features, i.e., mainly Layer-2 functionality such as switching and VLANs, but have to pay the surplus for all implemented features.

Unsuitable Feature Set Realization. Many feature designs and implementations are targeted for the ISP market rather than the IXP one due to its relatively small size. This starts with the standardization phase where the standardization committees are driven by ISPs and hardware vendors, as IXPs often lack resources to participate in all processes. One recent example is BGP FlowSpec, a method to communicate traffic filters via BGP from, e.g., an ISP customer to an ISP. It is "primarily designed to allow an AS to perform inbound filtering in their ingress routers of traffic that a given downstream AS wishes to drop." [115]. As such, hardware vendors implement FlowSpec filtering on the ingress path. In an IXP setting, this is unfeasible as IXPs handle peering traffic and tend to filter on egress [74] to protect member ports. Other examples include: statistics export limited to Layer-3/4, IXPs need Layer-2 information additionally; limited ingress/egress filtering on Layer-2 or Layer-3/4 exclusively, IXPs need Layer-2 and Layer-3/4 filtering to ensure platform stability and to implement value-added services.

Challenges of large Layer-2 domains. IXPs are realized via large Layer-2 domains. This comes with a number of challenges unknown to ISPs. These include broadcasting based protocols, like ARP and IPv6 Neighbor Discovery Protocol (NDP). Due to the large number of peers, large volumes of broadcast traffic can accumulate. This can overwhelm member routers, e.g., with information on MAC to IP mappings of routers from ASes that they do not peer with. Thus, Broadcast, Unknown Unicast, and Multicast (BUM) traffic should be dropped or at least rate limited in peering LANs.

Another challenge are Layer-2 loops—if two IXP ports are short-circuited in a way allowing Ethernet frames to be forwarded back into the IXP platform. Frames may be forwarded in an endless loop, which leads to overwhelmed backplanes in the edge switches and thus affects neighboring member ports on the same switch. Since Ethernet frames do not carry a time to live counter, an efficient detection is difficult. Therefore, it is necessary to filter which MAC address can send frames from which port and to block other traffic.

When an IXP grows to a multi-site setup, it needs routing—a Layer-3 feature—while maintaining the facade of a single Layer-2 domain. Often routing is realized via OSPF or IS-IS over MPLS. To maintain the Layer-2 domain and enable traffic engineering over the now present Layer-3, virtualization features are used, including VPLS, EVPN, and ECMP. This often goes along with using a vendor proprietary network management system. The setup introduces unnecessary cost and complexity—full routing is realized even though the setup is more or less static and all paths (including redundant ones) are known upfront, members join or leave only occasionally and their routers have well-known static MAC and IP addresses behind their dedicated port.

4.4 P4IX Requirements

Next, we outline requirements for a P4IX based on our experience operating a very large IXP. We divide them into operational and technical ones and assign them an identifier, i.e., technical requirement "TR1" and operational requirement "OC1". We later refer to these identifiers in the P4 pipeline concept (see Figure 4.4).

Among the overarching requirements are: (1) avoiding unnecessary complexity and implementing the bare minimum functionality for operating IXPs ranging from small single-edge ones to large multi-metro ones and (2) offering enough flexibility to implement value-added services. Today's services include dropping of unwanted traffic (blackholing [74]) and isolation via virtualization, e.g., peering with Cloud providers via dedicated peering LANs. Regarding possible services enabled by P4IX, we point to the work of Chiesa et al. [34]. These may include DDoS mitigation, enforcement of BGP policies (e.g., IRR/RPKI checks) in the data plane, and application specific peerings (e.g., video). All these services rely on the following building blocks: virtualization, admission control, traffic classification and filtering. Thus, P4IXes need to support these for a wide variety of packet header fields (Layer-2 to Layer-4).

P4IX Technical requirements. Firstly, Layer-2 traffic engineering capabilities similar to a WAN with MPLS or ECMP are required (TR1) for both single-metro and multi-metro setups (similar to [135, 134, 167]). However, we note that routes are almost static. Supporting multiple virtual peering LANs requires virtualization functionalities such as VLANs (TR2). Fine grained inspection of packet header fields is required to classify traffic, e.g., to avoid broadcast traffic, handle Layer-2 loops and BUM traffic (TR3). Note, TR2-3 are also needed to realize value-added services. ARP and NDP remain the protocols to ensure Layer-2 (Ethernet) connectivity. Thus, an IXP specific ARP/NDP handling mechanism which avoids broadcast storms is required (TR4). To enable members to purchase port capacities independent of their physical access bandwidths, rate limiting and Link Aggregation Groups (LAGs) are required (TR5). To enable monitoring and to generate statistical summaries, packet sampling is required (TR6). Lastly, no data plane software update should interrupt forwarding of traffic (TR7).

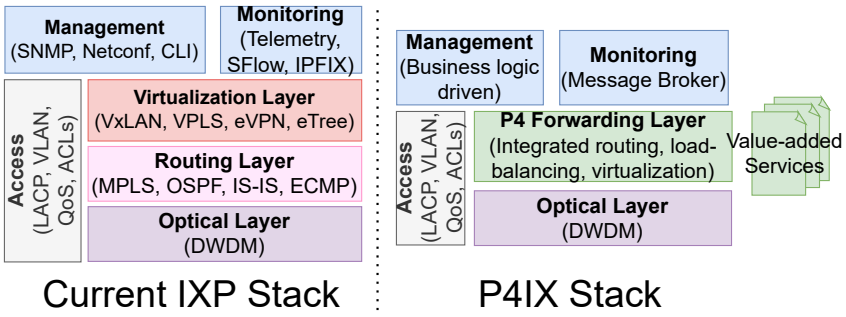

Figure 4.3: Traditional IXP setup versus P4IX concept.

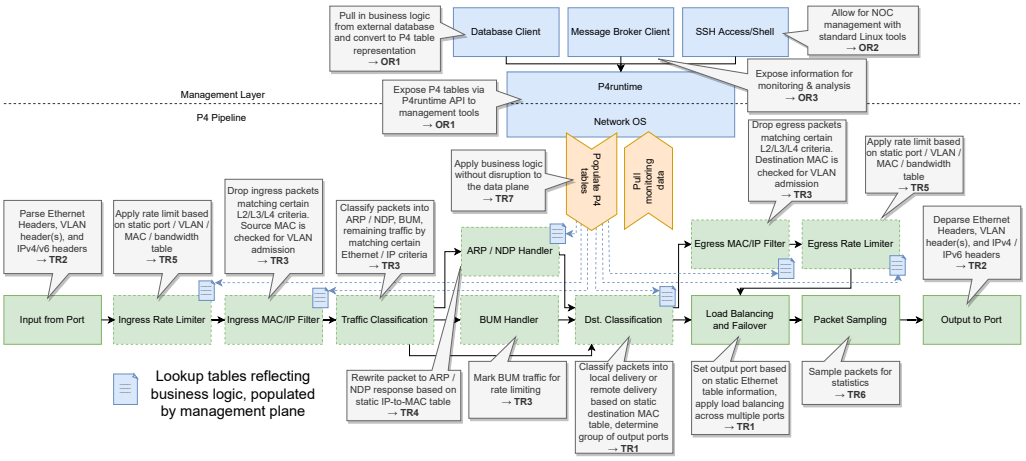

Figure 4.4: The P4 IXP model. The blue boxes indicate management-related solutions, green boxes show the parts of the P4 forwarding layer. Dashed boxes are not required for forwarding in the core, all boxes are required for edge routers.

P4IX Operational requirements. Here, we see a trade-off between degree of automation vs. capability of manual intervention for Network Operation Center (NOC) teams. The former is desirable to take full advantage of reducing operational cost. The latter may be required to troubleshoot problems. Thus, a P4IX concept requires a tight integration with IXP business logic, e.g., when a port is sold by the IXP provider, it should be automatically provisioned (OR1), but at the same time, there should be tooling to manually override automatic processes to be able to fix operational problems (OR2). For troubleshooting, monitoring information should be exposed to external tools (OR3).

4.5 The P4IX: Technical Concept

This section presents a technical concept for a P4IX. The main idea, see Figure 4.3, is to merge the virtualization and routing layer into a single P4 forwarding layer that relies on static routing. This reduces complexity, i.e., by removing all routing protocols. Furthermore, the P4IX concept does not introduce any changes to the optical underlay network, nor the platform's access technologies. Next, we discuss our concept in a bottom-up fashion.

Stage	Key	Value
Ingress & Egress Rate Limiter	Physical Port, MAC, VLAN	Bandwidth
Ingress & Egress MAC/IP Filter	Layer-2, Layer-3, Layer-4 Headers	Drop or pass
ARP / NDP Handler	IPv4 / IPv6-Address	MAC-Address
Dst. Classification	MAC-Address	Port Group

Table 4.1: Lookup table contents of the P4 pipeline stages.

4.5.1 P4 Forwarding Layer

Figure 4.4 outlines the P4 forwarding layer (bottom) together with its management layer (top) required for operation. The main interface

They interact via P4 lookup tables—the main interface for configuring the P4 pipeline. These are filled by the management layer and, then, used by the P4 pipeline at runtime.

We show and discuss the P4 pipeline from left to right. Packets enter the pipeline through the ingress port (*Input from port*) and leave through the output port (*Output to port*). When a packet enters the P4 pipeline, its headers have to be parsed (*Input from port*). A P4 parser graph is used to parse Ehternet, VLAN, and IPv4/v6 headers for further processing. Packets with invalid headers are detected by the parser and dropped immediately. This stage also separates header and payload, where the header is passed and transformed along the pipeline, while the payload stays until both are merged to be emitted via the output port (*Output to port*).

The stage *Ingress Rate Limiter* is responsible for rate limiting member traffic on ingress to the IXP's infrastructure based on a table that contains a mapping of ingress port, MAC, and VLAN tag to purchased capacity.

Next, the packet is matched against ingress filters to drop or rate limit traffic (*Ingress MAC/IP Filter*). This stage enforces basic Layer-2 network security by, e.g., tying each port to a single router MAC source address. It handles VLAN admission by matching source and destination MAC addresses as well as the VLAN tag against a table with whitelisting rules provided by the management layer. Any non-matching packets are dropped or shaped as required. Moreover, this stage is used to realize value-added services such as member or application specific peering LANs.

Packets passing the *Ingress MAC/IP Filter* are subject to a *Traffic Classification* stage. Here, each packet is tagged with one of three classes: (1) ARP / NDP traffic, (2) Broadcast, unknown Unicast or Multicast traffic (BUM), or (3) other traffic. This classification is then used to pass the packets to their next stage: *ARP / NDP Handler*, *BUM Handler*, or directly to the *Dst. Classification*. The *ARP / NDP Handler* is an IXP specific implementation for address resolution protocols. ARP and NDP requests are immediately transformed into replies using a MAC to IP mapping table generated by the management layer. This is feasible due to the relative static nature of IXP networks. This prevents ARP and NDP requests from flooding the IXP internal network. The *BUM Handler* rate limits BUM traffic to avoid BUM packets overwhelming the platform.

The next step in the pipeline determines the packet's rough destination (*Dst. Classification*). Two scenarios are possible: (1) *local* delivery, i.e., the destination interface is the current edge switch or (2) *remote* delivery, i.e., the packet needs to be routed to another edge

switch to reach its destination. As all members' MAC addresses are known in advance, this can be realized via local lookup tables.

Next, packets marked for *local* delivery are sent to the *Egress MAC/IP Filter* stage. This stage matches headers of various layers against egress filtering rules. These rules can be used to reflect user-defined blackholing/DDoS filtering measures (for reasons why this should be done on egress, see [74]). The next stage in this branch, the *Egress Rate Limiter* marks packets to be dropped if the member's rate limit is exceeded.

The *local* and *remote* paths merge again in the *Load Balancing and Failover* stage. However, they are handled separately. For *remote* delivery, there are typically multiple links available (due to redundancy requirements) for entering the IXP's core. Thus, this stage is responsible for distributing the packets among the available links according to metrics defined by the IXP operator, enabling, e.g., cost based routing or ECMP. In case of *local* delivery, packets are forwarded to the member while respecting LAGs. This stage is responsible for realizing failover mechanisms using, e.g., in-data-plane probing [135, 167], based on the management's layer holistic view of the network topology.

The *Packet Sampling* stage exports data for statistical and monitoring purposes. This data is traditionally generated from sampled packets, aggregated on a per-flow basis. In our pipeline, any stage that decides to drop a packet will mark this in the packet's metadata. It then continues to travel through the P4 pipeline until the egress parser stage. Thus, the sampling stage has full visibility of all packets entering the pipeline even if they are dropped, including the reason why the packet is dropped. This, as well as the full visibility of all Layer-2 to Layer-4 information, offers better visibility of IXP relevant data. Note, most traditional packet sampling protocols only export Layer-3/4 headers.

The final stage reads the packet metadata, deparses the respective headers for either remote or local delivery. If it is not marked drop it reassembles the packet and emits it via the output port.

4.5.2 Management and Monitoring

The above pipeline is based on pre-populated lookup tables, visualized as blue boxes in Figure 4.4, that reflect the IXP's functionality, i.e., ranging from member specific rate limits, LAG or VLAN configurations, to blackholing rules, as well as internal routing decisions and value-added service specific information. Thus, we next discuss options for how to populate these tables.

- The *traditional approach* is the most wide-spread management approach in today's networks. It implements embedded management and control on the device by providing CLI access and a limited set of configuration protocols (e.g., SNMP). The CLI is usually proprietary. Hardware vendors often provide an additional, proprietary network management system (NMS) that abstracts the CLI details under a GUI. This approach comes with vendor lock-in and due to its proprietary nature often leads to challenging automation.

- The *model-exposure approach* works by exposing a formal model of valid networking hardware configurations to the outside world. This allows validation of configurations or configuration changes upfront, which is highly beneficial for automation. After validation, changes are transferred to the hardware, usually in a transaction-based manner enabling rollbacks upon failure. Netconf/YANG [83, 14] provide a standardized that is currently reaching operative environments.

- The *external controller* approach was made popular by OpenFlow [18]. The networking hardware outsources all logic to external controller software, which instructs the hardware using a narrow interface to push forwarding decisions. All management functionality is implemented in the controller. This approach is very flexible, but also introduces complexity by adding a potentially large controller stack, which needs to be fail-safe for critical infrastructure to the system.

From the perspective of an IXP operator, mixing the traditional and the model-exposure approach is most promising since the external controller approach is a risky choice. The model-exposure approach allows for a tight integration with the IXP's business logic. It can be used to pull IXP member data from external data sources like ERP systems, generate configuration changes, validate them, and apply them to the data plane. This software-driven process enables a high degree of automation. At the same time, IXP NOC teams need to be able to troubleshoot networking hardware. These teams usually do not have software development skills and should, therefore, remain able to override any configuration change manually using a CLI until problems are fixed in the software stack. This aligns well with today's operational practices.

The model-exposure approach is achieved through two key software components that run on the host system of the P4 switch: (1) a database client that is connected to an external database containing all information related to the IXP business logic and (2) a data transformation pipeline executed upon every update of the database. The external database is not to be confused with a traditional SDN controller, as the P4 switch remains fully operational even if the database fails. The intelligence resides in the data transformation pipeline that generates P4 lookup table contents according to the IXP business logic and passes them to the P4 runtime, which in turn populates the P4 tables in the data plane accordingly. This is depicted with blue icons and dashed arrows in Figure 4.4. For a list of what information is stored in the lookup tables, we refer to Table 4.1.

Besides management, monitoring plays an important role. Currently, a large zoo of protocols is used, e.g., for the export of sampled packet data, there are three competing standards: SFlow, Netflow, Internet Protocol Flow Information Export (IPFIX) [194, 42, 44, 43]; other monitoring information such as interface information is exported via SNMP or Streaming Telemetry as promoted by OpenConfig[10]. Moreover, monitoring data needs to be distributed to a larger number of endpoints with different purposes, e.g., billing, network monitoring, statistics, DDoS mitigation, and data warehouses. The existing approaches are not well-suited for this scenario, as they are designed for a small number of data sinks. Consequently, often message brokers take on this job, i.e., there is a central infrastructure (e.g., a Kafka cluster) tasked with distributing measurement and monitoring data to different endpoints after decoding the necessary information from its wire format. With a P4IX architecture, this step can be skipped and a message broker client can run directly on the switch operating system (e.g., Linux) and push the necessary information directly to the message broker—notwithstanding that a standardized format is used for messaging to remain compatible with existing monitoring tools.

[10] `https://www.openconfig.net/`, last visited 12/1/2021

4.6 Discussion

The proposed P4IX is superior to the earlier proposed OpenFlow based SDX realizations (e.g., no reliance on external controllers, solutions for link failure, etc.). As such, it offers many advantages for IXP operators; however, moving an IXP to P4IX imposes significant changes, which introduces new challenges as well. In the following, we discuss advantages and disadvantages of the P4IX.

P4IX Advantages. The P4IX is based on standardized P4 which is supported by a wider range of hardware. Thus, it prevents vendor lock-in, which promises a considerable reduction in capital expenditures as well as a larger range of hardware. This is a considerable advantage given the highly varying requirements of various IXPs. In addition, the P4IX should reduce operational expenditures as the network becomes easier to manage and operate. On the one hand, technical complexity, e.g., routing protocols, can be eliminated, while on the other hand, tightly coupling the IXP's business logic with the data plane simplifies provisioning of new members and running value-added services. Moreover, it is now possible for IXPs to implement new functionalities themselves or tailor existing functionality to IXP-specific challenges. These can then be rolled out immediately. IXPs, thus, no longer need to wait until traditional hardware vendors release a requested feature. This enables tailored, potentially highly customized solutions independent from the rigid feature set or hardware constraints of ISP-oriented hardware. This promises to reduce the time-to-market, e.g., for value-added services. Moreover, the P4 implementations can be made open source, to fuel community driven collaboration between IXPs. IXPs can also share their implementations with their members so that these can review code and increase their trust in the platform's software.

P4IX Challenges. While a P4IX is promising, it does not come for free. A considerable shift in the IXP's development and management is required. So far, data plane development and, in particular, software development has not been the focus of IXP staff. This has been the responsibility of the hardware vendor. The same applies for data plane testing regarding compatibility with other hardware and for operational stability. Thus, we argue that, due to the size and structure of IXPs, data plane development and testing should be done collaboratively based on an open source foundation. While open source *control plane* projects such as the BIRD route server [208] are providing critical services at IXPs, *data plane* projects have not yet been set up. Moreover, to test the stability and compatibility of a production-ready P4IX is likely to require a substantial test lab.

Moreover, the hardware packaging of current P4 switches is not yet a good match for medium to large IXPs due to their port density requirement, i.e., for connecting hundreds to thousands of member ASes. Additional hardware is required to aggregate these links before being connected to the IXP's edge. This generates the need for additional rack space and increases the energy consumption. However, large and very large IXPs are the ones that may drive the development of a P4IX—they are the ones with the resources.

Besides the technical challenges, a P4IX also imposes organizational challenges. A higher degree of automation such as proposed in this work requires the retraining of the network engineers and NOC members at the IXP. In the mid-term, the CLI-oriented mindset of IXP staff needs to be transformed into a mindset organized around software development teams, i.e., release cycles, software engineering methods like SCRUM and development sprints.

4.7 Summary

In this chapter, we present the concept of a P4-based IXP that provides IXPs with a high degree of flexibility and allows to quickly react to unforeseeable events in the Internet. We provide a characterization of the IXP landscape, and find that IXPs in particular fall short with regard to what hardware vendors typically implement in the feature set of the forwarding devices. They do not focus on the very specific needs of IXP operators which underlines the need for an alternative solution. With the help of SDN—more precisely P4—we show how these specific needs can be reflected in a highly customizable IXP-P4 pipeline. By combining (1) the virtualization layer used for service realization and (2) the routing layer used for packet forwarding through the infrastructure, we created one unified P4-enabled forwarding layer. It addresses the need for IXP operators and allows for fine-grained adaptions and adjustments in operation as well as extensions that altogether can be deployed by writing code instead of relying on external hardware vendors with long-standing release cycles and the risk of having to replace the hardware.

The proposed P4-based IXP addresses IXP specific challenges and offers a high degree of flexibility to address them in rapidly short deployment cycles. It allows the Internet's core to handle unforeseeable demands and react to stress situations.

Chapter 5
The Internet Under Attacks

In this chapter, we now turn our attention to stress situations in the Internet with a digital cause. In contrast to the previously covered stress situation, e.g., coping with a global pandemic, cyberattacks originate from the digital world. They aim to impair the availability of online services. One example is the amplification DDoS attack. For background, see Section 2.4.

Amplification DDoSs attacks' traffic and harm are at an all-time high. To defend against such attacks, distributed attack mitigation platforms, such as traffic scrubbing centers that operate in peering locations, e.g., IXPs, have been deployed in the Internet over the years. These attack mitigation platforms apply sophisticated techniques to detect attacks and drop attack traffic locally, thus, act as sensors for attacks. However, it has not yet been systematically evaluated and reported to what extent coordination of these views by different platforms can lead to more effective mitigation of amplification DDoS attacks.

In this chapter, we ask the question: "Is it possible to mitigate more amplification attacks and drop more attack traffic when distributed attack mitigation platforms collaborate?" Our previously proposed concept is not suitable to answer this question since DDoS attacks cannot be solely solved by a P4-based single IXP. Hence, we collaborate with eleven IXPs from three different regions. These IXPs have more than 2,120 network members that exchange traffic at the rate of more than 11 Terabits per second.

We collect network data over six months and analyze more than 120k amplification DDoS attacks. To our surprise, more than 80% of the amplification DDoS attacks are not detected locally—supporting our point that a single IXP is not sufficient—although the majority of the attacks are visible by at least three IXPs. A closer investigation points to the shortcomings, such as the multi-protocol profile of modern amplification attacks, the duration of the attacks, and the difficulty of setting appropriate local attack traffic thresholds that will trigger mitigation. To overcome these limitations, we design and evaluate a collaborative architecture that allows participant mitigation platforms to exchange information about ongoing amplification attacks. Our evaluation shows that it is possible to collaboratively detect and mitigate the majority of attacks with a limited exchange of information and drop as much as 90% more attack traffic locally.

5.1 DDoS Mitigation in the Core of the Internet

As our commercial and social activity is increasingly moving online, due to the ongoing pandemic as well [92], cyberattacks are more frequent and devastating [179, 9, 31, 245]. By recent measures, the damage due to cyberattacks in 2020 alone is estimated to one Trillion USD [224], double than the damage in 2018.

Among the most popular cyberattacks are these that target online services. To generate voluminous attack traffic, attackers that are politically or commercially motivated, compromise computers around the globe. These so-called Distributed Denial of Service (DDoS) attacks are well orchestrated and typically exploit vulnerabilities of computing systems [190, 189]. In recent years, it is even possible to lease resources or compromised machines for attacks using booter services that are available in the public or dark market [50, 139]. Studies of DDoS attacks [9, 110, 245] have shown that attackers often first test the operation of their attack system by launching low volume attacks, or targeting low profile targets before launching the fully-fledged attack.

In terms of attack volume, recent studies have shown an exponential surge [142, 179]. Until 2012, the largest attack reported was less than 100 Gbps. In recent years, attacks with orders of magnitude higher traffic (up to 2 Tbps), have been reported, e.g., the 2018 memcached attack [2]. At the same time, attacks are becoming more sophisticated. Analysis of recent attacks shows that attackers can generate attacks with hundreds of millions of packets per second (Mpps) [144]. Thus, attackers are not only able to launch voluminous attacks but also attacks that require additional computation resources from defenders. The most successful of these attacks are *amplification attacks*, i.e., the attacker can create harm to the target that is up to 50,000 times higher than the original attack traffic the attacker generates by utilizing services like DNS or NTP as reflectors [210, 58]. If this trend continues, it is expected that in the next years attacks more than 10 Tbps and multiple Gpps will be reported. At this scale, no single infrastructure provider alone can defend them.

The response from the industry was to introduce DDoS attack detection and mitigation platforms deployed at various locations in the Internet [130, 128]. Among them, traffic scrubbing centers analyze incoming traffic and apply rules to detect DDoS. Detected attack traffic is then dropped locally [3, 103, 130, 238]. The required processing per packet or per flow processing for deep packet inspection increases the detection and mitigation cost of attacks, does not scale well, imposes performance penalties and is vulnerable to evasion tactics [129]. Other techniques, such as Remote Triggered Blackhole filtering [37] are more aggressive and scalable, but require the detection of attack from a separate system. Unfortunately, these coarse-grained mitigation techniques also drop legitimate traffic to the destination under attack, and, thereby cause collateral damage. FlowSpec allows for fine-grained mitigation. However, although it has adopted in intra-domain environments [29, 211], it has not been popular in inter-domain environments as it requires sharing of computational and network resources across independently administrated networks. More recently, finer-grained blackholing has been proposed to address the limitations [74]. Both traffic scrubbing centers and blackholing functionality are present and readily available in peering locations such as Internet Exchange Points (IXPs) [33, 104]. To the best of our knowledge this is not the case for FlowSpec. Today, it is well accepted that the previously mentioned DDoS mitigation techniques are effective. However, their performance has typically been evaluated at a single location [138, 73, 74, 228, 139]. For a small number of attacks, a previous work has utilized publicly available vantage points to infer the efficacy of Remote-triggered Blackholing [104].

IXP Code	#Members	Peak Traffic (Gbps)	Region	#Sampled Flows (Billions)
CE1	900+	9,000+	Central Europe	1,077.5
CE2	200+	150+	Central Europe	9.9
CE3	200+	150+	Central Europe	3.2
CE4	200+	100+	Central Europe	3.6
NA1	200+	800+	North America	78.8
NA2	75+	150+	North America	16.7
SE1	175+	400+	South Europe	30.5
SE2	75+	100+	South Europe	12.2
SE3	40+	10+	South Europe	2.2
SE4	30+	100+	South Europe	17.9
SE5	20+	50+	South Europe	2.0

Table 5.1: Statistics about the 11 IXPs in our study between April 27 to October 5, 2020.

In this chapter, we investigate whether coordination among attack detection and mitigation platform can be even more effective to: (1) detect and mitigate more amplification DDoS attacks, and (2) drop more amplification DDoS attack traffic locally that otherwise is carried to either be dropped later or to cause harm.

Our main observation is that the distributed nature in which reflectors are exploited for launching reflection DDoS attacks can be leveraged to realize better DDoS detection approaches. This way, infrastructures at different locations in the Internet can act as distributed "sensors" to better capture the global attack activity. Such sensors, in our case Internet Exchange Points, can collectively infer attack activity faster and potentially for even relatively small attacks. Thus, they can signal their peers about ongoing attacks.

The contributions of this work can be summarized as follows:

- We establish a collaboration with 11 Internet Exchange Points around the globe to utilize them as vantage points and collect traffic. Over a period of six months, we detect and analyze more than 120k amplifications DDoS attacks. Our analysis shows that more than 50% of the attacks are visible in more than three locations, in many cases in more than 5 locations.
- We show that more than 80% of the attacks that send traffic via one of our vantage points are not mitigated because local detection mechanisms aren't triggered (1) due to local attack traffic thresholds aren't exceeded or (2) due to the attack's multi-protocol profile remaining unseen at a single location.
- We show the critical role that infrastructures, like IXPs, located in the core of the Internet, can play in detecting and mitigating DDoS attacks. We show that around 45% of the reflectors' traffic is directly transferred from IXP members, and at least 30% of the attack targets are members of these IXPs.

5.2 Data Sets

For our study we leverage a distributed set of vantage points as well as network data and routing information. Our goal is to analyze all the available data to characterize recent amplification DDoS attacks and assess the potential benefits of collaborative DDoS detection and mitigation.

5.2.1 Vantage Points

We establish a collaboration with 11 Internet Exchange Points (IXPs) that operate at three regions around the world, namely Central Europe, South Europe, and North America. Three of these IXPs, namely `CE1`, `NA1`, and `SE1`, are the same IXPs that we introduced in 3.1.1. For an overview of the IXPs in our study we refer to Table 5.1.

All the IXPs we collaborate with are already offering DDoS mitigation solutions to their members. They offer BGP blackholing [73] or advanced blackholing [74] to block traffic to specific destinations or transport ports by dropping the traffic at the IXP, as a free service. IXPs are excellent locations for mitigation, as they have a large spare capacity. Thus, they can absorb large attacks at the scale of Tbps. All the IXPs we collaborate with have at least one scrubbing center [128] as member.

5.2.2 Flow Data

We get access to flow data collected at each of the 11 IXPs. Due to high volume, all these IXPs use the IPFIX [44, 43] that aggregates information per flow without storing the payload of the packets. For maintaining scalability, they sample packets at the rate of 1:10k.
Passive measurements. Collection of the flow data took place from 27th of April, 2020 to 5th of October, 2020. The total traffic of sampled data is 1.175 Petabytes, that corresponds to approximately 11,750 Petabytes of exchanged traffic in total (all 11 locations).
Active measurements. We collected flow data for self-initiated controlled DDoS attacks. The measurements took place between February 11th and 20th, 2021. During the self-attacks, a total of 4.6 TB was transmitted, resulting in about 340K sampled flows.

5.2.3 Metadata

IXP member lists. During the time of flow data collection, we have access to the list of members at each of the IXPs. The lists are updated every day, as new members are added (or removed) daily.

IXP route server BGP data. Moreover, we collect routing data from the route server at each IXP during the flow collection period. A route server [208] is a free service offered by all the 11 IXPs to their members. The IXP members have the option to announce prefixes to all the other members of the IXP. This service is very popular with more than 60% of all member networks using the IXP route server with an open peering policy. The IXP network members can announce their prefixes with only one BGP session, instead of establishing one session for each peer network. The Route Server also offers the option to announce prefixes to only some, none or all peers. We have access to both the input and (filtered) output at the 11 IXP's route servers. We analyze the output routes of the route server, i.e., the best path selected for route propagation to the peers. This allows us to

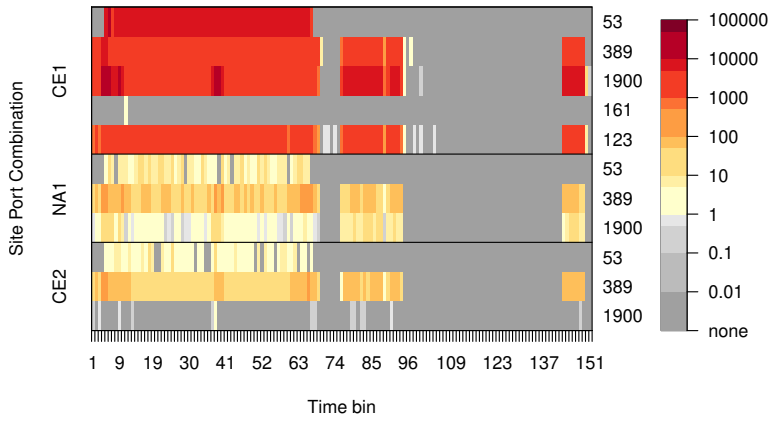

Figure 5.1: Attack against a large CDN on June 04, 2020. Each row corresponds to a different protocol used for the attack as observed at each IXP.

derive AS-distance information from the propagated BGP messages. For incoming traffic, however, the AS-distance correctness relies on the assumption of symmetric routing.

Internet routing registers. In addition, we also have access to public collector data sets. We used the Routing Assets Database (RADb) [198] to retrieve mappings from IP blocks to ASNs. This allows us to detect whether an IP address belongs to a peer of one of the IXPs. This further helps to calculate the distance between the IXP and the reflector or target respectively.

5.3 Anatomy of DDoS Attacks

In this section, we analyze real-world DDoS attacks in detail to assess their visibility across our vantage points. We approach this in two different ways. First, we have an exemplary look into a recent large scale DDoS attack that is well reported by the targeted infrastructure provider. Second, we analyze a set of self-attacks towards our measurement network connected to the IXP infrastructures we collaborate. We run various advertisement scenarios for our attacked IP space to observe how the distribution of DDoS attack traffic is affected.

5.3.1 A Recent Tbps Reflection Attack

We investigate an exemplary DDoS attack that is one of the largest ever reported in terms of attack traffic volume. The attack took place on June 4th, 2020 and targeted the CDN Akamai, which is member of all 11 IXPs we have data for. The peak attack traffic was reported to be 1.44 Tbps [126]. The attack is not only voluminous in terms of traffic, but also in terms of the number of packets. At the peak, around 385 million packets per second were generated. This makes the attack even more effective as it requires additional defense resources due to the processing of the very high number of packets, especially for traffic scrubbing centers. The attack is very sophisticated as nine different attack vectors are reported. These are multiple TCP and UDP specific attack vectors along with amplification attack vectors, i.e., SSDP, CLDAP and NTP. Our DDoS attack inference approach presented in § 5.4.1 successfully detected this event as an attack.

By analyzing the network flow data collected at the IXPs we confirm that the attack was visible at 3 of our vantage points. In Figure 5.1 we plot the reflector to target traffic volume for each IXP where the attack is visible. Note that these IXPs are located at different continents. The attack consists of 3 bursts and ends after a total duration of about 3 hours. The peak traffic observed at our vantage points totals at about 100 Gbps in terms of attack volume and at about 20 Mpps in terms of packets.

Five of the reported attack vectors are clearly visible in our data. The attack had an ON-OFF pattern, as also reported by Akamai [126]. Such patterns are common, as attackers try to avoid detection, and they also switch between attack sources and reflectors. Our data shows that the same target saw multiple smaller DDoS attacks one week earlier, using similar attack vectors. This can be considered a trial DDoS attack, a common pattern used to verify for example function of the reflectors. Furthermore, we observe the same target to be under attack multiple times with more than 1 Gbps of total attack volume across our whole period of observation, after the reported attack. All these attacks use the same aforementioned attack vectors. Thus, we conclude that DDoS attacks are indeed visible at different locations. However, the level of the attack traffic and amount of attack features visible at different locations may differ. As an example, the attack traffic at NA1 and CE2, is quite low compared to CE1. Furthermore, CE1 observes 5 attack vectors, but NA1 and CE2 only 3, respectively. Using DDoS detection practices in one location can lead to slower or no reaction, e.g., because the attack traffic volume is not large enough at a specific location, or the locally visible attack vectors remain under the radar of local defense mechanisms.

5.3.2 Self-Attacks

To gain ground truth for our study, the IXP operator attacks its own measurement infrastructure (for details see § 5.6) that is connected to 10 of the 11 IXP locations. During the attacks, the observable attack traffic is measured at these different vantage points. By adjusting the advertisement of the measurement network, it is possible to control the direct visibility of the DDoS target at the IXP locations and steer the traffic in different scenarios towards the network. To generate traffic samples from real-world DDoS attacks, a DDoS for hire service, also known as booter service, is utilized. The amplification protocols that are included in the self-attack experiment comprise DNS, NTP, SNMP, CLDAP, SOAP, SSDP, ARD. The attack traffic observed at the measurement infrastructure ranged from 100-600 Mbps for the protocols ARD and SSDP, more than 1 Gbps for SOAP, SNMP and up to 7 Gbps for NTP, DNS, CLDAP. The IXP's measurement infrastructure is attacked with the available amplification vectors with two different advertisement scenarios: (1) advertise measurement network only via the largest IXP, i.e., CE1 (unicast), and (2) advertise our measurement network at all locations (anycast). The self-attack traffic level observed at each location for the specific advertisement scenario is reported in Figure 5.2. To our surprise, we observe DDoS traffic towards the target at several vantage points even when advertising from a different location. When advertising the measurement network globally via anycast, we can see a stronger attraction of DDoS traffic locally especially for the locations in South Europe. This highlights the geographical distribution of DDoS traffic and the great potential for global mitigation efforts, even for networks that don't advertise their IP space via anycast.

Takeaway: Amplification DDoS attacks are globally distributed across networks, hence visible from different vantage points worldwide. At some, however, only with very low traffic rates that might hinder their local detection. Even for networks which don't

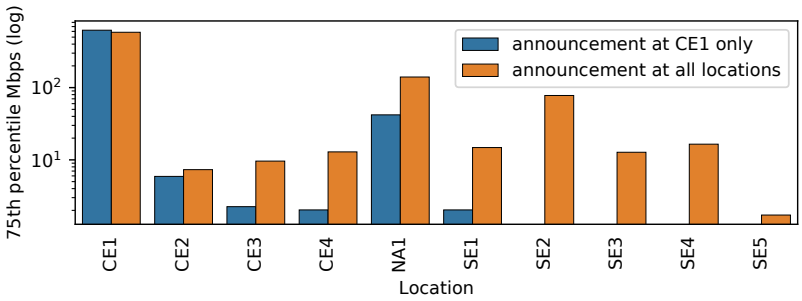

Figure 5.2: Attack traffic from self-attacks per location comparing two announcements scenarios (anycast and unicast).

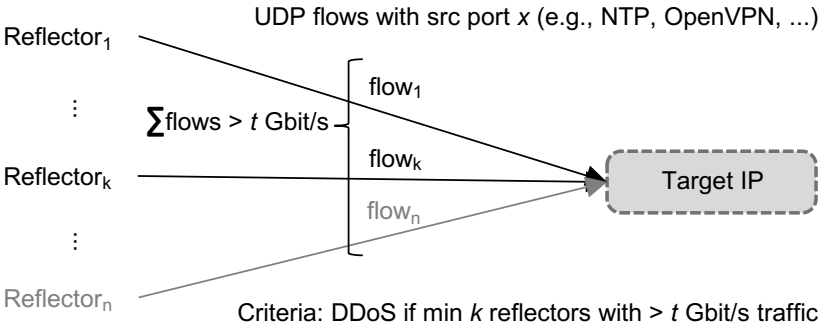

Figure 5.3: DDoS inference approach following [138].

employ a distributed infrastructure (e.g., anycast) attack traffic can be visible traversing different locations towards the target. The results show the great potential to more effectively detect and mitigate DDoS attacks.

5.4 Inference of DDoS Attacks

We focus on DDoS reflection attacks, that are responsible for some of the largest attacks known to date [5]. Their popularity and sophistication has increased the last years [108]. More frequently there are reports highlighting not only the traffic volume of DDoS attack, but also increasingly high numbers of packets per second [108].

5.4.1 Detecting DDoS Attacks in Flow Traces

Detection. To identify DDoS reflection attack traffic in the flow-level traces provided by the IXPs, we employ the approach proposed in [138] as shown in Figure 5.3. We consider an IPv4 address to be under attack, if its inbound traffic exceeds a threshold of $t = 1$ Gbps from more than $k = 10$ (reflector) IPs with the same source port (e.g., a known reflection protocol such as NTP). It is based on the assumption that it is unlikely for an Internet client to receive traffic from many sources with the same source port number (e.g., NTP) at a high traffic rate. We remark and will later show that typical DDoS attacks can generate much larger traffic volumes and involve reflectors in the number of 100s or even 1000s. Yet, prior work [138] has shown this filter to be capable to differentiate

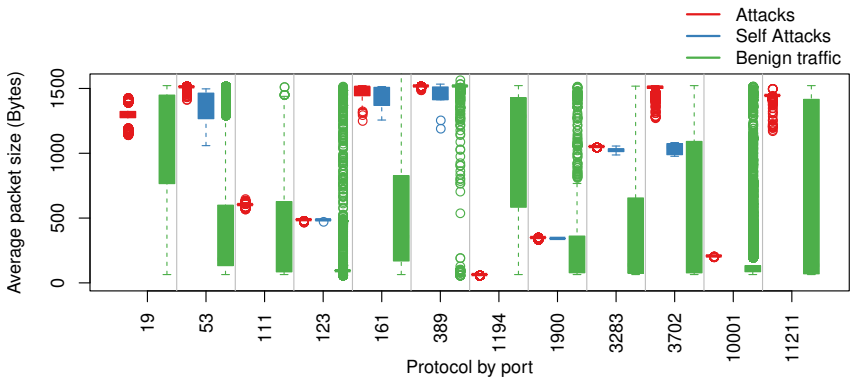

Figure 5.4: Validation of filtering approach. We compare the packet size characteristics of benign, attack traffic, and self-attack traffic for different protocol source ports.

between attack and benign traffic, an observation that we confirm in our validation. For our study, we focus on UDP-based amplification attacks.

We remark that the described filter approach [138] was proposed to be applied at a single site only. To be applicable in our multi-site scenario we extend it as follows. First, we define a flow as a septuple: (source (MAC address, IPv4 address, transport port), destination (MAC address, IPv4 address, transport port), IXP code). This ensures, that the same traffic flow traversing multiple IXPs will be captured as individual flows in our data to enable the later analysis of attack traffic visibility at different sites. Second, we define two variants of the detection threshold t. In the first variant (local threshold), the detection is applied to traffic from a single IXP only. In the second variant (global threshold), we detect a DDoS attack if the traffic sum exceeds t over all IXPs. We evaluate these thresholds in § 5.5.2.

Filtering. To further avoid false positive classifications, we filter the flow data for traffic having the source transport port set to the well-known port of popular DDoS reflection attack protocols. These (and the associated port number) are the following: Chargen (19), DNS (53), RPC (111), NTP (123), SNMP (161), CLDAP (389), OpenVPN (1194), SSDP (1900), ARMS (3283), WS-Discovery (3702), Device Discovery (10001), Memcached (11211). In addition, we take the packet size in account, as reported in [138]. By this, we populate a new, pre-filtered data set for attack detection purposes, consisting of 4 billion flows belonging to an average of 3TB of traffic exchanged data per day. In order to evaluate the accuracy of this filtering approach, we compare the average packet sizes per source port for benign traffic, attacks as defined by our filtering method and the recorded self-attacks as a ground truth. The results are shown in Figure 5.4: for any port we have data in the self-attacks, we observe a comparable packet size distribution, which deviates clearly from the benign traffic. In all other cases, we observe a clearly different packet size distribution between attack traffic using our filtering method and benign traffic. Our filtering method either selects traffic at the upper end of an Ethernet frame (high amplification factor, e.g., port 11211 for Memcached) or a with very small deviation around a characteristic packet size (e.g., 500 bytes for NTP monlists), which is expected and indicates a correct selection.

Attack Statistics. In Figure 5.5 we plot the number of attacks and the associated attack traffic volume for the amplifications attacks detected with our methodology. We notice a

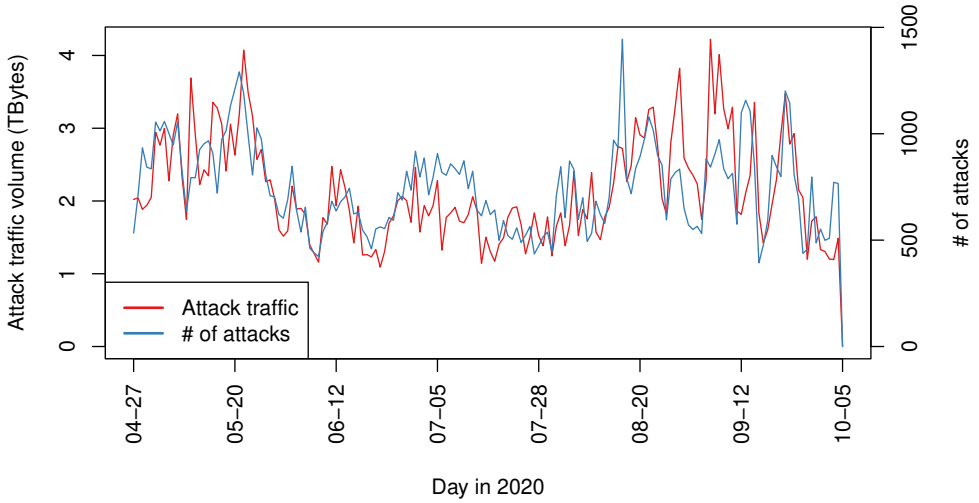

Figure 5.5: Daily number of observed DDoS amplification attacks and associated traffic volume during the period of our analysis.

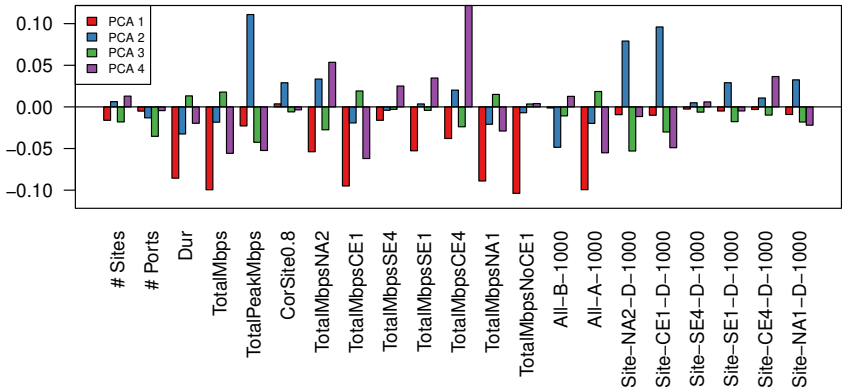

Figure 5.6: Principal Component Analysis (PCA): Contributions of different features to the rotation of the first 4 PCAs. This highlights that different IXP members add complementary information.

great variance among the total of 120k detected attacks. For most of the days, there are at least 500 attacks detected, whereas for some days, this number is about 300% higher. However, we did not notice any particular pattern, e.g., day of the week that has more attacks than others. However, we noticed that some of the services (ports) receive more traffic than others. The top ports (attack traffic volume) are: 123 (33%), 389 (30.8%), 53 (27%), 11211 (6%), and 1194 (1.2%). The other ports receive less than 1% of the attack traffic. The aggregated attack traffic volume that is exchanged in the 11 IXPs varies from 1 Terabyte to 4 Terabytes per day.

5.4.2 Validation: Attack vs. Benign Traffic

Features. To characterize benign and attack network flows, we derive 1,106 features from the flow-level traces. For an exhaustive list, we refer to Appendix A. These features include basic statistics like the duration of an attack or the overall as well as the peak

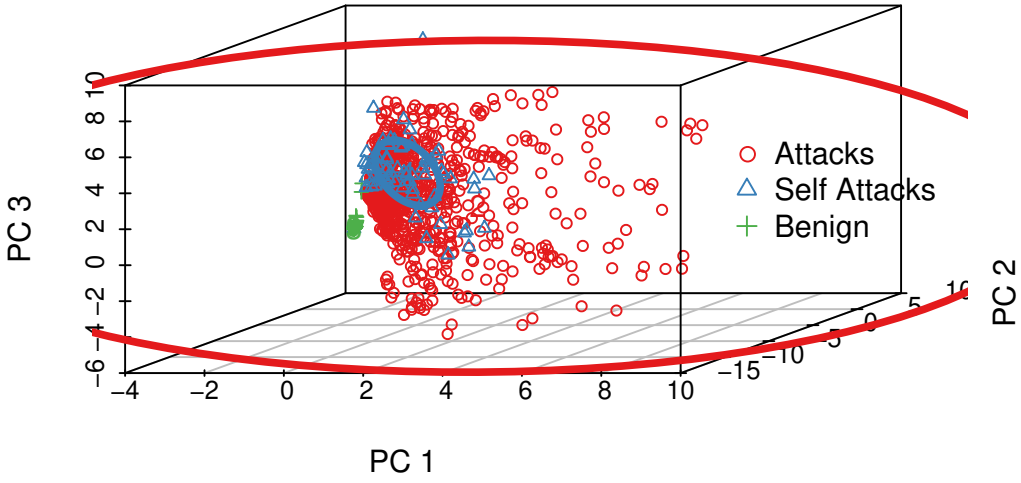

Figure 5.7: PCA: Projection of the attack, the self-attack, and benign traffic sample to the first three PCAs. The plot shows 1k samples for the attack and benign traffic, and an ellipsoid to show where the mass of the points fall.

traffic volume (in total, per transport-level protocol and per site). In addition, we consider the correlation of the attacks across different transport ports and across sites. Here, we compute the pair-wise correlation matrices for all features across time and then count how often the correlation exceeds the thresholds of 0.7, 0.8, and 0.9. Moreover, we perform local and global attack detection using different thresholds and include the time until attack detection as well as the traffic volume before, during, and after detection as additional features. As the overall traffic volume at the IXPs differs significantly, we also normalize the traffic per IXP site and compute the same feature sets for the normalized traffic volume. We remark that these features are of purely descriptive nature for our analysis and point to Figure 5.6 for a subset of the available features.

Attack vs. benign traffic. We validate the DDoS inference approach described in § 5.4.1 by analyzing the features for benign and attack traffic (i.e., traffic that matches our inference approach in § 5.4.1). We consider traffic within our unfiltered data set to be benign, if the destination IP address did not show up once using our detection mechanism. To reduce the dimensionality of the feature space, we apply a PCA. A PCA decomposition can be used to project a high-dimensional space to a lower-dimensional space by relying on the initial principle components. In effect it converts a set of values of M possibly correlated variables into a set of K uncorrelated variables, the PCAs. In that regard, PCA is a clustering algorithm for high-dimentional data. We find that a significant number of our features are correlated since the first 5 PCAs explain more than 25% of the variance and the first 50 more than 75% of the variance.

In Figure 5.7, we show the projection of our feature set to the first 3 PCA dimensions for both benign (cross) and attack (triangle) traffic. PC1, PC2, and PC3 are the three principal components. We also apply a k-means clustering to 8 clusters on the data and color the corresponding clusters. Note that Figure 5.7 is zoomed into the 0.01 to 0.99 quantiles for each dimension. We do so as there are a small number of outliers for each dimension (enclosed in the ellipsoid red envelop). We observe that benign and attack traffic can be visually clearly separated, which highlights that their flow-level characteristics in our feature space differ substantially. Moreover, this region only covers data points from three clusters: the red, green, and the blue one. The red and blue

ones only contain attacks while the green one contains mainly benign traffic samples. Note, the k-means clustering is a very simple clustering mechanism and with a more sophisticated mechanism it should be easily possible to separate attack from benign traffic samples.

Overall, this analysis suggests that the applied filter successfully detects DDoS attack traffic. Additionally, we manually inspected a random subset of the attack traffic to further support this finding. Thus, we use the DDoS attacks matching our inference approach as data set to inspect the distributed nature of DDoS attacks and illustrate our mitigation approach in the remainder of this chapter.

Relevant features per site. We next use the PCA to understand if the attack traffic features are homogeneous for the different IXPs. Therefore, we consider the contribution of the different features to the different PCAs. More precisely, we look at the rotation values. The rotation per feature and PCA captures which contribution the feature has to this specific PCA. In Figure 5.6 we show a bar plot of rotation values for the top features for the first four PCAs. This analysis shows, that the feature relevance differs per IXP. It is not only the volume that counts but also where the site is located, what type of member networks there are etc. This highlights that the different IXPs have complementary perspectives on the attack landscape. In the scope this chapter, this suggests that a cooperation of these IXPs in jointly detecting DDoS attacks by exchanging data is beneficial.

5.5 Detecting and Mitigating Thousands of DDoS Attacks

To understand the challenges and opportunities of combining the views of multiple vantage points to detect and mitigate attacks, we perform a detailed analysis of the more than 120k attacks we inferred with our detection method described in the previous section.

5.5.1 Challenges in Detecting DDoS Attacks

The detection and, thereby, also the mitigation of DDoS attacks is subject to challenges.

Detection Lag. Recent industry reports show that DDoS attacks are typically of short duration, i.e., less than one hour. For example, Cloudflare [46] reports that 90% of attacks last less than one hour. We observe similar characteristics in our data set (not shown): most of the attacks are relatively short-lived. Indeed, around 70% of the attacks have a duration of 10 minutes or less. 95% of the attacks lasted less than 50 minutes. The short duration of the attack traffic makes its detection challenging. A detection and mitigation approach that is too slow (e.g., by requiring longer sample periods for stable detection) will, thus, fail to detect a large bulk of the current DDoS attack landscape. By performing a collaborative DDoS detection proposed in this chapter, we show later that we can reduce the time required for detection and thereby increase the number of detected DDoS attacks.

Multi-Protocol Attacks. Another challenge is that most of the attacks do not rely on single transport port. Prior work has observed a tendency of DDoS attacks to utilize multiple attack vectors (i.e., amplification protocols) [46]. Thus, simple port-based blocking rules may not suffice in blocking DDoS traffic. In Figure 5.8, we show the distribution of the number of amplification protocols (ports) used to perform DDoS

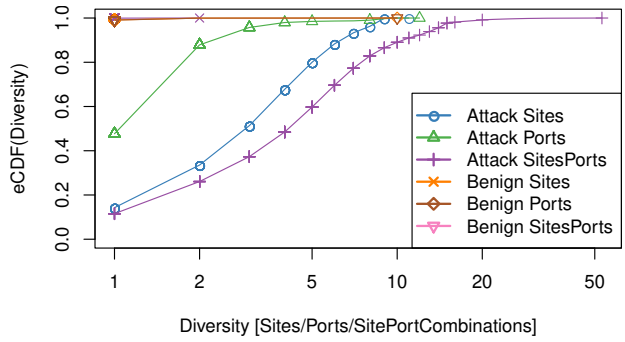

Figure 5.8: Diversity of ports and sites, and combination of both in attacks and benign data.

attacks (sites and combination discussed later). Our main finding is that most attacks involve 3 or more amplification protocols. More than half of the attacks in our data set use more than one amplification protocol. This holds both for short- and long-lived attacks.

Global vs. Local Thresholding. To detect amplification DDoS attacks the typical approach is to use local thresholds. These thresholds are only applied to the local traffic. This can be misleading as only a fraction of the attack traffic (below the local threshold) is routed via one location, but on aggregate the attack traffic yields a large DDoS. To show that this is quite often the case, in Figure 5.9 we plot the number of attacks detected using different local thresholds. Red bars annotate number of attacks visible at each IXP.

5.5.2 Opportunities in Detecting and Mitigating DDoS Attacks

Visibility of Attacks at Multiple Sites. A key observation in this chapter is that there exists visibility for the same DDoS attack at multiple sites. This is rooted in the fact that these DDoS attacks are executed by abusing a large set of reflectors distributed across many different networks. Thus, given inter-domain routing, the traffic paths from these reflectors to the attacked target can be expected to traverse many different networks. In Figure 5.8 we show the number of IXP sites at which the benign flows and DDoS attacks in our data sets are visible. In the same figure, we further show the distribution of the (IXP sites, amplification protocol/port) combination. For the benign data, we see most of the flows at a single site, using a single protocol. In contrast, 80% of the attacks are visible at more than one IXP, even if we further restrict this by amplification protocol. Given that the traffic volumes observed at each IXP site vary, attack detection at a single site alone is challenging. Yet, this result shows the opportunity in detecting and mitigating DDoS attacks: if IXPs unite to jointly detect DDoS attacks, the number of detectable attacks increases.

We notice that the majority of the attacks that are visible at the IXPs are often missed for both low and high local thresholds. Only, the very large IXP in our study, CE1, tends to only lose a small fraction, especially when the threshold is small. Our analysis shows that around 80% of the attacks are missed by a large majority of the IXPs, except the very large IXP. Indeed, the very large IXP's view contributes to the global view of the

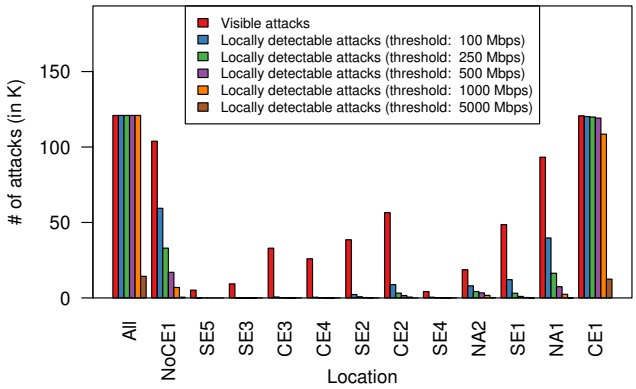

Figure 5.9: Number of attacks detected using local thresholds. Red bars annotate number of attacks visible at each IXP.

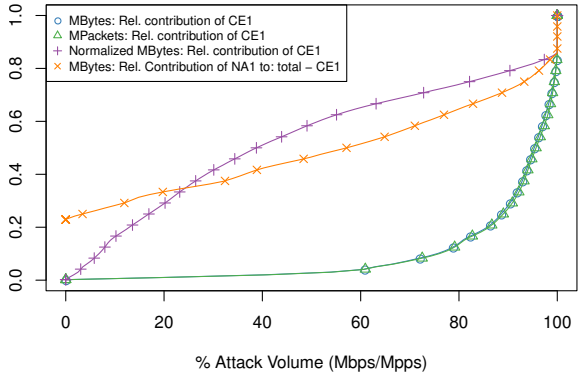

Figure 5.10: Visibility by IXPs on the attack traffic (bytes per second and packets per second).

ongoing DDoS attacks both in terms of bytes per second as well as packets per seconds as shown in Figure 5.10. Thus, although for very large IXPs local thresholds can be effective, such thresholds may not be sufficient for typical IXPs. Our study also shows that if local thresholds are set proportionally to the size of an IXP, with reference point a very large IXP, the detection of amplification DDoS attacks can be improved significantly, see Figure 5.11. However, the false positive rate may increase as well.

Potential Attack Traffic Savings. To estimate the potential attack traffic savings when information about the ongoing attack is shared, in Figure 5.12 (top) we compare the traffic that could have been detected and blocked at each IXP with a local or a global threshold. The difference is striking especially for the smaller sites and for various values of local and global threshold. In some cases the missed amplification DDoS attack traffic is close to 100% even for very low thresholds (100 Mbps) as shown in Figure 5.12 (bottom). It is worth noticing that high (local and even global) thresholds, e.g., 10 Gbps, may have also a negative effect as many of the amplification DDoS attacks do not send traffic at this rate.

Improved Reaction Time. A side benefit of using global information is that the amplification DDoS attack detection time is significantly improved. In Figure 5.13 we show that more than 80% of the attacks are detectable within 1 minute when the global threshold is

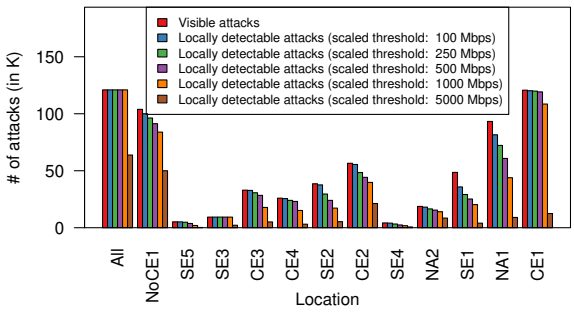

Figure 5.11: Number of attacks detected using local thresholds normalized by the average traffic volume of each IXP. Red bars annotate number of attacks visible at each IXP.

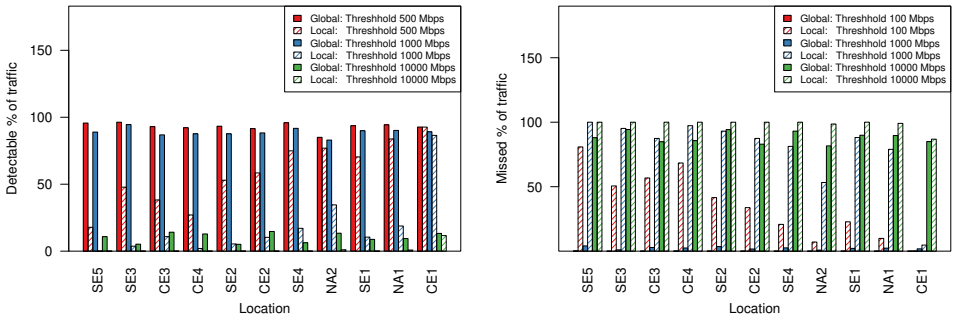

Figure 5.12: Traffic that could be detected and blocked with global information

100 Mbps. Around 70% (resp. 90%) of the attacks are detectable within 1 minute (resp. 5 minutes) with a global threshold of 250 Mbps. This allows to detect even short-lived attacks, i.e., the majority of the attacks that last for less than 10 minutes. Mitigation mechanisms are also more effective as they are activated earlier.

5.5.3 The Role of IXPs

Spare Capacity. IXPs are located in the core of the Internet. IXPs offer DDoS mitigation services, e.g., blackholing, as a free service to their members. Also, among their members are traffic scrubbing centers. Moreover, IXPs are peering infrastructures with very high capacity. To exemplify, the total capacity of the 11 IXPs we collaborate with is 65 Tbps, while their aggregated peak traffic of is around 11 Tbps. Thus, IXPs have spare capacity for absorbing and dropping even very large amplification attacks, at the scale of Tbps.

Proximity to Reflectors and Targets. To better understand the role that IXPs can play in defending against amplification DDoS attacks, in Figure 5.14 we plot the fraction of attack traffic that originates from reflectors with the relevant distance to IXPs in our study. To estimate the distance (in AS hops) from the reflector (IP) to an IXP we use routing information at the time of the attack, see § 5.2.3. Hop 1 refers to reflectors hosted in IXP members. Hop 0 refers to reflectors who's distance we could not estimate with our data. Recall, that the AS of IXPs is not visible in routing tables, thus, estimating the AS

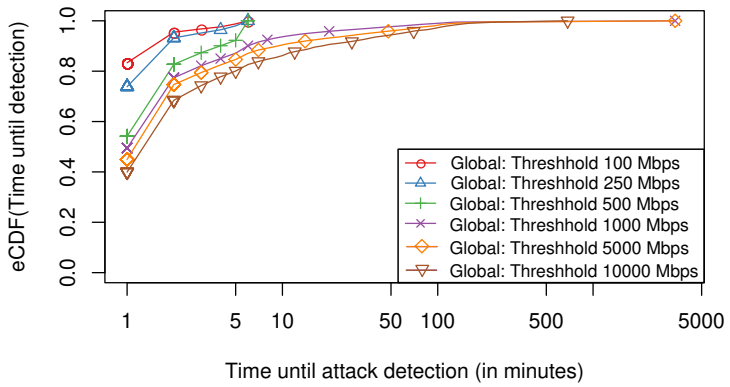

Figure 5.13: Time (in minutes) needed until an attack is detected in at least one site for different thresholds.

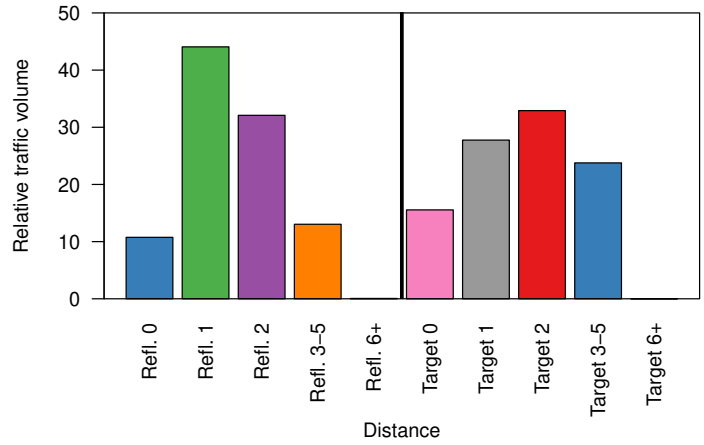

Figure 5.14: Fraction of attack traffic that originates from reflectors and target victim IPs with the relevant distance to the IXPs in our study. (Distance 1: corresponds to IXP members; Distance 0: we could not determine the distance.)

distance between reflectors and IXP is a complex task. More than 45% of the attack traffic originates from IXP members. This means that that by blocking the traffic at the IXP it is possible to drop attack traffic as close as possible to the reflector of the amplification DDoS. Another 30% of the attack traffic originates from networks that are two hops away from the IXPs, typically customers of the members of the IXP. Again, dropping this traffic will reduce significantly the attack traffic that is routed in the Internet as it stopped close to the source of the attack.

When we turn our attention to the targets of the attacks, see Figure 5.14, we also observe that a large fraction of the targets is relatively close to the IXPs. Around 30% of the amplification attack traffic targets IPs that are hosted in IXP members. This means that DDoS mitigation solutions can provide significant DDoS protection to IXP members.

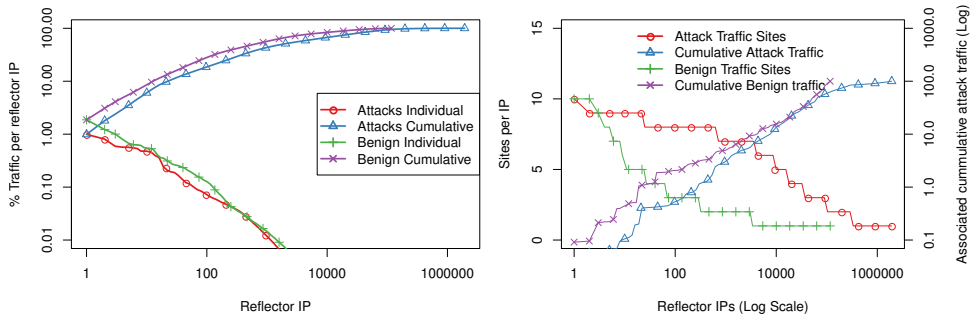

Figure 5.15: Attack traffic per reflector.

Moreover, by applying DDoS mitigation at the IXP, it is possible to reduce the AS-distance that attack traffic travels by one hop (see Hop 2 in Figure 5.14) for 35% of the attack traffic and 2-4 hops (see Hop 3-5) for 30% of the attack traffic.

Consolidation of Reflector and Target IPs. Another important observation derived by our analysis is that a relatively small number of reflector IPs are responsible for a large fraction of the attack traffic. In Figure 5.15 (top) we plot the attack traffic per reflector during for the 120k attacks we studied. Although there were more than 1.93 million reflector IPs identified, the top 1000 of them are responsible for about 40% of the attack traffic. This means that by blocking attack traffic from a relatively small number of reflector IPs yields significant reduction of the attack traffic. For the benign data, we see a similar image, in terms of relative traffic volume. In Figure 5.15 (bottom) we plot the number of our vantage points (sites) that observe reflector IPs as well as their associated attack traffic. Reflectors that are responsible for 50% of the attack traffic are visible at a minimum of three of our vantage points. This information can be shared across network infrastructures in a joined mitigation effort. In contrast, the benign data shows that less than 1% of the traffic is seen at 5 or more sites, and over 90% of the traffic is only visible at a single site. Moreover, it is noteworthy that the amount of source IP addresses involved in the benign data is much smaller than what we see for the attack traffic.

In Figure 5.16 we provide more insights on the consolidation of the reflectors with regards to the distance from the IXP. Again, a handful IPs are responsible for a large fraction of the attack traffic. For example, when we focus on the reflectors hosted in IXP members, some 200 IPs are responsible for more than 50% of the attack traffic that originates from direct members, which is around 22% of the total attack traffic.

5.6 Ethical Considerations

To comply with measurement ethics, we carefully design our study and take a number of measures that we describe next.

Traffic captures. Our study is based on traffic data that the IXPs regularly captured for operational purposes and are in compliance with legal requirements in the respective countries of operation. All traffic traces are aggregated at flow-level and thus do not contain any payload. Additionally, the data is processed and analyzed in-situ at the premise of the IXPs.

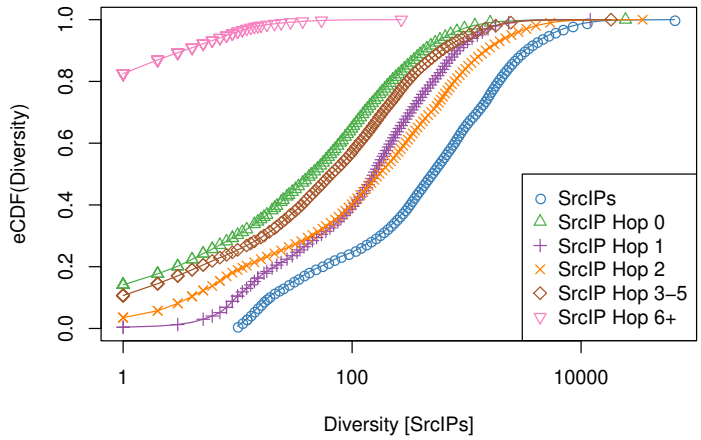

Figure 5.16: Attack traffic from reflectors based on the distance from IXPs that receive attack traffic. (Distance 1: IXP members; Distance 0: we could not determine the distance.)

Controlled self-attacks. Experience has shown that generating synthetic DDoS traffic in a real-world setup is hardly feasible. To obtain realistic traffic captures, we analyze traffic captures of self-attacks by following ethical guidelines and considerations that have been outlined for similar research studies [133]. Self-attacks run against a specially crafted autonomous system that belongs to the research infrastructure of one of the IXPs. The operator of the IXP executes the attack and several precautions are taken to limit potential negative effects of the attack. First, the IXP ensures that sufficient network bandwidth is available so that the likelihood of members being harmed by the targeted attack is minimized. Second, the IXP uses an experimental AS with no customer traffic and utilizes an unused /24 prefix that is allocated and announced only for the purpose of the experiment. While influence on external infrastructures (reflectors) cannot be completely avoided, the IXP captures the attack traffic to the infrastructure and continuously monitors the traffic sent by each reflector. The average traffic per reflector is typically between 500 kbps and 2 Mbps. The scope is limited by only purchasing the lowest possible low-volume attacks ($15) and further ensures no attack lasts longer than 5 minutes and the peak traffic is no more than 7 Gbps following recommendations from previous studies [133]. While contracting a booter service is a sensitive matter, the setup originates from a collaboration with law enforcement to study booter services and for operational tests. These include gaining insights into DDoS attack traffic from booters for ensuring operational safety (structure, link dimensioning, etc.), which is relevant at a national level where IXPs are considered critical infrastructure. During the self-attacks experiment, the IXP operator did not receive any complaints.

5.7 Related Work

The industry is currently investigating ways to exchange information among trusted parties to improve routing security towards a more resilient Internet. A global initiative backed by network operators, IXPs, content delivery networks, and cloud providers is MANRS [156]. "MANRS requires collaboration among participants and shared responsibility for the global Internet routing system" by sharing information for validation of

network announcements and registries, contact information for emergency situations, and anti-spoofing filters. DOTS [180] introduced requirements for enabling coordinated response to DDoS attacks. Our proposed DDoS Information Exchange Point can be used by participating partners to collectively fight against amplification DDoS attacks.

In the past, systems have been proposed to exchange information among networks to fight against DDoS, e.g., by using blockchain [209] or by introducing accountability to incentivize network operators to isolate sources of attacks in their networks [221]. Other proposed solutions are tailored to a small set of ISPs that are interconnecting with each other and have a relationship of customer-provider or peer [217]. Proposed systems enabled victims of DDoS to request attack monitoring and filtering on demand, and to pay upstream and remote ISPs for the services rendered [201]. Community efforts developed collaborative approaches to detect and neutralize botnets that participate in attacks [54, 141] and build collaborative IP blacklists [136, 164, 98]. These may suffer from shortcomings as they are not well maintained and sufficiently updated [165, 200].

At the national level, an anti-DDoS coalition have been formed. For example, in the Netherlands a national DDoS clearing house [52] is operational for collecting and sharing fingerprints of attacks and suitable mitigation rules among national network providers. Fingerprint extraction is done by dissecting pcaps of attacks, which may be shared through a DDoS database. The approach involves a considerable share of manual work and, to the best of the authors knowledge, there aren't any hard numbers on its efficiency. However, the project takes care of governance requirements and the legal implications of sharing sensitive data. Our approach is not a competitor of this project, but rather quantifies the potential of a distributed DDoS detection mechanism while extending the perspective to international vantage points.

More recent research on DDoS mitigation focused on full-blown scrubbing of traffic with programmable networking hardware, e.g., FPGAs [254] or P4 enabled switches [252]. These approaches aim more at applying fine grained filtering rules at scale to large amounts of traffic while staying flexible in adding, removing, or specifying new filters during operations. These approaches do not tackle distributed sensing of DDoS attacks nor do they tackle sensing at all and, thus, our proposed solution can complement and improve this new generation of DDoS mitigation platforms. The use of programmable networks was also suggested in [106] to enable verifiable in-network filtering for DDoS defense towards making IXPs or other involved infrastructures accountable in case of misbehavior.

5.8 Summary

DDoS attacks were first observed twenty years ago, but they are still one of the most serious threats. Amplification DDoS attacks have been repeatedly reported as both frequent and devastating, reaching 2 Tbps of attack traffic in recent years. In this chapter, we show that such amplification attacks are visible at multiple locations in the Internet. Unfortunately, the defense against such attacks is myopic and local today, and, thus, slow to react to attack and not effective, especially for short-lasting ones. We show that coordination in detecting and mitigating such attacks yields significant benefits, especially for smaller network infrastructures. In some cases, more than 80% more attacks and attack traffic can be detected and mitigated.

We also show that network infrastructures in the core of the Internet, such as IXPs, can drop attack traffic close to the location of the reflector, thus, reducing the distance that

attack traffic traverses only to either be dropped later or to create harm. We find that many networks can be relieved from carrying attack traffic. For more than 30% of the attacks, one network is relieved from carrying attack traffic. In another 30% of the attacks, exactly two networks are relieved. And finally, in more than 25% of the attacks, three or even more networks are relieved. This underlines the improved efficiency of DDoS attack mitigation in the core of the Internet using distributed visibility of ongoing global attacks.

The logic for collaboratively exchanging information about traffic characteristics can be implemented in P4. Moreover, the derived filters to drop the attack traffic can be integrated in the P4-based IXP concept introduced in Chapter 4 to create an IXP that further improves the resilience of the Internet.

Chapter 6
Meta-Telescope

The presented DDoS attack mitigation technique in Chapter 5 relies on an attack detection mechanism. So far, this mechanism is reactive; attack traffic is required to be observed at the collaborating IXPs to be detected. To improve this, we propose to use scanning activity as seen in telescopes. Unfortunately, the existing telescopes are limited in scope. Thus, we propose to use a meta-telescope.

Unsolicited traffic sent to advertised network space that does not host active services provides insights about misconfigurations as well as potentially malicious activities including the spread of Botnets, DDoS campaigns, and exploitation of vulnerabilities. Network telescopes have been used for many years to monitor such unsolicited traffic. Unfortunately, they are limited by the available address space for such tasks and, thus, limited to specific geographic and/or network regions. Even the institutions that dedicate a large address space to their Internet telescope [27, 166] are confined to what traffic is being sent to them. As attacks in the Internet typically have a highly distributed nature in terms of originating network type and continent, their preceding reconnaissance activity can barely be exhaustively collected by a single address block in a single network and country. To overcome this shortcoming, an Internet telescope that spans across multiple networks and countries is essential to gain information to help mitigate sophisticated attacks in the Internet.

In this chapter, we argue that telescopes do not need dedicated address space. Rather it suffices to focus on address space that is unlikely to be in use. Indeed, we observe that large parts of the advertised IPv4 address space are neither hosting users nor services. Thus, when such space is detected, the traffic sent to it is unsolicited background radiation. We refer to address space that capture such traffic as meta-telescopes.

By using central network vantage points we identify the largest and most distributed "meta-telescope" to date—consisting of more than 350k /24 blocks in more than 7k ASes. Using background radiation from the prefixes in the meta-telescope we highlight that unsolicited traffic differs by network/geographic region as well as by network type. Finally, we discuss our experiences and the challenges of operating a meta-telescope in the wild.

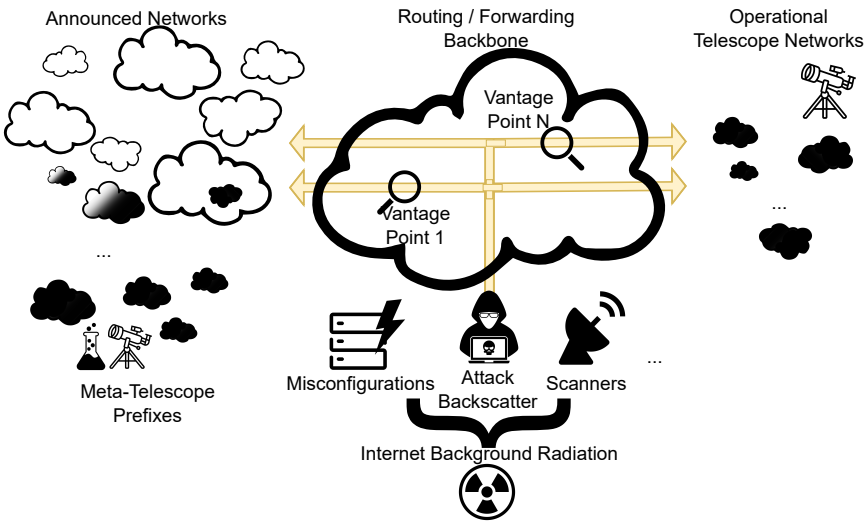

Figure 6.1: Sketch of the key ideas for identifying potential meta-telescope prefixes.

6.1 Intuition to Construct a Meta-Telescope

A telescope operator typically dedicates some of its allocated address space—thus one or a few prefixes—and tends to be limited to one geographic region and a few network locations. This address space has to be owned (or, at least, controlled) and advertised by the telescope operator. However, large and more distributed coverage is highly desirable, since certain IBR traffic components have been shown to be localized (with respect to the destination dark space) [207, 116]. The only known telescope that has been reported to be well distributed (spanning 1,300 networks) is operated by a major CDN provider and represents a variation of the original concept, since it leverages traffic reaching unused protocol ports on actually used (CDN) servers [207].

In this chapter, we introduce a novel concept to broadly capture IBR, which we call a "meta-telescope". A meta-telescope is based on the intuition that, with the availability of appropriate vantage points, one can (1) infer which address blocks on the Internet are unused and (2) capture traffic towards them—both without needing ownership or control of such address blocks. Our key observation is that a significant fraction of the Internet address space is indeed *advertised but unused* [206], i.e., it does not host users, servers, or other network equipment. These properties make such space ideal for monitoring unsolicited traffic destined to it that traverses accessible vantage points. From these intuitions, we develop and evaluate a methodology for identifying unlikely to be used Internet address space and build a "meta-telescope" that has very desirable properties, such as broad coverage of dark space both in terms of size and topological placement.

For an intuitive sketch of these key concepts see Figure 6.1: operational telescopes [27, 166] such as the ones shown on the right-hand side of the figure are dedicated network prefixes (black clouds) that are announced for the purpose of attracting IBR (bottom of the figure). The flow of the traffic throughout the Internet is shown as yellow arrows. At network vantage points it is possible to capture such traffic, e.g., as shown at Vantage Point 1. The advantage of monitoring in the middle of the Internet vs. at the edge of the Internet is that one may be able to observe traffic towards topologically and

geographically diverse destinations. The disadvantage is that one does not see all traffic towards each destination. On the left-hand side of the figure we show additional network prefixes (visualized as clouds). It is possible to use the vantage points to infer if any of these network prefixes originate any traffic. Using this criterion and auxiliary data sources, we show that it is possible to identify potentially-*dark* address blocks within prefixes (black clouds inside the white ones) or even entirely dark prefixes that do not originate any traffic (stand-alone black clouds). These black clouds represent the dark network blocks our meta-telescope monitors to capture IBR.

We propose, implement, and evaluate a pipeline of inference steps to apply to traffic data normally gathered by network monitors in order to identify *candidate* dark address blocks. We refer to them as meta-telescope prefixes. Meta-telescope prefixes can change over time and they are unlikely in the black lists of scanners or attackers. Moreover, the inference of meta-telescope prefixes helps us annotate flows that communicate with inactive address space. This enables us to scale up our ability to infer Internet background radiation as well as narrow down its originating networks at scale. The only requirement for our methodology is access to a network vantage point that regularly captures Internet traffic.

Our contributions can be summarized as follows:

- We develop and evaluate a methodology to identify globally advertised but unused address space to cumulatively contribute to a distributed meta-telescope. We refer to these prefixes as meta-telescope prefixes.
- Meta-telescope prefixes can be identified on demand according to various requirements regarding geographical footprint, network location, and address block size.
- Our analysis, using an Internet exchange point as vantage point, shows that in a single day more than 350k /24 IPv4 prefixes can be identified as meta-telescope prefixes. These are spread across more than 190 countries and 7,000 networks of various types, ranging from "eyeball" to data center to corporate. To the best of our knowledge, the cumulative size of meta-telescope prefixes we detect allows us to build the largest and most distributed telescope up to date.
- We comment on our experience in detecting and operating a meta-telescope in the wild. While spoofing can significantly affect our detection capabilities we show how to overcome this issue.
- We discuss how a meta-telescope can be used to shed light on scanning and other network activity across the Internet and geographical regions, as well as across types of networks.

6.2 Background and Related Work

For more than two decades, network telescope instrumentation—and more broadly, passively capturing and analyzing *unsolicited traffic* (IBR) at various vantage points—have allowed global visibility into a wide range of Internet phenomena: the automated spread of malicious software such as Internet worms or viruses [226, 174, 178, 175]; random spoofed source denial-of-service attacks [177]; large-scale botnet activities [67, 199]; macroscopic Internet blackouts due to natural disasters [64], network failures [12] and state censorship [68]; trends in IPv4 address space utilization [65, 66]; bugs and misconfigurations in popular applications [13], etc. Measurement and analysis of such

macroscopic phenomena are of key relevance for the security and reliability of the Internet infrastructure.

In the past, researchers have deployed dedicated infrastructure [250, 168, 59] for capturing traffic reaching large unused address blocks. However, due to the increasing scarcity and commercial value of IPv4 address space, the size of even the largest telescopes has been progressively eroding over the years. E.g., this is the case of the two largest telescopes broadly accessible to academic researchers: the UCSD Network Telescope [27]—once almost a /8 block—has been gradually shrinking as subnets are assigned by the owner, and recently saw a whole quarter of its addresses being sold [131]; the Merit Network telescope [166] has also progressively shrunk from originally a /8 to approximately the equivalent of a /13 due to steady sub-allocations.

For these reasons, researchers have also started developing ways to observe unsolicited traffic leveraging existing infrastructure. In 2012, Glatz et al. developed a scheme to dissect *one-way* (thus unsolicited) traffic observed in unsampled NetFlow records from the border routers of a regional academic backbone network. This traffic includes therefore packets destined to active hosts but towards ports where they do not host services or run clients. Unsampled traffic capture at medium-sized (or larger) networks is impractical and the infrastructure used in that study has been discontinued.

In 2019, Richter and Berger presented a "distributed telescope" approach leveraging existing logs of unsolicited packets blocked at the firewalls of $\approx$ 90,000 servers of a major CDN [207]. These servers are distributed over more than 1,300 networks and are live, offering services to end users. The distributed nature of this setup enabled the authors to uncover phenomena that cannot be captured by individual telescopes placed in only one location of the Internet topology. Specifically, they found evidence of local concentrations of unsolicited traffic, a phenomenon also recently observed by Hiesgen et al. [116] when comparing traffic from telescopes in Europe and the US.

The importance of widely distributed capture approaches and of sensor placement as a key factor in understanding and generalizing measurements from unused address space, were already highlighted in foundational studies from 20 years ago [177, 53, 168]—when large portions of unused IPv4 space were more accessible to researchers. The authors were able to deploy an "Internet Motion Sensor" (a hybrid telescope / honeynet-like infrastructure) consisting of 28 monitored blocks at 18 physical installations. Since then, researchers have rarely had access to largely distributed telescope infrastructure. In contrast, access to largely distributed honeynets is more common, but they represent infrastructure with different goals and characteristics—i.e., targeting specific classes of phenomena (e.g., malware, bruteforcing, exploits, etc.) and with each individual vantage point typically covering only one address or a small block [118].

The original approach we propose aims at (1) leveraging existing infrastructure (a meta-telescope operator does not need to own and allocate address space) while (2) enabling coverage that is both broad in size and diverse in terms of topological placement. In addition, (3) our approach offers the opportunity to identify untapped portions of unused IPv4 address space across the whole Internet. This feature also brings (4) another advantage: large individual telescopes tend to become notorious and their address blocks are often blacklisted by scanners and malicious actors [9]; by leveraging uncovered dark address blocks, our meta-telescope also promises to be more resistant to blacklisting.

IXP Code	#Members	Peak Traffic (Gbps)	Region	#Sampled Flows (Billions)
CE1	1,000+	12,000+	Central Europe	68.461
CE2	250+	150+	Central Europe	0.904
CE3	200+	150+	Central Europe	0.381
CE4	200+	150+	Central Europe	0.492
NA1	250+	1,000+	North America	8.471
NA2	125+	600+	North America	2.379
NA3	20+	10+	North America	0.031
NA4	20+	50+	North America	0.159
SE1	200+	1,000+	South Europe	2.807
SE2	10+	200+	South Europe	0.978
SE3	40+	50+	South Europe	0.23
SE4	40+	300+	South Europe	1,146
SE5	20+	10+	South Europe	0.179
SE6	30+	15+	South Europe	0.049

Table 6.1: IXPs: Basic statistics—week of April 24th 2023.

6.3 Data Sets

In this section we describe the vantage points and data sets used in our study. Specifically, Section 6.3.1 discusses the multiple vantage points we use to infer dark network blocks in the whole public IPv4 Internet and to characterize the IBR they receive. In our study we also use traffic data from three operational telescopes and an operational network that we were granted access to. We analyze specific properties of their traffic to inform several of our parameter choices in our inference methodology such as tuning of thresholds for average packet sizes (see Section 6.4). We describe these data in Section 6.3.2. Finally, Section 6.3.3 lists auxiliary data sets we used such as IP geolocation data or data that allow us to identify with certainty some active (i.e., non dark) subnets.

6.3.1 Network Vantage Points: IXP Sites

We partner with 14 Internet Exchange Points (IXPs) that have established Internet peering infrastructure in 3 regions of the world, i.e., North America, Central Europe, and South Europe. Three of these IXPs are the same as in our study in Chapter 3. These are CE1, NA1, and SE1. Eight further IXPs, namely CE2, CE3, CE4, NA2, SE2, SE3, SE4, and SE5, are the same as in our study in Chapter 5. Many of the large cloud providers and content delivery networks are members of the IXPs, as well as enterprise networks and regional or national eyeball networks [1, 33]. The size of the IXPs in our study varies in terms of number of members, i.e., peering networks, as well as peak traffic. For an overview of the IXPs in our study see Table 6.1.

For our study, we get once again access to network data collected at each of the 14 IXPs. The data is exported through the IPFIX protocol [44, 43] and contains aggregated packet header information about network flows of the IXPs. It does not contain any packet

Code	Location	Size (#/24s)	Daily /24 pkt count	Share of TCP traffic	Avg. IP pkt size (TCP)
TUS1	North America	1856	1.91M	93.82%	40.7B
TEU1	Central Europe	768	1.79M	90.38%	40.55B
TEU2	Central Europe	8	2.29M	79.5%	40.78B

Table 6.2: Operational telescopes: Basic statistics.

payloads. The flows are generated on a packet sampling. Collection of the flow data took place in the week from April 24th, 2023 to April 30th, 2023. The total amount of sampled data at all IXPs is 86,667 billion flows estimated to carry about 880 Petabytes of traffic.

6.3.2 Operational Telescopes

We obtained access to data from three network telescopes that are operated by three different organizations in three different countries. We use these telescopes to inform several of our parameter choices in our methodology (see Section 6.4), such as tuning of thresholds for average packet sizes. This type of data-driven indicators help us differentiate between "active" versus "dark" /24 networks. We also use these operational data sets to gather insights regarding top-targeted ports. We then juxtapose these observations with the ones obtained through our meta-telescope to shed light into spatial differences in scanning.

We analyze traffic for the week[11] April 24 – April 30, 2023 (see Table 6.2). We note that the TEU2 telescope only became operational during the course of our study. In addition, the network hosting this telescope is directly peering at ten of the IXPs and has transit connectivity via a tier-1 provider. The largest telescope in our study consists of 1,856 contiguous /24 subnets and each /24 subnet receives an average of 1.91 million packets per day. The share of TCP traffic is 93.8%. Similar results hold for the TEU1 telescope. The TEU2 telescope receives more UDP traffic than the other two and its share per /24 is also larger. Two of the telescopes receive traffic on any TCP/UDP port; for the TEU1 telescope ports 23 and 445 are blocked by their ingress router. It is noteworthy that some of the 768 /24 blocks in TEU1 are dynamically allocated to end users on a daily basis so that not all blocks are always actually dark. Table 6.2 also shows that the total packet count of IBR per /24 does not vary drastically. It typically lies to around 2 million packets per day per /24. (For TEU1 it is less due to some ports being blocked, as mentioned earlier.) We leverage this information in our inference methodology.

To better calibrate our inference approach to distinguish dark vs. active subnets, we also leverage traffic data from the same ISP that hosts the TUS1 telescope. Specifically, we use NetFlow records for the same week of April 24 – April 30, 2023. The ISP receives traffic for 26,079 unique /24 subnets, including those of the TUS1 telescope.

6.3.3 Auxiliary Data Sets

In addition to traffic data from IXP vantage points and operational telescopes/networks, we use a variety of data sets to either validate and supplement our inference pipeline or

[11] We verified that we find consistent values for each individual day.

analyze our results. For a comprehensive statement on ethical considerations, we refer to Section 5.6.

BGP Routing Data.

As a reference and starting point for our inferences, we use the portion of the IPv4 address space that is actually reachable through BGP. To obtain a complete list of prefixes announced in the global routing system, we use routing table (RIB) dumps from a large RouteViews collector (route-views4). For each day we consider, we combine the prefixes from all 12 RIB dumps available (RouteViews RIBs are dumped every 2 hours).

In addition, in the analysis of our results we characterize the portion of inferred dark address space of individual Autonomous Systems (ASes). To this end, we use CAIDA's prefix-to-AS mapping data set (published daily) from April 24, 2023 [25], which is based on RouteViews RIB dumps. We also use CAIDA's AS to Organization mapping data set [26] (published every few months) from April 11, 2023. This data set is generated using WHOIS information available from Regional and National Internet Registries to infer a mapping from AS numbers to the organizational entities that operate them.

Data on active IP addresses.

To partially validate our results and, afterwards, as a supplemental source of information, we use data from three measurement projects that confirms "liveness" of individual IP addresses: M-Lab's Network Diagnosis Tool (NDT), Censys, and ISI's Internet Address History. The NDT Speed Test is a single-stream performance measurement, initiated by users, of a connection's capacity for "bulk transport". We extract from the NDT "Unified Views" data the list of IPv4 addresses that perform speed tests and mark their /24 IPv4 prefixes as "used" on a daily basis for the week of April 24 – 30, 2023. The Censys Universal Internet [79] data set is generated by scanning the entire IPv4 address space on multiple ports and protocols on a daily basis. We identify roughly 4.8 million active /24 IPv4 prefixes using this data for the week April 24 – April 30, 2023. Finally, the history bits data from Internet Address History data set [236], generated on March 6, 2023, contains IPv4 addresses that respond to ICMP echo requests for scans starting from 2019. We aggregate data for the February, 2023 scan to generate a set of roughly 5.1 million /24 IPv4 prefixes, where we observe each prefix to contain at least 1 address which responded to the ICMP echo requests.

IP Geolocation and AS classification.

We geolocate—at a country level—the prefixes that our methodology identifies as advertised through BGP but unused using IP geolocation data from the Maxmind GeoLite2 data set [159] from April 25, 2023. To classify ASes into business categories, we use the IPInfo "IP to Company" commercial database from April 23, 2023.

6.4 Methodology

In this section we describe our methodology and inference pipeline for identifying candidate meta-telescope prefixes. Since this study is the first attempt at developing a "meta-telescope" (i.e., a first-of-its-kind analysis) both our filter logic as well as the

thresholds are chosen conservatively to ensure that we can identify meta-telescope prefixes with a high confidence while keeping the number of false positives low. Unless stated otherwise, all numbers in this section are for all 14 IXPs and the 24-hour period from April 24th, 2023 00:00 UTC through April 25th, 2023 00:00 UTC.

6.4.1 Telescope Traffic Analysis

A key intuition of our inference methodology is that traffic towards dark address blocks (IBR) is likely to present different characteristics from traffic towards populated address blocks (i.e., hosting services, clients, etc.). We indeed find that the size of the vast majority of IBR packets tends to be minimal (i.e., typically just made of the IP and TCP header). To derive a packet size threshold to use in our inference pipeline, we perform a detailed analysis of this property by leveraging traffic data from the operational telescopes and the ISP described in Section 6.3.2.

In the last column of Table 6.2 we report the average IP packet size we observe for TCP packets captured at each of the three operational telescopes: in all three cases, the average value is smaller than 41 bytes (and greater than the minimum size of 40 bytes, that is, 20 bytes each for IP header and the TCP header respectively). We also verify that this property is consistently present when we separately aggregate packets by /24 block. We find that *at least* 93% of all TCP packets destined to the telescopes have a size of 40 bytes. We note that 40 bytes correspond to a typical `TCP-SYN` packet, i.e., a packet with no IP or TCP options set nor any payload content. In addition, we see a step at 48 bytes. This is again a typical size of `TCP-SYN` packet with one option. Large amounts of packets attempting to open a TCP connection are often originated by malware trying to compromise new hosts, malicious actors performing network reconnaissance, or even benign research scanners [80, 7]. Therefore, we speculate that utilizing the TCP packet size as an indicator or "fingerprint" for distinguishing between "dark" and "live" traffic is a useful filtering step in our inference methodology.

To confirm this hypothesis, we look at (both directions of) the traffic traversing a production network containing both dark and active subnets. Specifically, we analyze NetFlow traffic data captured at the border routers of the same ISP that operates the `TUS1` telescope (described in Section 6.3.2). Looking at traffic destined to the ISP, we find 26,079 unique /24 subnets that receive traffic (including the /24 subnets of the `TUS1` telescope). On the other hand, only 7,923 /24 subnets out of the 26,079 ones are seen to be originating traffic within that ISP. Hence, for the week of interest, there are 18,151 (i.e., 26,079 - 7,923) *dark subnets*—including the 1,856 /24 subnets dedicated to the `TUS1` telescope—that by definition are not used. To identify *active subnets* we focus on the 7,923 networks that we see activity from. To filter out networks that may be considered active because of spoofed traffic, we impose the conservative constraint that a network is considered active only if we have observed at least 10 million packets originating from that subnet during the course of the full week (based on our observation of the distribution of per subnet packet counts). With this constraint at hand, the number of active /24 subnets identified drops to 5,835.

Now that we have obtained "labels" for active and dark subnets, we can select and tune a classification criterion to optimally distinguish between the two classes. For this work, we evaluated two features (see Table 6.3): (1) median packet size and (2) average packet size destined to a /24 subnet. The classification rule simply checks whether the median (average) packet size destined at a /24 subnet is less or equal than a threshold of N bytes. If so, that /24 subnet is considered to be "dark". Otherwise, it is classified as "active".

Median Packet Size	Classified Dark, but are Active	Classified Active, but are Dark	Classified Dark, and are Dark	Classified Active, and are Active	F1-score
Threshold (bytes)	False Positive Rate	False Negative Rate	True Positive Rate	True Negative Rate	
40	6.96%	0.39%	99.61%	93.04%	98.70%
42	8.00%	0.39%	99.61%	92.00%	98.54%
44	22.59%	0.09%	99.91%	77.41%	96.45%
46	22.64%	0.09%	99.91%	77.36%	96.44%

Average Packet Size	Classified Dark, but are Active	Classified Active, but are Dark	Classified Dark, and are Dark	Classified Active, and are Active	F1-score
Threshold (bytes)	False Positive Rate	False Negative Rate	True Positive Rate	True Negative Rate	
40	0.00%	99.10%	0.90%	100.00%	1.83%
42	0.53%	56.83%	43.17%	99.47%	60.25%
44	**0.87%**	**0.41%**	**99.59%**	**99.13%**	**99.65%**
46	1.08%	0.27%	99.73%	98.92%	99.69%

Table 6.3: Tuning the packet-size based fingerprint that allows distinguishing between "active" versus "dark" /24 subnets using ISP data.

We experimented with various threshold choices for N, and the sensitivity analysis of Table 6.3 showcases our results. A good classification outcome is reached when using the *average packet size* criterion with a threshold of 44 bytes. This setting achieves very high accuracy and, importantly, a very low false positive rate. This value provides the second-best F1-score[12]. Nevertheless, between the two best thresholds (44 versus 46 bytes average packet size), we opt to use the 44 bytes threshold due to its lower false positive rate.

6.4.2 Inference Pipeline

Our inference pipeline consists of seven steps (Figure 6.2) to exclude address blocks that are unfit (e.g., reserved space) or whose traffic does not exhibit typical IBR characteristics. At a high level, we seek the following characteristics for the meta-telescope prefixes: (1) they are routed and advertised but not reserved for special purposes [56]; (2) they do not originate any traffic; (3) the size of their incoming TCP packets is small, i.e., with an average smaller than 44 bytes (recall 6.3.2); and (4) they do not receive too much traffic.

We start with the roughly 6 million IPv4 destination /24 subnets that are included in the IXP traffic data set and apply the following filtering steps:

1. **TCP Traffic.** As noted earlier, `TCP SYN` packets are very common in IBR and UDP is very noisy. As such, we remove any subnet that does not receive any TCP traffic. About 300k /24 subnets are filtered out in this step.

2. **Average Packet Size.** Moreover, as `TCP SYN` packets dominate background radiation TCP traffic, we remove any subnet that receives TCP traffic with an average packet size above 44 bytes. This filter removes about 800k /24 subnets from the data set.

3. **Source Address Unseen.** Next, we require that potential meta-telescope prefixes do not originate any traffic. Hence, we remove any subnets from which we observe traffic. Some 100k subnets are removed in this step. While this is a straightforward filter, it is susceptible to source address spoofing, thus also removing subnets that do not actually generate traffic (but whose addresses were used in spoofed packets). Yet, the likelihood that the sampled IXP traffic includes scanning packet towards an address and in the same period spoofed packets "from" that address is rather low. However, as we increase the duration of the data set, this issue has to be addressed, since the likelihood of including spoofed packets increases.

[12]The F1-score measures the overall classification accuracy, and is defined as $F_1 := 2\text{tp}/(2\text{tp} + \text{fp} + \text{fn})$, where tp denotes the true positives, fp the false positives and fn the false negative results.

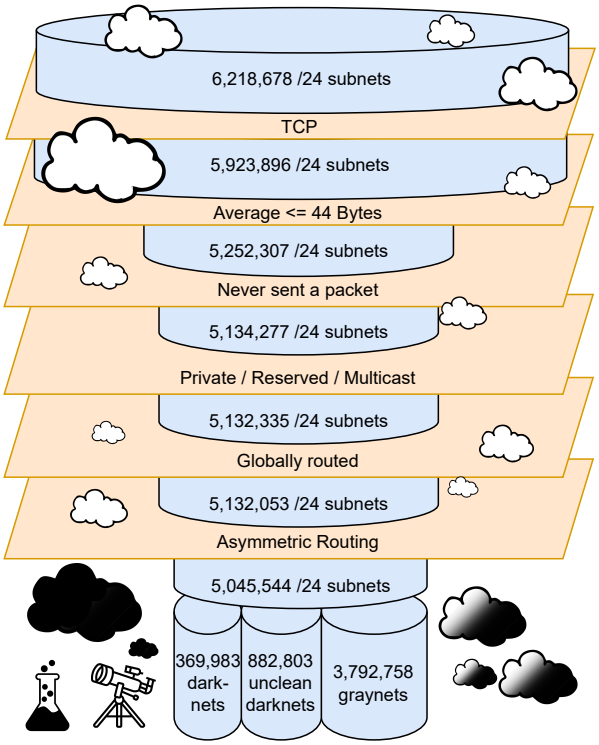

Figure 6.2: Illustration of inference pipeline: # of /24 blocks—all IXPs and 24th April 2023.

4. **Private / Multicast / Reserved.** Telescopes must be reachable in the public Internet. Hence, we remove all IP blocks from the IXP data set that are inside private IPv4 blocks, multicast IPv4 blocks, or reserved IPv4 blocks. About 2 thousand /24 subnets are removed by this filter.

5. **Globally Routed.** Our next filter ensures that candidate dark prefixes should not only be located inside a routed address block, but also need to be publicly announced. Here, we rely on daily snapshots from RouteViews [235] and remove any /24 subnet that is not inside an announced prefix as seen by RouteViews. About 300 subnets are hereby removed.

6. **Asymmetric Routes.** Another challenge is due to asymmetric routing in the Internet. For example, Content Delivery Networks (CDNs) often receive lots of `TCP ACK` packets from an IP, but send their data to them via another path not visible at any of the IXPs. These packets will typically be 40 bytes long similarly to the `TCP SYN` packets targeted by our average-packet-size filter. We leverage the fact that IBR is limited compared to production traffic [197] and apply a conservative volume-based filter of 1.7M packets on average per day per /24. This step filters out 90k /24 blocks.

7. **Classification.** Finally, we classify all /24 blocks into three classes: (1) dark (i.e., meta-telescope prefix), (2) unclean darknets, and (3) graynets. For a block of IP addresses to be a meta-telescope prefix, all IPv4 addresses have to survive the above filter steps. Unclean blocks are those that have at least one IP surviving the filter steps and at least one IP that did not survive the filters but did not originate traffic. If one

IP inside a block of IPs originates traffic in which another IP survived our filters, we consider this a graynet.

After the last filtering step, we are left with 5M /24 blocks. Our classification pipeline labels about 370k of them as dark (potential meta-telescope prefixes), 880k as unclean, and 3,8M as gray.

6.4.3 Evaluation

We evaluate the results of our inference in three ways: (1) We verify its ability to identify the address space used by known telescopes; (2) We compare port count statistics from the traffic we observe towards our inferred dark prefixes against traffic observed at operational telescopes; (3) We use three data sets of host activity/responsiveness to identify (a lower bound for) false positives in our inferences. In addition, we use results from this last comparison to further refine and finalize our list of meta-telescope prefixes.

To assess the effectiveness of our inferences, we check if we can infer as dark the address space of the three operational telescopes we have been granted access to (see Section 6.3.2). Table 6.4 summarizes our findings. Using 7 days of data from `CE1` we can infer as dark 31.15% and 87.5% of the address space of `TEU1` and `TEU2`, respectively. The remaining address space is inferred as either unclean or gray due to spoofing or we did not see any traffic to it. We cannot find the address space of the `TUS1` telescope, as it is not visible at this IXP. However, using data from all IXPs we can infer as dark 23.5% of `TUS1`'s address space in a single day and 77% of it in a week. Notably, `TEU1` contained 503 active /24 blocks on the day that we report on and, thus, the 38 inferred blocks correspond to 14.3% of the unused space in that telescope.

Overall, we find that our conservative approach of using a threshold of 1.7M packets per /24 per day eliminates quite some /24s of the telescopes. Indeed, it might not necessarily be the ideal choice, as `TEU2` receives more than 2.2M packets on average. Furthermore, since the telescope's address space is directly announced at 10 of the vantage points, its traffic is well observed. That is why we are unable to infer a single /24 inside `TEU2` even with data for several days. Still, even with this conservative threshold, we can infer a large amount of dark space while not having to deal with many false positives (as we show at the end of this section).

Code	Size (/24s)	#Inferred meta-telescope prefixes			
		1 day		7 days	
		`CE1`	All	`CE1`	All
`TUS1`	1856	0	437	0	1424
`TEU1`	768	38	33	262	247
`TEU2`	8	0	0	7	7

Table 6.4: Meta-telescope coverage of IPv4 address space of the operational telescopes for 1 and 7 days.

To further underline our ability to detect potentially-dark address space, in Figure 6.3 we plot a Hilbert curve of an IPv4 address block which contains the address blocks of a network telescope. Every pixel corresponds to a /24 block of the IPv4 space. The colorized pixels correspond to inferred dark blocks. Uncolored pixels correspond to blocks without data or that are inferred as unclean or gray. The boundaries of the

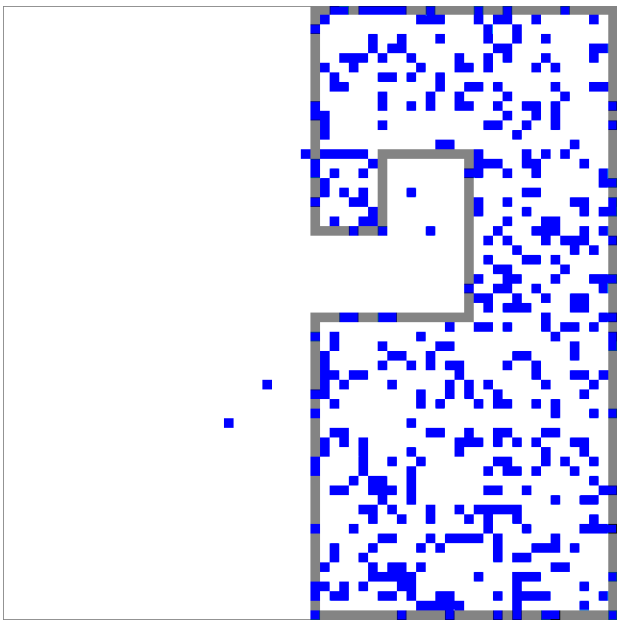

Figure 6.3: Hilbert curve of IPv4 address space colored by meta-telescope annotated with a gray box for the address space of an operational telescope.

operational telescope are marked in gray and we clearly see that almost all blue pixels correctly fall within this area. That there are a few, i.e., 5, outside is also not surprising, since not all of the address space outside this block must necessarily be in use.

We then compare port statistics of packets towards the inferred dark prefixes against those from operational telescopes. Table 6.5 shows statistics we extract for the top-10 TCP ports by analyzing raw PCAP data collected from the three telescopes. We observe a high degree of similarity between the three sites but also some notable differences. Ports 22, 80, and 443 are in the top list of all three. However, port 6379, a top-5 port in `TUS1` and `TEU2`, does not appear in `TEU1`. This highlights the importance of studying scanning traffic destined to multiple, distributed vantage points, a significant advantage of our meta-telescope approach. When comparing these numbers against those from traffic we observe at the IXP vantage points towards the blocks we infer as dark, we find a perfect overlap for the top ports, namely 22, 23, 80, 443, and 8080.

The above results confirm our ability to identify dark blocks (true positives) and suggest that overall the traffic towards these prefixes is consistent with high-level properties of IBR. To complete our evaluation, we lastly assess the presence of false positives (i.e., blocks inferred as dark whereas they are active) to the extent that publicly available data sets allow for. The indications of activity we use are: (1) hosts replying on any transport port upon contacting as reported by Censys [79], (2) end-users performing speed tests of their Internet connectivity as reported by NDT [163], and (3) hosts replying to ICMP echo requests as reported by ISI [236]. We aggregate this data at /24-block granularity and we compare it with our initial list of 369,983 blocks we infer as dark. We find that 51,337 out of these 369,983 /24 blocks have been active, i.e., 13.9% of our inferred-dark /24 blocks are (at least for some IP addresses) active. This (positive) result shows there is still significant room for improvement to our pipeline in terms of false positives[13]. However,

[13]Especially considering that these three data sets can only provide a lower bound for active networks, i.e.,

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	23	22	23
#2	6379	80	22
#3	22	443	80
#4	80	8080	6379
#5	443	3389	445
#6	8080	5555	25565
#7	25565	60023	443
#8	5555	81	8080
#9	3389	8443	8090
#10	60023	2375	3389

Table 6.5: Operational telescopes: Top 10 ports in descending order by popularity (April 24-30, 2023).

we can apply such active networks ground-truth data to further filter our inferences and obtain a more accurate final set of meta-telescope prefixes. In the remainder of the chapter our analysis is based on this final set of prefixes. Table 6.6 summarizes our results in terms of inferred dark /24 blocks after applying this final correction (by individual vantage points and overall).

6.4.4 Limitations

The IXP vantage points used in our study impose some limitations to our results.

Sampling. The IXP data set is generated based on packet samples. To gain statistical significance a lot of packets may be required, which is one reason for some of the variations in our results, e.g., when looking at data from different days. Indeed, high-volume DDoS attacks are much easier to capture than low volume background radiation. However, the large amount of addresses being scanned and the possibility to extend our inference to arbitrarily long time frames, help overcome this limitation. We further discuss the impact of sampling in Section 6.6

Routing. Another limitation is that we can only see traffic routed via the vantage point. Note that IXP customers may use alternative routes for part of their traffic, which, for example, restricts the visibility of some of the smaller IXPs.

Asymmetric Routing. Given the prevalence of asymmetric routing, one cannot presume that the observation of traffic in one direction implies that the traffic in the reverse direction is routed over the same path. Hence an IP address block that appears dark, may not be so, as the traffic might be routed via a different path. We tackle this challenge in two ways. First, we only consider networks that receive less than 1.7M packets per day. Second, we eliminate all IPs that appear to be active in any of our auxiliary data sets.

Locality. Given that IXPs pursue the motto: "keep local data local" [33], the data set may be prone to geographical bias. However, (especially) the larger IXPs offer many services that are attractive also to peers from other regions, including ease of connectivity and remote peering [30]. Moreover, hypergiants [21], including cloud providers, often peer

they do not necessarily identify all /24 blocks on the Internet that have at least one active IPv4 address.

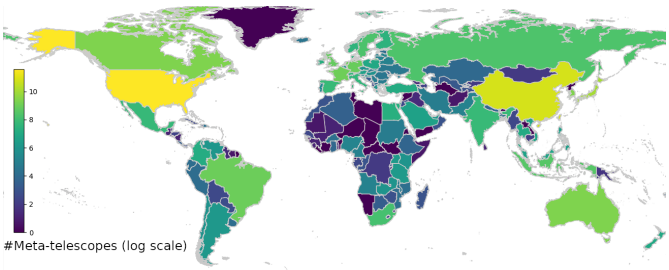

Figure 6.4: World map colored according to the number of /24 blocks in meta-telescope prefixes (logarithmic scale).

at IXPs.

Spoofing. Unfortunately, spoofing, e.g., sending packets with an incorrect source IP address, is quite common in the Internet [152]. Spoofing impacts our methodology, since these packets may use source IP addresses of potential meta-telescope prefixes and, thus, will disqualify the whole block. As we expand our data set, e.g., by expanding the observation period or by adding more vantage points, we likely look at more spoofed traffic. As such, the amount of inferred meta-telescope space decreases. To counteract this behavior, we will allow for a small number of potentially spoofed packets, see Section 6.6.

6.5 Meta-Telescope Properties

In this section, we analyze the properties of the meta-telescope we built. For this analysis, we use the ipinfo and pfx2as data sets (see Section 6.3) to determine the geographic location of the meta-telescope prefixes and to which ASes they belong. This analysis helps us understand the basic properties of the meta-telescope as well as obtain insights about the geographic distribution of the inferred prefixes and their network types.

Basic Properties

Using our methodology we are able to identify a very large number of /24 blocks as meta-telescope prefixes, namely up to 318,646 /24s in a single day using all vantage points. These are originated by more than 7,000 ASes, which are in turn located in almost 200 different countries (see Table 6.6).

We find variance in the results depending on the vantage point. Even though some of the vantage points have a very limited visibility on Internet traffic, they are still useful in helping us to identify more than 250 meta-telescope prefixes in various regions of the world while larger vantage points are able to infer thousands of meta-telescope prefixes around the globe. Note that, when combining multiple vantage points, we obtain a smaller number of prefixes compared to the largest individual contributors (`CE1` and `NA1`). This is because more information about individual /24s is provided to our filtering criteria, which are also designed to be conservative and prioritize low false positives.

Using the inferred meta-telescope prefixes from all IXPs, in Figure 6.4 we show a visualization of the distribution of the meta-telescope prefixes across the world in a logarithmic scale.

IXP	#Inferred meta-telescope prefixes	#ASes	#Countries
CE1	397,000	8,529	201
CE2	21,340	1,597	124
CE3	61,607	3,982	173
CE4	2,178	455	84
NA1	395,585	8,960	198
NA2	12,489	919	102
NA3	262	128	17
NA4	1,054	299	74
SE1	34,222	2,269	152
SE2	56,638	2,078	132
SE3	3,782	729	97
SE4	43,573	2,431	152
SE5	1,949	667	104
SE6	270	104	33
All	318,646	7,195	194

Table 6.6: Overview of the meta-telescope prefixes we identify (individual vantage point and overall).

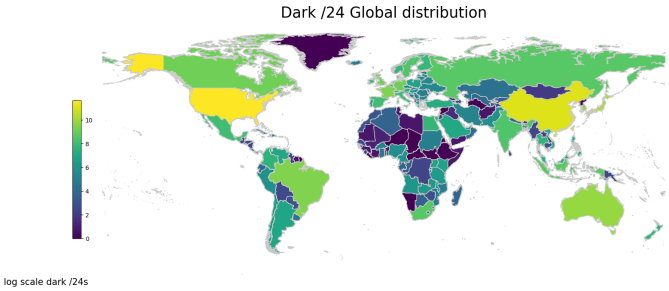

Figure 6.5: World map colored according to the number of /24s meta-telescope blocks as seen by CE1 (logarithmic scale).

For comparison, we show the maps for IXPs CE1 (Figure 6.5), NA1 Figure 6.6, and all IXPs (Figure 6.7) We note the different color scaling for NA1 as the number of inferred meta-telescope prefixes is higher for this vantage point. We find that all vantage points infer meta-telescope prefixes in all regions of the world. Once again, most inferred meta-telescope prefixes are located in the USA. Vantage point CE1 has the best visibility of meta-telescope prefixes that belong to China. We see that the inferred meta-telescope prefixes are located in IP spaces belonging to almost every country. This includes even small countries that we usually do not have network telescope insights for, given that none of the operational telescopes that we are aware have presence in such countries.

Most meta-telescope prefixes are located in the USA according to our geographic mapping. We find that NA1 which has the "best" visibility of the traffic in that region is able to infer the largest number of meta-telescope prefixes in the USA. By adding information from other vantage points, the number of inferred meta-telescope prefixes in the USA decreases slightly. This is likely due to the impact of spoofed traffic. That the USA is dominating is likely due to the fact that a large fraction of address space was allocated to

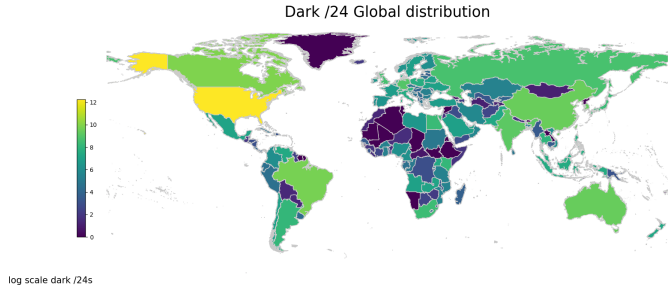

Figure 6.6: World map colored according to the number of /24s meta-telescope blocks as seen by NA1 (logarithmic scale).

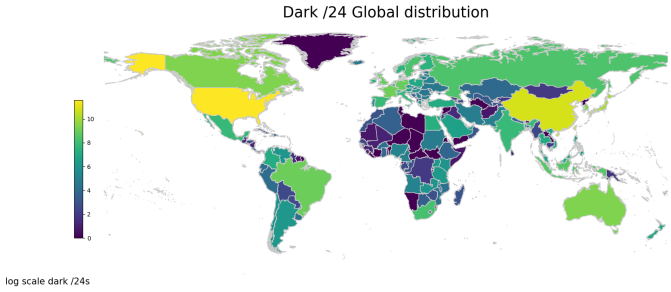

Figure 6.7: World map colored according to the number of /24s meta-telescope blocks as seen by all sites combined (logarithmic scale).

US-organizations in the early days of the Internet. Most of these large, i.e., mainly /8 blocks, seem to remain primarily unused. To our surprise, the country ranked second by the number of inferred meta-telescope prefixes is China. Visibility into prefixes within China is one example that highlights the benefits of the proposed methodology, since this can allow researchers to study scanning and other macroscopic activities destined to usually unobserved address space. Regions that are still not well covered include central Africa, some middle eastern regions, and North Korea. To overcome this aspect, one might need vantage points closer to these regions.

Meta-Telescope Prefixes: Examples

Next, we look at some example meta-telescope prefixes. The first is about a /9 block; the second is about a large network telescope subnet we are aware of.

In Figure 6.8, we plot the Hilbert Curve of a /8 where an inferred /9 meta-telescope prefix is inferred using data from three different vantage points, namely (1) CE1, (2) NA1, and (3) all vantage points combined. Here, colored pixels refer to /24 address blocks of inferred meta-telescope prefixes, whereas white pixels correspond to /24 blocks that are either gray or unclean or for which we have no data. We find that CE1 has great visibility of the /9 block in the right half of the Hilbert Curve. However, the visibility of the left half is not as good. This is partially due to it containing some unannounced space and partially due to lack of traffic.

The NA1 vantage point does not infer a single /24 block of the right /9 block as meta-telescope prefixes. However, it identifies the blocks inside the left /9 as meta-telescope

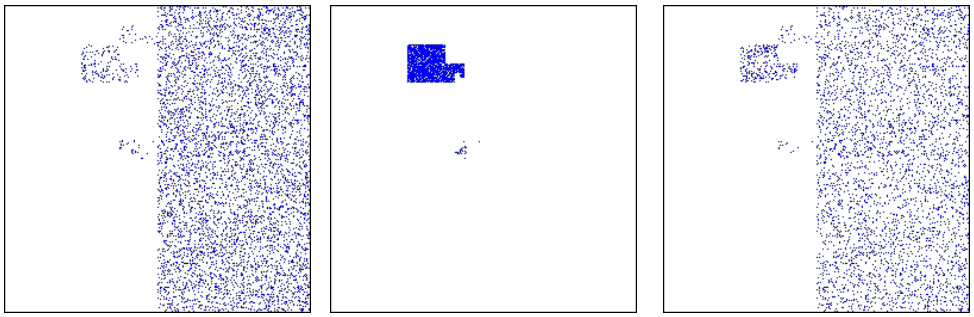

(a) CE1 (b) NA1 (c) All sites

Figure 6.8: Hilbert Curves of a /8 with colored meta-telescope prefixes.

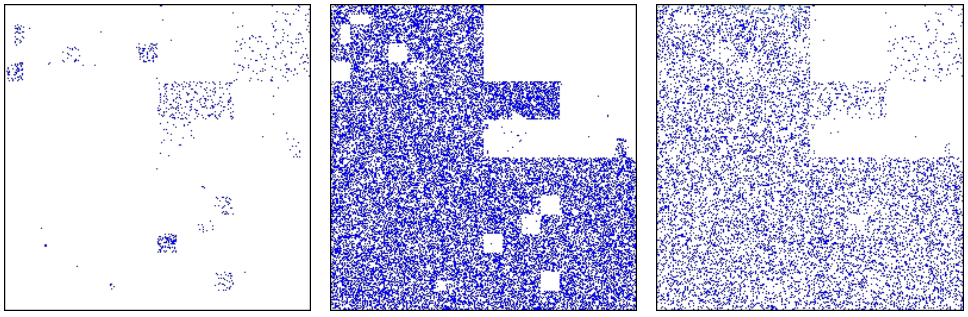

(a) CE1 (b) NA1 (c) All sites

Figure 6.9: Hilbert Curves of a /8 which contains a known telescope with colored meta-telescope prefixes.

prefixes. These coincide precisely with the ones from CE1. When combining the data from all vantage points, the /9 is again visible, however with a slightly lower density. The latter is likely due to the added noise of spoofing which causes some /24 address blocks to be classified as gray. Note, we do not know whether the /9 on the right side or the /14 on the left side are operational telescopes or if they just happen to be unused IPv4 space.

Next, we take a closer look at another /8. We selected it since we know that it contains an operational telescope. Figure 6.9 shows the corresponding Hilbert Curves in which we observe the address space of the telescope in the upper and lower left, and the lower right quarters. The upper right quarter does not belong to the telescope. Unlike before, CE1 is unable to infer many of the actual telescope prefixes as meta-telescope prefixes, hence the number of colored pixels in Figure 6.9(a) is small. In Figure 6.9(b), we show the meta-telescope prefixes inferred by vantage point NA1. Here, we find that many meta-telescope prefixes are inferred correctly and we can clearly see the boundaries of the telescope's address space.

Still, some sizable blocks inside the telescope space remain undetected (white) using data from NA1. However, for some of these blocks, the CE1 vantage point has visibility. When integrating data from all vantage points, our inferred meta-telescope prefixes match the known operational telescope prefixes, see Figure 6.9(c). One reason for the need of multiple vantage points has to do with the different route announcements, i.e.,

World Region	Total	ISP	Enterprise	Education	Data Center
All	318,559	158,262	56,598	79,206	24,493
North America	119,919	30,756	28,407	44,729	16,027
South America	10,680	9,492	849	99	240
Europe	58,990	34,284	11,105	10,323	3,278
Asia	106,411	68,180	12,255	21,958	4,018
Africa	9,411	7,458	1,401	318	234
Oceania	12,373	7,726	2,437	1,741	469
International	729	344	125	38	222

Table 6.7: Number of meta-telescope /24 prefixes using the union data set, per type and continent.

using more specifics, and preferences of specific neighbors. This results in different route propagation which can lead to some blind spots at certain vantage points, e.g., `NA1` or `CE1`. This example highlights that large as well as small vantage points can add substantial value to the overall visibility. Furthermore, by combining data from multiple vantage points one can increase visibility, but the trade-off is that the spoofing concern becomes more challenging.

Meta-Telescope Prefixes: Network Types

Next, we focus on studying the network type of the ASes that host the inferred meta-telescope prefixes. We rely on the classification offered by the "ipinfo" data set, and consider the following network types: ISP, Enterprise, Education, and Data Center. Table 6.7 shows the resulting numbers of meta-telescope prefixes per network type and geographic region for all vantage points. This highlights that we are able to identify meta-telescope prefixes in all network types in every region of the world, ranging from a few dozens to tens of thousands. Thus, our methodology gives us access to meta-telescope prefixes inside ISP, Enterprise, and Data Center networks which up to now has been a rather difficult task for researchers. This result underlines the novel contribution of our methodology.

The results confirm that North America hosts the largest share of inferred meta-telescope prefixes. In addition, most meta-telescope prefixes are located inside ISP networks rather than educational networks as is common for most operational telescope networks. The second largest number of meta-telescope prefixes is located in address ranges of educational institutions. The third most ones are inside Enterprise networks, while the smallest number belongs to Data Center networks. Africa is not covered as well. Somewhat surprisingly, this is also true for South America. However, the likely explanation is that we do not have an IXP vantage point within South America. The row labeled as "International" refers to prefixes that cannot be mapped to only one region (a small number of prefixes).

Meta-Telescope: Prefix Coverage

Next, we examine what fraction of an address space is inferred as meta-telescope space regardless of the AS that hosts it. For this task, we focus on large prefixes that are advertised, and can be observed via the RouteViews data set. Specifically, we look at advertised prefixes ranging from /8 to /16 and calculate for each the *prefix index*, i.e., the

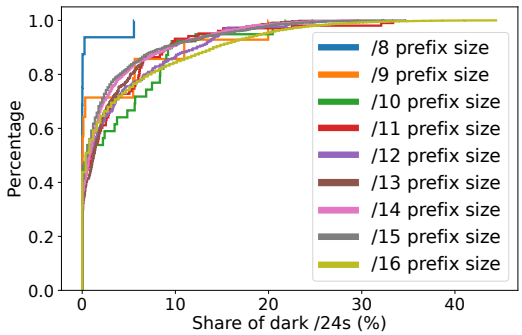

Figure 6.10: Prefix index: ECDF for different prefix sizes.

portion of /24 blocks that are identified as meta-telescope prefixes within their covering prefix. In Figure 6.10 we show an ECDF of the prefix index for each prefix size under consideration.

We find that a surprisingly large share of meta-telescope /24 prefixes is found in such large announced prefixes. For example, more than 6.6% of all /8 BGP announcements have more than 5% meta-telescope address space. This number even rises to roughly 20% for 12% of all /9 announcements. For smaller announced prefixes, i.e., /10 to /15, we find that the share of meta-telescope address space decreases slightly to roughly 10% for most of them. However, we find that for a few /16 announcements, the share of meta-telescope address space is larger than 40%.

We also check if the fraction of meta-telescope space changes with network type, see Figure 6.11 in the Appendix. We find that there are only small differences, with one exception: data center networks tend to have a smaller fraction of meta-telescope space. This is likely due to the fact that data centers have emerged in times where IPv4 address space was already relatively scarce. Looking by continent (see Figure 6.12 in the Appendix), EU followed by AF have the least share, which is again consistent with IPv4 address scarcity.

Meta-Telescope Prefix Distribution

We also check if the fraction of meta-telescope space changes with network type. We find that there are only small differences with one exception. Figure 6.11 shows that Data Center networks tend to have a smaller fraction of meta-telescope space. Looking by continent, see Figure 6.12, EU followed by AF have the least share which is again consistent with IPv4 address scarcity.

6.6 Meta-Telescope Challenges

Next, we comment on the challenges we encountered while aiming to infer meta-telescope prefixes. The first challenge relates to (packet-based) flow sampling, which results in a high variability in the results depending on the day and the vantage point. The second challenge relates to spoofing, which can significantly impact the results. The third challenge relates to sampling. Lastly, we comment on challenges based on our experience in operating a meta-telescope.

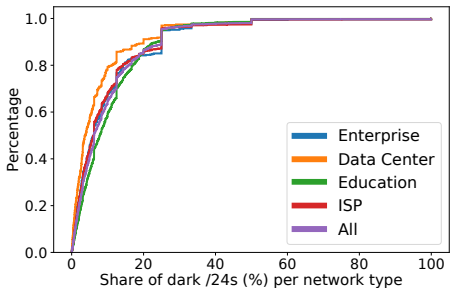

Figure 6.11: Network type index: ECDF for different network types.

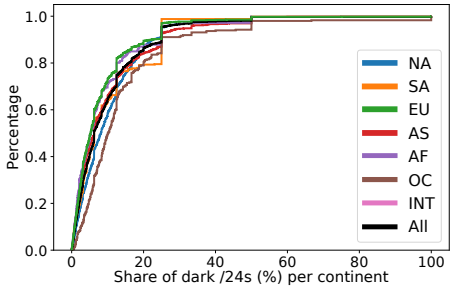

Figure 6.12: Continent index: ECDF for different world regions.

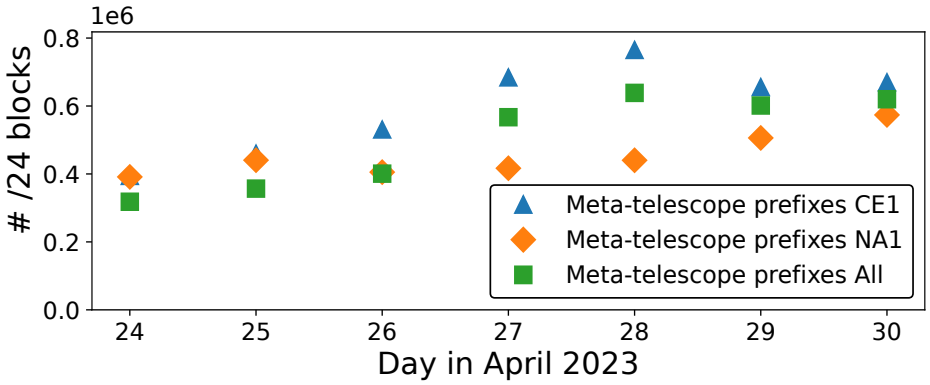

Figure 6.13: Number of daily meta-telescope prefixes for `CE1`, `NA1`, and all.

Meta-Telescope Prefixes Variability

To highlight the diurnal variability, Figure 6.13 depicts the number meta-telescope prefixes for each of the seven days considered in our study. Using data from the vantage point `CE1` from April 24th, 2023, we are able to infer 397,000 /24 blocks. In contrast, using data from the same vantage point four days later, we infer roughly twice as many meta-telescope prefixes. The same kind of variability can be observed for the other vantage points during the week. Another trend across all vantage points is that we are able to infer more meta-telescope prefixes during weekends. One possible explanation is that enterprise or educational networks do not have any major activity outside of the

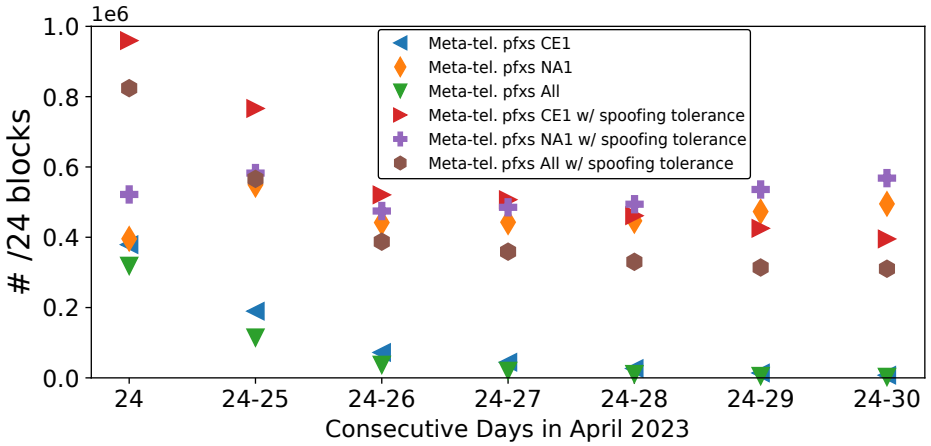

Figure 6.14: The effect of spoofing to the number of meta-telescope prefixes for `CE1`, `NA1`, and All.

working hours. Hence, the effect of sampling becomes less apparent due to the lower traffic volume processed at each vantage point.

Overall, one has to consider multiple trade-offs when inferring meta-telescope prefixes. Should one aim at obtaining stable prefixes, it is better to check if a prefix is among the meta-telescope prefixes identified in multiple days. Moreover, we recommend employing the approach of inferring meta-telescope prefixes on a daily basis also to account for the dynamic nature of Internet routing and address space utilization. Namely, routing changes and/or changes in the use of the address space can alter the observed behavior of network prefixes.

Effect of Spoofing

Spoofing refers to packets with fake source IP addresses. Thus, if we encounter a packet with a spoofed source IP address from a possible meta-telescope prefix range, we may erroneously classify it as an invalid meta-telescope prefix range. Note that spoofing is quite common in the Internet and we have validated that at all of our vantage points we regularly observe spoofed traffic from the three operational telescopes we had access to. We are certain that traffic originating from any operational telescope is spoofed traffic, since these "dark" IP spaces do not initiate, interact with, nor respond to any network traffic.

Spoofed packets negatively effect our inference capability for identifying meta-telescope prefixes. Per our methodology in Section 6.4, (1) we filter out any destination address that is seen sending traffic, and (2) we classify any /24 block as a graynet if it contains at least one "sending" address. In fact, due to the latter step, even a few spoofed packets from a single address could remove a full /24 block from our candidate meta-telescope prefixes.

Given that actors that resort to spoofing activities are typically selecting IP sources across routed and unrouted address space we can leverage this characteristic to better cope with spoofing [148]. The key intuition here is that one could observe the spoofing activities that occur within *unrouted* IP space, and use this information to obtain a "baseline" for spoofing behavior (similarly to [65]). We examined traffic from known unrouted IP space to identify how many spoofed packets one should expect and adjust the filters in our

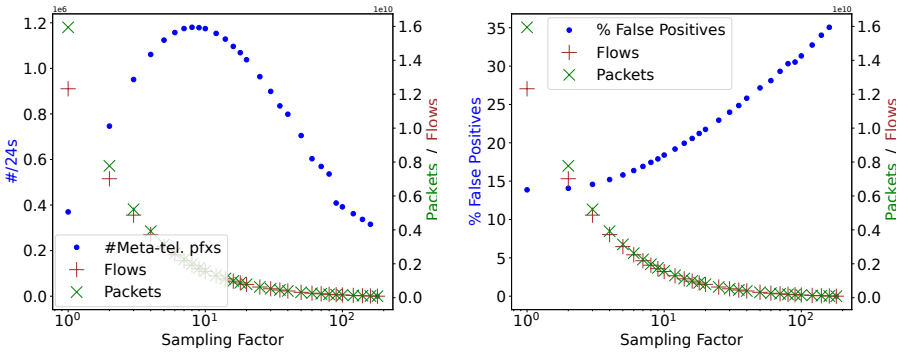

(a) Number of /24 meta-telescope prefixes

(b) Percent of false positives

Figure 6.15: Effect of performing the meta-telescope prefix inference on various sub-sampled data sets.

methodology. As a result, we manage to obtain a tolerance for packets to be seen sourced "from" a /24 block without mistakenly tagging that block as a graynet. We calculate this tolerance for each vantage point and each time frame. Hence, we can adapt to events that may cause an increase in the number of spoofed packets. Concretely, we calculate the 99.99th percentile of packets seen per /24 block inside 2 unrouted /8 blocks per day. We allow this many packets to be sourced by any meta-telescope prefix on that day. For most sites and a single day this threshold is zero packets; for some other sites, it is one or two packets. For seven days the threshold can rise up to four packets per day.

To highlight the impact of spoofing, Figure 6.14 shows how the number of meta-telescope prefixes that we infer decreases when we add more days of data. The number of blocks for all sites decreases from 350k to 4k. If we add the spoofing tolerance the numbers change from above 800k to 400k. Similar observations hold for `CE1`.

The vantage point `NA1` does not seem to suffer that much from spoofing as the number of inferred meta-telescope prefixes does not decrease substantially. In fact, for some days the inferred meta-telescope prefixes increase. Nevertheless, the spoofing tolerance (which is very small for `NA1`) is still meaningful and allows us to reach higher numbers of inferred meta-telescope prefixes.

Effect of Sampling

Since our analysis is applied to a data set consisting of sampled data, we want to better understand the impact of sampling to our inference method. To this end, we create additional data sets with different sub-sampling rates from the original data set of all 14 IXPs from April 24th, 2023. As it is not possible to get data from the IXPs using lower sampling rates, i.e., inspecting more packets to create the flow data, we use higher sampling rates, i.e., investigating a lower number of packets, to get a view on the effect of sampling. For a sub-sampling factor of 2, we only consider every second packet in the IXP data. For a factor of 3, only every third, and so on. Figure 6.15 shows the results when applying our inference method to various sub-sampled portions, starting with the original data set, e.g., a sampling factor of 1. We plot the number of flows and number of packets in the respective sub-sampled data sets. In Figure 6.15(a), we

plot the absolute number of inferred /24 meta-telescope prefixes on the sub-sampled data sets. We see that using higher sampling rates, the inference method first identifies more /24 meta-telescope prefixes, which is likely linked to spoofing becoming less represented in the data set and, hence, eliminating fewer /24 blocks that are actually meta-telescope prefixes. However, when considering only every 100th packet of the original data set, the inference method becomes blind to many parts of the Internet and starts to infer fewer meta-telescope prefixes. Finally, when reaching a sub-sampling factor of 180, our inference method can no longer identify a single meta-telescope prefix. In Figure 6.15(b), we plot the share of false positives of the analyses performed on the respective sub-sample data sets. We find that the rate of false positives is monotonously increasing when using higher sub-sampling rates. Overall, there seems to be a sweet-spot of sampling rates to use, if the data set is prone to spoofing. However, we conclude that using lower sampling rates will provide the most reliable data set of meta-telescope prefixes.

Summary of Challenges

Overall, under a reasonable sampling rate, false positives (i.e., identifying an active prefix as meta-telescope prefix) are by design not a concern, as we elaborated in Section 6.4. However, when the sampling rate is very high, false positives are a concern, since there are not enough samples to conclude with high confidence which part of the address space is not active. By increasing the period of our study, see Section 6.6, and due to spoofing, see Section 6.6, the false negative rate may increase, i.e., non-active prefixes may not be identified as meta-telescope prefixes. The size of vantage points also plays an important role. As we elaborated in Section 6.5, smaller vantage points typically have a higher false negative rate, i.e., less prefixes are identified as meta-telescope prefixes, compared to larger vantage points. The vantage point's physical proximity to the allocated unused space region also plays a role. Typically, the false negative rate is lower for allocated unused space in the same region (continent) of a vantage point, as shown in Section 6.5.

6.7 Meta-Telescope Insights

In this section, we discuss some insights distilled by the inferred meta-telescope to underscore the new abilities unleashed by this new measurement approach. Specifically, we examine scanning activities targeting different geographic regions and different network types.

Targeted Ports by Geographic Region

Using the inferred meta-telescope prefixes, we study traffic destined to TCP ports to shed light into services / applications contacted by nefarious actors. This analysis offers insights to cybersecurity analysts about ongoing security incidents, such as, e.g., scans for exploitable ports (e.g., ssh or telnet ports targeted by Mirai botnet variants [9]), reconnaissance activities for vulnerabilities on ports such as 3389 (Microsoft remote desktop services), randomly spoofed DDoS attacks (observed in our inferred meta-telescope as "backscatter" traffic [177]).

We start by inspecting the distribution of the most popular destination ports. We first compile the list of top-targeted ports for all meta-telescope prefixes of each region. We

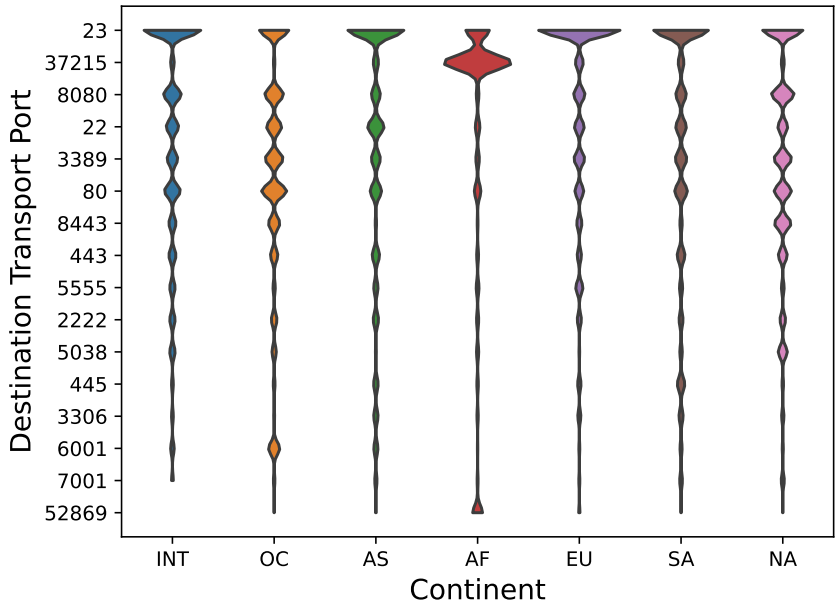

Figure 6.16: Bean plot of activity for the top 16 destination ports in meta-telescope traffic per world region.

then join these lists to get an aggregate list of top-16 ports. This list includes ports often probed by scanners, e.g., ports 23 (telnet) and port 2222 (used by derivatives of the Mirai botnet [9]), and popular services, e.g., port 80 / 443 (HTTP / HTTPS).

In Figure 6.16 we use bean plots to visualize the distribution of traffic per top-16 ports as observed by the meta-telescope prefixes. A bean plot allows for visual comparison of univariate data between groups. Here, the groups are the world regions and the values are the port popularity. The ports are ordered by the total popularity in descending order. Overall, we find that port 23 dominates in all regions except OC and AF. Port 37215 is among the top-10 due to its popularity in AF. This port is used for attacks on Huawai HC532 routers. This port is not aimed so aggressively in other regions. Another port that is mainly popular in AF is 52869, which seems to be associated with "Satori", i.e., a Mirai variant that scans both ports 37215 and 52869 to spread. So traffic to port 37215 seems to serve two purposes but most likely it is associated with the Satori botnet. We also see popular Web ports, namely ports 8080, 80, 443, 8443. However, somewhat surprisingly, 8080 is observed to be the most popular one. However, given that network administrators are more prone to securing services listening to TCP/80, adversaries may have adjusted their strategies to look for alternative HTTP ports. We also note that 8080 is another port targeted by the Mirai botnet. Figure 6.16 shows the port activities relative to the total activity within the specific region. We plot the port activities in all regions relative to the overall traffic share in all meta-telescope prefixes in Figure 6.17. It highlights that SA, OC, and INT receive a very small share of the overall traffic. Figure 6.17 also puts into global perspective the activity against port 37215 which is observed to dominate the AF region.

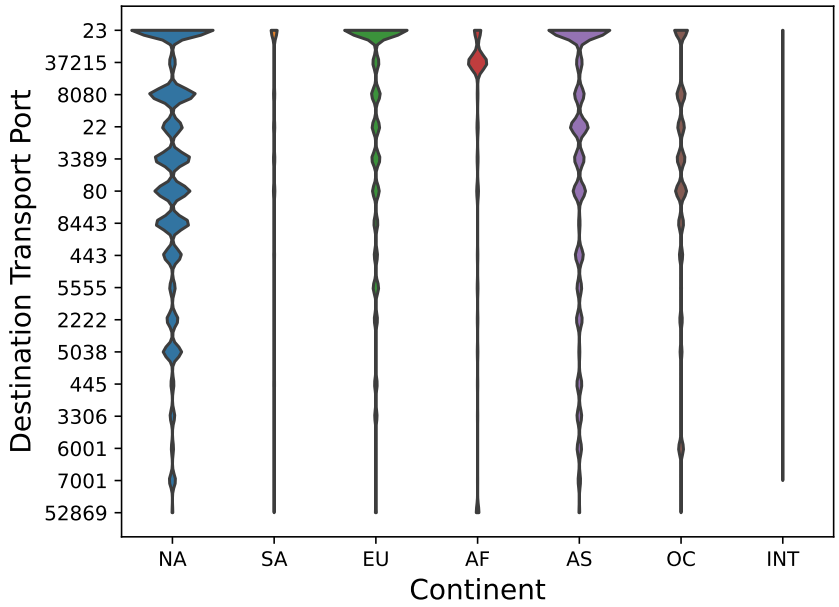

Figure 6.17: Bean plot of activity for the top 16 destination ports in meta-telescope traffic per world region relative to overall traffic.

Targeted Ports by Network Type

Next, we check if there are interesting observations when we categorize traffic by the network type that the meta-telescope prefix is located in. We again compile top-lists per network type and then obtain their union. This yields a total of 12 top destination ports, see Figure 6.18. Compared with the lists of the previous section, ports 52869, 6001, 7001, and 3306 are now excluded. No new ports are identified. The ports amiss are ports that are only popular in some of the regions; e.g., port 3306 is mainly popular in AF and NA, port 52869 in AF, port 6001 in OC, and 7001 in NA. Port 23 is again the most popular one. Port 37215 is not that prevalent anymore in a single group, but has contributions in all sub-categories. Still, ISPs contribute the most and these are mainly located in AF. Overall, we see that the popularity of the top 6–8 ports after removing port 23 is roughly similar. Also, note that the port rankings can significantly vary across categories; e.g., port 80 is more popular within data centers and educational networks, and less preferred when it comes to ISP-based meta-telescope. A likely explanation is that scanners are trying to find unprotected Web servers within data centers. In a similar manner, activity against port 5038—a MLDB database port—is higher within data centers compared to ISP-based or enterprise networks.

Targeted Ports by Network Type in NA and EU

Given the significant differences in port activities by region and by network type we now take a closer look at the two regions with the largest meta-telescope prefix space, namely NA and EU. Hereby, we separate the meta-telescope prefixes by network type,

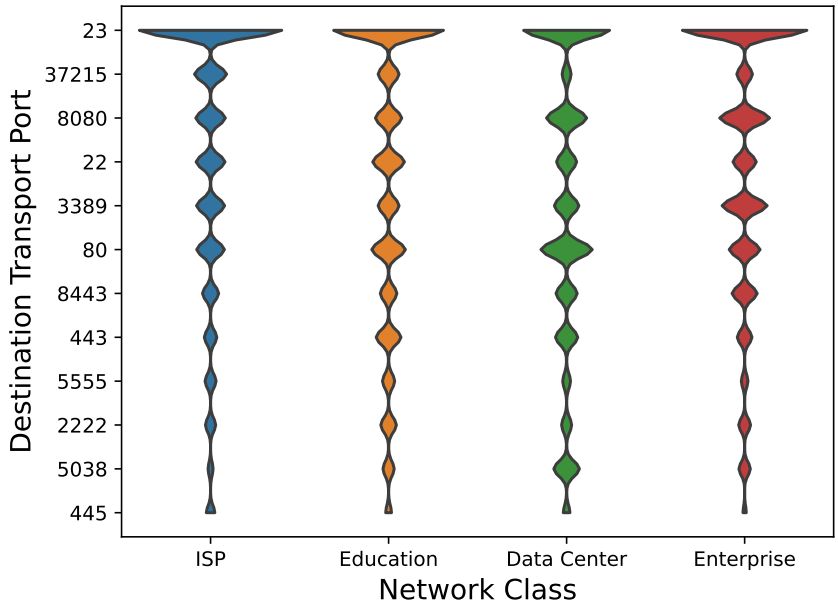

Figure 6.18: Bean plot of activity for the top 12 destination ports in meta-telescope traffic per network type.

see Figure 6.19 and Figure 6.20.

For NA we again see that multiple ports are popular. We also see that port 80 is of particular interest in data centers and educational networks. The same is true for the database port 5038. For Enterprise and ISPs networks port 3389 stands out. Port 23 is still the most prominent port. For EU port 23 dominates by far. At the same time, similar observations with the NA hold. Interestingly, port rankings differ among NA and EU scanned prefixes. Hereby, we note that the differences in the percentages for ranks 4–8 are minimal. Still, port 37215 moves to rank 6 for EU and is no longer in the top 10 for NA. In addition, port 7001 gets introduced into the NA region.

6.8 Ethical Considerations

Traffic Captures and Data Products. Our study is based on traffic statistics and data that the IXPs may gather for operational purposes and are in compliance with legal requirements in the respective countries of operation. All traffic traces are aggregated at the flow level and do not contain any payload. Additionally, the data is processed and analyzed in-situ at the IXP premises and all analyses can potentially happen in an online manner using only aggregate information at the /24 subnet. We utilize the flow data available at the IXP to extract two data products: (1) the list of network /24 prefixes that, using the methodology proposed in this chapter, are inferred to be "unused"—we refer to the set of these prefixes as the meta-telescope; (2) traffic flow traces (which are a subset of the overall IXP flow data sets) destined to the meta-telescope's prefixes that could be used to shed light into Internet-wide scanning activities, malware campaigns, etc.

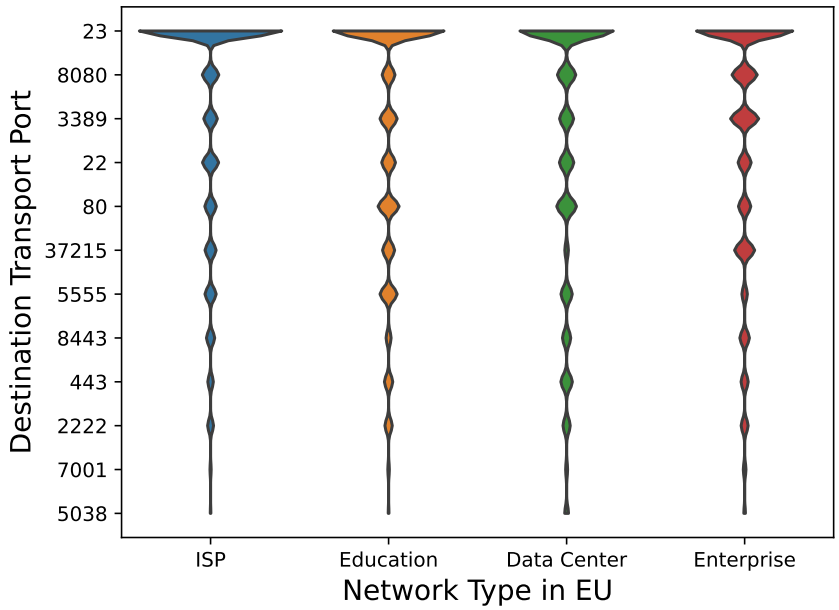

Figure 6.19: Bean plot of activity for the top 12 destination ports in meta-telescope traffic per network type for regions EU.

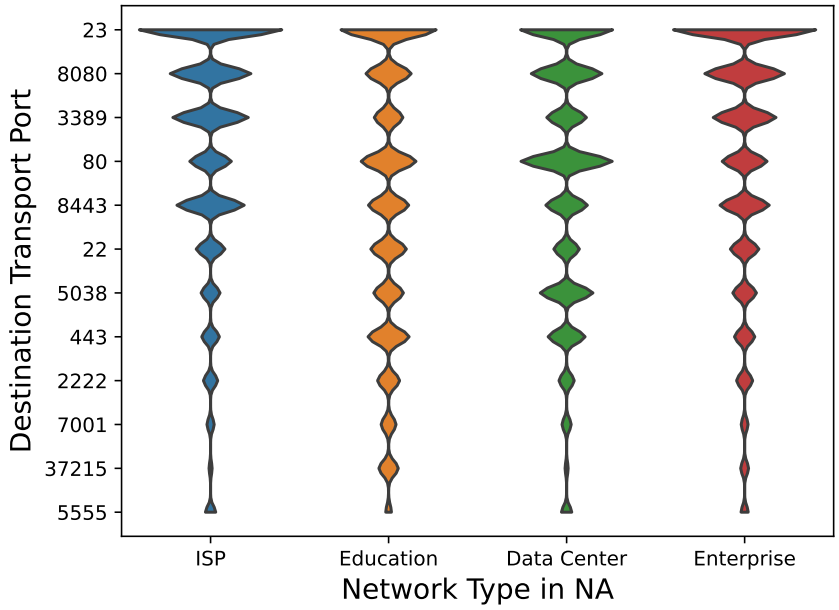

Figure 6.20: Bean plot of activity for the top 12 destination ports in meta-telescope traffic per network type for regions NA.

Leveraging Inferred Unused Prefixes for a Meta-telescope. The key idea of our study is to identify prefixes that can serve as telescopes for a vantage point. One may argue that these prefixes are not controlled by the vantage point operator (i.e., the IXP in our case). However, the traffic passes the vantage point and uses the resources of the vantage point. Therefore, this traffic does not differ from any of the other traffic at the vantage point and the vantage point operator has an operational interest in monitoring such traffic. E.g., it is not atypical for operators to use appliances that process the sampled traffic data for upstream analysis, traffic insights, DDoS mitigation and other threat analysis. These appliances monitor traffic that transits the IXPs, which do not own either end of the communication, and offer insights that could be used for traffic engineering purposes, capacity planning, peering arrangements, routing policies, etc. These insights help the operators operate their networks in a safe manner and offer quality service to their customers. This chapter proposes a new approach that fits into the same scope of data analysis and allows the vantage point operators to leverage the traffic traces they already process to obtain a novel product, namely the meta-telescope and its data products. These new data products can then be utilized to identify "unwanted" traffic, i.e., background radiation traffic that is inherently malicious/suspicious, since destined to subnets that are seemingly unused. Once the meta-telescope prefixes are identified, only a small number of IP /24 subnets needs to be further monitored (about 5%).

Note that the distilled data products are not meant to be broadly or publicly shared. Although the vantage point operators may elect to share aggregate meta-data (e.g., targeted ports per region or per country) with the community, the main utility of the extracted data sets lies within the operator itself. For instance, the operator may utilize the data sets to enhance the cybersecurity posture of their customers by informing them of suspicious IPs originating from the customer prefixes and destined to the inferred meta-telescope prefixes. Similarly, the meta-telescope operator may share the threat intelligence extracted from the meta-telescope prefixes with appropriate authorities, e.g., Computer Emergency Response Team (CERT) organizations, to inform them about the onset of new malicious activities or nefarious scanning campaigns. To minimize any potential reputation harm that may arise from Type I errors (i.e., false positives) when detecting unused subnets, we followed a very conservative approach in the calibration of the parameters of our methodology (as we had discussed earlier). To further minimize these risks, one could complement the meta-telescope insights with the observations from operational/real network telescopes (such as the three we employ in our study) and other existing threat intelligence data sets.

Operational Telescopes. Our study is based on data that the operational telescopes regularly capture for operational purposes and are again in compliance with the legal requirements in the respective countries. All data is processed and analyzed in-situ at the premise of the telescope operator. The telescopes receive only unidirectional traffic destined to the unused address space, and we do not probe or interact with any of the source IPs sending traffic to the telescope.

Routing and Active Measurement Data Sets. Our study is based on data made accessible by the respective data provider either via research license or via public domain access. None of the used data was created specifically for the purpose of this study.

6.9 Discussion

Dealing with Spoofing. Our experience in detecting meta-telescope prefixes shows that spoofing has a significant effect. Indeed, the number of detected meta-telescope prefixes reduces over time. Even with the spoofing tolerance technique, some missed meta-telescope prefixes can be detected. An alternative approach is to exclude network flows by networks that do not implement spoofing filters, namely BCP 38 [93]. Projects such as Spoofer [152] maintain a list of networks that have not adopted BCP 38. Another approach is to exclude flows originating from IPs that do not belong to a peering network's network cone [153] from our analysis network. We are aware that the network cone is not always accurate. However, this approach will reduce the spoofed traffic significantly.

The Vantage Point Effect. Our methodology can be applied to network flows collected at any vantage point. For this study, we presented an in-detail analysis of the meta-telescope prefixes that can be inferred when using Internet exchange points of different size in terms of number of members and traffic and from different regions. However, shortcomings such as routing visibility, asymmetric routing, sampled traffic, spoofing, and members' demographics may limit the number of meta-telescope prefixes we discover. Network flows captured at large Internet service providers do not suffer from asymmetric routing. In this case, the routing visibility is less of a concern, and many of them already have BCP 38 implemented. Some ISPs also collect network flow information at a high sampling rate. For all the above reasons, there is the potential to detect even a higher number of meta-telescope prefixes when analyzing ISP data.

Meta-telescope Information as a Service. Internet exchange points and service providers that implement our methodology can detect information about meta-telescope prefixes, and they can also infer which of their peers or customers send traffic to them. They can offer this information as an opt-in service to their customers to make them aware that traffic that originates from their network has a meta-telescope prefix as destination, helping their customers to make traffic engineering and filtering decisions. Our results show that the set of meta-telescope prefixes is quite stable for a couple of days. However, the set of meta-telescope prefixes will vary when the observation window increases in duration and traffic conditions change rapidly, e.g., unused space is allocated to hosts. We argue that additional vantage points and regular measurements to detect meta-telescope prefixes are needed before meta-telescope information as a service will be able to handle spontaneous prefix allocation and traffic changes.

Federated Meta-telescope. The detection of meta-telescope prefixes can inform and improve the operation of more customers than those of a single IXP or ISP. The detection can be shared among trusted parties to detect meta-telescope prefixes with higher accuracy collectively. It is also possible to develop a standard to enable operators to opt-in to the measurements, e.g., a BGP community or embedment into RPKI that marks announced but unused space. Encoding known only to involved parties will help in keeping this tagging hidden so the prefixes will not be excluded from scanners and attackers. Getting this information into operation has the potential to impact network operators' operations significantly. The research community can also benefit and contribute to this effort.

IPv6 Meta-telescope. In this chapter, we focus on the detection of IPv4 meta-telescope prefixes. IPv6 address space is much larger and less active. In addition, IPv6 address assignment varies per network and vendor. As IPv6 traffic increases, it is important also to detect meta-telescope IPv6 prefixes. Given the vastness of the IPv6 space, our filtering pipeline would likely need adjustments. The lack of complete and reliable hit

lists [101] and archives of active measurements for IPv6 further complicate the detection of meta-telescope IPv6 prefixes.

6.10 Summary

For decades network telescopes have been used to collect and analyze unsolicited traffic to detect misconfigurations or malicious activities including the spread of Botnets, DDoS campaigns, and exploitation of vulnerabilities. A limitation when deploying a network telescope is that its visibility is limited to the scanning and attack activity received in the announced prefixes, which are typically limited to one regional location. In this chapter, we develop and evaluate a methodology to detect advertised but unused space worldwide and potentially originated by any organization. This approach can provide insights similar to those of collectively operating network telescopes in all portions of this space, which we refer to as meta-telescope prefixes.

By utilizing traffic flows collected in the core of the Internet, i.e., at Internet exchange points, we show that it is possible to detect more than 350k /24 IPv4 address blocks, in 7k ASes and 190 countries, which can be used in a meta-telescope. The size of this meta-telescope is by far larger than any other operational telescope, and it is also highly distributed and observing (unsolicited) traffic towards networks of different types. These features allow us to answer measurement questions for unsolicited traffic arriving in networks at different regions and types, without the need to own, advertise, and dedicate address space for a telescope, and the operational overhead of running it. We also comment on our experience in detecting meta-telescope prefixes. While spoofing can significantly reduce the detection of meta-telescope prefixes, we propose ways to overcome this issue.

Chapter 7
Conclusion

In this thesis, we tackle four research questions, with the goal of investigating how the Internet reacts to analog and digital stress factors. While studying these, we derive, realize, and evaluate possible solutions that have the potential to improve the resilience of the Internet.

7.1 Summary

We study the reaction of the Internet to the COVID-19 pandemic in Chapter 3. This helps us to answer the first research question:

(1) How does the core of the Internet react to unprecedented traffic volume?

We find that the traffic volume increased by 15-20% almost within a week— while overall still modest, this constitutes a large increase within this short time. However, despite this surge, we observe that the Internet infrastructure can handle the new volume, as most traffic shifts occur outside of traditional peak hours. When looking directly at the traffic sources, it turns out that, while hypergiants still contribute a significant fraction of traffic, we see (1) a higher increase in traffic of non-hypergiants, and (2) traffic increases in applications that people use when at home, such as Web conferencing, VPN, and gaming. While many networks see increased traffic demands, in particular, those providing services to residential users, academic networks experience major overall decreases. Yet, in these networks, we can observe substantial increases when considering applications associated with remote working and lecturing.

We find that the reaction time to events is crucial for the resilience of the Internet. We, therefore, in Chapter 4, tackle our second research question:

(2) Can novel programming paradigms improve network configuration flexibility?

We present the P4IX, a technical concept for a generic P4 packet processing pipeline for IXPs. The P4IX concept is built upon a comprehensive requirements analysis: we

characterize the IXP landscape and provide first-hand insights of a large IXP operator (more than 1000 well-distributed ports). Moreover, we use our insights to critically discuss the P4IX from an operational, technical, and organizational perspective.

Next, we move to stress situations from the digital world. We measure the benefit for DDoS attack mitigation of collaboratively exchanging traffic characteristics attacks between IXPs in Chapter 5. This provides us the means to tackle our third research question:

(3) Can the collaborative exchange of traffic characteristics benefit the distributed mitigation of DDoS attacks?

We collaborate with eleven IXPs that operate in three different regions. These IXPs have more than 2,120 network members that exchange traffic at the rate of more than 11 Terabits per second. We collect network data over six months and analyze more than 120k amplification DDoS attacks. To our surprise, more than 80% of the amplification DDoS attacks are not detected locally, although the majority of the attacks are visible by at least three IXPs. A closer investigation points to the shortcomings, such as the multi-protocol profile of modern amplification attacks, the duration of the attacks, and the difficulty of setting appropriate local attack traffic thresholds that will trigger mitigation. Our evaluation shows that it is possible to collaboratively detect and mitigate the majority of attacks with limited exchange of information and drop as much as 90% more attack traffic earlier in the network.

Finally, in Chapter 6, we introduce the concept of meta-telescopes to capture unsolicited Internet traffic. A meta-telescope is based on the intuition that, with the availability of appropriate vantage points, one can (1) infer which address blocks on the Internet are unused and (2) capture traffic towards them—both without having control of such address blocks. This helps us to answer our last research question:

(4) Can we distill traffic anomaly trends by focusing on inactive destinations?

We develop and evaluate a methodology for identifying unlikely-to-be-used Internet address space and build a meta-telescope. Such a meta-telescope identifies and captures unsolicited traffic to more than 350k /24 blocks in more than 7k ASes. Through the analysis of background radiation towards these networks, we also highlight that unsolicited traffic differs by destination network/geographic region as well as by network type. Finally, we discuss our experience and challenges when operating a meta-telescope in the wild.

7.2 Future Directions

In this section, we discuss possible future directions of our work.

Implementation of Reactive Internet Peering Infrastructure. The goal we started with is to improve the reactive capabilities of Internet peering infrastructure. As such, we propose a flexible concept for a P4-based IXP in Chapter 4. This provides us with the opportunity to realize services, e.g., blackholing, traffic isolation, enforcement of BGP policies, the exchange of information about traffic characteristics as presented in Chapter 5, and the inference of meta-telescope prefixes as presented in Chapter 6. Integration of these services in an implementation of the P4-based IXP concept is left for future work.

IPv6. This thesis focuses on IPv4 traffic. Yet, IPv6 traffic is increasing both in terms of traffic volume and connectivity. Thus, it becomes increasingly important to study IPv6 traffic. We now discuss how our methodologies from Chapter 5 and Chapter 6 can be applied to IPv6 traffic.

We expect that no technical adjustments are necessary to collaboratively exchange IPv6 traffic characteristics between IXPs to improve the distributed mitigation of DDoS attacks. However, the benefits are likely relatively small compared to IPv4 as there are not that many attacks using IPv6 yet. The reason for this is that most abandoned hosts are more likely to have no IPv6 address configured. Furthermore, it is significantly more difficult for attackers to locate hosts via IPv6 scanning due to the largeness of the IPv6 address space. Nevertheless, every detected and mitigated attack is an improvement for the resilience of the Internet.

With regard to our inference of unused IP address space, some adjustments have to be made. Unsolicited IPv6 traffic has different characteristics in terms of distribution across the IPv6 address space due to its vastness [59]. Moreover, for the elimination of false positives, we need reliable hit lists [101] as well as archives of active measurements. The lack of those complicates the detection of meta-telescope IPv6 prefixes. Still, we expect that our approach of inferring unused address spaces to work well with IPv6 traffic. We leave the inference of unused IPv6 address space to future work.

ISPs as Vantage Points. In this thesis, we focus mostly on IXPs as vantage points. In Chapter 3, we use data from an ISP network and derive valuable insights. Indeed, ISP networks are important vantage points in the Internet and provide us with complementary data to the IXP vantage points. Hence, we plan to apply our analyses from Chapter 5 and Chapter 6 to ISP networks in future work. We now discuss how our methodologies can be applied to ISP networks.

The methodology for exchanging information about traffic characteristics is generic enough to be applicable to ISP networks and potentially other types of networks. However, some adjustments are again necessary. For example, our distance calculation relies on information from the IXP route server which does not exist in ISP networks. An appropriate data source can be the RIB of the routers in the ISP network.

Regarding the inference of unused IPv4 address space, we expect that the presented methodology can be applied to ISP networks without any adjustments.

Other Stressful Factors. Amplification reflection DDoS attacks are not the only kind of DDoS attacks. The detection mechanism presented in Chapter 5 can be tweaked to focus on other types of DDoS attacks as well. For example, the 1 Gb/s threshold can be replaced with a 1 Mp/s threshold, and the data set can be expanded to other transport protocols. Using this mechanism, we expect to detect attacks that belong to, e.g., botnet attacks, and attacks that are based on overwhelming a host by sending a huge number of requests. We expect that changing these thresholds still allows us to detect amplification reflection DDoS attacks since they generate many packets. This can help us to detect and mitigate a larger set of DDoS attacks. However, it may increase the number of false positives and, thus, requires more future work.

Regarding further stress situations from the analog world, we always have to be prepared to analyze and learn from them. However, we hope that no such situation will arise any time soon.

Appendix A
DDoS Attack Feature Set

For analyzing our flow data, we define 1106 features. The following list shows how they can be divided into feature classes, see Tables A.1, A.2, A.3, A.4, and A.5. Some classes contain multiple features as they are parameterized by IXP and/or by port. Therefore, we also add how many features they contribute for clarity.

The sites are `CE1`, `CE2`, `CE3`, `CE4`, `SE1`, `SE2`, `SE3`, `SE4`, `SE5`, `NA1` and `NA2`. When normalizing by size, we are multiplying the traffic by the relative average traffic volume to the biggest IXP. The source transport protocols are 19 (chargen), 53 (DNS), 111 (RPC), 123 (NTP), 161 (SNMP), 389 (CLDAP), 1194 (OpenVPN), 1900 (SSDP), 3283 (ARMS), 3702 (WS-Discovery), 10001 (Device Discovery), 11211 (Memcached). The thresholds are 100 Mbps, 250 Mbps, 500 Mbps 1 Gbps, 2.5 Gbps, 5 Gbps and 10 Gbps. Every time bin comprises one minute of traffic.

Feature Class	#Features	Description
Sites	1	Number of sites involved in the attack
Ports	1	Number of source transport ports involved in the attack
SitesPorts	1	Sum of source transport ports seen at the sites, where the attack is visible
Dur	1	Total duration of the attack in minutes
DurAttack	1	Duration in minutes where the attack volume is greater than t (In our study: 1 Gbps)
TotalMbps	1	Volume of the attack in Mbps, summed across all sites and all source transport ports

Table A.1: List of features used for the PCA.

Feature Class	#Features	Description
TotalMbpsAttack	1	Volume of the attack in Mbps, summed across all sites and all source transport ports, while the volume is greater than t
TotalPeakMbps	1	Peak of the attack volume in Mbps, summed across all sites and all source transport ports
Peak Mbps	1	Peak of the attack volume in Mbps, single site, single source transport port
TotalMbpsCE1	1	Sum of the attack traffic across all source transport ports in Mbps, seen at site CE1
TotalMbpsAttackCE1	1	Sum of the attack volume across all source transport ports in Mbps, seen at site CE1 while exceeding t
TotalPeakMbpsCE1	1	Peak attack volume across all source transport ports, seen at site CE1, in Mbps
PeakMbpsCE1	1	Peak attack volume of a single source transport port, seen at site CE1, in Mbps
TotalMbpsNoCE1	1	Volume of the attack in Mbps, seen at all sites but CE1, all source transport ports
TotalMbpsAttackNoCE1	1	Volume of the attack in Mbps, seen at all sites but CE1, all source transport ports while exceeding t
TotalPeakMbpsNoCE1	1	Peak volume of the attack in Mbps, seen at all sites but CE1, across all source transport ports
PeakMbpsNoCE1	1	Peak volume of the attack in Mbps, seen at all sites but CE1, across a single transport port
Cor[Site | Port]{0.7,0.8,0.9}	6	Counter for correlation of the attack between sites and source transport ports, respectively, being greater than .7, .8, .9, respectively per minute.
TotalMbps[IXP*]	11	Volume of the attack in Mbps, as seen at the 11 sites, all source transport ports, respectively
TotalMbps[PORT*]	12	Volume of the attack in Mbps, summed across all sites, for each of the 12 source transport ports in our study
PeakMbps[IXP*]	11	Peak volume of the attack in Mbps, as seen at the 11 sites, respectively, single source transport port

Table A.2: List of features used for the PCA (continued (1)).

Feature Class	#Features	Description
PeakMbps[PORT*]	12	Peak volume of the attack in Mbps, summed across all sites, for each of the 12 source transport ports in our study
TotalMpps	1	Sum of packets transmitted for the attack across all sites, all source transport protocols, in Mpps
TotalMppsAttack	1	Sum of packets transmitted for the attack across all, all source transport ports, sites, while exceeding t, in Mpps
TotalPeakMpps	1	Peak of packets transmitted for the attack, summed across all sites, all source transport ports, in Mpps
PeakMpps	1	Peak of packets transmitted for the attack at any site, single transport port, in Mpps
TotalMpps[IXP*]	11	Sum of packets transmitted across all source transport ports, at the 11 sites, respectively
TotalMpps[PORT*]	12	Sum of packets transmitted at all sites, for each of the 12 source transport protocols in our study
TotalMbpsNorm	1	Volume of the attack, summed across all source transport ports and all sites, normalized by their size
TotalMbpsAttackNorm	1	Volume of the attack in Mbps, summed across all source transport ports, all sites, normalized by their size, while exceeding t
TotalPeakMbpsNorm	1	Peak of the attack volume in Mbps, summed across allsource transport ports, all sites, normalized by their size
PeakMbpsNorm	1	Peak of the attack volume in Mbps, single source transport port, at a single site, normalized by their size
TotalMbpsNormNoCE1	1	Volume of the attack in Mbps, all source transport ports, seen at all sites but CE1, normalized by their size
TotalMbpsAttackNormNoCE1	1	Volume of the attack in Mbps, all source transport ports, seen at all sites but CE1, normalized by their size, while exceeding t
TotalPeakMbpsNormNoCE1	1	Peak volume of the attack, summed all source transport ports, seen at all sites but CE1, normalized by their size
PeakMbpsNormNoCE1	1	Peak volume of the attack, single source transport ports, seen at all sites but CE1, normalized by their size

Table A.3: List of features used for the PCA (continued (2)).

Feature Class	#Features	Description
TotalMbpsNorm[IXP*]	11	Volume of the attack in Mbps, all source transport ports, as seen at the 11 sites, normalized by their size respectively
PeakMbpsNorm[IXP*]	11	Peak volume of the attack in Mbps, single source transport port, as seen at the 11 sites, normalized by their size, respectively
Allthresh-Before -[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest volume of a single site, before the respective threshold was exceeded
Allthresh-Detect -[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest volume of a single site, while the respective threshold is exceeded
Allthresh-After -[THRESHHOLD*]	7	Volume of traffic across all single source transport ports that belong to an attack,greatest volume of a single site, after therespective threshold is no longer exceeded
Allthresh-Time -[THRESHHOLD*]	7	Amount of time bins for which the attack volume across all source transport ports, greatest of a single site, exceeded the respective threshold
Allthreshnorm-Before -[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest of a single site, normalized by its size, before the respective threshold was exceeded
Allthreshnorm-Detect -[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest of a single site, normalized by its size, while the respective threshold is exceeded
Allthreshnorm-After -[THRESHHOLD*]	7	Volume of traffic across all source transport ports that belong to an attack, greatest of a singnle site, normalized by its size, after the respective threshold is no longer exceeded

Table A.4: List of features used for the PCA (continued (3)).

Feature Class	#Features	Description
Allthreshnorm-Time -[THRESHHOLD*]	7	Amount of time bins for which the attack volume across all source transport ports, greatest of a single site, normalized by its size, exceeded the respective threshold
SiteThresh-[IXP*]-Before	77	Volume of the attack, for every site respectively, single source transport port, before exceeding the respective threshold
SiteThresh-[IXP*]-After -[THRESHHOLD*]	77	Volume of the attack, for every site respectively, single source transport port, after the respective threshold is no longer exceeded
SiteThresh-[IXP*]-Detect -[THRESHHOLD*]	77	Volume of the attack, for every site respectively, single source transport port, while exceeding the respective threshold
SiteThresh-[IXP*]-Time -[THRESHHOLD*]	77	Amount of time bins, for every site respectively, for every threshold, single source transport port, before exceeding the respective threshold
GlobalThresh-[IXP*]-Before -[THRESHHOLD*]	77	Volume of the attack, adding all site's volume to every site respectively, all source transport ports, before exceeding the respective threshold
GlobalThresh-[IXP*]-After -[THRESHHOLD*]	77	Volume of the attack, adding all site's volume to every site respectively, all source transport ports, after the respective threshold is no longer exceeded
GlobalThresh-[IXP*]-Detect -[THRESHHOLD*]	77	Volume of the attack, adding all site's volume to every site respectively, all source transport ports, while exceeding the respective threshold
GlobalThresh-[IXP*]-Time -[THRESHHOLD*]	77	Amount of time bins, when adding all site's volume to the respective site, for every threshold, all source transport ports, while exceeding the respective threshold
SiteThreshNorm-[IXP*]-Before -[THRESHHOLD*]	77	Volume of the attack, for every site, normalized by its size, single source transport port, before exceeding the respective threshold
SiteThreshNorm-[IXP*]-After -[THRESHHOLD*]	77	Volume of the attack, for every site respectively, normalized by its size, single source transport port, after the respective threshold is no longer exceeded

Table A.5: List of features used for the PCA (continued (4)).

Feature Class	#Features	Description
SiteThreshNorm-[IXP*]-After -[THRESHHOLD*]	77	Volume of the attack, for every site respectively, normalized by its size, single source transport port, after the respective threshold is no longer exceeded
SiteThreshNorm-[IXP*]-Detect -[THRESHHOLD*]	77	Volume of the attack, for every site respectively, normalized by its size, single source transport port, while exceeding the respective threshold
SiteThreshNorm-[IXP*]-Time -[THRESHHOLD*]	77	Amount of time bins, for every site respectively, normalized by its size, for every threshold, single source transport port, before exceeding the respective threshold
Total	1106	

Table A.6: List of features used for the PCA (continued (5)).

List of Abbreviations

ARP Address Resolution Protocol 42

AS Autonomous System 7

ASN Autonomous System Number 7

BEREC Body of European Regulators for Electronic Communications 9

BGP Border Gateway Protocol 1

BUM Broadcast, Unknown Unicast, and Multicast 45

CAPEX Capital Expenditures 10

CDN Content Delivery Network 9

COVID-19 Coronavirus Disease 2019 11

DDoS Distributed Denial of Service 1

DFN Deutsches Forschungsnetz (German Research Network) 10

DNS Domain Name System 15

DoS Denial of Service 1

Euro-IX European Internet Exchange Association 9

IBR Internet Backround Radiation 14

IETF Internet Engineering Task Force 12

IP Internet Protocol 5

IPFIX Internet Protocol Flow Information Export 50

IPv4 Internet Protocol Version 4 7

IPv6 Internet Protocol Version 6 7

ISP Internet Service Provider 9

IXP Internet Exchange Point 2

LAN Local Area Network 10

MAC Media Access Control 11

List of Figures

List of Tables

Bibliography

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

[2] Akamai. Web Attack Report Shows Hospitality Industry Under Siege From Botnets, 2018. Available at: `https://www.ir.akamai.com/news-releases/news-release-details/akamai-state-internet-security-summer-2018-web-attack-report` Last accessed: 2023-11-30.

[3] Akamai. Prolexic Technologies by Akamai, 2021. Available at: `https://www.akamai.com/us/en/products/security/prolexic-solutions.jsp` Last accessed: 2023-11-30.

[4] Akamai. DDoS Protection: Reference Architecture, 2023. Available at: `https://www.akamai.com/resources/reference-architecture/ddos-protection` Last accessed: 2023-10-19.

[5] Amazon. AWS Shield Threat Landscape Report, 2020. Available at: `https://aws-shield-tlr.s3.amazonaws.com/2020-Q1_AWS_Shield_TLR.pdf`, Last accessed: 2023-11-30.

[6] Amazon. Overview of Amazon Web Services, 2023. Available at: `https://docs.aws.amazon.com/pdfs/whitepapers/latest/aws-overview/aws-overview.pdf` Last accessed: 2023-09-28.

[7] A. Anand, M. Kallitsis, J. Sippe, and A. Dainotti. Aggressive Internet-Wide Scanners: Network Impact and Longitudinal Characterization. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*, 2023.

[8] G. Antichi, I. Castro, M. Chiesa, E. L. Fernandes, R. Lapeyrade, D. Kopp, J. H. Han, M. Bruyere, C. Dietzel, M. Gusat, A. W. Moore, P. Owezarski, S. Uhlig, and M. Canini. ENDEAVOUR: A Scalable SDN Architecture For Real-World IXPs. *IEEE Journal on Selected Areas in Communications*, 2017.

[9] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou. Understanding the Mirai Botnet. In *Proceedings of the USENIX Security Symposium*, 2017.

[10] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram. BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects. Rfc, Internet Engineering Task Force, 2023. Available at: `https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification` Last accessed: 2023-10-11.

[11] N. Bayat, K. Mahajan, S. Denton, V. Misra, and D. Rubenstein. Down for Failure: Active Power Status Monitoring. *Future Generation Computer Systems*, 2021.

[12] K. Benson, A. Dainotti, K. C. Claffy, and E. Aben. Gaining Insight into AS-Level Outages Through Analysis of Internet Background Radiation. In *Proceedings of the Network Traffic Measurement and Analysis Conference*, 2013.

[13] K. Benson, A. Dainotti, K. C. Claffy, A. Snoeren, and M. Kallitsis. Leveraging Internet Background Radiation for Opportunistic Network Analysis. In *Proceedings of the ACM Internet Measurement Conference*, 2015.

[14] M. Bjorklund. The YANG 1.1 Data Modeling Language. RFC 7950, Internet Engineering Task Force, 2016. Available at: `http://www.rfc-editor.org/rfc/rfc7950.txt` Last accessed: 2023-07-20.

[15] J. Blendin, F. Bendfeldt, I. Poese, B. Koldehofe, and O. Hohlfeld. Dissecting Apple's Meta-CDN During an iOS Update. In *Proceedings of the ACM Internet Measurement Conference*, 2018.

[16] Boardcomm. What Is BGP and Why Is It Important?, 2023. Available at: `https://academy.broadcom.com/blog/network-operations/what-is-bgp-and-why-is-it-important` Last accessed: 2023-10-11.

[17] R. Bogutz, Y. Pradkin, and J. Heidemann. Identifying Important Internet Outages. In *IEEE Big Data*, 2019.

[18] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, and G. Varghese. P4: Programming Protocol-Independent Packet Processors. *ACM Computer Communication Review*, 2014.

[19] P. Bosshart, G. Gibb, H. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz. Forwarding Metamorphosis: Fast Programmable Match-action Processing in Hardware for SDN. *ACM Computer Communication Review*, 2013.

[20] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig. A Hypergiant's View of the Internet. *ACM Computer Communication Review*, 2017.

[21] T. Böttger, F. Cuadrado, and S. Uhlig. Looking for Hypergiants in PeeringDB. *ACM Computer Communication Review*, 2018.

[22] T. Böttger, G. Ibrahim, and B. Vallis. How the Internet Reacted to COVID-19 — A Perspective from Facebook's Edge Network. In *Proceedings of the ACM Internet Measurement Conference*, 2020.

[23] M. Bruyere, G. Antichi, E. L. Fernandes, R. Lapeyrade, S. Uhlig, P. Owezarski, A. W. Moore, and I. Castro. Rethinking IXPs' Architecture in the Age of SDN. *IEEE Journal on Selected Areas in Communications*, 2018.

[24] C. Partridge, P. Barford, D. D. Clark, S. Donelan, V. Paxson, J. Rexford, and M. K. Vernon. *The Internet Under Crisis Conditions: Learning from September 11*. National Academies Press, 2003.

[25] CAIDA. RouteViews Prefix to AS Mappings Dataset for IPv4 and IPv6, 2023. Available at: `https://www.caida.org/catalog/datasets/routeviews-prefix2as/` Last accessed: 2023-12-09.

[26] CAIDA. The CAIDA UCSD AS to Organization Mapping Data Set, 2023-04-11, 2023. Available at: `https://www.caida.org/catalog/datasets/as-organizations/` Last accessed: 2023-12-09.

[27] CAIDA. The UCSD Network Telescope, 2023. Available at: https://www.caida.org/projects/network_telescope/ Last accessed: 2023-12-09.

[28] M. Candela, V. Luconi, and A. Vecchio. Impact of the COVID-19 Pandemic on the Internet Latency: A Large-Scale Study. *Computer Networks*, 2020.

[29] K. Carriello. Arm Yourself Against DDoS Attacks: Using BGP Flow Specification for Advanced Mitigation Architectures. Available at: https://forum.ix.br/files/apresentacao/arquivo/131/04%2012%20%2009%2030%20%20kleber.pdf Last accessed: 2023-12-09, 2017.

[30] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois. Remote Peering: More Peering without Internet Flattening. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*, 2014.

[31] O. Çetin, C. Gañán, L. Altena, T. Kasama, D. Inoue, K. Tamiya, Y. Tie, K. Yoshioka, and M. van Eeten. Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai. In *Proceedings of the Network and Distributed System Security Symposium*, 2019.

[32] Central Intelligence Agency. The World Factbook: Field Listing — Internet Users, 2023. Available at: https://www.cia.gov/the-world-factbook/field/internet-users/ Last accessed: 2023-10-11.

[33] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is More to IXPs Than Meets the Eye. *ACM Computer Communication Review*, 2013.

[34] M. Chiesa, C. Dietzel, G. Antichi, M. Bruyere, I. Castro, M.. Gusat, T. King, A. W. Moore, T. D. Nguyen, and P. Owezarski. Inter-Domain Networking Innovation on Steroids: Empowering IXPs with SDN Capabilities. *IEEE Communications Magazine*, 2016.

[35] Y. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan. Are We One Hop Away from a Better Internet? In *Proceedings of the ACM SIGCOMM HotNets*, 2015.

[36] K. Cho, C. Pelsser, R. Bush, and Y. Won. The Japan Earthquake: The Impact on Traffic and Routing Observed by a Local ISP. In *ACM CoNEXT Special Workshop on the Internet and Disasters*, 2011.

[37] Cisco. Remotely Triggered Black Hole Filtering - Destination Based and Source Based, 2005. Cisco White Paper, available at: http://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf Last accessed: 2023-12-10.

[38] Cisco. Broadband Network Gateway Overview, 2012. https://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r5-2/bng/configuration/guide/b-bng-cg52xasr9k/b-bng-cg52xasr9k_chapter_010.pdf Last accessed: 2023-12-10.

[39] Cisco. Introduction to Cisco IOS NetFlow - A Technical Overview, 2012. http://www.service-desk.co/white_papers/cisco_netflow.pdf Last accessed: 2023-12-10.

[40] Cisco. Cisco Annual Internet Report (2018–2023) White Paper, 2020. Available at: https://www.cisco.com/c/en/us/solutions/collateral/executive-perspectives/annual-internet-report/white-paper-c11-741490.html Last accessed: 2023-10-11.

[41] Cisco. Network Requirements for Webex Teams Services, 2020. Available at: https://help.webex.com/en-us/WBX000028782/Network-Requirements-for-Webex-Teams-Services Last accessed: 2023-12-10.

[42] B. Claise. Cisco Systems NetFlow Services Export Version 9. RFC 3954, Internet Engineering Task Force, 2004. Available at: http://www.rfc-editor.org/rfc/rfc3954.txt Last accessed: 2023-07-20.

[43] B. Claise and B. Trammell. Information Model for IP Flow Information Export (IPFIX). RFC 7012, Internet Engineering Task Force, 2013. Available at: http://www.rfc-editor.org/rfc/rfc7012.txt Last accessed: 2023-07-20.

[44] B. Claise, B. Trammell, and P. Aitken. Specification of the IPFIX Protocol for the Exchange of Flow Information. RFC 7011, Internet Engineering Task Force, 2013. Available at: http://www.rfc-editor.org/rfc/rfc7011.txt Last accessed: 2023-07-20.

[45] Cloudflare. Understanding and mitigating NTP-based DDoS attacks, 2014. Available at: https://blog.cloudflare.com/understanding-and-mitigating-ntp-based-ddos-attacks/ Last accessed: 2023-10-03.

[46] Cloudflare. DDoS attack trends for 2021 Q1, 2021. Available at: https://blog.cloudflare.com/ddos-attack-trends-for-2021-q1/ Last accessed: 2023-12-10.

[47] Cloudflare. Main components: Autonomous edge, 2023. Available at https://developers.cloudflare.com/ddos-protection/about/components/#autonomous-edge Last accessed: 2023-10-19.

[48] Cloudflare. Memcached DDoS Attack, 2023. Available at: https://www.cloudflare.com/learning/ddos/memcached-ddos-attack/ Last accessed: 2023-11-01.

[49] Cloudflare. One year of war in Ukraine: Internet trends, attacks, and resilience, 2023. Available at: https://blog.cloudflare.com/one-year-of-war-in-ukraine/ Last accessed: 2023-10-17.

[50] B. Collier, D. R. Thomas, R. Clayton, and A. Hutchings. Booting the Booters: Evaluating the Effects of Police Interventions. In *Proceedings of the ACM Internet Measurement Conference*, 2019.

[51] Comcast. COVID-19 Network Update, 2020. Available at: https://corporate.comcast.com/covid-19/network/may-20-2020 Last accessed: 2023-12-10.

[52] No More DDoS Coallition Consortium. National Anti-DDoS-Coalition, 2021. Available at: https://www.nomoreddos.org/en/ Last accessed: 2023-12-11.

[53] E. Cooke, M. Bailey, Z. M. Mao, D. Watson, F. Jahanian, and D. McPherson. Toward Understanding Distributed Blackhole Placement. In *Proceedings of the 2004 ACM Workshop on Rapid Malcode*, 2004.

[54] E. Cooke, F. Jahanian, and D. McPherson. The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets. In *Workshop on Steps to Reducing Unwanted Traffic on the Internet*, 2005.

[55] Nippon Telegraph & Telephone Corporation. Tier 1 ISPs: What They Are and Why They Are Important, 2006. Available at: `https://www.gin.ntt.net/wp-content/uploads/2020/01/IDC_Tier1_ISPs.pdf` Last accessed: 2023-09-28.

[56] M. Cotton, N. Vegoda, R. Bonica, and B. Haberman. Special-Purpose IP Address Registries. RFC 6890, Internet Engineering Task Force, 2013. Available at: `http://www.rfc-editor.org/rfc/rfc6890.txt` Last accessed: 2023-07-20.

[57] Cybersecurity & Infrastructure Security Agency. NTP Amplification Attacks Using CVE-2013-5211, 2023. Available at: `https://www.cisa.gov/news-events/alerts/2014/01/13/ntp-amplification-attacks-using-cve-2013-5211` Last accessed: 2023-10-03.

[58] J. Czyz, M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *Proceedings of the ACM Internet Measurement Conference*, 2014.

[59] J. Czyz, K. Lady, S. G. Miller, M. Bailey, M. Kallitsis, and M. Karir. Understanding IPv6 Internet Background Radiation. In *Proceedings of the ACM Internet Measurement Conference*, 2013.

[60] M .V. B. da Silva, A. S. Jacobs, R. J. Pfitscher, and L. Z. Granville. IDEAFIX: Identifying Elephant Flows in P4-Based IXP Networks. In *IEEE Global Communications Conference*, 2018.

[61] M. V. B. da Silva, A. S. Jacobs, R. J. Pfitscher, and L. Z. Granville. Predicting Elephant Flows in Internet Exchange Point Programmable Networks. In *IEEE Conference on Advanced Information Networking and Applications*, 2019.

[62] A. da Silveira Ilha, Â. C. Lapolli, J. A. Marques, and L. P. Gaspary. Euclid: A Fully In-Network, P4-Based Approach for Real-Time DDoS Attack Detection and Mitigation. *IEEE Transactions on Network and Service Management*, 2020.

[63] Daimler Trucks. Daimler trucks, 2023. Available at: `https://www.daimlertruck.com/en` Last accessed: 2023-11-16.

[64] A. Dainotti, R. Amman, E. Aben, and K. C. Claffy. Extracting Benefit from Harm: Using Malware Pollution to Analyze the Impact of Political and Geophysical Events on the Internet. *ACM Computer Communication Review*, 2012.

[65] A. Dainotti, K. Benson, A. King, K. C. Claffy, M. Kallitsis, E. Glatz, and X. Dimitropoulos. Estimating Internet Address Space Usage Through Passive Measurements. *ACM Computer Communication Review*, 2013.

[66] A. Dainotti, K. Benson, A. King, B. Huffaker, E. Glatz, X. Dimitropoulos, P. Richter, A. Finamore, and A. Snoeren. Lost in Space: Improving Inference of IPv4 Address Space Utilization. *IEEE Journal on Selected Areas in Communicatoin*, 2016.

[67] A. Dainotti, A. King, K. C. Claffy, F. Papale, and A. Pescapè. Analysis of a "/0" Stealth Scan from a Botnet. In *Proceedings of the ACM Internet Measurement Conference*, 2012.

[68] A. Dainotti, C. Squarcella, E. Aben, K. C. Claffy, M. Chiesa, M. Russo, and A. Pescape. Analysis of Country-wide Internet Outages Caused by Censorship. In *Proceedings of the ACM Internet Measurement Conference*, 2011.

[69] Danga Interactive. Memcached, 2018. Available at: `https://memcached.org/` Last accessed: 2023-10-11.

[70] DE-CIX. DE-CIX Virtual Get-together - Focus Middle East & Asia, 2020. Available at: `https://www.youtube.com/watch?v=DfPt10aopns` Last accessed: 2023-12-10.

[71] S. Deering and R. Hinden. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460, Internet Engineering Task Force, 1998. Available at: `http://www.rfc-editor.org/rfc/rfc2460.txt` Last accessed: 2023-11-21.

[72] DFN. German National Research and Education Network: COVID-19 Newsticker. Available at: `https://www.dfn.de/alle-meldungen-aus-dem-newsticker-zur-covid-19-pandemie/`, 2020.

[73] C. Dietzel, A. Feldmann, and T. King. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *Proceedings of the Passive and Active Measurement Conference*, 2016.

[74] C. Dietzel, M. Wichtlhuber, G. Smaragdakis, and A. Feldmann. Stellar: Network Attack Mitigation Using Advanced Blackholing. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*, 2018.

[75] X. Dimitropoulos, P. Hurley, A. Kind, and M. Stoecklin. On the 95-percentile billing method. In *Passive and Active Network Measurement*, 2009.

[76] M. Dischinger, M. Marcon, S. Guha, K. Gummadi, R. Mahajan, and S. Saroiu. Glasnost: Enabling End Users to Detect Traffic Differentiation. In *Proceedings of the USENIX Networked Systems Design and Implementation*, 2010.

[77] F. Dobrian, A. Awan, D. Joseph, A. Ganjam, J. Zhan, V. Sekar, I. Stoica, and H. Zhang. Understanding the Impact of Video Quality on User Engagement. In *Proceedings of the ACM SIGCOMM*, 2011.

[78] J. Dong, M. Chen, and A. Suryanarayana. Subcodes for BGP Finite State Machine Error. RFC 6608, Internet Engineering Task Force, 2012. Available at: `http://www.rfc-editor.org/rfc/rfc6608.txt` Last accessed: 2023-11-15.

[79] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman. A Search Engine Backed by Internet-Wide Scanning. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2015.

[80] Z. Durumeric, M. Bailey, and J. A. Halderman. An Internet-Wide View of Internet-Wide Scanning. In *Proceedings of the USENIX Security Symposium*, 2014.

[81] Z. Durumeric, E. Wustrow, and J. A. Halderman. ZMap: Fast Internet-Wide Scanning and its Security Applications. In *Proceedings of the USENIX Security Symposium*, 2013.

[82] W. Eddy. Transmission Control Protocol (TCP). RFC 9293, Internet Engineering Task Force, 2022. Available at: `http://www.rfc-editor.org/rfc/rfc9293.txt` Last accessed: 2023-10-11.

[83] R. Enns, M. Bjorklund, J. Schoenwaelder, and A. Biermann. Network Configuration Protocol (NETCONF). Request for Comments 6241, Internet Engineering Task Force, 2011. Available at: http://www.rfc-editor.org/rfc/rfc6241.txt Last accessed: 2023-07-20.

[84] European Commission. Commission and European Regulators Calls on Streaming Services, Operators and Users to Prevent Network Congestion, 2020. Available at: https://digital-strategy.ec.europa.eu/en/news/commission-and-european-regulators-calls-streaming-services-operators-and-users-prevent-network Last accessed: 2023-12-10.

[85] European Internet Exchange Association. BEREC Preliminary Assessment of the Underlying Assumptions of Payments from Large CAPs to ISPs, 2023. Available at: https://www.berec.europa.eu/system/files/2022-10/BEREC%20BoR%20%2822%29%20137%20BEREC_preliminary-assessment-payments-CAPs-to-ISPs_0.pdf Last accessed: 2023-10-03.

[86] European Internet Exchange Association. OBJECT: Fair Share Debate and Potential Impact of SPNP on European IXPs and Internet Ecosystem, 2023. Available at: https://www.euro-ix.net/media/filer_public/1a/e4/1ae40d86-95ea-460a-920d-3b335c2439d4/spnp_impact_on_ixps_-_final.pdf Last accessed: 2023-10-03.

[87] European Parliament. Network Cost Contribution Debate, 2023. Available at: https://www.europarl.europa.eu/RegData/etudes/ATAG/2023/745710/EPRS_ATA(2023)745710_EN.pdf Last accessed: 2023-10-03.

[88] Fastly. A New Architecture of the Modern Internet, 2023. Available at: https://www.fastly.com/network-map/ Last accessed: 2023-09-28.

[89] T. Favale, F. Soro, M. Trevisan, I. Drago, and M. Mellia. Campus Traffic and e-Learning during COVID-19 Pandemic. *Computer Networks*, 2020.

[90] Federal Communication Commission. The Most Important Part of the Internet You've Probably Never Heard Of, 2023. Available at: https://www.fcc.gov/news-events/notes/2023/08/02/most-important-part-internet-youve-probably-never-heard Last accessed: 2023-10-11.

[91] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *Proceedings of the ACM Internet Measurement Conference*, 2020.

[92] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. A Year in Lockdown: How the Waves of COVID-19 Impact Internet Traffic. *Communications of the ACM*, 2021.

[93] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827, Internet Engineering Task Force, 2000. Available at: http://www.rfc-editor.org/rfc/rfc2827.txt Last accessed: 2023-12-11.

[94] Forbes. Netflix Starts To Lift Its Coronavirus Streaming Restrictions, 2020. Available at: https://www.forbes.com/sites/johnarcher/2020/05/12/netflix-starts-to-lift-its-coronavirus-streaming-restrictions/#7bcba5bf4738 Last accessed: 2023-12-11.

[95] Internet Engineering Task Force. Introduction to the Internet Engineering Task Force, 2023. Available at: https://www.ietf.org/about/introduction/ Last accessed: 2023-10-19.

[96] Internet Engineering Task Force. Requests for Comments, 2023. Available at: https://www.ietf.org/about/introduction/#rfcs Last accessed: 2023-10-19.

[97] Deutsches Forschungsnetz. The National Research and Eduction Network, 2023. Available at: https://www.dfn.de/en/network/ Last accessed: 2023-09-28.

[98] J. Freudiger, E. De Cristofaro, and A. Brito. Controlled Data Sharing for Collaborative Predictive Blacklisting. In *Proceedings of the Conference on Detection of Intrusions and Malware & Vulnerability Assessment*, 2015.

[99] Riot Games. League of Legends: Troubleshooting Connection Issues, 2020. Available at: https://support-leagueoflegends.riotgames.com/hc/en-us/articles/201752664-Troubleshooting-Connection-Issues Last accessed: 2023-12-11.

[100] J. L. Garcia-Dorado, A. Finamore, M. Mellia, M. Meo, and M. Munafo. Characterization of ISP Traffic: Trends, User Habits, and Access Technology Impact. *IEEE Transactions on Network and Service Management*, 2012.

[101] O. Gasser, Q. Scheitle, P. Foremski, Q. Lone, M. Korczyński, S. D. Strowes, L. Hendriks, and G. Carle. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proceedings of the ACM Internet Measurement Conference*, 2018.

[102] T. A. Ghebreyesus. WHO Director-General's Opening Remarks at the Media Briefing on COVID-19, 2021. Available at: https://www.who.int/director-general/speeches/detail/who-director-general-s-opening-remarks-at-the-media-briefing-on-covid-19---11-march-2020 Last accessed: 2023-10-02.

[103] D. Gillman, Y. Lin, B. Maggs, and R. K. Sitaraman. Protecting Websites from Attack with Secure Delivery Networks. *IEEE Computer Magazine*, 2015.

[104] V. Giotsas, G. Smaragdakis, C. Dietzel, P. Richter, A. Feldmann, and A. Berger. Inferring BGP Blackholing Activity in the Internet. In *Proceedings of the ACM Internet Measurement Conference*, 2017.

[105] J. Gogna. What Is BGP, and Why Does the Internet Depend on It?, 2021. Available at: https://www.howtogeek.com/760059/what-is-bgp/ Last accessed: 2023-10-11.

[106] D. Gong, M. Tran, S. Shinde, H. Jin, V. Sekar, P. Saxena, and M. S. Kang. Practical Verifiable In-network Filtering for DDoS Defense. In *IEEE ICDCS*, 2019.

[107] Google. COVID-19 Community Mobility Report, 2020. Available at: https://www.google.com/covid19/mobility/ Last accessed: 2023-12-11.

[108] Google. Exponential Growth in DDoS Attack Volumes, 2020. Available at: https://cloud.google.com/blog/products/identity-security/identifying-and-protecting-against-the-largest-ddos-attacks Last accessed: 2023-12-11.

[109] Google. Keeping Our Network Infrastructure Strong amid COVID-19, 2020. https://www.blog.google/inside-google/infrastructure/keeping-our-network-infrastructure-strong-amid-covid-19/ Last accessed: 2023-12-11.

[110] H. Griffioen, K. Oosthoek, P. van der Knaap, and C. Doerr. Scan, Test, Execute: Adversarial Tactics in Amplification DDoS Attacks. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2021.

[111] Deutsche Börse Group. Deutsche Börse AG, 2023. Available at: https://www.deutsche-boerse.com/dbg-en/ Last accessed: 2023-11-16.

[112] A. Gupta, N. Feamster, and L. Vanbever. Authorizing Network Control at Software Defined Internet Exchange Points. In *Proceedings of the Symposium on SDN Research*, 2016.

[113] A. Gupta, R. MacDavid, R. Birkner, M. Canini, N. Feamster, J. Rexford, and L. Vanbever. An Industrial-Scale Software Defined Internet Exchange Point. In *Proceedings of the USENIX Annual Technical Conference*, 2016.

[114] A. Gupta, L. Vanbever, M. Shahbaz, S. P.Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. SDX: A Software Defined Internet Exchange. In *Proceedings of the ACM SIGCOMM*, 2014.

[115] S. Hares, R. Raszuk, D. McPherson, and M. Bacher. Dissemination of Flow Specification Rules. Request for Comments 8955, Internet Engineering Task Force, 2020. Available at: http://www.rfc-editor.org/rfc/rfc8955.txt Last accessed: 2023-07-20.

[116] R. Hiesgen, M. Nawrocki, A. King, A. Dainotti, T. Schmidt, and M. Wählisch. Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope. In *Proceedings of the USENIX Security Symposium*, 2022.

[117] David Hilbert. *Über die stetige Abbildung einer Linie auf ein Flächenstück*. Springer, 1891. (German).

[118] Honeynet Project. Know Your Enemy: Honeynets, 1999. Available at: https://www.honeynet.org/ Last accessed: 2023-12-11.

[119] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. In *Proceedings of the ACM SIGCOMM*, 2013.

[120] Information Sharing and Analysis Center. Distributed Denial of Service (DDoS) Attacks, 2021. Available at: https://www.aha.org/system/files/media/file/2021/03/distributed-denial-of-service-ddos-attacks-march-2021.pdf Last accessed: 2023-10-10.

[121] Information Sciences Institute. Internet Protocol version 4. RFC 791, Internet Engineering Task Force, 1981. Available at: http://www.rfc-editor.org/rfc/rfc791.txt Last accessed: 2023-11-21.

[122] Information Sciences Institute. Transmission Control Protocol. RFC 793, Internet Engineering Task Force, 1981. Available at: http://www.rfc-editor.org/rfc/rfc793.txt Last accessed: 2023-10-11.

[123] Intel. Let your networks soar with intel tofino expandable architecture, 2022. Available at: https://www.intel.com/content/dam/www/central-libraries/us/en/documents/2022-09/intel-tofino-expandable-architecture-paper.pdf Last accessed: 2023-10-19.

[124] International Telecommunication Union. Press Release: New 'State of Broadband' Report Carns of Stark Inequalities Laid Bare by COVID-19 Crisis, 2020. Available at: https://www.itu.int/en/mediacentre/Pages/PR20-2020-broadband-commission.aspx Last accessed: 2023-12-11.

[125] International Telecommunication Union. Individuals Using the Internet, 2023. Available at: https://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx Last accessed: 2023-10-11.

[126] L. Jakober. Akamai Mitigates Sophisticated 1.44 Tbps and 385 Mpps DDoS Attack, 2020. Akamai Blog, https://blogs.akamai.com/2020/06/akamai-mitigates-sophisticated-144-tbps-and-385-mpps-ddos-attack.html Last accessed: 2023-12-11.

[127] Laurent Joncheray. A Simple Active Attack Against TCP. In *Proceedings of the USENIX Security Symposium*, 1995.

[128] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti. Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem. In *Proceedings of the ACM Internet Measurement Conference*, 2017.

[129] M. Jonker and A. Sperotto. Measuring Exposure in DDoS Protection Services. In *IEEE Network and Service Management*, 2017.

[130] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras. Measuring the Adoption of DDoS Protection Services. In *Proceedings of the ACM Internet Measurement Conference*, 2016.

[131] B. Kantor, P. Karn, K. C. Claffy, J. Gilmore, H. Magnuski, B. Garbee, S. Hansen, B. Horne, J. Ricketts, J. Traschewski, and P. Vixie. AMPRNet, 2019. Available at: https://web.archive.org/web/20190719144558/https://www.ampr.org/amprnet/ Last accessed: 2023-12-11.

[132] T. Karagiannis, D. Papagiannaki, and M. Faloutsos. BLINC: Multilevel Traffic Classification in the Dark. In *Proceedings of the ACM SIGCOMM*, 2005.

[133] M. Karami, Y. Park, and D. McCoy. Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services. In *Proceedings of the World Wide Web Conferences*, 2016.

[134] N. Katta, A. Ghag, M. Hira, I. Keslassy, A. Bergman, C. Kim, and J. Rexford. Clove: Congestion-Aware Load Balancing at the Virtual Edge. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*, 2017.

[135] N. Katta, M. Hira, C. Kim, A. Sivaraman, and J. Rexford. HULA: Scalable Load Balancing Using Programmable Data Planes. In *Proceedings of the Symposium on SDN Research*, 2016.

[136] S. Katti, B. Krishnamurthy, and D. Katabi. Collaborating Against Common Enemies. In *Proceedings of the ACM Internet Measurement Conference*, 2005.

[137] L. A. D. Knob, R. P. Esteves, L. Z. Granville, and L. M. R. Tarouco. Mitigating Elephant Flows in SDN-Based IXP Networks. In *IEEE Symposium on Computers and Communications*, 2017.

[138] D. Kopp, C. Dietzel, and O. Hohlfeld. DDoS Never Dies? An IXP Perspective on DDoS Amplification Attacks. In *Proceedings of the Passive and Active Measurement Conference*, 2021.

[139] D. Kopp, J. Santanna, M. Wichtlhuber, O. Hohlfeld, I. Poese, and C. Dietzel. DDoS Hide & Seek: On the Effectiveness of a Booter Services Takedown. In *Proceedings of the ACM Internet Measurement Conference*, 2019.

[140] S. S. Krishnan and R. K. Sitaraman. Video Stream Quality Impacts Viewer Behavior: Inferring Causality using Quasi-Experimental Designs. In *Proceedings of the ACM Internet Measurement Conference*, 2012.

[141] M. Kührer, T. Hupperich, C. Rossow, and T. Holz. Exit from Cell? Reducing the Impact of Amplification DDoS Attacks. In *Proceedings of the USENIX Security Symposium*, 2014.

[142] C. Labovitz. Internet Traffic 2009-2019, 2020. Slides available at: `https://www.lacnic.net/innovaportal/file/4016/1/lacnog-internet-traffic-2009-2019.pdf` and recording available at: `https://www.youtube.com/watch?v=jGnVcCQUCdk` Last accessed: 2023-07-20.

[143] C. Labovitz. Pandemic Impact on Global Internet Traffic, 2020. NANOG 79. Slides available: `https://storage.googleapis.com/site-media-prod/meetings/NANOG79/2208/20200601_Labovitz_Effects_Of_Covid-19_v1.pdf` Last accessed: 2023-10-17.

[144] C. Labovitz. Tracing Volumetric DDoS to its Booter / IPHM Origins, 2021. NANOG 82. Slides available: `https://storage.googleapis.com/site-media-prod/meetings/NANOG82Virtual/2368/20210611_Labovitz_Tracing_Ddos_End-To-End_v1.pdf` Last accessed: 2023-12-11.

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

[146] A. Lakhina, K. Papagiannaki, M. Crovella, C. Diot, E. D. Kolaczyk, and N. Taft. Structural Analysis of Network Traffic Flows. In *Proceedings of the ACM SIGMETRICS*, 2004.

[147] T. Leighton. Can the Internet Keep Up with the Surge in Demand? Available at: `https://blogs.akamai.com/2020/04/can-the-internet-keep-up-with-the-surge-in-demand.html`, 2020.

[148] F. Lichtblau, F. Streibelt, T. Krüger, P. Richter, and A. Feldmann. Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses. In *Proceedings of the ACM Internet Measurement Conference*, 2017.

[149] K. Lougheed and Y. Rekhter. Border Gateway Protocol (BGP). RFC 1105, Internet Engineering Task Force, 1989. Available at: `http://www.rfc-editor.org/rfc/rfc1105.txt` Last accessed: 2023-10-11.

[150] K. Lougheed and Y. Rekhter. Border Gateway Protocol (BGP). RFC 1163, Internet Engineering Task Force, 1990. Available at: http://www.rfc-editor.org/rfc/rfc1163.txt Last accessed: 2023-10-11.

[151] K. Lougheed and Y. Rekhter. Border Gateway Protocol 3 (BGP-3). RFC 1267, Internet Engineering Task Force, 1991. Available at: http://www.rfc-editor.org/rfc/rfc1267.txt Last accessed: 2023-10-11.

[152] M. Luckie, R. Beverly, R. Koga, K. Keys, J. A. Kroll, and K. C. Claffy. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2019.

[153] M. Luckie, B. Huffaker, K. C. Claffy, A. Dhamdhere, and V. Giotsas. AS Relationships, Customer Cones, and Validation. In *Proceedings of the ACM Internet Measurement Conference*, 2013.

[154] A. Lutu, D. Perino, M. Bagnulo, E. Frias-Martinez, and J. Khangosstar. A Characterization of the COVID-19 Pandemic Impact on a Mobile Network Operator Traffic. In *Proceedings of the ACM Internet Measurement Conference*, 2020.

[155] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *Proceedings of the ACM Internet Measurement Conference*, 2009.

[156] MANRS Initiative. Mutually Agreed Norms for Routing Security, 2021. Available at: https://www.manrs.org/ Last accessed: 2023-12-11.

[157] L. F. C. Martins, Í. Cunha, and D. Guedes. An SDN-based Framework for Managing Internet Exchange Points. In *IEEE Symposium on Computers and Communications*, 2018.

[158] J. Mauch, J. Snijders, and G. Hankins. Default External BGP (EBGP) Route Propagation Behavior without Policies. RFC 8212, Internet Engineering Task Force, 2017. Available at: http://www.rfc-editor.org/rfc/rfc8212.txt Last accessed: 2023-11-15.

[159] Maxmind GeoLite2. GeoIP2 and GeoLite City and Country Databases, 2023. Available at: https://www.maxmind.com Last accessed: 2023-12-09.

[160] M. McKeay. Parts of a Whole: Effect of COVID-19 on US Internet Traffic, 2020. Available at: https://blogs.akamai.com/sitr/2020/04/parts-of-a-whole-effect-of-covid-19-on-us-internet-traffic.html Last accessed: 2023-12-11.

[161] M. McKeay. The Building Wave of Internet Traffic. Available at: https://blogs.akamai.com/sitr/2020/04/the-building-wave-of-internet-traffic.html Last accessed: 2023-12-11, 2020.

[162] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: Enabling Innovation in Campus Networks. *ACM Computer Communication Review*, 2008.

[163] Measurement Lab. The M-Lab NDT Data Set 2023-04-24 – 2023-04-30, 2023. Available at: https://measurementlab.net/tests/ndt Last accessed: 2023-12-09.

[164] L. Melis, G. Danezis, and E. De Cristofaro. Efficient Private Statistics with Succinct Sketches. In *Proceedings of the Network and Distributed System Security Symposium*, 2016.

[165] L. Melis, A. Pyrgelis, and E. De. Cristofaro. On Collaborative Predictive Blacklisting. *ACM Computer Communication Review*, 2018.

[166] Merit Network, Inc. ORION: Observatory for Cyber-Risk Insights and Outages of Networks. Available at: `https://www.merit.edu/initiatives/orion-network-telescope/` Last accessed: 2023-12-11, 2023.

[167] R. Miao, H. Zeng, C. Kim, J. Lee, and M. Yu. Silkroad: Making Stateful Layer-4 Load Balancing Fast and Cheap Using Switching ASICs. In *Proceedings of the ACM SIGCOMM*, 2017.

[168] B. Michael, C. Evan, J. Farnam, N. Jose, and W. David. The Internet Motion Sensor - A Distributed Blackhole Monitoring System. In *Proceedings of the Network and Distributed System Security Symposium*, 2005.

[169] G. Michaelson. How do We Count ASNs?, 2021. Available at: `https://blog.apnic.net/2021/03/19/how-do-we-count-asns/` Last accessed: 2023-10-02.

[170] Microsoft. Prepare Your Organization's Network for Microsoft Teams. `https://docs.microsoft.com/en-us/microsoftteams/prepare-network`, 2020.

[171] 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.

[172] Microsoft. Microsoft Azure Cloud Platform, 2023. Available at: `https://azure.microsoft.com/en-us` Last accessed: 2023-09-28.

[173] N. Mishra, S. Pandya, C. Patel, N. Cholli, K. Modi, P. Mehta, M. Chopade, S. Patel, and K. Kotecha. Memcached: An Experimental Study of DDoS Attacks for the Wellbeing of IoT Applications. *Sensors*, 2021.

[174] D. Moore and C. Shannon. The Spread of the Witty Worm. *IEEE Symposium on Security and Privacy*, 2005.

[175] D. Moore, C. Shannon, and J. Brown. Code-Red: A Case Study on the Spread and Victims of an Internet Worm. In *Proceedings of the ACM Internet Measurement Conference*, 2002.

[176] D. Moore, C. Shannon, G. Voelker, and S. Savage. Network Telescopes: Technical Report. *CAIDA*, 2004.

[177] D. Moore, G. Voelker, and S. Savage. Inferring Internet Denial-of-Service Activity. In *Proceedings of the USENIX Security Symposium*, Washington, D.C., 2001.

[178] D. Moore, V.Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the Slammer Worm. *IEEE Symposium on Security and Privacy*, 2003.

[179] C. Morales. NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us, 2018. Available at: https://www.netscout.com/blog/asert/netscout-arbor-confirms-17-tbps-ddos-attack-terabit-attack-era Last accessed: 2023-12-11.

[180] A. Mortensen, T. Reddy, and R. Moskowitz. DDoS Open Threat Signaling (DOTS) Requirements. RFC 8612, Internet Engineering Task Force, 2019. Available at: http://www.rfc-editor.org/rfc/rfc8612.txt Last accessed: 2023-07-20.

[181] H. Mostafaei, D. Kumar, G. Lospoto, M. Chiesa, and G. Di Battista. DeSI: A Decentralized Software-Defined Network Architecture for Internet eXchange Points. *IEEE Transactions on Network Science and Engineering*, 2021.

[182] Mozilla Foundation. Public Suffix List, 2020. https://publicsuffix.org/ Last accessed: 2023-12-11.

[183] F. Musumeci, A. C. Fidanci, F. Paolucci, F. Cugini, and M. Tornatore. Machine-Learning-Enabled DDoS Attacks Detection in P4 Programmable Networks. *Journal of Network and Systems Management*, 2022.

[184] RIPE NCC. RIPE Database Query, 2020. Available at: https://apps.db.ripe.net/db-web-ui/query Last accessed: 2023-12-11.

[185] Netflix. Reducing Netflix Traffic Where it's Needed While Maintaining the Member Experience, 2020. Available at: https://media.netflix.com/en/company-blog/reducing-netflix-traffic-where-its-needed Last accessed: 2023-12-11.

[186] Netnod. Netnod Reference 23-019, 2023. Available at: https://www.netnod.se/sites/default/files/2023-05/Future%20of%20Connectivity%20Netnod%20comments.pdf Last accessed: 2023-10-03.

[187] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-Performance Internet Applications. *SIGOPS Operating Systems Review*, 2010.

[188] Oberlo. How Many People Use the Internet?, 2023. Available at: https://www.oberlo.com/statistics/how-many-people-use-internet Last accessed: 2023-10-11.

[189] E. Osterweil, A. Stavrou, and L. Zhang. 21 Years of Distributed Denial-of-Service: A Call to Action. *IEEE Computer Magazine*, 2020.

[190] E. Osterweil, A. Stavrou, and L. Zhang. 21 Years of Distributed Denial-of-Service: Current State of Affairs. *IEEE Computer Magazine*, 2020.

[191] J. S. Otto, M. Sánchez, D. Choffnes, F. Bustamante, and G. Siganos. On Blind Mice and the Elephant: Understanding the Network Impact of a Large Distributed System. In *Proceedings of the ACM SIGCOMM*, 2011.

[192] R. Padmanabhan, A. Schulman, D. Levin, and N. Spring. Residential Links Under the Weather. *Proceedings of the ACM SIGCOMM*, 2019.

[193] PeeringDB. The Interconnection Database, 2020. https://www.peeringdb.com Last accessed: 2023-12-11.

[194] P. Phaal, S. Panchen, and N. McKee. InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks. RFC 3176, Internet Engineering Task Force, 2001. Available at: http://www.rfc-editor.org/rfc/rfc3176.txt Last accessed: 2023-07-20.

[195] S. Previdi, K. Talaulikar, C. Filsfils, K. Patel, S. Ray, and J. Dong. BGP-LS Extensions for Segment Routing BGP Egress Peer Engineering. RFC 9086, Internet Engineering Task Force, 2021. Available at: http://www.rfc-editor.org/rfc/rfc9086.txt Last accessed: 2023-11-15.

[196] E. Pujol, I. Poese, J. Zerwas, G. Smaragdakis, and A. Feldmann. Steering Hyper-Giants' Traffic at Scale. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*, 2019.

[197] L. Quan, J. Heidemann, and Y. Pradkin. Trinocular: Understanding Internet Reliability Through Adaptive Probing. In *Proceedings of the ACM SIGCOMM*, 2013.

[198] RADb. The Internet Routing Registry, 2021. Available at: https://www.radb.net Last accessed: 2023-12-11.

[199] E. Raftopoulos, E. Glatz, X. Dimitropoulos, and A. Dainotti. How Dangerous Is Internet Scanning? A Measurement Study of the Aftermath of an Internet-Wide Scan. In *Proceedings of the Network Traffic Measurement and Analysis Conference*, 2015.

[200] S. Ramanathan, A. Hossain, J. Mirkovic, M. Yu, and S. Afroz. Quantifying the Impact of Blocklisting in the Age of Address Reuse. In *Proceedings of the ACM Internet Measurement Conference*, 2020.

[201] S. Ramanathan, J. Mirkovic, M. Yu, and Y. Zhang. SENSS Against Volumetric DDoS Attacks. In *Proceedings of the Annual Computer Security Applications Conference*, 2018.

[202] REDImadrid. REDImadrid. Available at: https://www.redimadrid.es/ Last accessed: 2023-12-11, 2023.

[203] Y. Rekhter and T. Li. Border Gateway Protocol 4 (BGP-4). RFC 1771, Internet Engineering Task Force, 1995. Available at: http://www.rfc-editor.org/rfc/rfc1771.txt Last accessed: 2023-10-11.

[204] Y. Rekhter, T. Li, and S. Hares. Border Gateway Protocol 4 (BGP-4). RFC 4271, Internet Engineering Task Force, 2006. Available at: http://www.rfc-editor.org/rfc/rfc4271.txt Last accessed: 2023-10-11.

[205] M. H. Ribeiro, K. Gligoric, M. Peyrard, F. Lemmerich, M. Strohmaier, and R. West. Sudden Attention Shifts on Wikipedia Following COVID-19 Mobility Restrictions. In *Proceedings of the AAAI Conference on Web and Social Media*, 2021.

[206] P. Richter, M. Allman, R. Bush, and V. Paxson. A Primer on IPv4 Scarcity. *ACM Computer Communication Review*, 2015.

[207] P. Richter and A. Berger. Scanning the Scanners: Sensing the Internet from a Massively Distributed Network Telescope. In *Proceedings of the ACM Internet Measurement Conference*, 2019.

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

[209] B. Rodrigues, T. Bocek, A. Lareida, D. Hausheer, S. Rafati, and B. Stiller. A Blockchain-Based Architecture for Collaborative DDoS Mitigation with Smart Contracts. In *AIMS*, 2017.

[210] C. Rossow. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proceedings of the Network and Distributed System Security Symposium*, 2014.

[211] J. Ryburn. DDoS Mitigation Using BGP Flowspec, 2015. NANOG 63. Recording available: `https://www.youtube.com/watch?v=ttDUoDf6xzM&list=PLO8DR5ZGla8hhQXL_9_IRcw4HA9yu3JEC&index=15` Last accessed: 2023-12-11.

[212] Sachsen-Gigabit. Sachsen-Gigabit, 2023. Available at: `https://www.sachsen-gigabit.de/wps/portal/giga/cms/menu_main/sachsengigabit` Last accessed: 2023-11-16.

[213] B. Sanghani. COVID-19 & IXPs, 2020. RIPE 80, `https://ripe80.ripe.net/wp-content/uploads/presentations/27-ripe80-covid-ixp-1.pdf` Last accessed: 2023-12-11.

[214] B. Schilz and R. Maunier. Experience on Deploying a New Remote PoP During COVID-19 Restriction, 2020. RIPE 80, `https://ripe80.ripe.net/wp-content/uploads/presentations/26-Volterra-Ripe-connect-presentation.pdf` Last accessed: 2023-12-11.

[215] B. Schlinker, H. Kim, T. Cui, E. Katz-Bassett, H. V. Madhyastha, I. Cunha, J. Quinn, S. Hasan, P. Lapukhov, and H. Zeng. Engineering Egress with Edge Fabric: Steering Oceans of Content to the World. In *Proceedings of the ACM SIGCOMM*, 2017.

[216] C. Schuba, I. Krsul, M. Kuhn, E. Spafford, A. Sundaram, and D. Zamboni. Analysis of a Denial of Service Attack on TCP. In *IEEE Symposium on Security and Privacy*, 1997.

[217] V. Sekar, N. Duffield, O. Spatscheck, K. van der Merwe, and H. Zhang. LADS: Large-scale Automated DDoS Detection System. In *Proceedings of the USENIX Security Symposium*, 2006.

[218] P. Sermpezis, V. Kotronis, A. Dainotti, and X. Dimitropoulos. A Survey among Network Operators on BGP Prefix Hijacking. *ACM Computer Communication Review*, 2018.

[219] M. Z. Shafiq, L. Ji, A. X. Liu, and J. Wang. Characterizing and Modeling Internet Traffic Dynamics of Cellular Devices. In *Proceedings of the ACM SIGMETRICS*, 2011.

[220] SiliconAngle. Ukrainian Websites Knocked Offline in Massive DDoS Attack, 2022. Available at: `https://siliconangle.com/2022/02/23/ukrainian-websites-knocked-offline-massive-ddos-attack/` Last accessed: 2023-10-09.

[221] D. R. Simon, S. Agarwal, and D. A. Maltz. AS-Based Accountability as a Cost-effective DDoS Defense. In *Proceedings of the First Workshop on Hot Topics in Understanding Botnets*, 2007.

[222] K. Singh and A. Singh. Memcached DDoS Exploits: Operations, Vulnerabilities, Preventions and Mitigations. In *IEEE Conference on Computing, Communication and Security*, 2018.

[223] A. Sivanathan, H. H. Gharakheili, and V. Sivaraman. Can We Classify an IoT Device using TCP Port Scan? In *IEEE Conference on Information and Automation for Sustainability*, 2018.

[224] Z. M. Smith, E. Lostri, and J. A. Lewis. The Hidden Costs of Cybercrime on Government, 2020. Available at: https://www.mcafee.com/blogs/other-blogs/executive-perspectives/the-hidden-costs-of-cybercrime-on-government/ Last accessed: 2023-12-11.

[225] J. Snijders. Internet Network Operations During Pandemics, 2020. Available at: https://www.youtube.com/watch?v=tFeVlzBxICc Last accessed: 2020-09-05.

[226] S. Staniford, D. Moore, V. Paxson, and N. Weaver. The Top Speed of Flash Worms. In *ACM Workshop on Rapid Malcode*, 2004.

[227] Statista. Number of Internet and Social Media Users Worldwide as of July 2023, 2023. Available at: https://www.statista.com/statistics/617136/digital-population-worldwide/ Last accessed: 2023-10-11.

[228] K. Subramani, R. Perdisci, and M. Konte. IXmon: Detecting and Analyzing DRDoS Attacks at Internet Exchange Points. *ArXiv*, 2020. Available at: https://arxiv.org/pdf/2006.12555.pdf Last accessed: 2023-12-11.

[229] Tele2. Tele2, 2023. Available at: https://www.tele2.com/ Last accessed: 2023-11-16.

[230] Telegeography. State of the Network: Updates on COVID-19, 2020. Available at: https://www2.telegeography.com/network-impact Last accessed: 2023-12-11.

[231] Deutsche Telekom. Deutsche Telekom Netz, 2023. Available at: https://www.telekom.de/netz Last accessed: 2023-11-16 (German).

[232] The Register. Ukraine Hit by DDoS Attacks, Russia Deploys Malware, 2022. Available at: https://www.theregister.com/2022/02/23/ukraine_ddos_russia_malware/ Last accessed: 2023-10-09.

[233] C. Timberg. Your Internet is Working. Thank These Cold War-Era Pioneers Who Designed it to Handle Almost Anything last accessed: 2023-12-11. Available at: 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.

[234] M. Trevisan, D. Giordano, I. Drago, M. Mellia, and M. Munafo. Five Years at the Edge: Watching Internet from the ISP Network. In *Proceedings of the ACM Conference on Emerging Networking Experiments and Technologies*, 2018.

[235] University of Oregon. University of Oregon RouteViews Project, 2023. Available at: http://www.routeviews.org/routeviews/ Last accessed: 2023-12-09.

[236] USC/LANDER Project. Internet Addresses IPv4 Response History Dataset, 2023. Available at: https://ant.isi.edu/datasets/index.html Last accessed: 2023-12-09.

[237] P. Vervier, O. Thonnard, and M. Dacier. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In *Proceedings of the Network and Distributed System Security Symposium*, 2015.

[238] T. Vissers, T. Van Goethem, W. Joosen, and N. Nikiforakis. Maneuvering around clouds: Bypassing cloud-based security providers. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2015.

[239] F. von Bornstaedt, M. Roettgermann, I. Korthals, F. T. Johansen, and H. Lonsethagen. "The Sending Party Network Pays": A First Step Towards End-to-End Quality of Service. In *IEEE Conference on Intelligence in Next Generation Networks*, 2011.

[240] A. V. Vu, J. Hughes, I. Pete, B. Collier, Y. T. Chua, I. Shumailov, and A. Hutchings. Turning Up the Dial: the Evolution of a Cybercrime Market Through Set-up, Stable, and COVID-19 Eras. In *Proceedings of the ACM Internet Measurement Conference*, 2020.

[241] D. Wagner, D. Kopp, M. Wichtlhuber, C. Dietzel, O. Hohlfeld, G. Smaragdakis, and A. Feldmann. United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale. In *Proceedings of the ACM Conference on Computer and Communications Security*, 2021.

[242] D. Wagner, S. A. Ranadive, H. Griffioen, M. Kallitsis, A. Dainotti, G. Smaragdakis, and A. Feldmann. How to Operate a Meta-Telescope in your Spare Time. In *Proceedings of the ACM Internet Measurement Conference*, 2023.

[243] D. Wagner, M. Wichtlhuber, C. Dietzel, J. Blendin, and A. Feldmann. P4IX: A Concept for P4 Programmable Data Planes at IXPs. In *Proceedings of the ACM SIGCOMM Workshop on Future of Internet Routing & Addressing*, 2022.

[244] M. Wählisch, O. Maennel, and T. C. Schmidt. Towards Detecting BGP Route Hijacking Using the RPKI. *ACM Computer Communication Review*, 2012.

[245] A. Welzel, C. Rossow, and H. Bos. On Measuring the Impact of DDoS Botnets. In *Proceedings of the European Workshop on System Security*, 2014.

[246] Wirtschafts- und Sozialwissenschaftliches Institut. Deutlicher Anstieg: 24 Prozent der Erwerbstätigen arbeiten aktuell vorwiegend oder ausschließlich im Homeoffice , 2021. Available at: `https://www.boeckler.de/pdf/pm_wsi_2021_02_16.pdf` Last accessed: 2023-12-11 (German).

[247] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz. P4P: Provider Portal for Applications. In *Proceedings of the ACM SIGCOMM*, 2008.

[248] Q. Xu, J. Huang, Z. Wang, F. Qian, A. Gerber, and Z. M. Mao. Cellular Data Network Infrastructure Characterization and Implication on Mobile Content Placement. In *Proceedings of the ACM SIGMETRICS*, 2011.

[249] K-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat. Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering. In *Proceedings of the ACM SIGCOMM*, 2017.

[250] V. Yegneswaran, P. Barford, and D. Plonka. *On the Design and Use of Internet Sinks for Network Abuse Monitoring*. Springer, 2004.

[251] C. Zakaria, A. Trivedi, M. Chee, P. Shenoy, and R. Balan. Analyzing the Impact of Covid-19 Control Policies on Campus Occupancy and Mobility via Passive WiFi Sensing. *Transactions on Spatial Algorithms and Systems*, 2020.

[252] M. Zhang, G. Li, S. Wang, C. Liu, A. Chen, H. Hu, G. Gu, Q. Li, M. Xu, and J. Wu. Poseidon: Mitigating Volumetric DDoS Attacks with Programmable Switches. In *Proceedings of the Network and Distributed System Security Symposium*, 2020.

[253] X. Zhang, H. Hsiao, G. Hasker, H. Chan, A. Perrig, and D. G. Andersen. SCION: Scalability, Control, and Isolation on Next-Generation Networks. In *IEEE Symposium on Security and Privacy*, 2011.

[254] Z. Zhao, H. Sadok, N. Atre, J.C. How, V. Sekar, and J. Sherry. Achieving 100Gbps Intrusion Prevention on a Single Server. In *Proceedings of the USENIX Operating Systems Design and Implementation*, 2020.