Technische Universität Berlin
Fakultät für Elektrotechnik und Informatik
Lehrstuhl für Intelligente Netze
und Management Verteilter Systeme

Empirical Analysis of the Effects and the Mitigation of IPv4 Address Exhaustion

vorgelegt von

M.Sc. Philipp Richter
geboren in Berlin

von der Fakultät IV – Elektrotechnik und Informatik
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Naturwissenschaften
- Dr. rer. nat. -

genehmigte Dissertation

Promotionsausschuss:

Vorsitzender:	Prof. Dr.-Ing. Sebastian Möller, Technische Universität Berlin
Gutachterin:	Prof. Anja Feldmann, Ph.D., Technische Universität Berlin
Gutachter:	Prof. Vern Paxson, Ph.D., University of California, Berkeley
Gutachter:	Prof. Steve Uhlig, Ph.D., Queen Mary University of London

Tag der wissenschaftlichen Aussprache: 2. August 2017

Berlin 2017

Abstract

IP addresses are essential resources for communication over the Internet. In IP version 4, an address is represented by 32 bits in the IPv4 header; hence there is a finite pool of roughly 4B addresses available. The Internet now faces a fundamental resource scarcity problem: The exhaustion of the available IPv4 address space. In 2011, the Internet Assigned Numbers Authority (IANA) depleted its pool of available IPv4 addresses. IPv4 scarcity is now reality.

In the subsequent years, IPv4 address scarcity has started to put substantial economic pressure on the networks that form the Internet. The pools of available IPv4 addresses are mostly depleted and today network operators have to find new ways to satisfy their ongoing demand for IPv4 addresses. Mitigating IPv4 scarcity is not optional, but mandatory: Networks facing address shortage have to take action in order to be able to accommodate additional subscribers and customers. Thus, if not confronted, IPv4 scarcity has the potential to hinder further growth of the Internet. Addressing is a collective and global effort, and interconnectivity among networks forms the very basis of the Internet. At the same time, the decentralized nature of the Internet and independent decisions by its different stakeholders create a complex and opaque problem space.

Different approaches to mitigating IPv4 address scarcity in the short and in the long term exist. The solution space includes increasing the utilization of already available IPv4 address resources, introducing IPv4 address multiplexing techniques, incrementally transitioning to IPv6 (but maintaining IPv4 compatibility), or purchasing IPv4 address space on address markets. Individual network operators make independent decisions on which approach—or combination of approaches—to pursue. Each option has different ramifications, benefits, and consequences for the individual networks and their connected end users. At the same time, the increasing and disparate deployment of different mitigation techniques and the coexistence of two addressing protocols (IPv4 and IPv6) adds heterogeneity to the network and hence changes the overall connectivity and communication structure of the Internet. In spite of the pressing relevance of the topic, we lack a comprehensive understanding of IPv4 address exhaustion and its effects, as well as of the pervasiveness, impact, and ramifications of the different mitigation strategies.

This dissertation provides a systematic empirical analysis of the phenomenon of IPv4 address space exhaustion, its effect on the Internet as a whole, and its stakeholders individually. We first provide an empirical lay-of-the-land of the history and the current status of the IPv4 address space and illuminate the interplay between policy and governance decisions and address use. We then develop techniques that allow us to measure the potential, the pervasiveness, and the ramifications of the individual mitigation strategies that network operators may choose to pursue. In particular, we measure global IPv4 address activity patterns, which allows us to both study exhaustion effects on address activity and to assess the potential for increasing the utilization of the IPv4 address space. We develop tools to detect Carrier-Grade NAT (CGN) presence on the Internet at scale, and identify dominant properties of deployed CGN instances and their respective impact. Lastly, we examine aspects of IPv6 adoption, where we measure inter-domain connectivity and traffic carried over IPv4 and IPv6, and interactions between IPv4 and IPv6 traffic in detail, allowing us to pinpoint barriers and challenges for IPv6 adoption.

We strive to both illuminate the broader impact of address exhaustion on the Internet and its structure as well as to provide practical insights to support the ongoing process of mitigating address scarcity. Our results can serve as a basis for network operators and policymakers to make informed decisions on how to approach IPv4 address exhaustion. They may also inform future measurement studies and the design of operational systems, which need to adapt to this increasingly complex environment.

Zusammenfassung

IP-Adressen stellen essentielle Ressourcen für die Kommunikation über das Internet dar. In Version 4 des IP-Protokolls sind 32 Bit für IP-Adressen reserviert, was in einer Maximalanzahl von etwa 4 Milliarden verfügbaren Adressen resultiert. Das Internet ist nun mit einer fundamentalen Ressourcenknappheit konfrontiert: Der Erschöpfung des global verfügbaren IPv4-Adressraums. Im Jahr 2011 hat die Internet Assigned Numbers Authority (IANA) ihre Adressreserven erschöpft.

In den folgenden Jahren hat Adressknappheit zu substantiellem ökonomischen Druck auf die Netze des Internets geführt. Reserven verfügbarer IPv4-Adressen sind weitgehend erschöpft und Netzbetreiber müssen neue Methoden finden, um ihren fortlaufenden Adressbedarf zu decken. Mitigation von IPv4-Adressknappheit ist unumgänglich: Netzbetreiber müssen aktiv werden, um zusätzliche Kunden aufnehmen zu können. Wird IPv4-Adressknappheit also nicht konfrontiert, kann sie potentiell das weitere Wachstum des Internets behindern. Adressierung ist ein kollektives und globales Unterfangen, gleichzeitig führen jedoch die dezentrale Struktur des Internets und die unabhängigen Entscheidungen der unterschiedlichen Interessenvertreter zu einer komplexen und undurchsichtigen Problematik.

Verschiedene Ansätze zur kurzfristigen und langfristigen Mitigation von IPv4-Adressknappheit existieren. Die Optionen reichen von Effizienzsteigerung bereits vorhandener Adressblöcke, über Einführung von IPv4-Multiplexingtechnologien, inkrementeller Transition zu IPv6 (bei weitergehender IPv4-Kompatibilität), bis hin zum Erwerb zusätzlicher Adressblöcke auf Adressmärkten. Individuelle Netzbetreiber treffen unabhängige Entscheidungen bezüglich des präferierten Ansatzes, oder der Kombination von Ansätzen. Jede Option hat unterschiedliche Vor- und Nachteile und Konsequenzen für die individuellen Netze und deren Endnutzer. Gleichzeitig verstärkt die zunehmende und ungleichartige Anwendung unterschiedlicher Mitigationstechnologien die Heterogenität des Internets und verändert die Struktur des Internets. Ungeachtet der Relevanz der beschriebenen Thematik existieren erstaunlich wenig Erkenntnisse über IPv4-Adressknappheit und ihre Folgeeffekte, als auch über die Verbreitung und die Auswirkungen der unterschiedlichen Mitigationstechnologien.

Diese Dissertation bietet eine systematische empirische Analyse des Phänomens der Adressknappheit des Internets. Hierbei werden sowohl Effekte auf das Internet generell, als auch auf seine individuellen Teilhaber analysiert. Zunächst analysieren wir die aktuelle und historische Situation und Entwicklung des IPv4-Adressraums und studieren das Zwischenspiel zwischen Adressvergabepraktiken und tatsächlicher Adressverwendung. Im Folgenden entwickeln wir Technologien und Analysen, welche uns erlauben, das Potential, die Verbreitung, und die Folgen individueller Mitigationsstrategien zu erörtern. Im Speziellen messen wir Aktivitätsmuster des globalen IPv4-Adressraums, um sowohl die Effekte der Adressknappheit zu erörtern, als auch das Potential für effizientere Adressnutzung aufzuzeigen. Wir entwickeln Methoden zur Detektierung von Carrier-Grade NAT (CGN) im Internet, und zeigen die Eigenschaften und Auswirkungen der identifizierten Instanzen auf. Schlussendlich analysieren wir einige Aspekte der IPv6 Adoption, wobei wir uns sowohl auf Inter-Domain Konnektivität und Datenverkehr fokussieren, als auch auf die Interaktionen zwischen IPv4- und IPv6-Datenverkehr. Hier zeigen wir Barrieren und Herausforderungen im Zuge der Adoption von IPv6 auf.

Diese Arbeit zielt darauf ab, sowohl den weitgehenden Effekt von Adressknappheit auf das Internet aufzuzeigen, als auch praktisch relevante Ergebnisse zu liefern, welche den fortlaufenden Prozess der Mitigation von IPv4-Adressknappheit unterstützen können. Unsere Ergebnisse können sowohl Netzbetreibern als auch Kontrollorganen helfen, informierte Entscheidungen bezüglich der Mitigation von IPv4-Adressknappheit zu treffen. Darüber hinaus geben wir Einblicke in die Eigenschaften des sich weiterentwickelnden Internets bei gleichzeitiger IPv4-Verknappung, was den Entwurf und die Entwicklung zukünftiger Studien und operativer Systeme unterstützen kann.

Acknowledgments

The last years were certainly some of the most intense, fascinating, and demanding of my life. I had the chance to learn from some of the sharpest, most determined, and most supportive people I have ever met. I owe you my deepest gratitude.

Anja Feldmann, thank you for your continuous support over all these years. I will never forget writing research papers with you at 3AM on a Saturday night. I doubt that many advisors are willing to go as far as you to make your students flourish and succeed. I am grateful for having a person as strong as you as my advisor. Georgios Smaragdakis, my sincere thanks goes to you for your unparalleled support and help in all these years. From day one until graduation, you were on hand with help and advice. A true friend. Walter Willinger, thank you very much for your support, particularly in the beginning of this endeavor. Nikolaos Chatzis, I thank you for your support. I would have never gotten into Internet Measurement without my former Bachelor advisor and now dissertation committee member: Steve Uhlig. Steve, what did you get me into? A huge thank you at this point.

I still vividly remember my first visit at ICSI in Berkeley in 2013. My two internships at ICSI laid the groundwork for this dissertation. Vern Paxson, I am grateful for your advice, guidance, and leadership on this journey. Your determination and vision for the big picture, but also attention to detail and rigor, continuously brought out the best in me. I have learned a lot from you. Mark Allman, I admire your straightforward and sometimes sarcastic way. It has been a great pleasure to work with and learn from you. Randy Bush, your experience and help has proven invaluable to this work. Narseo Vallina-Rodriguez, I thank you for your contribution to this work. I have rarely encountered persons that are as invested in a project as you are. Florian Wohlfahrt, I thank you for your contribution. I am grateful to the ICSI staff, particularly Jaci Considine and Maria Eugenia Quintana, for arranging my two internships at ICSI in 2013 and 2014.

Arthur Berger, when you approached me in Vancouver in 2014 for the possibility of an internship in Akamai, I immediately felt that we would have a productive and great future together. I am deeply grateful for your guidance and support. Working with and learning from you has been a great pleasure. David Plonka, your sharp mind and perspectives on things always impressed me and I truly enjoy working with you. My gratitude goes to the Custom Analytics group in Akamai, particularly Matt Olson, Keung Chi Ng, and Steve Hoey, for your support prior to, during and after my internship in Akamai.

During the last years, I spent a lot of days (and nights) in the office. Enric Pujol, I am very grateful for having had you as my office mate and as collaborator for many years. We shared intense moments and in the end, we got *something* done. Matthias Rost, I thank you for your support and for the refreshing breaks in your office. Niklas Semmler, I truly enjoyed having you as office mate. I am grateful to our administrators, in particular Christian Struck, for making various exotic experimental setups possible.

I would not have succeeded without the constant support of my family and my friends, particularly Kathrin, Mascha, and Moritz. Your support in the last years was priceless. Last but certainly not least, I want to thank you, Anne, for your tireless support in the many years from me starting my studies in computer science a decade ago up until this very day. Thank you.

I am grateful for partial support for my studies from the DFG/German Research Foundation (Leibniz Prize funds, FKZ FE 570/4-1) and from the US National Science Foundation via research grant CNS-1111672.

List of Publications

Parts of this thesis are based on the content included in the set of published papers listed below. These papers have been co-authored with other researchers. All my co-authors are acknowledged here. I thank all of them for their highly valuable contribution.

International Conferences

Understanding the Share of IPv6 Traffic in a Dual-Stack ISP.
ENRIC PUJOL, PHILIPP RICHTER AND ANJA FELDMANN.
Proceedings of the Passive and Active Measurement Conference (PAM), 2017.

Beyond Counting: New Perspectives on the Active IPv4 Address Space.
PHILIPP RICHTER, GEORGIOS SMARAGDAKIS, DAVID PLONKA AND ARTHUR BERGER.
Proceedings of the ACM Internet Measurement Conference (IMC), 2016.
★ Awarded “Best Paper”.

A Multi-perspective Analysis of Carrier-Grade NAT Deployment.
PHILIPP RICHTER, FLORIAN WOHLFART, NARSEO VALLINA-RODRIGUEZ, MARK ALLMAN, RANDY BUSH, ANJA FELDMANN, CHRISTIAN KREIBICH, NICK WEAVER AND VERN PAXSON.
Proceedings of the ACM Internet Measurement Conference (IMC), 2016.
★ Awarded with the 2017 “IRTF Applied Networking Research Prize”.

Distilling the Internet’s Application Mix from Packet-Sampled Traffic.
PHILIPP RICHTER, NIKOLAOS CHATZIS, GEORGIOS SMARAGDAKIS, ANJA FELDMANN AND WALTER WILLINGER.
Proceedings of the Passive and Active Measurement Conference (PAM), 2015.

Peering at Peerings: On the Role of IXP Route Servers.
PHILIPP RICHTER, GEORGIOS SMARAGDAKIS, NIKOLAOS CHATZIS, JAN BOETTGER, ANJA FELDMANN AND WALTER WILLINGER.
Proceedings of the ACM Internet Measurement Conference (IMC), 2014.

International Journals

Lost in Space: Improving Inference of IPv4 Address Space Utilization.
ALBERTO DAINOTTI, KARYN BENSON, ALISTAIR KING, BRADLEY HUFFAKER, EDUARD GLATZ, XENOFONTAS DIMITROPOULOS, PHILIPP RICHTER, ALESSANDRO FINAMORE AND ALEX C. SNOEREN.
IEEE Journal on Selected Areas in Communications, 34(6), 2016.

A Primer on IPv4 Scarcity.
PHILIPP RICHTER, MARK ALLMAN, RANDY BUSH AND VERN PAXSON.
ACM Computer Communication Review, 45(2), 2015.
★ Selected for “Best of CCR 2015”.

Contents

1 Introduction

The Internet's design philosophy has facilitated enormous, rapid, and de-centralized growth of the Internet, from a specialized research facility to a massive network of global importance. As of 2016, the Internet had more than 3.5 billion users worldwide [271] and emerged as one of the key drivers for human progress and prosperity. The proliferation of the Internet has promoted enormous economic opportunity, fueled distributed innovation, education, information exchange, and developed into an indispensable pillar of modern society. To accomplish this, the Internet depends on the availability of a set of protocols that allow universal exchange of data across heterogeneous networks. Despite the Internet's rapid expansion and critical importance, the core protocols supporting the networks' fundamental functions have seen very little change over time.

One of these core functions is global addressing. As originally designed, the Internet architecture calls for IP (Internet Protocol) addresses to uniquely identify devices. In IP version 4 [223], which is the dominant addressing protocol in today's Internet, an address is represented by 32 bits in the IPv4 header. Hence there is a finite pool of roughly 4B addresses available.[1] When IPv4 was introduced in 1981, the Internet was a small networked system connecting a few dozens of networks, and 4B of addresses seemed plenty. What followed was unprecedented growth of the Internet which went along with a dramatic increase in the number of connected devices and users. In 2011, the IANA (Internet Assigned Numbers Authority), which governs global IP address assignments, exhausted its pool of globally available IPv4 addresses. This event was a watershed moment for the Internet: We ran out of IPv4 addresses, one of the Internet's key resources.

The predicament of IPv4 address exhaustion was long foreseen. The Internet community realized the possibility of exhaustion early on, with the IAB (Internet Advisory Board)[2] discussing alternatives, mitigation strategies, and possible successors already in 1991, describing the looming issue of IPv4 address shortage as "clear and present danger to the future successful growth of the worldwide Internet" [154]. In 1998, the successor of IPv4, IPv6 [103] was standardized. Still, as of 2017, the majority of Internet traffic is carried over IPv4, and only a minority of hosts can communicate over IPv6 [46, 96, 130]. Despite substantial efforts within the Internet Engineering Task Force (IETF) and within the network

[1] Fewer than 4B addresses are usable in practice, which we discuss in Chapter 2.

[2] IAB initially referred to "Internet Activities Board", which was later renamed to "Internet Advisory Board".

operator community to promote a smooth transition, IPv4 scarcity commenced before we reached substantial levels of IPv6 deployment.

Addressing on the Internet is a collective and global effort, and interconnectivity among networks forms the very basis of the Internet. Networks facing IPv4 address scarcity must ensure continuous interconnectivity with the rest of the Internet at all costs, to provide service to their customers. In turn, networks wanting to transition over to IPv6 are dependent on the rest of the IPv4 Internet transitioning as well, including networks with large supplies of unused IPv4 addresses. Thus, the incentive structure on how to approach and solve the problem of IPv4 address exhaustion is highly ambiguous. The need for global coordination and governance of the Internet and its address resources, limited measures to enforce such governance, and various business interests and resource availability of the different stakeholders of the Internet result in a complex and opaque tussle.

IPv4 address scarcity has the potential to hinder ongoing growth of the Internet and has started to put substantial economic pressure on ISPs (Internet Service Providers). IPv4 addresses are—in most regions of the world—exhausted and network operators have to find new ways to satisfy their ongoing demand for addresses. Mitigating IPv4 scarcity is not optional, but mandatory: ISPs facing address shortage *have* to take action, in order to be able to accommodate additional subscribers, customers, and devices. The difficulty to get additional IPv4 address space has led ISPs to pursue various alternatives to mitigate their individual address scarcity issues. The solution space ranges from increasing the utilization of already available IPv4 address resources, applying IPv4 address multiplexing techniques, incrementally transitioning to IPv6 (but maintaining IPv4 compatibility), purchasing IPv4 address space on address markets, or a combination thereof. All approaches come with direct costs for the involved networks. The choice of which approach to adapt when and the connected necessary investments have a direct impact on the economic success or non-success of individual ISPs. Thus, network operators now have to make business-critical decisions on which mitigation strategy to adapt when.

The coexistence of two addressing protocols (IPv4 and IPv6) and the increasing and disparate deployment of different approaches to mitigate IPv4 address scarcity in different networks also affects fundamental connectivity and communication patterns of the Internet. Mechanisms that multiplex the currently usable IPv4 address space further erode the Internet's end-to-end principle and we now face a scenario in which the usage profile of individual IP addresses is highly dynamic and can vary vastly. This poses new challenges for application developers, for a range of operational systems, and for Internet measurements in general. Applications need to be able to function in such an increasingly complex environment. Systems that rely on the notion of an IP address, e.g., host reputation systems, geolocation systems, and systems to measure Internet reliability need to adapt to this challenging situation. Internet measurements that track growth and spread of the Internet need to be revised to account for a situation with increasing complexity of IPv4 activity.

The exhaustion of the IPv4 address space has far-reaching consequences for the Internet as a whole and its stakeholders, including ISPs, content providers, regulators, and end users, individually. Yet, in spite of its relevance, IPv4 address exhaustion has received comparably little in the way of systematic empirical assessment.

Goal of this Dissertation

The goal of this dissertation is to provide a systematic analysis of the phenomenon of IPv4 address space exhaustion and its effect on the Internet as a whole, and its individual stakeholders. We strive to both yield practical insights to inform the ongoing process of mitigating IPv4 address scarcity as well as to assess the broader impact of address scarcity on the Internet, its properties, and its structure. The scope of this dissertation involves both assessment of the historical development and the current status of the IPv4 address space and its management as well as measurement, analysis, and evaluation of the different options that individual networks may choose from to cope with address scarcity. Our results can aid network operators and policymakers to make informed decisions on how to approach IPv4 address scarcity. They also provide insight into properties of the evolving Internet in the face of IPv4 exhaustion, which may inform future measurement studies and the design of operational systems.

The remainder of this chapter is structured as follows: We introduce the Internet Protocol Suite and the role of the Internet Protocol in Section 1.1. We introduce the problem of IPv4 address shortage in the face of a growing Internet in Section 1.2, and outline the solution space in Section 1.3. We highlight the contributions of this dissertation in Section 1.4 and outline the contents of this dissertation in Section 1.5.

1.1 The Internet Protocol Suite

The Internet is a layered system. It consists of a large number of interconnected packet networks that support communication between host computers using the Internet protocols. To communicate, hosts must implement some or all of the protocols that together comprise the Internet protocol suite. The interplay of the individual protocols eventually enables end-to-end data communication.

The Internet protocol suite can be expressed using the abstraction of the TCP/IP model, shown in Figure 1.1. It consists of four layers.[3] Each layer provides a certain functionality to the above layer and relies on the functionality of the layer below.

- **Application Layer:** In the top layer, application-layer protocols facilitate the exchange of data for specific applications, and are almost exclusively implemented in software on end hosts communicating with each other. Today, we see a large diversity of application-layer protocols. The HTTP protocol, enabling the Web, is the most prominent application-layer protocol in today's Internet [233]. Emails are delivered using the SMTP protocol, and locally fetched from clients using the POP3(S) or IMAP(S) protocol. Other protocols, such as the BitTorrent protocol [3], allow Peer-to-Peer file-sharing between end hosts.
- **Transport Layer:** The transport layer transports data between two communicating applications. The two most common transport layer protocols are TCP (Transmission Control Protocol [224]) and UDP (User Datagram Protocol [220]). TCP and UDP provide the abstraction of a *socket* to the application layer, which allows applications to exchange data with remotely located applications. TCP provides a *reliable bytestream* between two end hosts, whereas UDP provides a connectionless and unreliable service. The transport layer functions are also implemented mostly in software in end hosts communicating with each other.
- **Internet Layer:** The task of the Internet layer is to carry packets of data from the source to a destination in the Internet. The core function here is both the determination of a *route* between source and destination as well as the actual forwarding of packets. The network layer provides a unified abstraction of network end hosts, using IP addresses, and thus allows heterogeneous networks to become interconnected. Network layer functionality is typically implemented in a mixture of soft- and hardware.

[3] Several different notions of the TCP/IP model exist in the literature. Some representations feature and discuss five layers, splitting the *Network Interface* layer into a *Link Layer* and a *Physical Hardware Layer*, e.g., [172, 266]. In this dissertation, we adapt the notion presented in [92], presenting four layers and noting the *physical hardware* as an additional layer.

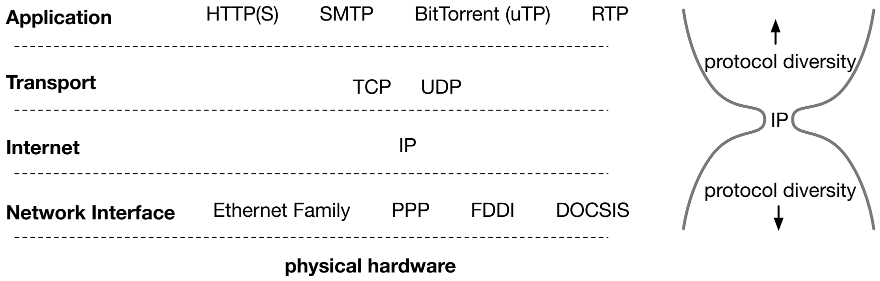

Figure 1.1: TCP/IP model for data transmission on the Internet with example protocols.

- **Network Interface Layer (Link Layer):** Protocols in this layer facilitate exchange of data frames from one network element to an adjacent network element. This involves both the physical transmission of bits, dependent on the transmission medium (e.g., copper wire, fiber, or radio frequency) as well as medium access control mechanisms for shared mediums (e.g., in the case of RF), low-level buffering and, sometimes, reliable frame delivery. Link layer functionality is typically implemented in hardware, e.g., within network interface cards in hosts, routers, and switches.

An application developer does not need to implement networking functionalities, but can use the abstraction of a socket, as provided by the transport layer. The implementation of the transport layer (e.g., TCP) within the host's operating system takes care of assembly/disassembly of a byte stream into packets, but does not need to route packets to their destinations. The actual transmission and routing of packets happens solely within the Internet layer, which in turn relies on the network interface layer for sending data to the next adjacent node. Hence, transport and application protocols are implemented and operate within the end hosts. The network interface layer, on the other end of the spectrum, operates within a certain network segment, e.g., within a home network or within an ISP. Data on its way from a source to a destination typically traverses multiple network segments, carried by various link layer protocols over various media (fiber, copper, RF, etc.). The Internet layer is the only layer in the protocol stack that requires truly global coordination across all networks between the two hosts that exchange data, since its task is *global* addressing and routing of packets.

The Internet Layer: The Narrow Waist of the Protocol Stack

The Internet layer glues the networks that comprise the Internet together. It carries a multitude of traffic from various applications from the upper layers, and IP packets in turn are carried over a multitude of different links with different characteristics on the lower layer. Thus, the Internet layer unifies both heterogeneous links and networks (lower layers) as well as a diverse set of applications (upper layers). The Internet layer protocol (IP) penetrates *all* parts of the Internet and must be spoken by all devices on the Internet that work on the Internet layer. This includes every end host, router, and middlebox (that accesses the network layer) on the path. It provides a simple and unified abstraction of the underlying networks and provides no guarantees for reliable data delivery. Today, only two Internet layer protocols are in use: IPv4 and IPv6.

- **Internet Protocol Version 4 (IPv4):** IPv4 is the initial IP protocol and was standardized in 1981 [223]. The IPv4 address space is 32 bits in size, allowing for a theoretical maximum of 4.29B unique IP addresses. Addresses are typically written in decimal notation, where each byte boundary is denoted with a dot, e.g., *130.149.0.1*. Some 592M IPv4 addresses are in a reserved state, e.g., for multicast, private use, and future use [93]. This leaves roughly 3.7B unique IPv4 addresses that are globally routable, i.e., can be assigned to end hosts to communicate via the Internet.

- **Internet Protocol Version 6 (IPv6):** IPv6, the designated successor of IPv4, was introduced in 1998 [103] and provides a vastly larger address space. IPv6 addresses are of 128-bit length, allowing for a theoretical maximum of 2^{128} or $3.4 * 10^{38}$ unique addresses. IPv6 addresses are represented in 8 groups of 16 bits each, where each group is written as four hexadecimal digits and the groups separated by colons. Groups that only consist of zeros can be omitted in presentation view. An example of such an address is *2001:638:809:ff1f::2:1*. The lower 64 bits of an IPv6 address are referred to as the interface identifier, and the upper 64 bits are referred to as the network portion of the address. This reduces the number of available IPv6 addresses to 2^{64}, or $1.8 * 10^{19}$ addresses.

Over the course of the last decades, the Internet has seen enormous innovation both in terms of the available applications as well as in terms of networking technology that enabled massive capacity, bandwidth, and latency improvements. Its applications have evolved from simple text-based protocols, such as USENET [141], to applications delivering rich media content including audio and video in real-time. This innovation primarily took place in the application layer, fueled by exponentially growing computational capabilities of end hosts. At the same time, the capacity, bandwidth, and latency of the Internet vastly improved, mostly due to the availability of cheap and fast networking hardware and advances in methods to transmit data at high rates via fiber, copper, and RF links (e.g., the success of the Ethernet family [150]). As of today, we see a large number of different protocols in the application layer as well as an increasing number of protocols and hardware in the network interface layer. The center of the protocol stack however, the Internet layer, consists of only two protocols, which leads to the "hourglass shape" of the Internet Protocol Suite, as shown in Figure 1.1.

1.2 Internet Growth and Address Shortage

When IPv4 was introduced in 1981, the possible number of 4.3 billion unique IP addresses seemed vast, given that at that time only several dozen networks were interconnected on the Internet [221]. Scarcity was not seen as a looming issue, and hardware constraints, such as limited and expensive memory, led to a standardization of 32-bit long address fields in the IP header.

In the following 35 years, the Internet grew at unprecedented and unforeseen rates. With the decommissioning of the NSFNET backbone in 1995 [199] and the growth of the commercial Internet in the following years, the number of connected networks, devices, and users inflated dramatically. Figure 1.2 shows estimates of the International Telecommunication Union (ITU) on the number of Internet users and broadband subscriptions. The number of Internet users grew from one billion in 2005 up to 3.7 billion in 2016. During the same time, the number of fixed broadband subscriptions grew by a factor of 4 and mobile broadband subscriptions grew by a factor of 14 between 2007 and 2016. Notably, 2015 marks the first year in which the number of mobile subscriptions exceeds the estimated number of Internet users. This is due to an increasing number of users with multiple broadband subscriptions.

As of 2016, the ITU reports more than 3.6B mobile broadband subscriptions and more than 880M fixed broadband subscriptions worldwide [271]. Besides a steadily growing user base, the number of connected *devices* can be expected to grow even faster due to implementation of networking capabilities in an increasing number of devices, ranging from toys, smart meters, home appliances, to vehicles and industrial appliances: the "Internet of Things". Estimations for the number of connected devices range as high as predicting more than 8B devices connected via the Internet as of 2016, which could exceed 20B devices by 2020 [122].

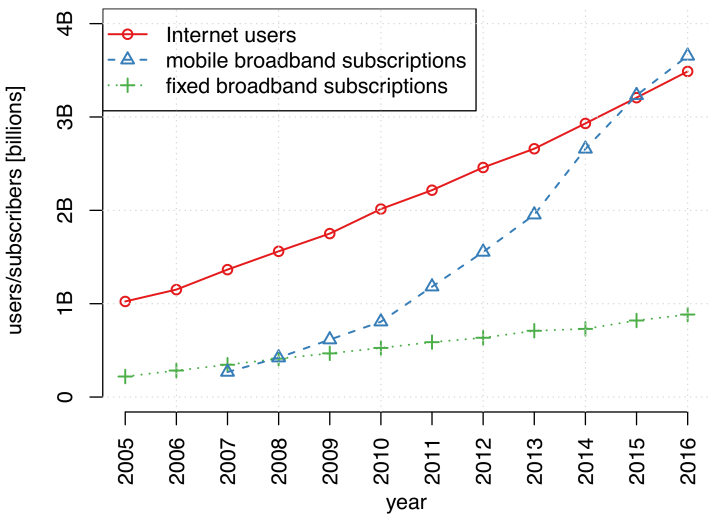

Figure 1.2: Number of Internet users and broadband subscriptions over time (ITU data [271]).

The IPv4 address space provides a theoretical maximum of 4.29B unique addresses, of which 3.7B are globally routable. Assuming perfect utilization of the available address space, one IP address per subscription, and sound estimations by the ITU, we would have reached the limit of routable IPv4 addresses in 2015 (per Figure 1.2). However, utilizing every single IPv4 address is impossible, since the network routing system cannot keep enough state to deal with each individual address and therefore aggregates addresses into blocks. Addresses need to be globally unique and are assigned in the form of network blocks by the IANA and the respective Regional Internet Registries (RIRs).[4] As of 2011, the IANA exhausted its address block reserves. In the subsequent years, the four largest RIRs exhausted their regional address pools as well. As of 2017, only ISPs residing in the African region can still receive IPv4 addresses from their registry.

1.3 Overcoming IPv4 Scarcity

The size of the IPv4 address space has proven insufficient to provide a unique IP address for every connected device on the Internet. The Internet community can now choose from a set of possible options to mitigate IPv4 scarcity, both in the short term as well as in the long term. More precisely, network operators facing shortage of IPv4 addresses can choose to follow one, or a combination of, three approaches to cope with their ongoing demand for IPv4 address space: *(i)* more efficiently use the available IPv4 address space, *(ii)* multiplex IPv4 address space using address sharing techniques such as Carrier-Grade NAT (CGN), and/or *(iii)* transition to IPv6.

Approach (i): **Use IPv4 address space more efficiently:** The IANA and four out of five of the Regional Internet Registries have depleted their address pools for allocation. However, not all allocated IP addresses are in active use. Addresses were given out generously in the early years of IPv4, and there was little incentive to conserve and efficiently use the allocated address blocks. As of 2017, large portions of the IPv4 address space remain unrouted, and are thus not in public use. Recent studies (including our own) find utilization levels for the routed address space at around 50% to 60% [100, 283]. Hence, significant usable IPv4 address space remains. Increasing the utilization efficiency of the available IPv4 address space requires both technical as well as governance measures. Technical measures

[4] We describe the framework in place to assign IP addresses in detail in Section 2.

include first the detection of possibly un- or underutilized address ranges and then consolidation (by renumbering) of address assignments within networks. Governance measures include the adaption of address management policies that allow re-assignment of already allocated address blocks, e.g., through the establishment of policies for address transfer markets.

Approach (ii): **Multiplex IPv4 address space with Carrier-Grade NAT:** Alternatively, or additionally, networks can get by with many fewer addresses by multiplexing. That is, end hosts are assigned internal IP addresses, which are then translated onto fewer publicly routed IPv4 addresses, making simultaneous use of a single IP address by multiple end hosts possible. Today, numerous approaches to perform address sharing at scale are available and are already in use by large ISPs. However, widespread use of NAT raises concerns about eroding end-to-end connectivity and semantics, which could directly affect end-users' connectivity and has the potential to limit the functionality of some applications. Moreover, large-scale NAT deployment raises concerns by law enforcement agencies due to the erosion of attribution of IP addresses to end-users [108,114], and has the potential to affect IP address-based measurement systems like host reputation and geolocation systems.

Approach (iii): **Transition to IPv6:** IPv6, the successor to IPv4, extends the routable address space by several orders of magnitude. IPv6 reflects the natural long-term solution to the address scarcity problem. However, transitioning the Internet over to IPv6 presents a mammoth task, since it involves replacing the very central protocol of the Internet protocol suite. The transition towards an IPv6 Internet is ongoing, and has been ongoing for almost two decades, fueled by intense efforts to promote adoption within the networking community. Broad and widespread IPv6 adoption was intended to happen long before IPv4 address scarcity commenced [66, 94]. However, the fraction of both IPv6-enabled networks as well as native IPv6 traffic on the Internet remains comparably small—adoption of IPv6 remains problematic and only slowly increases. As of January 2017, Google reports some 15% of clients accessing Google to be IPv6 enabled, with adoption rates as high as 48% in Belgium, around 30% in the US and Germany, and increasing support in other European countries. Nonetheless, the per-host adoption rate still ranges at or below 1% for most countries, including China, Russia, as well as European countries such as Spain [130]. IPv6 is by itself not compatible with IPv4 and requires complex transition mechanisms to ensure compatibility between the IPv4 and IPv6 Internet (e.g., [205, 261]).

Individual network operators make independent decisions on which approach, or which combination of approaches, to pursue. The feasibility of each alternative or combination of approaches depends on the individual requirements and access to resources of individual networks. Networks with large supplies of lightly used IPv4 address space might decide to consolidate their address space usage to accommodate more hosts, or to become sellers in the IPv4 marketplace. Other networks might opt to purchase additional IPv4 addresses or to deploy Carrier-Grade NAT, if their network properties allow for it. Networks might opt to adopt IPv6, in addition to IPv4, in combination with IPv4-prolonging mechanisms, or to postpone IPv6 adoption until more widespread deployment.

Each approach comes with different benefits and costs for the respective network operator, and results in different ramifications for end users and the broader topology and structure of the Internet. From a research perspective, several issues arise: What is the potential that each of these options have? Which option is preferred by what type of networks and why? What are the ramifications for networks deploying them and their customers? What technologies will have what impact on the Internet, its topology and its measurability? Will we find ourselves in a long-term situation in which IPv4, IPv6 and technologies like CGN operate in parallel? What will be the corresponding impact on the Internet? Following up on the presented solution space, this dissertation seeks to answer some of these questions.

1.4 Contributions

The goal of this dissertation is to measure the effects and the mitigation of IPv4 address exhaustion. We first provide an empirical lay-of-the-land of the history and the current status of the IPv4 address space. We then study the pervasiveness and the ramifications of the three possible mitigation strategies: *(i)* more efficiently use the available IPv4 address space, *(ii)* multiplex IPv4 address space using address sharing techniques such as Carrier-Grade NAT (CGN), and/or *(iii)* transition to IPv6.

In particular, we make the following four contributions:

1) **Systematic framing of IPv4 address exhaustion**
We develop a systematic framing of the fundamentals of the IPv4 address exhaustion phenomena and connected issues. We study how the current ecosystem of IPv4 address space has evolved since the standardization of IPv4, leading to the complex and opaque scenario we face today. We outline the evolution in address space management and its effects on address space usage patterns, identifying key factors of the scarcity issue. We characterize the solution space to overcome scarcity and open the perspective of address blocks as virtual resources, highlighting issues such as differentiation between address blocks, resource certification, and issues arising when transferring address space.

2) **Analysis of IPv4 address space activity and utilization**
We perform a detailed study of Internet-wide IPv4 address activity, using different techniques and vantage points. We are able to identify and attribute address activity patterns to network restructurings, user behaviors, and, in particular, various address assignment practices. Drawing upon our metrics of spatio-temporal address utilization, traffic volume, and estimates of relative host counts, we illuminate how IPv4 exhaustion manifests itself in the various regions of the world. We pinpoint, and quantify the prevalence of, underlying addressing mechanisms and their effect on activity and utilization. Based on our metrics, we provide upper bounds for potential utilization increase.

3) **Analysis of Carrier-Grade NAT deployment**
We present the first broad and systematic study of the deployment and the behavior of Carrier-Grade NAT instances in the Internet. We develop a methodology to detect the existence of hosts behind CGNs by extracting non-routable IP addresses from peer lists we obtain by crawling the BitTorrent DHT. We complement this approach with improvements to the Netalyzr [170] troubleshooting service, enabling us to determine a range of indicators of CGN presence as well as detailed insights into properties of CGNs. Combining our data sources we illustrate the scope of CGN deployment in the Internet, and report on characteristics of deployed CGNs and their effects on end users' connectivity.

4) **Analysis of IPv4/IPv6 connectivity and traffic**
We study several important aspects of IPv6 adoption in the Internet. We provide an analysis of the control- and data planes at two Internet Exchange Points (IXPs), emphasizing IPv4 and IPv6 differences and recent developments. To identify individual traffic components, we devise a methodology to classify the application mix from sparsely sampled traffic traces. We then assess the interplay between available IPv4 and IPv6 connectivity and actual traffic exchange in a residential network. Our analyses allow us to highlight the status of IPv6 deployment, and to assess pitfalls when measuring adoption. Moreover, we identify barriers, opportunities, and challenges for ongoing IPv6 adoption.

1.5 Overview and Structure

The structure of this dissertation closely follows our aforementioned contributions. In this section, we present a more detailed overview and put the contents and contributions pertaining to each chapter into context.

Chapter 2: A Brief History of the IPv4 Address Space

In Chapter 2, we study how the management and governance of the IPv4 address space has evolved over time and we assess the eventual impact of governance decisions on address consumption and routing. Our observations lead to a set of insights and challenges that arise, given that IPv4 addresses are now scarce virtual resources. Having a comprehensive understanding of "all things concerning IPv4 addresses" serves as an essential foundation for our assessment of IPv4 address exhaustion. It also allows us to draw conclusions on how policies affect IPv4 address exhaustion over time, what effects these developments have on the status quo, and what new issues arise.

Obtaining an accurate picture of the IPv4 address space proves difficult: The Internet rapidly transformed from a research project into a network of global importance. Institutions and frameworks in charge of allocating IP addresses had to be established on-the-fly. Pressured by increasing IPv4 scarcity concerns and unprecedented growth of the Internet, policies to manage IP address space were introduced. Facing IPv4 exhaustion and growth in address markets, policies needed to be adapted, challenged by legal constraints, and agreed upon by different stakeholders. Many aspects of this transition are poorly documented and the resulting legal and technical framework is complex and opaque. To shed light into this complex space, we survey, analyze, and combine historic and current policy documents, empirical data on past and present address allocations, address transfers, and data from the global routing table.

Our contributions can be summarized as follows:

- We study the evolution of 35 years of management of the IPv4 address space. Compiling and interpreting the available historic documents, we discuss the various institutions that govern address space assignments past and present and highlight differences and salient technical developments that influence these processes.
- We augment our survey of management practices with empirical data on address allocations by the various institutions as well as data from the global routing table. Our empirical assessment of address allocations allows us to study the impact of management decisions on the reality of IPv4 address space consumption and use.
- We discuss IPv4 addresses as virtual resources. We elaborate on current issues of resource certification of IP addresses and differentiation across IP addresses. We study the actual developments in the emerging IPv4 address transfer market, where we first introduce the various policies introduced by the governing institutions and connected challenges. We then provide empirical data on the number of transferred IPv4 address blocks in the various regions of the world.

This comprehensive understanding of the interplay between management and policy aspects of IPv4 address resources, and their development over time serves as the foundation for our analysis. Hereby, we highlight the political component of IPv4 address exhaustion, which directly affects network operators in need of additional IPv4 address space.

Chapter 3: Address Activity in the Wake of Exhaustion

In Chapter 3, we analyze global IPv4 address space activity. We are interested in understanding what portions of the IPv4 address space are indeed in active use, to what degree individual address blocks are used, and if address exhaustion manifests itself in usage patterns. Empirical data on IPv4 address space utilization forms not only the basis for policymaking when it comes to address transfers or reassignments, but also more fundamentally helps to understand the scope of the problem and the effects of exhaustion on the properties of the broader Internet. For individual network operators, methods to track, quantify, and measure address space activity help to assess the potential for address space consolidation and to precisely identify regions with little utilization. For Internet measurements and operational systems, e.g., host reputation systems, a detailed understanding of address activity, and how it changes in the face of IPv4 scarcity, serves as a fundamental basis.

Yet, measuring address activity at scale poses a challenge, since no single method or vantage point can comprehensively capture address activity. Today, we face a situation in which individual addresses and address ranges vary in their periods of activity and in that activity's nature and volume. As a result, questions about the number of IP addresses active at a point of time, let alone their usage characteristics, are difficult to answer. Detailed measurements and metrics that capture this diversity are necessary.

Our contributions can be summarized as follows:

- We assess the visibility of different methods and datasets into global IPv4 address activity, considering different aggregation levels. We then leverage server logs detailing the number of client requests of a major global Content Distribution Network (CDN) to study address activity. We find that our CDN logs allow for unprecedented visibility into the global IPv4 space on a per-IP granularity. Our CDN logs illuminate a detailed picture of world-wide activity from 1.2B unique IPv4 addresses contacting the CDN, which is unseen in other data sources.
- We analyze the historic evolution of IPv4 address activity in terms of active IPv4 addresses, finding the effect of IPv4 address exhaustion clearly reflected in our measurements. We study dynamics of the IPv4 address population on different timescales, finding that the pool of active IPv4 addresses is in constant flux and that most dynamics are the result of changes within networks and are largely hidden from the global routing table. We identify a variety of address block activity patterns and are able to attribute them to network restructuring, user behaviors, and, in particular, various address assignment practices.
- We introduce metrics that allow us to quantify prevalent addressing practices at scale and to study current address space utilization on a detailed level, allowing us to identify additional utilization potential in the IPv4 address space. We augment our binary address activity metrics with corresponding traffic volumes and relative host count metrics to study the relationship between address scarcity and traffic dynamics. Combining our different metrics of address activity, we then derive Internet-wide demographics of the active IPv4 address space.

Our study of IPv4 address activity provides unprecedented visibility into the current status of the IPv4 address space. It hence serves our goal to foster the discussion on how to mitigate IPv4 scarcity, by uncovering the potential for increasing the utilization of the current IPv4 address space, and by showing how address scarcity affects some fundamental properties of the broader Internet. Our metrics can be used to track IPv4 activity on an ongoing basis.

Chapter 4: Carrier-Grade NAT to the Rescue

In Chapter 4, we develop techniques that allow us to detect the presence of Carrier-Grade NAT (CGN) devices, study their prevalence in the Internet, their properties, associated challenges, and their effect both on end users and the broader Internet. CGNs actively prolong the lifetime of IPv4 and hence CGN deployment is a controversial topic [145], both politically and technically. As a result, ISPs typically do not publicly advertise their CGN deployment, resulting in a lack of available data. Data on CGN deployment is crucial both to track how networks cope with IPv4 address scarcity as well as to understand the direct ramifications that CGN deployment has for the affected end users. However, NATs operate transparently, which makes their detection challenging. It is even more challenging to distill the properties of these devices and their potential impact on users' connectivity. To shed light on CGN deployment in the Internet, we develop suitable tools and apply them on a broad scale.

Our contributions can be summarized as follows:

- We develop a technique to detect CGN deployment at scale by leveraging information gathered by crawling the BitTorrent Distributed Hash Table (DHT). With our technique, we are able to detect CGN deployment in remote ASes without the need for customized measurements inside these networks. We complement these measurements with data and detection methods using the Netalyzr [170] platform to detect CGN presence using instrumented devices within the respective ISPs.
- We provide a network-wide view of CGN deployment in today's Internet. Our measurements cover more than 60% of ISPs that connect end-users to the Internet. We present statistics on global CGN deployment, per geographical region and per the type of ISP. Our measurements present a first overview of the pervasiveness of CGN deployment in today's Internet. We enrich our findings with a survey that we conducted among network operators.
- We extend our detection techniques with tailored measurements leveraging our two vantage points, which allow us to study properties of the identified CGN deployments. In particular, we study IP address and port allocation and mappings, the degree of resource sharing, and topological properties of CGNs. Our measurements allow us to study the real-world effect on the connectivity of end users behind CGN, and to identify challenges that ISPs face when deciding to deploy CGN.

The techniques we develop prove efficient not only in detecting CGN deployment, but also to analyze the salient properties of CGNs. Our analysis uncovers unexpected broad deployment of CGN across many networks, which should both interest ISPs as well as regulators. We also identify pitfalls and ramifications that this mitigation technique has for end users and for ISPs wanting to implement it, and their impact on the Internet topology.

Chapter 5: The Shift Towards IPv6

In Chapter 5, we present measurements that add to our understanding of the current status of IPv6 deployment in the Internet. Broad transition to IPv6 was intended to happen many years ago [66, 94]. Yet, this transition did not happen as quickly as anticipated and thus we now face a situation with concurrent IPv4 and IPv6 deployments. The Internet is a running system and replacing the very core protocol (IP) presents an ongoing and highly complex challenge. Our analysis identifies and highlights some of the challenges faced during this transition. Just finding suitable vantage points and adoption metrics to track the degree of actual IPv6 deployment and use becomes increasingly complicated. Tracking IPv6 deployment and identifying challenges and barriers for its adoption, as well as for the migration of traffic from IPv4 over to IPv6, are of critical importance both for network operators wanting to transition to IPv6 and for the Internet community at large. Our contributions can be summarized as follows:

- We develop and apply techniques to reconstruct the control plane over IPv4 and IPv6 at two Internet Exchange Points (IXPs), where hundreds of networks interconnect. We study the options that networks have to interconnect at such IXPs and assess some discrepancies of the IPv4 and IPv6 control plane. We then take traffic into account and contrast our connectivity-based findings with the actual amount of traffic carried on peering links. While our findings suggest increasing IPv6 inter-domain connectivity, our traffic analysis cautions that connectivity is only part of the puzzle.
- We develop a methodology to classify the application-layer protocol of the exchanged traffic at our IXP. Traffic classification proves difficult here, since our available dataset consists of sparse packet samples. Applying our method to the exchanged traffic, we find that while the aggregated application mix seems consistent with widely reported statistics, each individual peering link carries a distinct set of applications. Given the heterogeneity of traffic carried over individual peering links, we can expect disparate shifts to IPv6 for each individual link.
- We then shift our focus and study traffic carried over IPv4 and IPv6 in a residential network. This vantage point gives us the ability to precisely discern what portions of the traffic *are* and *could be* exchanged over either IPv4, IPv6, or both protocols. Our method tags subscribers and their respective traffic flows to be either IPv4 or IPv6 capable as well as with the protocol they are carried over. With our analysis, we draw a detailed picture of barriers that prevent traffic from being carried over IPv6 as well as of the potential share of traffic that could immediately be shifted over to IPv6, once service providers offer their content over IPv6.

Our analyses allow us to comment on the current status of IPv6 deployment and pitfalls when measuring adoption, and identify barriers and potential for adoption, and resulting challenges.

2
A Brief History of the IPv4 Address Space

In this chapter, we survey the 35 years of the community's management and use of IP address space, and the resulting challenges that the community now faces in light of IPv4 exhaustion. Certainly, the exhaustion of the IPv4 address space is a multi-faceted problem. It does not only present a technical issue, but is heavily affected and influenced by management and governance decisions. When IPv4 was standardized, scarcity was seen as a minor issue, IPv4 addresses did not pose a valuable commodity, and management was informal. Yet, with the proliferation of the Internet, more formal and distributed frameworks for managing IP address allocations were established. With the prospect of exhaustion, these management bodies now faces yet another challenge: Managing allocation of what is now a scarce, and yet seemingly fungible, good. While scientists and engineers often ignore such "soft" issues, policies and regulations ultimately shape what we can deploy in production.

We first discuss the evolution of the relevant policy structures and organizations in place to manage the IPv4 address space. We then shift our focus over to the demand side and complement our view with empirical data on address allocations and data from the global routing table. We next correlate the different management practices with their impact on address consumption and address use, as seen from the global routing table. Having this empirical lay of the land of the IPv4 address space, we then discuss IPv4 addresses as virtual resources. The various legal frameworks in place over the last decades created a situation in which different address blocks come with different policies and are hence subject to different regulations. We highlight the current developments in the Internet Governance space and study regulatory aspects of IPv4 address markets. Again, we complement our observations with empirical data on address space transfers.

2.1 Evolution of Address Management

From its standardization in 1981 [223] until now, the management of IP addresses has undergone drastic changes. The changes were mainly a result of the evolution of the Internet from a research network to a global commercial network and the corresponding need to establish international frameworks to manage its critical resources. We elaborate on this evolution in three time phases: The *Early Registration*

Phase starting with the arrival of IPv4, the *Needs-based Provision* Phase leading to the modern registry framework, and the recently entered *Depletion and Exhaustion* Phase.

2.1.1 First Phase: Early Registration

Initially, address blocks were allocated quite informally, with Jon Postel serving as the "czar" personally attending to each allocation. Postel periodically re-published RFCs enumerating the current address assignments ("please contact Jon to receive a number assignment") [221]. At that time, addresses block allocations came in one of three *classes*: class A networks (2^{24} addresses), class B (2^{16}), and class C (2^{8}). Classful addressing required a *network identifier* of one of these distinct types, meaning that an operator requesting significantly more addresses than provided by a particular threshold would instead be allocated a larger class network. Given the coarse-grained nature of the differences between these classes, this policy led to heavy internal fragmentation and thus waste of address space.

Early (1981) in the Internet's evolution, parties had already registered 43 class A networks, allocating in total more than 700M addresses [221]—vastly larger than the number of hosts actually connected at that time.[1] While scarcity in address blocks was not mentioned as a looming issue, the notion of different sizes of networks (A, B and C) suggests early recognition of the finite nature of network address blocks and the need for some sort of stewardship when parceling them out to different parties. The responsibility for the management of address space led to formalizing the notion of the IANA (first mentioned in IETF documents in 1990 [231]), and, in the same timeframe Solensky, drawing upon allocation statistics, predicted IPv4 address exhaustion in the late '90s [262].

2.1.2 Second Phase: Needs-based Provision

The need for a more distributed and parsimonious framework to allocate IP addresses—shaping the modern registry structure—appeared at least as early as 1990 [85], with further refinements in 1992 and 1993 [126]. The discussion at that time included the need to distribute the administration of IP address blocks to regional registries, covering distinct geographic regions to better serve the respective local community—consciously fragmenting the registry. In addition, classless inter-domain routing (CIDR)[2] and private address space[3] arose in 1993–4 to further conserve publicly routable address space.

The modern framework of Regional Internet Registries (RIRs), established in the years between 1992 and 2005, was very specific that conservation of address space was a primary goal [142]. Five RIRs emerged, run as non-profit organizations: RIPE for Europe in 1992, APNIC for the Asia-Pacific in 1993, ARIN for the North-Americans in 1997, LACNIC for Latin America in 2002, and AfriNIC for Africa in 2005.

The RIRs manage the distribution of IP address resources, each according to their local policies. Policies within the RIRs are created using a community process; for details of the process for each RIR, see [28,41,48,174,242]. For the most part, anyone can submit an RIR policy proposal which then undergoes an open discussion and review process, usually carried out on mailing lists as well as in working group

[1] Address registration statistics in terms of number of blocks and block holders varied heavily among the first published RFCs.

[2] CIDR [121] supported routing and forwarding on bit-aligned, as opposed to the previous byte-aligned, variable-length prefixes. CIDR denotes prefixes as a combination of an IP address and a corresponding network mask, such as *1.1.2.0/23* specifying a network with 2^{9} IP addresses that share their top 23 bits. Introducing CIDR required significant network restructuring efforts as well as changes to routing protocols and hardware (see, for example, [120]).

[3] Reserved address blocks not globally routable, and thus usable concurrently within multiple networks as long as the given hosts do not require globally reachable IP addresses [229].

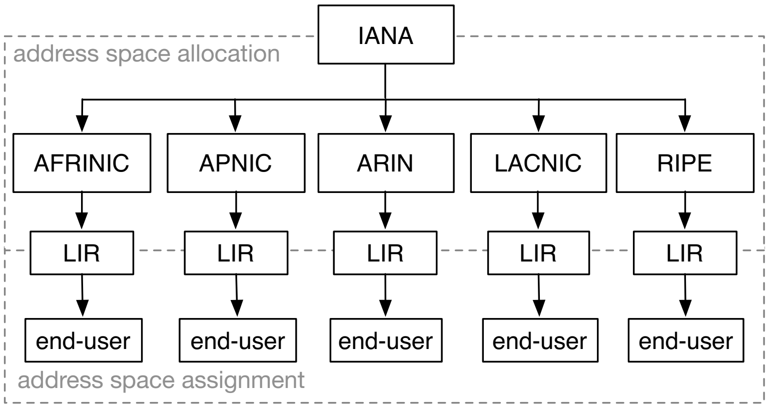

Figure 2.1: Regional Internet Registry system.

and policy meetings. Adopting a proposal requires the community to reach a degree of consensus as reflected in these discussions.

We sketch the structure of the RIR framework in Figure 2.1. The IANA serves as the parent organization, *allocating* large free address blocks (/8, i.e., 2^{24} addresses, granularity) to an RIR once their regional free pool reaches a low threshold level. The RIRs then further allocate subsets of these address blocks to their members, the so-called LIRs (Local Internet Registries), which are mainly ISPs. The LIRs then *assign* address blocks to either smaller ISPs or for their own infrastructure. Thus, the allocation of a block reserves it for (future) use, while the assignment of parts of an allocation puts that subset into use.[4] ISPs decide for themselves whether to become LIRs—meaning entering a direct contractual relationship with the respective RIR—or to rely upon their upstream provider to assign address space to them.[5]

During the needs-based provision phase, one of the key principles was that receivers of address space (LIRs) must *justify* their need for the address blocks they receive, though some RIRs no longer require this in some contexts (e.g., RIPE for “last /8” allocations—see below). LIRs requesting new allocations had to provide documentation showing a sufficient *utilization rate* of prior allocations, namely that a given proportion of prior allocations were assigned to end-users as well as documentation of the intended use of new allocations. RIRs might also request more detailed information, such as how many and what type of hosts were connected to assigned subnets. LIRs passed these policies on to their end-customers. For example, if a customer of a transit provider required blocks of IP addresses, they had to fill out corresponding LIR-specific forms detailing the intended use of that block (e.g., [209]).

The global nature of the Internet raises the question of when an organization is supposed to be served by a specific geographic region. Whether or not a company can become an LIR under a specific RIR is not explicitly stated, but is usually determined by the registered address of a company. However, there also are organizations with multiple subsidiaries as members of—and holding address resources from—multiple RIRs [148]. While address blocks are theoretically assigned and “used” by organizations operating inside the region of the allocating RIR, current policies are inconsistent regarding explicit constraints on the geographic region of an address block’s actual use in the sense of where connected devices reside.[6]

[4] The APNIC and LACNIC regions also have *National Internet Registries* (NIRs), which act as intermediaries between the RIR and the LIR to serve specific countries. For example, JPNIC does so for Japan.

[5] Under some circumstances, RIRs can also assign address space directly to end users—so-called provider independent (PI) address space. Such assignments usually arise due to the user’s need to connect to multiple upstream providers (multi-homing), and thus requiring independent address space. For more details, see for example § 4.2 in the ARIN NRPM [49] or the RIPE policy documents [249]. For a practical guide for operators, see [67].

[6] ARIN has a policy proposal to explicitly allow out-of-region use [50], and a RIPE official stated that RIPE permits out-of-region use, assuming that the address blocks originate at some point from within the RIPE region (e.g., by a router at

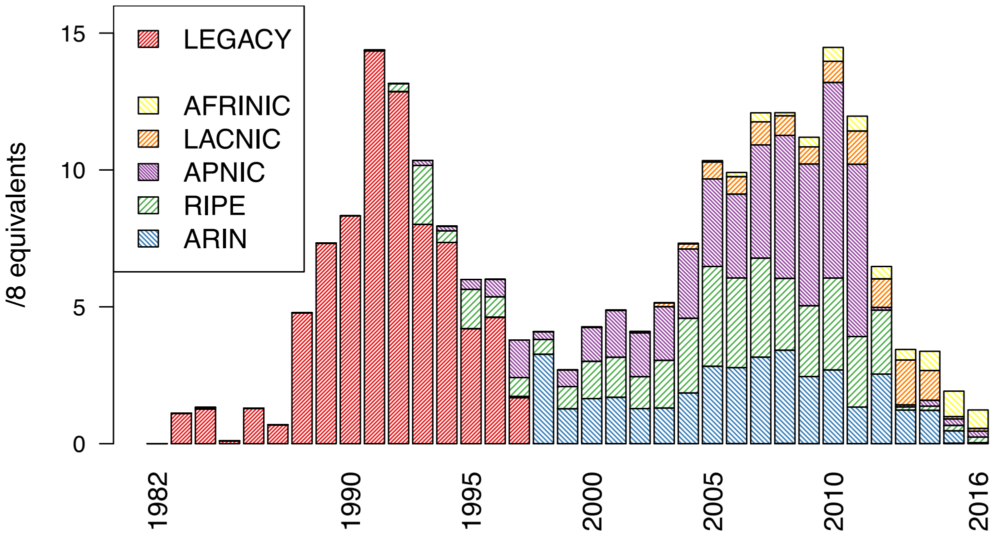

Figure 2.2: Yearly allocations of IPv4 address blocks.

2.1.3 Third Phase: Depletion and Exhaustion

The five RIR communities agreed to a policy regarding address block allocation upon the onset of exhaustion, which ICANN—the international body responsible for the IANA function—ratified in 2009 [147]. The policy dictated that when the IANA's IPv4 free pool reached five remaining /8 blocks, the IANA would distribute these blocks simultaneously and equally to the five RIRs. In February 2011, the IANA allocated its last five free /8 address in accordance with the policy, one to each RIR [208]. After that point, from a global perspective the pool of available IPv4 addresses was fully depleted.

Once the RIRs started to allocate from this last block from the IANA, the "last /8" policies introduced by each RIR went into effect (e.g., APNIC's per [38]), imposing more restrictive allocation policies to further conserve this final address block and to allow new market entrants to still receive a last allocation, e.g., to implement IPv4-to-IPv6 transition mechanisms. Thus, LIRs could receive a single (small) allocation from this block. This transition occurred in April 2011 for APNIC, in September 2012 for RIPE, and in June 2014 for LACNIC, upon the exhaustion of their respective free pools. ARIN exhausted its pool in September 2015, while AFRINIC's pool should last until 2018 [124].[7] LIRs in need of address space now need to find other means of obtaining address space.

2.2 Evolution of Address Block Allocation

Per the above, almost all of the free IP address blocks have been distributed. We can group today's address blocks into three categories: (i) blocks given out prior to the RIRs' existence, termed *legacy* address space;[8] (ii) blocks given out during the era of the RIRs, termed *allocated* address blocks; and (iii) *reserved* address blocks, such as those set aside for multicast and private addressing.

a European Internet Exchange Point) [239]. Numbering resources under the stewardship of LACNIC must be distributed among organizations legally constituted within its service region, and mainly serving networks and services operating in this region. The AFRINIC community, on the other hand, has discussed explicitly limiting out-of-region use to prevent possible exploitation of their IP address resources from operators in other regions [27].

[7] We set the exhaustion date to when the RIRs started to allocate from their last /8, consistent with [124].

[8] LACNIC (and possibly AFRINIC) uses the date of ARIN's inception as their "legacy" threshold, not their own formation, as they would otherwise be unable to apply their policies to addresses that predate their formation.

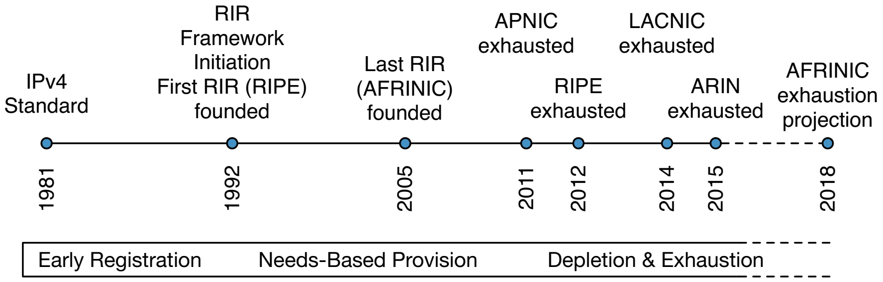

Figure 2.3: Evolution of address management.

Figure 2.3 shows a timeline of the most significant events in the evolution of address block allocation. One cannot pinpoint the transition between the above-mentioned phases precisely: the RIRs were founded years apart, hence ISPs in some regions received legacy address space for a longer period than in other regions. ARIN, for example, began in 1997, whereas RIPE was founded in 1992. Thus, address space holders in the European region received allocated address blocks earlier while holders in North America were still receiving legacy blocks. The transition between phase 2 and 3 is ongoing as of 2017, as one RIR (AFRINIC) still has unexhausted free pools.

In the remainder of this section we present an empirical lay-of-the-land of the state of these allocations.

2.2.1 History of Address Block Allocations

The IPv4 address space consists of 2^{32} possible addresses, an equivalent of 256 /8 address blocks. Of these 256 /8 blocks, 35.3 are reserved by the IETF, e.g., for multicast, private use, and future use. This leaves 220.7 /8s worth routable address space.

In the following, we present a historical view on IPv4 address consumption from an RIR allocation point-of-view. We rely on allocation files provided by the RIRs [207]. Figure 2.2 shows the address blocks given out by the registries over the years as well as those given out prior to the existence of the modern RIR framework (shown as LEGACY).

Two peaks in address consumption are quite visible: The first occurs in the "Early Registration Phase" in the late 80's and early 90's. As discussed in the previous section, address space conservation was not yet a primary concern, and classful allocations resulted in heavy internal fragmentation of address space. The allocation rate drastically decreased in subsequent years, as address space conservation was implemented by the RIRs. Address consumption rates in the late '90s and early 2000s suggested IPv4 address exhaustion would not happen before 2020. The second peak, starting in the mid-2000s was dominated by allocations in the APNIC region, and compromised more than 50% of all allocations in 2010 and 2011. After the exhaustion of the IANA free pool in 2011, a rapid decline in further allocations in 2012 is quite visible. Currently, fewer than 6 /8 equivalents are available for distribution by the RIRs.

The responsibility for the administration of legacy address blocks was transferred to ARIN upon its inception in 1997 [200]. ARIN subsequently re-distributed some of these legacy blocks to the various other RIRs for respective holders located outside the ARIN region. This happened in the course of the ERX (Early Registrations Transfer) project [238]. Yet, most legacy address space is still administered by ARIN, a symptom of North America's dominance of the early Internet.

	handed out /8s	**of which legacy /8s**	**available /8s**
ARIN	100.2	∼ 64.9	0.35
RIPE	48.1	∼ 11.9	0.78
APNIC	51.8	∼ 4.4	0.43
LACNIC	11.1	∼ 0.6	0.02
AFRINIC	5.9	∼ 0.02	1.27
total	**217.0**	**∼ 81.8**	**2.9**
% of routable	**98.3%**	**∼ 37.1%**	**1.3%**

Table 2.1: Address space statistics (January 2017).

Table 2.1 gives an overview of the distribution of the address space among the RIRs (in January 2017). The first column is the number of /8 equivalents, as listed in the allocation files of the RIRs. The second column is an estimate of how much address space is legacy (given out in Phase 1) for each RIR.[9] The last column shows the number of /8s per RIR that are available for allocation. We observe that close to 97% of the IPv4 address space has already been allocated, with less than 3% available for further allocation. Some address blocks are in a reserved state (e.g., for temporal assignments for Internet experiments or conferences), and thus neither available nor handed-out. The heavy allocation rates in the last years prior to exhaustion mainly reflect heavy consumption in the APNIC region. This could reflect a degree of hoarding, but might simply reflect booming Internet deployment in Asia.

2.2.2 History of Routing

In the last section we outlined how the management of IP addresses evolved over time. From a pure allocation perspective, the address space is now close to fully exhausted. One important question is the degree to which allocation reflects actual use. We can consider this in two parts: (1) the degree to which elements of allocated blocks are routed, and thus potentially in use; (2) the degree to which addresses within routed blocks are in fact used. In this chapter, we assess the first of these, as we can much more readily obtain insight into it (via the global routing table as publicly available from the RouteViews project) than we can for the second consideration, which we will study in Chapter 3.

Figure 2.4 shows the number of routed address blocks (expressed as /8 equivalents) over the last 16 years, along with the cumulative total of allocations made by the RIRs. We see that by 1997 more than 25% of the routable address space was advertised, which gradually increased to over 70% in January 2014. While there is an increasing trend in the '00s, in the last two years the rate has been fairly stagnant, perhaps reflecting address exhaustion. It should be noted that the growth of the Internet in its early prime, starting in 1997, used some 50% of the available address space, while the 25% routed prior to that time is likely due to classful allocations and rather lax allocation policies.

Figure 2.5 shows the evolution of the routed address space from 1997 until 2014, by plotting for each /8 the fraction of routed addresses, ranging from white (no address blocks advertised) to black (all address blocks advertised), with the various ranges annotated according to their address types.[10]

[9] For ARIN, we consider all address blocks handed out prior to December 1997 as legacy. For the other RIRs, we consider all address blocks transferred as part of the ERX project as legacy, in addition to blocks *25/8*, *51/8*, *53/8* and *57/8* for RIPE and *43/8* for APNIC. Some of these blocks may have been voluntarily returned or otherwise changed their status. Thus, the number of legacy blocks only serves to give a sense of the landscape.

[10] A few /8 *legacy* block ranges of former class A networks were not given out, and are thus allocated. In addition, some smaller address blocks in the former class B range were allocated by the RIRs, hence the notation "mainly" in the figure.

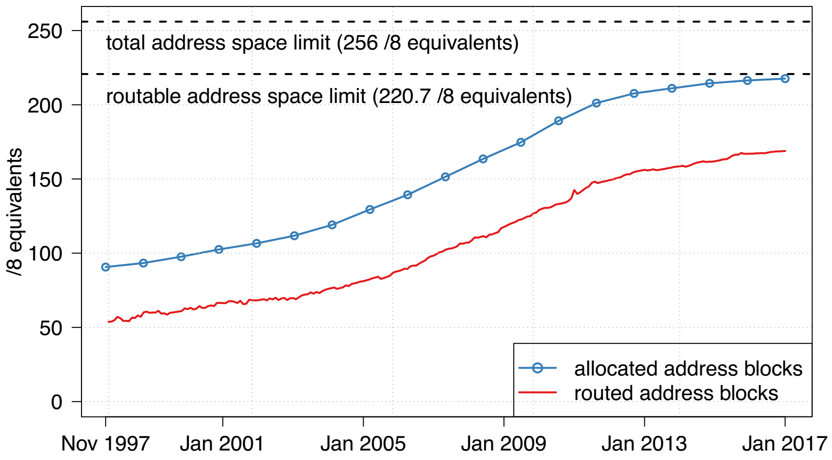

Figure 2.4: Allocated and routed address blocks.

The most striking observation from this plot is that the use of address blocks is very unevenly distributed. Address ranges assigned prior to the existence of the RIRs, the legacy ranges, exhibit much fewer routed address blocks, whereas the RIR-allocated ranges show a gradually increasing and consistent routing pattern. We see that the measures taken in Phase 2, namely the delegation of finer-grained address blocks (CIDR), together with the address conservation principles of the RIRs, indeed had noticeable effect. Hence, efficient address management greatly improved the utilization of address space, but did not enhance utilization in legacy ranges outside of their scope of operation. Today, address blocks in the legacy range have the greatest supply of free and usable address space. In fact, as of February 2015 more than 90% of the allocated address space is routed but only some 50% of legacy address space.

The caveat when using routing tables to reason about the utilization of address blocks is that, while it gives an indication of address space use (clearly visible here), a routed address block does not necessarily mean that it is in active use. Recent estimates range from 47% to 60% [98, 100, 283] of routed /24 address blocks that are actually used, meaning that they are actively engaged in communication. Actual use of address blocks can be measured actively (e.g., probing every IP address with a *ping*) or by relying on passive measurements (e.g., identifying those parts of the address space that actively engage in communication—emitting traffic). We return back to this issue and study IPv4 address space utilization in detail in Chapter 3. While an address block being routed does not imply its actual use, unrouted address blocks, on the other hand, might be in private use for interconnecting networks not publicly reachable.

Hence, while the IP address space is close to fully exhausted, from an allocation perspective, scarcity seems to be less of an issue from a purely technical perspective (e.g., routing). While it requires further work to quantify "efficient use", we can clearly see significant differences between legacy address space and allocated address space.

2.3 IP Addresses as a Resource

IP addresses are virtual resources. In this section, we elaborate associated issues.

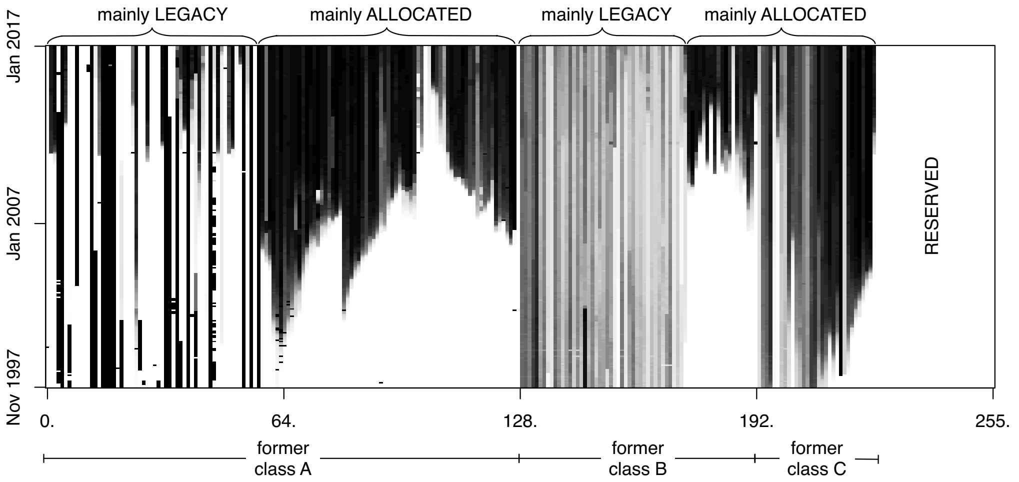

Figure 2.5: Evolution of the distribution of routed address blocks over the total address space.

2.3.1 Addresses: All The Same, Only Different

At first, one might consider IP addresses as a fully homogeneous (fungible) resource, but in fact not all addresses have equivalent properties. First, the size of a given address block governs its routability. Larger address blocks are less likely to be filtered by other operators, and can be de-aggregated into smaller entities, allowing networks to better engineer their route announcements. In addition, an address block comes with *history*: for example, a block previously used by spammers will more likely be found on blacklists, limiting one dimension of its usability. Finally, the properties of address blocks differ depending on their allocation standing and any associated policy restrictions, as noted in the next section.

The case of allocated address blocks

Allocated address blocks given out by the RIRs (Phase 2) are contractually constrained—in a more or less explicit way—as not constituting the *property* of the respective holder. ARIN, RIPE and AFRINIC have explicit "no property" statements in the documents a receiver of address space must agree to [25, 57, 244], while LACNIC and APNIC have more implicit statements in their contracts, not mentioning ownership or property by name. LACNIC states that it can withdraw address blocks from holders [176] and APNIC states that it [only] hands out resources on a "license basis" [42]. The RIRs apply different policies for address space they give out, both with regard to the requirement to document how address space is used as well as with regard to transferability of address blocks. Hence, for RIR-allocated address blocks, the holder will generally have to agree to policies and eventual policy changes as imposed by the respective RIR. Thus, the region associated with an address block directly affects the policies that govern it and thus also its value.

With respect to the possibility of RIRs unilaterally *reclaiming* unused address space from LIRs, the policy documents differ. ARIN clearly rules out unilateral reclamation in its current RSA [57]. APNIC does not mention this possibility by name in its documents, but states that *"If an allocation or assignment becomes invalid then the address space must be returned to the appropriate IR"* [42]. AFRINIC states the possibility of *"revocation or withholding of the service supplied"* [25], and RIPE that it might deregister resources if members fail to comply with their policies [244]. We are not aware of any cases of a unilateral reclamation of allocated address space to date, aside from those where an address holder went defunct without successor.

The case of legacy address blocks

Legacy address blocks, on the other hand, are not in general governed by contractual requirements imposed by any RIR. A noteworthy point with regard to IP addresses as resources is the ongoing discussion whether IP addresses can be considered property or not [253]. Per Figure 2.5, much of today's unrouted address space is legacy, and thus not considered to be subject to current RIR policy.

The RIRs do maintain the registry databases and the anchors for reverse DNS mappings for legacy blocks. However, the attitude of the RIRs towards holders of legacy resources varies. In the course of the last decades, the RIRs—mainly ARIN [55]—started several initiatives to contact holders of legacy address space with the goal of establishing some contractual agreements between the holder and the RIR. As the documentation of legacy allocations is often poor (e.g., outdated information), many holders of legacy resources might not even be approachable. ARIN offers LRSAs (Legacy Registration Services Agreement) [56] to holders of legacy address space in their region. LRSAs establish a more formal relationship between the address holder and ARIN, contain an explicit "no property" clause, and also contractually obligate the legacy holder to ARIN's policies, including the policy for transfer to other entities (or when the holder requests additional address space from ARIN). In late 2007 ARIN sent out more than 18K letters to legacy holders [55]. Their data shows that as of 3 years later, fewer than 1,000 LRSAs were in turn requested by the holders, and LRSAs cover less than 15% of the legacy address space in the ARIN region [53]. One address broker publicly suggests to legacy holders to not sign such LRSAs [157]. Another ARIN document states "*All of the IP address space that ARIN administers, including legacy space, is subject to ARIN policy*" [60]. RIPE, on the other hand, adopted a proposal in February 2014 to offer registration services to holders of legacy address space and not impose particular regulations on transfers of registered legacy address blocks [246].

Regarding the possibility of *reclaiming* unused legacy address blocks, ARIN states that it will not attempt to unilaterally reclaim legacy address space [54]. APNIC and RIPE ran initiatives to contact holders of *legacy* address blocks to recover address space [43, 247] but left the decision up to the respective holder. In case of the RIPE initiative, 400 holders were contacted of which 16 returned address space to RIPE. However, there are prominent examples of voluntarily returned legacy address blocks, such as Stanford University voluntarily returning its /8 legacy address block in 2000 [78] as well as some other organizations [149].

A meeting convened by ICANN in 2012 informally addressed issues related to legacy address resources [148]. The discussion involved representatives from the RIRs, network operators holding legacy and non-legacy address resources, and address brokers. On one hand, it was argued that legacy resources by their nature do not differ from other IP address blocks, and should thus be subject to the same policies. On the other hand, holders of legacy address space argued that *grandfathering* applies—meaning that as legacy address space was given out prior to RIR policies, they are not subject to any policies subsequently created by RIRs.

Hence, the open question with regard to legacy holders is whether they are bound to the terms of the registry that currently provides registration services to them—in a more general way, whether they hold ownership rights for their addresses or not.

2.3.2 Resource Certification and Enforcement

In the case of IP addresses, no global system exists to either authoritatively verify the ownership of a given address block nor to prevent the usurping of address blocks by illegitimate users. Inter-domain routing as instantiated by BGP does not itself provide any mechanisms to ensure routing only by a block's legitimate holder. While the community readily recognizes BGP's lack of security features, including its inability to authenticate routes, a large body of research and accompanying deployment efforts has done little to change this situation in productive environments (see [77] and references therein).

The RIRs publicize the mapping of address spaces to their respective holders via registry databases (WHOIS), which can be queried publicly, and by delegating the respective reverse-DNS zones (`.in-addr.arpa`) to authoritative nameservers specified by the address holders. This latter enables the holders to specify PTR records for IP addresses in the respective namespace (not a fundamental requirement or hallmark of ownership, but certainly operationally useful). Neither of these mechanisms provide sufficient information to directly validate (or invalidate) route advertisements, such as by authoritatively indicating the origin AS.[11] Thus, the administrative management of address space is largely decoupled from its actual use. The degree to which a prefix is usable by some entity—and which entities have the capability to use it—simply depends on how far a route advertisement for the given prefix propagates, which directly translates into how many hosts on the Internet can interact with hosts in the given address block.

The propagation or non-propagation of prefix advertisements depends on the route filtering performed by the border routers of ISPs. To configure these filter settings, the community has established routing registries (IRR), where network operators can register route objects to express prefix ownership in the form of prefix-AS mappings [77]. The various IRR databases are managed by several independent organizations, including ISPs, RIRs and others [158]. However, not all address space is registered in some registry (only around 50% according to [269]) and information in these registries is known to be significantly inaccurate [168]. Many IRRs allow their participants to introduce essentially any route object without further validation [270]. Complications with the IRR can again result in ISPs not filtering advertisements from their peers using IRR information at all [104]. There are well-known cases of erroneous IP address block advertisements, be it hijacking of address blocks by spammers [228] or advertisements caused by misconfigurations. As an example, a Pakistani ISP erroneously advertised a prefix belonging to YouTube in 2008, resulting in an extensive global outage for that service [201].

The Internet Engineering Task Force (IETF) has developed a solution to this problem based on the RPKI (Resource Public Key Infrastructure) [179]. The basic function of the RPKI is to provide cryptographically verifiable attestations to address space and AS number allocations using a X.509 based hierarchy. RPKI uses the IANA and the RIRs as trust anchors, which give out certificates for resources they manage. Currently, RPKI services are offered by the RIRs as a free opt-in service only to their members.[12] Based on the RPKI database, routers can verify that an AS advertising a specific prefix is in fact authorized to do so, which is referred to as RPKI-based origin validation [74]. This only prevents accidental advertisements and is not intended to prevent malicious attacks, as the full AS path is not validated but only the origin of the path [75].[13] RPKI is supported by current routers from Cisco, Juniper,

[11] The ARIN WHOIS database recently started to provide a field for the origin AS, but the field is often unset and prominent cases of inconsistencies exist [212].

[12] ARIN requires legacy resource holders to sign an LRSA in order to be eligible to register their resources in the RPKI. Moreover, ARIN requires any operators wanting to *use* the ARIN RPKI data to sign a Terms of Service Agreement that includes an indemnification clause [58].

[13] To overcome this, AS-Path validation is necessary [180].

and Alcatel-Lucent, yet as of January 2017 only about 10% of the routable address space is covered by RPKI and by far the largest share of that address space is in the RIPE region [245].

While the problem of securing the advertisement of a prefix by only the respective holder is well-known and many approaches have been proposed over the years, little has changed in productive environments. Faced with the increasing scarcity of IP addresses (and the corresponding increasing value of addresses as resources), a functional scheme for certifying resources will be a key requirement in the near future in order to prevent illegitimate address space use.

2.3.3 Address Markets

Given that free address pools are now mostly exhausted and that demand for IPv4 address space will likely continue to grow (at least until significantly broader IPv6 deployment), address space transfers arise as a natural step necessary to further distribute address space to those networks that need it. In light of the issues discussed above—namely the fragmentation of addresses into legacy and non-legacy address blocks subject to varying RIR-policies, and connected ownership discussions as well as the lack of widely adopted resource certification mechanisms—the landscape of such address transfers is at best murky. Network operators have already started buying and selling address blocks under varying conditions, as we outline in the following. This resulted in the emergence of several *address brokers* (e.g., [23, 155, 156]); companies that assist network operators wanting to buy or sell address space. Eventually, the RIRs learned to encourage the use of address brokers to mediate transactions within the strict confines of RIR policies.

RIR Transfer Policies

Today, four[14] out of the five RIRs allow address space transfers among their members [44, 49, 175, 249]. In addition, ARIN, RIPE and APNIC offer *transfer listing services* [40, 59, 240], where network operators can list address blocks they want to sell and express the need for certain amounts of address space they want to buy. These services aim to help interested parties to come together, but use of them is not mandatory. RIPE publicizes aggregated statistics for address space requests and offerings, listing fewer than one million available addresses, and more than 50 million requested addresses, as of January 2017 [240]. These listings do not include any prices, as the negotiations remain entirely at the discretion of the respective parties.

We observe a striking difference between how the RIRs perceive their roles when it comes to conducting transfers under their policies. Except for RIPE, the RIRs still require the receiving party of a transfer to *justify* their need for more address space according to their already established policies. For example, APNIC requires transfer recipients to document use rates for past allocations as well as detailed plans for the use of transferred resources [44], while ARIN states that recipients must demonstrate the need for up to a 24-month supply following their established policies [49].

RIPE—as of February 2014—removed all "justification of need" clauses from their policies. Address space can be transferred from any member to any other member without the need to make statements of how the transferred addresses will be used by the recipient. The proposal [248] argued that address conservation will be in the interest of the members themselves (to not waste address space). With regard to concerns about possible address hoarding by wealthy LIRs it states that "*markets* [for other

[14] AFRINIC states the possibility of transfers between LIRs [26], but prohibits any such transfers in their LSA unless they arise due to Mergers & Acquisitions [25].

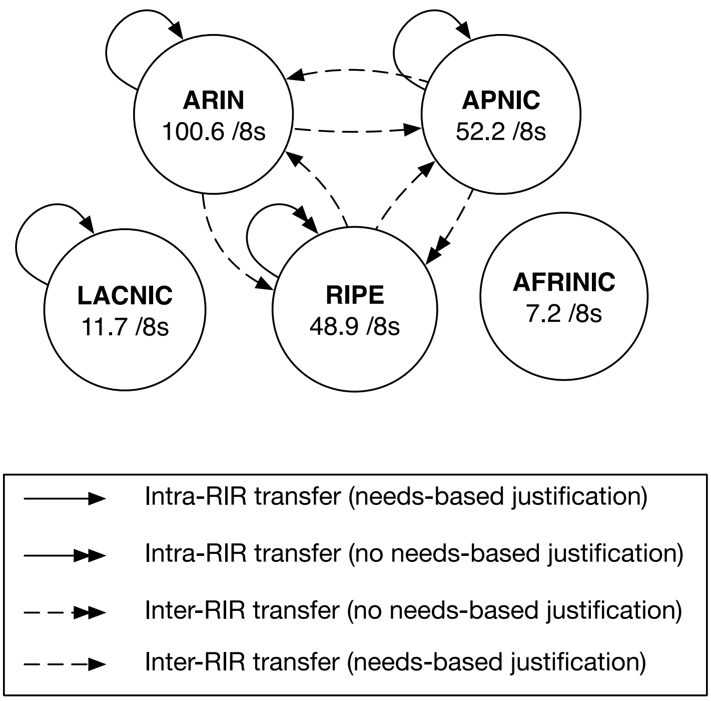
ARIN
100.6 /8s
APNIC
52.2 /8s
LACNIC
11.7 /8s
RIPE
48.9 /8s
AFRINIC
7.2 /8s
Intra-RIR transfer (needs-based justification)
Intra-RIR transfer (no needs-based justification)
Inter-RIR transfer (no needs-based justification)
Inter-RIR transfer (needs-based justification)

Figure 2.6: Current address space transfer policies of the RIRs and the administered address space.

commodity goods] *function well and in a competitive manner, and there is no reason why the trade of IPv4 addresses will be any different*".

As of January 2017, *Inter-RIR Transfers*—i.e., transfers between address holders in different regions—are possible between the ARIN, APNIC, and RIPE regions. ARIN explicitly requires justification of need on the receiving side of a transfer—even if the recipient is located in a different region [52]. RIPE subsequently introduced a policy that requires needs-based justification as a prerequisite to allow incoming transfers from the ARIN region [243]. Figure 2.6 summarizes the transfer policies in place by the RIRs along with the address space they administer.

Another scenario in which address block transfers happen—and happened long before modern transfer policies were established—is due to *Mergers & Acquisitions*. In this case, address blocks are part of the assets of a company. Since the related contracts are often confidential, these transfers are not publicly listed by the RIRs—with the exception of APNIC, which requires full disclosure of the involved parties and publicly lists the corresponding address blocks [39]. The RIR's documents make no explicit statements about the justification of need for the transferred allocations. ARIN only states that the transferred resources will be subject to ARIN policies [49], while APNIC states that it will "review the status" of the allocations, requiring full disclosure of all allocations held by the "entities in question". If that is not provided, APNIC will "require that they be returned" [44].

RIR Transfer Statistics

Figure 2.7 shows monthly aggregates of address blocks that were transferred under the previously introduced RIR transfer policies. In the years between 2010 and 2014, only comparably few address blocks were transferred, as shown in Figure 2.7(a). During this period, not all RIRs exhausted their pools yet and address space could still be obtained from the RIRs following their regular processes. Starting in 2014, however, we observe a steep increase in the number of monthly address transfers, and through

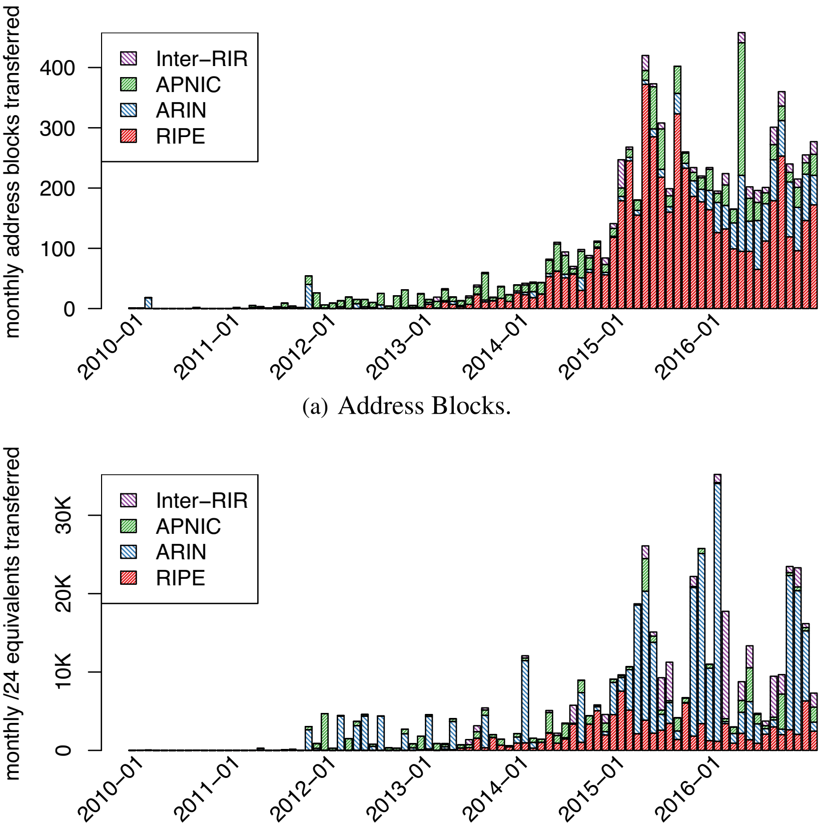

(a) Address Blocks.

(b) /24 Equivalents.

Figure 2.7: Monthly address block transfers per the respective RIR policy.

January 2017 already more than 8,500 address blocks were transferred. In terms of address blocks (i.e., transactions), most transfers happen in the RIPE region. Indeed, we see that the number of transfers steeply increased since RIPE allowed intra-RIR transfers without the need for justification in 2014. Figure 2.7(b) shows the number of transferred addresses (as /24 equivalents) on a monthly basis. Contrary to the number of transactions (i.e., address blocks), we see that the majority of addresses are transferred in the ARIN region. Transfers in the ARIN region typically correspond to larger address blocks, when compared to other RIRs, such as RIPE.

Growing numbers of address transfers, both in terms of address blocks and /24 equivalents show that there is ongoing demand for IPv4 address space. Network operators are indeed willing to turn to the emerging address market to acquire additional IPv4 address resources. Still, the number of transferred IPv4 addresses is lower compared to pre-exhaustion allocation rates. This observation suggests that network operators likely also apply other techniques, besides acquiring addresses on markets, to mitigate their IPv4 scarcity issues. Additionally, some address space transfers might not undergo the official RIR processes, and are thus missing from these statistics.

Transfers Outside the RIRs

Given that neither the legal nor the technical aspects of address space transfers are under the full control of the current RIR framework, parties can also conduct transfers separately from the RIRs. To the extent that these occur, a definitive determination of the party possessing a given allocation becomes more difficult because the RIRs no longer possess accurate records.

Even though ARIN states that legacy holders are subject to ARIN policies, recent transfers, such as the well-known sale of more than 660K IP addresses from the Nortel bankruptcy to Microsoft, have raised concerns whether they complied with proper ARIN transfer policy. Mueller et al. [197] state that while ARIN was formally involved in the transfer, likely no needs-based evaluation was performed on the receiving side, and that ARIN's intervention boiled down to a *"face-saving exercise"*. As the relationship of legacy holders towards the RIRs is not entirely clear, one IP address trader has suggested that legacy address holders in the ARIN region could de-register their address space there and re-register it with a different RIR, such as RIPE [157]. Doing so would effectively allow inter-region transfers from ARIN to RIPE without undergoing any transfer process. But currently there is no process to de-register from an RIR.

Aside from transfers that were formally noticed by the RIRs (such as the above example), address transfers can also happen without the involvement of any registry at all. While address space can be of various types (*legacy*, allocated to a holder by an RIR, assigned by a holder to an end-user, PI-assigned directly from the RIR to an end-user), bound to various contractual limitations, not much prevents any party from unofficially transferring an address block to another entity. This is known as a *"black market"* transfer. This possibility stems from the decoupled nature of address block management and actual address block use. If RIRs do not acknowledge such transfers, registry information becomes in turn inaccurate and incomplete, making the attribution of address blocks to their respective holders difficult.

In the simplest terms, we can view a transfer as simply an address block—or parts of it—formerly in use by some entity A now being used by some entity B, possibly outside the purview of any RIR regulation. If the routing of the concerned address block is possible after the transfer (it is not filtered by networks), and (to a lesser degree) the corresponding reverse-DNS zones become under the control of the receiver (e.g., by subdelegation of reverse-DNS zones by the previous owner), the transfer would be successful.

It is unclear whether it is even feasible to detect the occurrence of such transfers. Livadariu et al. attempted to detect such transfers by looking for changes in routing origins over time [182, 183]. One difficulty here is that transferred address blocks are not necessarily routed before they are transferred. Indeed, prior routing might be unlikely, as unrouted address space is likely also unused and thus more likely to be transferred. Also, whether such a transfer would be reflected in the reverse-DNS is unclear, as NS records might simply not be changed and PTR records might be unchanged or switched off. Shifts in traffic, latency changes or geographical changes might be due to transfers but also due to restructurings within a company.

Thus, defining the boundaries of what exactly an address transfer is and what it is not is not straightforward. It is likely that the official RIR transfer policies only cover a fraction of the total address transfers occurring in various instantiations of the above scenarios. While transfers undergoing the RIRs policies are publicly listed [39, 51, 241] and quantifiable, the number of address transfers outside this framework is unknown and requires further research.

2.4 Chapter Summary

In this chapter, we studied the entanglement of policies and technology for IPv4 addresses. Our historic analysis of the management structure reveals that address management and address policies directly affect the degree of address consumption and eventual utilization. Address conservation was not a primary goal in the early days of the Internet, and close to 40% of the entire IPv4 address space was handed out before the establishment of the global registry framework. The RIR framework introduced stricter rules for subsequent address allocations to better conserve address resources. These policies have a measurable effect on address utilization in the global routing table, where we find that *legacy* address blocks are more likely unrouted, in contrast to addresses subsequently assigned by an RIR. Early decisions and address allocation policies are still clearly visible many years later. Thus, Internet governance decisions have a direct and far-reaching impact.

Today's situation in which IPv4 addresses are valuable commodities raises several new and unsolved challenges, most prominently the need for resource certification as well as the need for policies to regulate the emerging market for IPv4 addresses. The RIRs and some members of the community actively promote resource certification based on the RPKI, and yet many network operators do not use or enforce RPKI-based validation. If IPv4 address scarcity intensifies in the future, the inability to validate the holder of an address block has the potential to add increasing instability to the Internet's routing system. Besides technical challenges, policy issues regarding the management and regulation of emerging address markets compound the situation. The current registry framework is split across five institutions (RIRs), and each RIR applies individual policies for address resources transferred within or across regions. Thus, individual addresses blocks now come with different policies and restrictions on their transferability, depending on the region they are registered in. Additionally, the RIRs inherited management obligations for address space given out before the framework was even established. Holders of such *legacy* address blocks are not necessarily subject to the same policies that apply to subsequently allocated address space.

Since IP addresses are an inherently global resource, it is questionable whether the distributed registry framework can cope with the looming issues and provide sufficient resource liquidity. Future scenarios for the management could include a more competitive environment among RIRs, or tighter cooperation or even re-centralization of the five regional registries. While it is unclear whether some IP address block holders have ownership rights for their IP addresses, secondary markets already exist to facilitate their exchange. However, the uncertainties associated with address space transfers—both the legal status of legacy address blocks and the varying policies among RIRs when it comes to such transfers—will also complicate how pricing develops. Different prices for address resources in the different regions have the potential to create more pronounced scarcity in some regions and to create market entry barriers for ISPs in need of address space. As a result, increasing numbers of transfers outside the RIR processes are possible. Such transfers, in turn, can result in less accurate registration data of the RIRs, further complicating the management situation of the IPv4 address space.

The complex policy space makes it increasingly difficult for network operators to make decisions about when to apply which mitigation strategy to overcome IPv4 address scarcity. Despite the technical and legal challenges in this tussle, we see an increasing number of listed IPv4 address transfers. Thus, IPv4 address scarcity is real and networks are willing to pay for additional IPv4 addresses. While purchasing IPv4 addresses is evidently a viable option for many network operators, the number of monthly transferred addresses is lower compared to pre-exhaustion allocation rates. This observation suggests that networks likely also use other means to mitigate their individual scarcity issues, which we will follow up on in the remainder if this dissertation.

3 Address Activity in the Wake of Exhaustion

In the last chapter we laid the groundwork for our understanding the current status of the IPv4 address space, both from a management as well as from a routing perspective. In this chapter, we measure and analyze *activity* of the IPv4 address space. Measuring IP address space activity has caught the attention of the research community for some time, often to assess the current state and expansion of the Internet [79, 96, 99, 226, 283]. We primarily focus on assessing the current degree of utilization of the IPv4 address space and the potential for utilization increase, and we study how the exhaustion of the IPv4 address space manifests itself in global address activity.

ISPs now need to find open-ended ways to accommodate the needs for IPv4 connectivity of their customers, e.g., by increasing the utilization efficiency of their respective address blocks. Policy makers need to establish regulatory guidelines for the emerging marketplace for IPv4 address space. Identifying regions of the address space that show only comparably little utilization—and pinpointing the root causes for that—can help networks to consolidate their IPv4 address space and possibly free up IPv4 address reserves. Identifying inactive portions of the address space, in turn, helps to determine possible sellers of address space. More fundamentally, having a detailed picture of IPv4 address space utilization helps to quantify the current scarcity problem and puts our findings from the routing table from the last chapter into proper perspective. A detailed understanding of address activity also serves as a foundation for security-critical systems that rely on the notion of IP addresses, e.g., for client reputation [37, 136], as well as for systems that rely on active IP addresses to perform measurements, e.g., geolocation systems [131, 166, 278] and network troubleshooting systems [115, 167, 226].

Recent studies that present Internet-wide statistics on IPv4 address space utilization either measure or estimate the total number of active IPv4 addresses [283] and address blocks [99] in the Internet by relying on a number of diverse data sources. However, the total number of active IPv4 addresses and blocks only partially captures address space utilization. Moreover, with the exhaustion in allocation of IPv4 addresses, the situation will likely be changing and will reflect the independent decisions of network operators, where the administration of the IP address space is under the control of the respective administrative domain (Autonomous System, of which about 51K can be found in the global routing table, as of 2015). Varying resource demands and operational practices, as well as available supply of free and unused IP address space, blurs the notion of an "active" IP address. Today, we face a situation in which individual addresses and address ranges vary in their periods of activity and in that activity's nature and volume. For example, dynamic addressing, network reconfigurations, and users' schedules dramatically affect periods of activity. Traffic characteristics and volumes of active address blocks range widely, from lightly used addresses and sparsely-populated blocks to proxy gateways connecting thousands of devices to the Internet.

As a result, questions about the number of IP addresses active at a point of time, let alone their usage characteristics, are still difficult, if not impossible, to answer. This problem becomes even more difficult when characterizing address space usage over time, since we face the problem of choosing the right time granularity to observe such activity. Detailed measurements of address space activity helps to elucidate the current state of the IPv4 address space exhaustion, and has practical use and implications for ISPs and regulators. ISPs now need to make business-critical decisions such as how to adapt their address assignment practices in order to maximize the utilization of their available address resources. Regulators currently have to rely on estimations and predictions when introducing new policies that will ultimately affect what will be deployed in practice.

In this chapter, we provide an unprecedented, detailed, and longitudinal view of IPv4 address space activity. We first discuss methods to study IPv4 address activity, and provide a comparison of the visibility into the IPv4 space from various vantage points. We then study address activity as seen through the lens of a large commercial CDN that serves almost 3 trillion requests per day. This unique vantage point enables us to measure Internet-wide IPv4 address activity at the granularity of individual IP addresses, over a period that spans a full year. Our study provides a number of insights on the state and growth of the Internet in the face of increasing IPv4 scarcity.

The findings of this chapter can be summarized as follows:

(i) We assess the visibility of different methods and datasets into global IPv4 address activity, considering different aggregation levels. While we find disparities across datasets, we find that our CDN logs allow for unprecedented visibility into the global IPv4 space on a per-IP granularity. Our CDN logs draw a detailed picture of world-wide activity from 1.2B unique IPv4 addresses contacting the CDN.

(ii) We find that, after years of constant linear growth, the total number of active IPv4 addresses has stagnated since 2014. We also find that state-of-the-art active measurement campaigns miss up to 40% of the hosts that contact the CDN.

(iii) We show that despite the stagnation in the number of active IPv4 addresses in 2014, the *set* of active addresses is far from constant. In fact, over the course of a year, more than 25% of the active IP address pool changes. Most client networks contribute, with varying degrees, to this "address churn" and this churn is barely visible in the global routing table.

(iv) We identify a variety of address block activity patterns, and attribute them to network restructuring, user behaviors, and various address assignment practices. Based on our observations, we introduce metrics that allow us to quantify prevalent addressing practices at scale and comment on additional utilization potential within these already active address blocks.

(v) We augment our address activity metrics with corresponding traffic volumes and relative host counts, which we derive from *HTTP User-Agent* samples, observing a trend of increasing traffic for addresses already heavily trafficked. Combining our three key metrics of address activity, we then derive Internet-wide demographics of the active IPv4 address space and discuss the broader implications that our study has towards enhancement of current operational and measurement practices.

3.1 Measuring Address Activity

In this section, we first introduce the various methods that have been used in the past to measure and capture IP address space activity. To this end, we discuss active and passive approaches used in related work. We then present a visibility comparison of datasets gathered from 7 data sources, including data from 4 passive vantage points and 3 active measurement campaigns.

3.1.1 Methods to Measure Address Activity

Active Measurements

A popular way of assessing IP address activity is by actively probing IP addresses (IPs), e.g., with ICMP (Internet Control Message Protocol [222]) echo queries. Heidemann et al. presented a census of IP address activity by systematically probing the allocated IPv4 address space with ICMP echo requests as early as 2008 [138], which was followed by studies that also capture aspects of network management, e.g., diurnal activity patterns [79, 227] and Internet reliability [226]. Recent improvements in active scanning techniques were introduced by Durumeric et al. in ZMap [111], that enable scanning of the entire IPv4 address space within less than one hour or even in less than 5 minutes [24]: a milestone in worldwide active measurement.

Note that a reply from an IP address does not necessarily indicate that a unique host is indeed active or even exists; tarpits [35], firewalls, and other middleboxes might send replies to probe traffic destined to other IP addresses, or even entire IP address ranges. Also, active measurements cannot capture activity at all timescales, as a reply might be dependent on many factors [227, 257]. It is also common that network administrators and home routers block ICMP traffic, thus, active measurements are not always successful in detecting active address blocks [100]. Advanced active measurement techniques that scan specific ports can be used to increase the detection success of active IPs [111].

Passive Measurements

Dainotti et al. [99] used passive measurements of packet captures and network flow summaries recorded at three passive vantage points and found 3.9M active /24 blocks from passive measurements in 2012. In Section 3.1.2 we present a comparison of the used datasets, and additional vantage points, measured in 2013 [100]. Studying address activity only at the /24 level might be misleading, as the utilization within /24 blocks can vary widely. To our best knowledge, only one related piece of work, by Zander et al. [283], estimates the number of active IPv4 addresses (in contrast to address blocks). They combine data from nine different passively captured datasets and two active datasets. Their active data sources consist of ICMP echo scans and TCP SYN scans on port 80. Their passive data sources include Wikipedia page edit histories, lists of potential spam senders, addresses of clients tested by Measurement Lab [8], web clients participating in an IPv6 readiness test, server logs of game clients connecting to Valve's Steam online gaming platform, as well as NetFlow records from border routers at Swinbourne University of Technology and Caltech. Over the course of their measurement period (end of 2011 until June 2014), they measured activity from 740 million unique IPv4 addresses in 5.9M /24 address blocks. They use a statistical capture/recapture model to account for invisible addresses and estimate the total number of active IPv4 addresses to be 1.2 billion (6.3M /24 address blocks) as of 2014.

Other Related Work

A number of studies proposed techniques to identify dynamically assigned IPv4 addresses and uncover their dynamics. Xie et al. [280] introduced a novel method, UDmap, that takes, as input, user-login traces (e-mail logins in their study) and identifies the dynamic IPv4 addresses by associating the unique login information of each user with the set of IPs it utilizes. They concluded that address dynamics exhibit a large variation across networks, ranging from hours to several days. Jin et al. [160] proposed and evaluated a technique to identify static and dynamic address blocks based on distinct traffic activity patterns of static and dynamic addresses, when countering outside scanning traffic. Moura et al. [196] proposed an active ICMP-based method to scan the addresses of an ISP in search of blocks that rely on dynamic host configuration protocol (DHCP) to dynamically assign IPv4 addresses to users and also to estimate DHCP churn rates. Padmanabhan et al. [213] used data gathered from RIPE Atlas probes to analyze the frequency and events associated with address assignment changes. They found that a number of ISPs around the world periodically reassign addresses after a fixed period, often a multiple of 24 hours. They also found that some address changes are correlated with network and power outages occurring at customer premises equipment.

Plonka and Berger count active World-Wide Web (WWW) client addresses by passive measurement and develop temporal and spatial address classification methods [218]. Their work has similarities to ours here in its use of CDN server logs (in fact, the same logs we utilize) and in its spatio-temporal approach, but differs in that they study only IPv6 addresses.

3.1.2 Vantage Point Visibility

Both active and passive measurements of address activity come with inherent biases. Active measurements do not have a direct topological bias (i.e., an IP address can be actively probed independently of its topological location), but may hide significant portions of address activity, since not every active host replies to probes. Moreover, machines performing active measurements might accrue reputation, which might result in blacklisting and thus further limited visibility. Active measurements can only estimate address activity based on replies to active queries and can hence not provide an "in situ" notion of address activity. Passive measurements do allow for measuring "in situ" address activity as observed in the Internet, without the need to actively send probing traffic. However, passive measurements come with an inherent topological bias, since there is no single vantage point that can capture all traffic exchanged on the Internet and hence illuminate the activity of *all* active IP addresses.

To gauge the bias and disparity when relying on different vantage points to measure IPv4 address activity, we compare the visibility in terms of active IPv4 /24 address blocks, as observed from 4 different passive vantage points, and 3 datasets gathered by actively probing the entire IPv4 address space. Each of the vantage points retains traffic data in different formats and thus requires different filtering approaches for use in a census. In this dissertation, we rely on these datasets for the sole purpose of studying disparities in their visibility. The details of the various filtering steps are described in [100]. All measurements were collected in the period between July and September 2013.

Passive Datasets

SWITCH: We collected NetFlow records from all the border routers of SWITCH, a national academic backbone network serving 46 single-homed universities and research institutes in Switzerland [265].

The monitored address range of SWITCH contains 2.2 million IP addresses, which correspond to a continuous block slightly larger than a /11.

R-ISP: We collected per-flow logs from a vantage point monitoring traffic of about 25,000 residential ADSL customers of a major European ISP [123]. The vantage point is instrumented to run Tstat, an open source passive traffic flow analyzer [117] that stores transport-level statistics of bidirectional flows.

UCSD-NT: We collected full packet traces from the /8 network telescope operated at the University of California San Diego [13]. Network telescopes, also called darknets, passively collect unsolicited traffic—resulting from scans, misconfigurations, bugs, and backscatter from denial of service attacks—sent to routed regions of the address space that do not contain any hosts.

IXP: We collected packet samples from one of the largest IXPs in the world, which is located in Europe and interconnected some 500 networks which exchanged more than 400PB monthly at the time this measurement was taken [29].

For all passive vantage points, we developed methods to detect and remove spoofed traffic. The details of these methods are outlined in [100] and we will only give an overview of the methods in the following. In the case of **R-ISP**, we only count IP addresses, for which we see bidirectional flows. In the case of **SWITCH**, we only count bidirectional TCP flows with at least 5 packets and 80 bytes. For **UCSD-NT**, we manually identified and removed large-scale spoofing events and removed TCP packets with no flags set and UDP packets without payload. For the **IXP** dataset, we set thresholds for the minimum number of packets and their corresponding packet size sent to or from specific /24 address blocks. We adjusted and validated our methods by comparing the seen-as-active portion of the address space against unrouted address space, as well as routed address ranges within **UCSD-NT** and **SWITCH**, which do not emit any traffic. After applying our anti spoofing heuristics, we find that every vantage point detects less than 0.05% of unrouted /24 address blocks as active, and less than 0.04% of the dark /24 address blocks as active.

Active Datasets

ISI: We used the ISI Internet Census dataset *it55w-20130723* [1], obtained by probing the routed IPv4 address space with ICMP echo requests[1] and retaining only those probes that received an ICMP echo reply from an address that matched the one probed (as recommended [152]).

HTTP: We extracted IP addresses from logs of Project Sonar's HTTP (TCP port 80) scan of the entire IPv4 address space on October 29, 2013 [137]. For each /24 block, we stored how many IP addresses responded to an HTTP GET query from the scan.

ARK-TTL: We processed ICMP traceroutes performed by CAIDA's Archipelago to each /24 in the routed IPv4 address space between July and September 2013 [146]. Specifically, we extracted the ICMP Time Exceeded replies sent by hops along the traceroute path.

Disparity of IPv4 Address Space Visibility across Vantage Points

Table 3.1 shows the number of /24 address blocks that were detected as active by each passive vantage point and active probing technique, as well as the number of address blocks that were unique to each vantage point. The third column shows the number of /24s observed in the data set that were not also

[1] We did not use reverse DNS PTR scans of the IPv4 space for the same reasons articulated in [138], namely that many active IP addresses lack DNS mappings, and many unused IP addresses still have (obsolete) DNS mappings.

Dataset	# /24s	# Unique /24s within active/passive	# Unique /24s among active + passive
Active			
ISI	4,589,213	1,319,283	398,334
HTTP	3,161,064	189,831	76,189
ARK-TTL	1,627,363	40,284	24,533
All Active	4,837,056		
Passive			
SWITCH	3,599,380	147,220	54,905
UCSD-NT	3,149,944	61,443	24,134
R-ISP	3,797,273	176,721	59,278
IXP	3,090,645	195,328	55,155
All Passive	4,468,096		
Total	5,306,935		

Table 3.1: Visibility of /24 address blocks and intersection across different vantage points and probing techniques.

observed in the (top) other active data sets or (bottom) other passive data sets; the fourth column is the number of /24s observed that were not observed in any other data set. The final total is the number of /24s from which we measure activity. Every vantage point detects a significant unique set of active IPv4 address blocks. Note that the active ICMP ping scan (**ISI**) shows the largest number of unique active /24 address blocks, when compared to the other vantage points, which could hint at the limited coverage of our passive vantage points. On the other hand, our passive vantage points detected some 850K /24 address blocks as active that were invisible in **ISI**.

Our results highlight that there exists no single vantage point or method that can capture all IP address activity on the Internet; the combination of the 7 data sources shows activity for 5.3M /24 address blocks and significant unique contributions. Our measurements underline the difficulty of finding appropriate data sources to capture and analyze address activity at scale. With this observation in mind, we next introduce another dataset, which shows both broad coverage (i.e., exceeding the number of /24 address blocks that were visible in our earlier study and in related work) as well as granularity (i.e., it allows us to study address activity on a per-IP level).

3.2 A CDN as an Observatory

Next, we introduce the dataset that we use to study detailed IPv4 address activity of Web clients, as seen from a major CDN. To assess the visibility of our dataset, we provide a comparison of our passive IP address activity logs with active probing and provide a geographic breakdown of address visibility.

Server logs of one of the world's largest CDNs form the foundation for this study. In the year 2015, the CDN operated more than 200,000 servers in 120 countries and 1,450 networks, serving content to end-users worldwide. The CDN serves close to 3 trillion HTTP requests on a daily basis. The content does not only include websites of its customers, but also a multitude of embedded Web objects (e.g., advertisements), smartphone application content, video streaming, and software updates. In the following, we do not differentiate between the individual content types, but refer to them as Web objects. Each time a client fetches a Web object from a CDN edge server, the server creates a log entry, which is then processed and aggregated through a distributed data collection framework. After processing, we have access to the exact number of requests ("hits") issued by each single IP address. In this work,

Description	**IP addresses** total	avg.	**/24 blocks** total	avg.	**ASes** total	avg.
Daily: 08/17/15 - 12/06/15	975M	655M	5.9M	5.1M	50.7K	47.9K
Weekly: Jan - Dec 2015	1.2B	790M	6.5M	5.3M	53.3K	47.8K

Table 3.2: Datasets: Totals and averages of active IPv4 addresses observed by the CDN per snapshot.

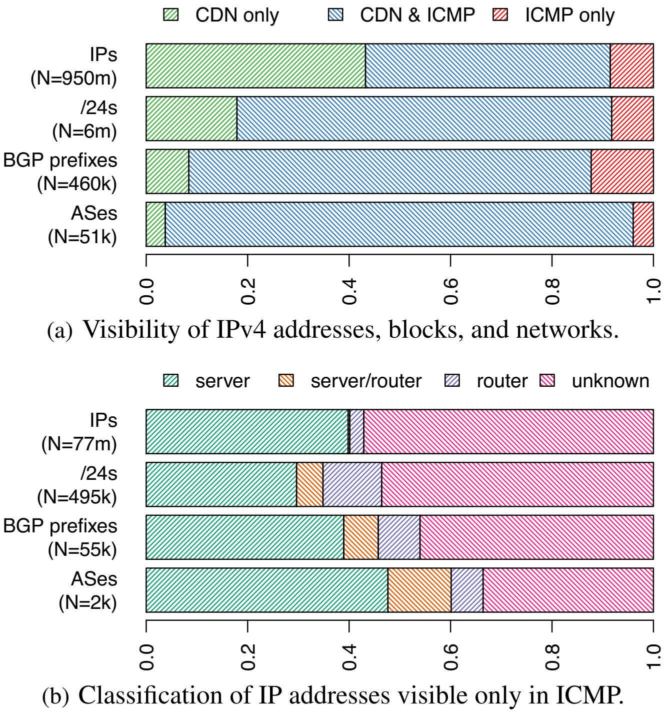

(a) Visibility of IPv4 addresses, blocks, and networks.

(b) Classification of IP addresses visible only in ICMP.

Figure 3.1: Visibility into the IPv4 address space of the CDN compared with active measurements (Oct. 2015).

we rely on two datasets, which are shown in Table 3.2. For the year-long dataset, we have weekly aggregates of all IP addresses and for the daily dataset, we cover a period of 4 months. In the following, we refer to an IP address as *active* if the CDN handled a request from that IP address in the given time interval. Correspondingly, we refer to an IP address as *inactive* if there was not such a request. Here, requests refer to successful WWW transactions, i.e., an IP address will only be associated with a request if the client initiated a successful TCP and HTTP(S) connection and successfully fetched an object. Therefore, address activity is *evident* from our log dataset and a major advantage compared to other passive measurements. The second advantage of our dataset is its *granularity*, both space and time-wise. The logs contain numbers of requests on a per-IP level, illuminating a detailed picture of address activity.

To assess the view from our vantage point, we next compare the set of addresses visible from the CDN to those which replies to ICMP queries. For this, we use the aggregated counts of CDN-observed active IP addresses and compare them to the union of all IP addresses that were seen in 8 ICMP scans performed using ZMap [111] and made publicly available by *scans.io* [6].[2] Figure 3.1(a) shows this comparison where the green bars are entities seen by CDN but not ZMap, the blue bars are entities seen by both

[2] We chose to show the comparison for October 2015 because the largest number of ICMP scans is available for this month.

the CDN and in ZMap, and the red bars are entities seen only by ZMap. As illustrated in Figure 3.1(a), over 40% of the 950 million IPv4 addresses show activity in the CDN logs but do not appear to be active from ICMP probes. This difference is likely mainly attributable to hosts that sit behind NAT gateways [236] and firewalls that do not permit replies to external requests via ICMP or to hosts that respond only intermittently. While this pitfall is well-known [100, 115], we are not aware of any prior studies that quantify this effect at large scale. This incongruity is less pronounced when aggregating the address space to /24 prefixes and ASes.[3] For routed prefixes and ASes, the number of (in)visible units is comparable for both methods, with ICMP outnumbering the CDN for the case of prefixes. Thus, measuring address space activity on a per-prefix or even per-AS level, active measurements provide a significant coverage. On the per-IP level, however, active measurements miss significant activity. We acknowledge that there is a bias in favor of the CDN logs with respect to WWW clients since we compare an entire month worth of CDN logs against 8 snapshots of ICMP scans, which will, naturally, not capture hosts that are active for only short periods of time.

3.2.1 Non-Web Activity

Despite the fact that much WWW content is hosted on the CDN platform and all the successful connections reported, our dataset has at least two limitations: *(i)* the platform typically does not receive requests from Internet "infrastructure" such as routers and servers (though some routers and servers do obtain software updates from the WWW, and some servers obtain content from the WWW to, in turn, complete requests from their clients) and *(ii)* an IP may be assigned to a user who did not interact with the CDN platform. To assess these, we next compare the portion of IPs that do reply to an ICMP request but are not present in our CDN dataset, i.e., the red bars on the right in Figure 3.1(a). While this is roughly only 8% of the IPv4 addresses in the combined CDN/ICMP dataset, we are interested in examining them further.

Figure 3.1(b) shows a classification of these IP addresses, prefixes, and ASes. Here, we use additional data to identify servers and router infrastructure. To identify servers, we rely on additional data gathered by ZMap [111] and made publicly available on *scans.io* [6]: IP addresses that replied to server connection requests using HTTP(S), SMTP, IMAP(S) or POP3(S). To identify router IP addresses, we use one month worth of the Ark [80] dataset and extracted all router IP addresses that appeared on any of the traceroutes (N=490M), i.e., they replied with an ICMP TTL Exceeded error. Close to half of the addresses that did not connect to the CDN, indeed, can be attributed to server or infrastructure IP addresses. This fraction increases when aggregating to prefixes and ASes. We also note, however, that about half of these IP addresses did not show any server or infrastructure activity. These IP addresses might be (a) serving infrastructure that is not present in the Ark dataset, or infrastructure running other protocols than those probed by ZMap, or (b) practically unused IP addresses, or (c) active IP addresses that simply do not connect to the CDN.

3.2.2 Geographical View

To gauge the geographic coverage of our dataset, as well as how it compares to active probing techniques regionally, we next dissect our dataset into geographic regions and countries. To accomplish this, we use allocation data provided by the RIRs [207] to assign regions and countries to each IP address.[4]

[3] Here, we count a prefix/AS as active if we see activity from at least one IP address within the respective prefix/AS.

[4] We acknowledge that the exact geographic location of an IP address does not necessarily correspond to the country where the IP address was registered. We chose this dataset because it is publicly available, and we believe that this data is sufficient to highlight regional characteristics for the purpose of this study.

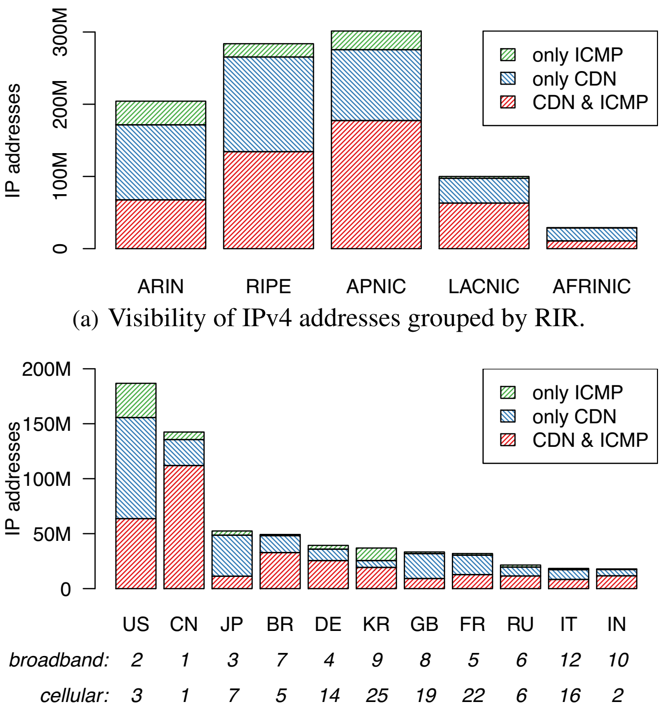

(a) Visibility of IPv4 addresses grouped by RIR.

(b) Top countries, annotated with their rank of broadband and cellular subscribers (ITU, 2015).

Figure 3.2: IP address activity by geographic region.

Figure 3.2(a) shows for each RIR the number of IP addresses that were visible both in the CDN dataset as well as responded to ICMP (bottom red bars), the number of IP addresses that were only visible in the CDN logs (middle blue bars), as well as those IP addresses that were invisible in our logs, but appeared in ICMP scanning campaigns (green bars on top). We observe that the CDN logs provide substantial additional visibility in all regions. When put in relation to the total number of active IP addresses per region, this effect is particularly pronounced in the African region, where the CDN logs increase the number of visible active IP addresses by more than 150%.

In Figure 3.2(b), we show the partition of addresses for the top countries in terms of the number of addresses seen in the CDN logs and the ICMP scans. In addition, we annotate each country with its rank of fixed broadband and cellular subscribers, based on ITU data [153]. Here, we see that the top countries ranked by broadband subscribers are also the top countries visible in the CDN logs. Thus, the coverage of our dataset largely agrees with ITU estimates on global Internet subscribers. This effect is much less-pronounced when ranking countries per cellular subscribers, perhaps because most cellular networks deploy Carrier-Grade NAT (which we will study in Chapter 4), blurring the relationship of subscribers to IP addresses. Secondly, we also see that the fraction of ICMP-responding IP addresses varies heavily per country. In China, for example, we find that close to 80% of the IP addresses do respond to ICMP requests, whereas in Japan only about 25% of the IP addresses reply to ICMP requests. An observation to keep in mind when, e.g., using active measurement techniques to reason about Internet penetration or address space utilization in specific parts of the world.

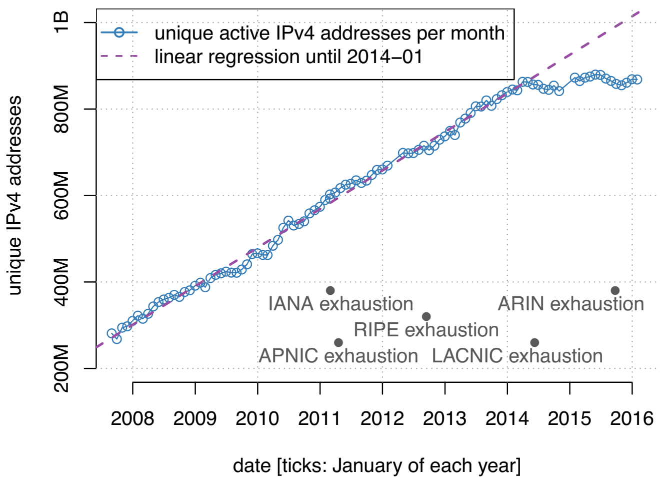

Figure 3.3: Unique active IPv4 addresses observed monthly by a large CDN.

3.3 Rethinking Address Activity

The study of the Internet's growth has attracted the interest of the research community since its early days. One fundamental dimension of this growth is the utilization of the available address space. As originally envisioned, every device on the Internet needs a globally unique IP address to be part of the Internet. Thus, the number of active addresses is a natural metric to track growth of the Internet. Figure 3.3 shows the number of monthly total active IPv4 addresses, as seen by a large commercial CDN.[5] For many years, we observe a linear growth in terms of active IPv4 addresses. In 2014 however, the number of active IP addresses stagnated and we see a relatively constant number of IPv4 addresses fetching content from the CDN.[6] This observation underlines a fundamental point in the history of the Internet, as observed through the lens of the CDN: The growth of active IPv4 addresses has subsided.

IPv4 address space scarcity has recently come to the full attention of the research and operations communities, as four out of the five Regional Internet Registries that manage global IP address assignments have exhausted their available IPv4 address space [232] as of 2015. Figure 3.3 is annotated with the respective exhaustion dates for each RIR. The prospect of exhaustion fueled intense discussions about how to ensure unhindered growth of the Internet by introducing technical as well as political measures to satisfy the ongoing demand, until we reach sufficient IPv6 adoption [96].[7] A fundamental problem, however, is that getting an accurate and detailed picture of the current state of IPv4 address space activity, and how this activity evolves over time, is quite difficult due to the Internet's decentralized structure. Past studies (per our discussion in Section 3.1.1) typically relied on active or passive measurements to enumerate active addresses and blocks. Given that we have now entered a period of stagnation, we argue that a sole enumeration of active IPv4 addresses does not draw a sufficiently accurate picture of address space utilization.

[5] Note: The values in Figure 3.3 are about 5% greater than those reported in Akamai's State of the Internet Report, [31], as the latter restricts to those addresses for which bandwidth is measured, which is also discussed in that report. As the present work is not concerned with bandwidth, we omit this condition.

[6] The number of unique IPv4 addresses contacting the CloudFlare CDN is also stagnant in the timerange between July 2015 and July 2016, ranging at $\approx$ 800M unique IPv4 addresses monthly [187].

[7] In this work, we exclusively focus on IPv4. We note that IPv6 address activity grew significantly during the year of 2015. The number of weekly active /64 IPv6 prefixes (union of active IPv6 prefixes per week) grew from 200M to more than 400M from September 2014 to September 2015. However, we emphasize that IPv6 /64 prefix counts are not directly comparable to IPv4 address counts. For more details on IPv6 client activity seen from the CDN, we refer readers to Plonka and Berger's work [218].

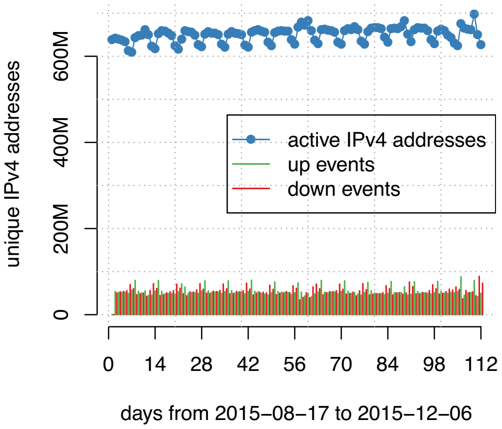

Figure 3.4: Daily active IPv4 addresses and *up/down* events.

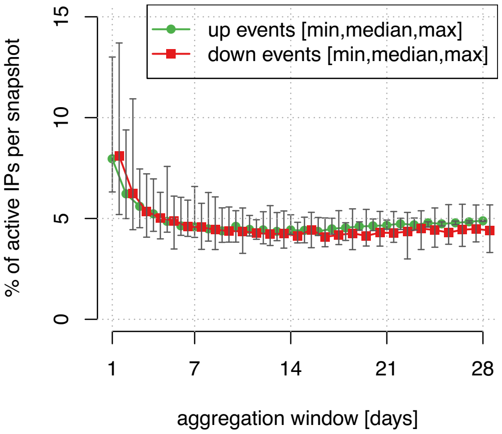

Figure 3.5: Median *up/down* events between subsequent snapshots for different aggregation windows (union of active IPs within window).

In the following, we seek to gain a better understanding of address space activity and utilization. In particular, we explore the following questions:

Q1 How differently does activity manifest itself at different timescales? What are the long- and short-term dynamics of IPv4 address space utilization? (Section 3.4)

Q2 Precisely, what operational practices contribute to these dynamics and which knobs could be adjusted to improve utilization? (Section 3.5)

Q3 Can we extend our understanding of address utilization when taking traffic volume, and measures of the number of connected hosts into account? (Section 3.6)

Q4 Can we extract meaningful address space utilization demographics when combining our various metrics of address space activity? How do these demographics compare for different regions? (Section 3.7)

3.4 Macroscopic View of Activity

In this section, we study IPv4 address activity on a broad scale. In particular we focus on how many addresses our vantage point observes as well as how consistent the set of active IPs is over time. Then, we focus on spatial properties of the observed dynamics and compare our observations with what is visible from the global routing table.

3.4.1 Volatility of Address Activity

To assess address activity over time, we show in Figure 3.4 the daily number of unique IPv4 addresses that contact the CDN over the course of 16 weeks. We observe about 650M unique active IPv4 addresses on a daily basis, and less on weekend days. Although Figure 3.4 shows a relatively constant *number* of active IPv4 addresses, the *set* of addresses can vary. To capture changes in the population of active addresses, we define an *up* event if an address is not seen in a given window of time, e.g., a day or 7 days, but then is seen in the subsequent window. Likewise a *down* event occurs if an address is seen in a given window of time, but not seen in a subsequent window. Figure 3.4 shows an average of 55M daily *up* events, likewise for *down* events. Hence, each day we see 55M addresses showing activity that were not active the day before. Another 55M addresses are active that day, but not on the next day.

We next assess whether this churn appears only on short timescales (i.e., due to short-term inactivity of certain IP addresses) and disappears on longer timescales, e.g., when comparing subsequent weeks to each other as opposed to days. We hypothesize that—if the day-to-day dynamics are the result of short-term inactivity—the churn in active IP addresses decays to zero when comparing larger time windows. In Figure 3.5, we hence partition the 112 days of Figure 3.4 into non-overlapping windows, of a given size. For a window size of 7 days, for example, there would be 16 windows, or snapshots. In each window, we note the union of all active IP addresses. Then, for window i and $i + 1$ we compute the percentage of addresses that had an *up* event as 100 × (the number of addresses in window $i + 1$ that are not present in window i) divided by the number of addresses in window $i + 1$. Hence, for a window size of 7 days, we obtain 15 such percentages. We then note the minimum, median, and maximum of these percentages. We do the analogous computation for *down* events. In Figure 3.5, the two red and green points at $x = 1$ on the x-axis show the min, median, and max of the percentage of addresses that had up/down events on a daily basis, corresponding to Figure 3.4. On an average day, about 8% of the active addresses "come," another 8% "go." We see that the maximum values for *up/down* events are as high as 14%, reflecting changes from weekdays to weekends and vice versa. The red/green dots at $x = 7$ show these statistics when we aggregate our dataset into weeks and compare subsequent weeks. The interesting observation from this figure is that, while churn is more apparent on short timescales (particularly for window sizes for 1 and 2 days, related to day-of-the-week effects), the dynamics in *up/down* events do not decay to zero for higher aggregates. Indeed, we observe that the churn level for aggregates larger than 7 days remains constant at roughly 5%. Thus, whichever aggregation level we choose (days, weeks, months), the set of active IP addresses is in constant change, both on short, as well as on long time scales.

To highlight the long-term effects, Figure 3.6 shows, weekly, the number of newly appearing and disappearing IP addresses as compared to the first week of 2015. That is, for each week in 2015 (x-axis), we show the number of addresses that were *not* active in the first week (positive y-axis, *appear*), but in the given week and also the number of IP addresses that were active in the first week, but *not* in the given week (negative y-axis, *disappear*). In fact, the set of active addresses has changed by as much as 25% over the course of 2015.

3.4.2 Volatility Across Networks and Prefixes

Having seen that the active portion of the IPv4 address space is highly volatile in nature, we next study topological and spatial features of the observed dynamics. In particular, we study *(i)* if networks contribute similar levels of churn, *(ii)* the size of *up/down* events, in terms of prefixes and *(iii)* if this churn is also reflected in the global routing table.

A network view of churn: In Figure 3.7, per Autonomous System (AS), we show the median percentage of IP addresses with an *up* event for each snapshot. That is, we partition the set of addresses into ASes, and we repeat the calculation of Figure 3.5 for addresses in each AS, and obtain a median percentage (calculated over the different snapshots) for each AS. Figure 3.7 shows the CDF of these medians. We only consider ASes for which we saw at least 1K active IP addresses during our observation period and we only show *up* events; the CDF for *down* events is similar. The takeaway from this figure is that highly dynamic IP address activity is not a phenomenon limited to a small number

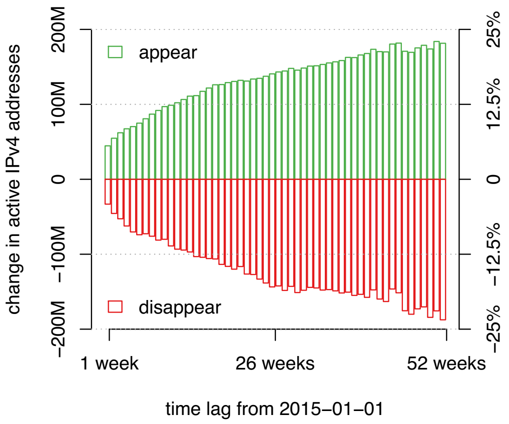

Figure 3.6: Difference in active IPv4 addresses compared to first snapshot of our period (weekly).

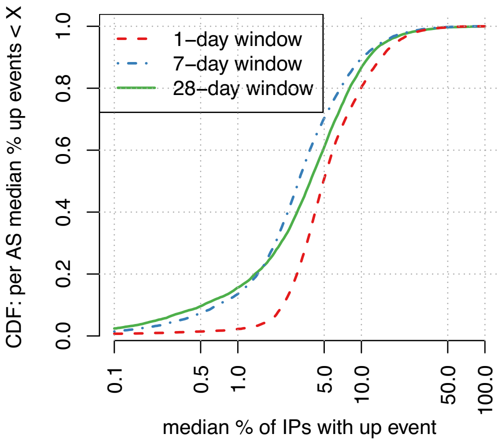

Figure 3.7: CDF: Median % of *up* events per AS and snapshot (only ASes with $>$ 1,000 active IPs, N = 8.6K).

of ASes - rather, about 10% to 20% (depending on window size) of the ASes have a 10% or higher median percentage of IPs with an *up* event. Only about half of the ASes have a churn rate below 5%. We observe similar churn rates for different aggregation windows, with a slight decrease in volatility for some ASes at higher aggregation levels. Thus, churn is a ubiquitous phenomenon, which we observe for a large number of networks.

A prefix view of churn: So far, we have considered *up* and *down* events on a per-address basis (for different time window sizes). Next, we are interested in whether *up* and *down* events really only affect single addresses, or rather entire address ranges. In particular, we are interested in entire prefixes that have been inactive and then some or all of the addresses become active, which we expect would likely indicate network operator actions as opposed to independent, individual user behavior.

To accomplish this, for each per-address *up* event, we find the smallest prefix mask m (where a smaller mask corresponds to a prefix that contains more addresses) in which all addresses either had an *up* event or showed no activity in both snapshots.Figure 3.8 shows a histogram of the fraction of per-address *up* events, for a given window size, where we assign each *up* event to its tagged prefix mask m (the histogram for *down* events looks similar). For example, for a window size of 1 day, more than 70% of the per-address *up* events are associated only with a mask $\geq$ /31, indicating that these dynamics typically only affect individual IP addresses.

For larger aggregates (e.g., 28-days), we still see more than a third of the *up* events in the $\geq$ /31 range, however we also observe some *up* events spanning larger ranges of addresses, with more than 38% of month-to-month *up* events affecting larger address blocks with a mask $\leq$ /24. Thus, a key observation when studying churn across different time aggregates is that a significant proportion of long-term events (38% on a month-to-month aggregation) affect entire prefix masks $\leq$ /24, some of them as large as an entire /16 prefix. These “bulky” events hint towards changes in address assignment practice (e.g., network restructurings), as opposed to churn caused by individual ON/OFF activity of a single IP address. While this is an expected property and holds for some portion of the month-to-month churn, we also notice that this certainly does not hold for all events on larger timescales. In fact, even on a month-to-month scale, more than 36% of the events only affect prefixes of size /31 or even /32, i.e., single IP addresses.

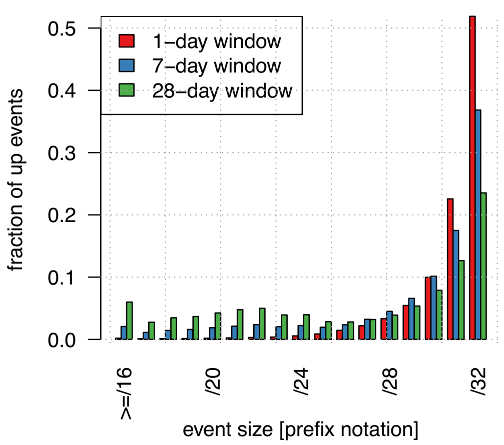

Figure 3.8: Size distribution of *up* events for different time ranges.

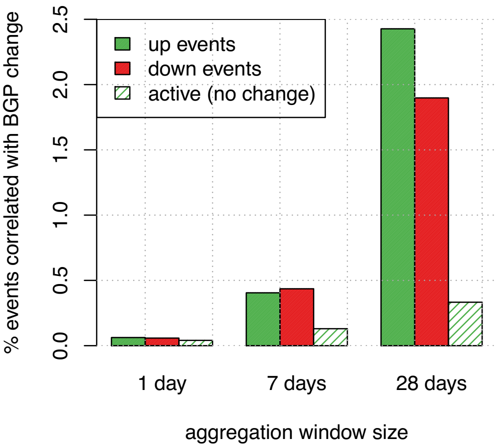

Figure 3.9: Percentages of *up/down* events that go together with a change in BGP. The y-axis ranges between at 0% and 2.5%.

A routing table view of churn: Given that the active IP address population changes by about 25% over the course of a year (per Figure 3.6), we next study whether these dynamics are also reflected in the global Border Gateway Protocol (BGP) routing table. To assess this question, we associate each IP address with its origin AS using daily snapshots of the global routing table.[8] We show in Figure 3.9 the fraction of *up/down* events that go together with a BGP change. Here, we consider both route announcements, withdrawals, as well as origin AS changes, as a "BGP change" event. The green bars show the percentage of *up* events that go together with a BGP change, and the red bars show the percentage of *down* events. In addition, we also plot the fraction of steadily active (no *up/down* event) IP addresses and for what fraction of them we observe changes in the routing table. While we can clearly see that *(i)* IP addresses with *up/down* events are much more likely to correlate with events in the routing table when compared to steadily active addresses and *(ii)* that, on higher aggregation levels, *up/down* events are more likely to correlate with BGP changes, reflecting network changes, we find that *(iii)* only a tiny minority of these events are visible in the global routing table (less than 2.5% for monthly aggregation levels). Thus, the vast majority of volatility in IP address activity is entirely hidden from the global routing table.

3.4.3 Volatility During One Year

Next, we study those IP addresses that were first inactive for a long period and then became active, along with IP addresses that showed activity but then went inactive. We pick the first two months of our observation period (January, February 2015) and the last two months of our observation period (November, December 2015), where we take the union of all active IP addresses that were seen within each snapshot. We then compare the two snapshots, and also the associated BGP activity. Table 3.3 summarizes our results. Continuing the trend shown in Figure 3.8, that churn becomes bulkier on longer time scales, we observe that more than half of the events (65% and 54%, respectively) affected entire address blocks, and are, thus, more likely to be caused by operational changes. However, another large chunk of long-term volatility affects smaller aggregates, down to single IP addresses. The main result

[8] We rely on daily snapshots from a RouteViews collector in AS6539. For larger window sizes, we determine the origin AS for a given IP address using a majority vote of all daily IP-to-AS mappings within the respective time window.

	appear	*disappear*
total	139M	129M
entire /24 prefix affected	65%	54%
BGP no change	87.1%	90.4%
BGP origin change	3.3%	7.1%
BGP announce/withdraw	9.6%	2.5%

Table 3.3: IP addresses that *appeared/disappeared* comparing Jan/Feb 2015 and Nov/Dec 2015, percentage of those IP addresses where the entire containing /24 prefix *appeared/disappeared*, and corresponding BGP changes.

in Figure 3.9, that only a small minority of these events coincide with BGP changes, also pertains at the year-long time scale in Table 3.3. In fact, most of these IP addresses were—and are—still routed by the same AS.

More than 30K ASes announce IP addresses that show long-term volatility in our dataset without any change in BGP configuration. The top 10 ASes, in terms of IP addresses of the class that *appear* or *disappear*, contribute about 30% of the total addresses in each class. These top 10 ASes include major ISPs connecting both residential and cellular mobile users. In fact, we find that ASes contributing the most IP addresses to the *appear* class are also those ASes contributing the most addresses to the *disappear* class. Focusing on our two sets of top 10 ASes, we find 7 of those contributing to the *appear* class are also among the top 10 contributing to the *disappear* class. Thus, while contributing large number of IP addresses with high volatility, the total number of active IP addresses for these ASes varied only marginally, in the order of a few percent. Hence, we can attribute the majority of long-term volatility to AS-internal dynamics, as opposed to, e.g., ASes starting to route and use newly allocated or purchased address blocks.

3.5 Microscopic View of Activity

Given observations of churn in the active IPv4 address space, we now drill down into their root causes. Network operators can assign and use IP addresses in a multitude of different ways. Indeed, many factors contribute to how a network operator assigns IPs to client hosts, e.g., address pool size, lease time, client population, type of clients (enterprise or residential), or other operational practices (static/dynamic address assignment). Thus, it is challenging to characterize the IP assignment strategies within a single network, let alone an entire address space.

To offer a glimpse of how activity typically manifests in different areas of the IPv4 address space, in Figure 3.10, we show examples of activity patterns in four address blocks. Here, we examine specific /24 prefixes, which allows us to present a spatio-temporal view of activity in address-level detail. To generate these plots, we rely on our 4 months' worth of daily IP address activity (x-axis). We then align all IP addresses within the selected /24 on the y-axis in increasing order. Having this "activity matrix" in place, we plot a red point for each day on which a given address was active. With these examples in mind, we introduce two root causes for churn in address activity:

Regular activity patterns: Address assignment practice. The four examples in Figure 3.10 show strong differences in daily address activity. While we see a non-uniform, light utilization in Figure 3.10(a), with a day-of-week pattern for few active addresses, we see heavier utilization in Figures 3.10(b), 3.10(c), and 3.10(d), with a variety of activity patterns involving dynamic assignment from

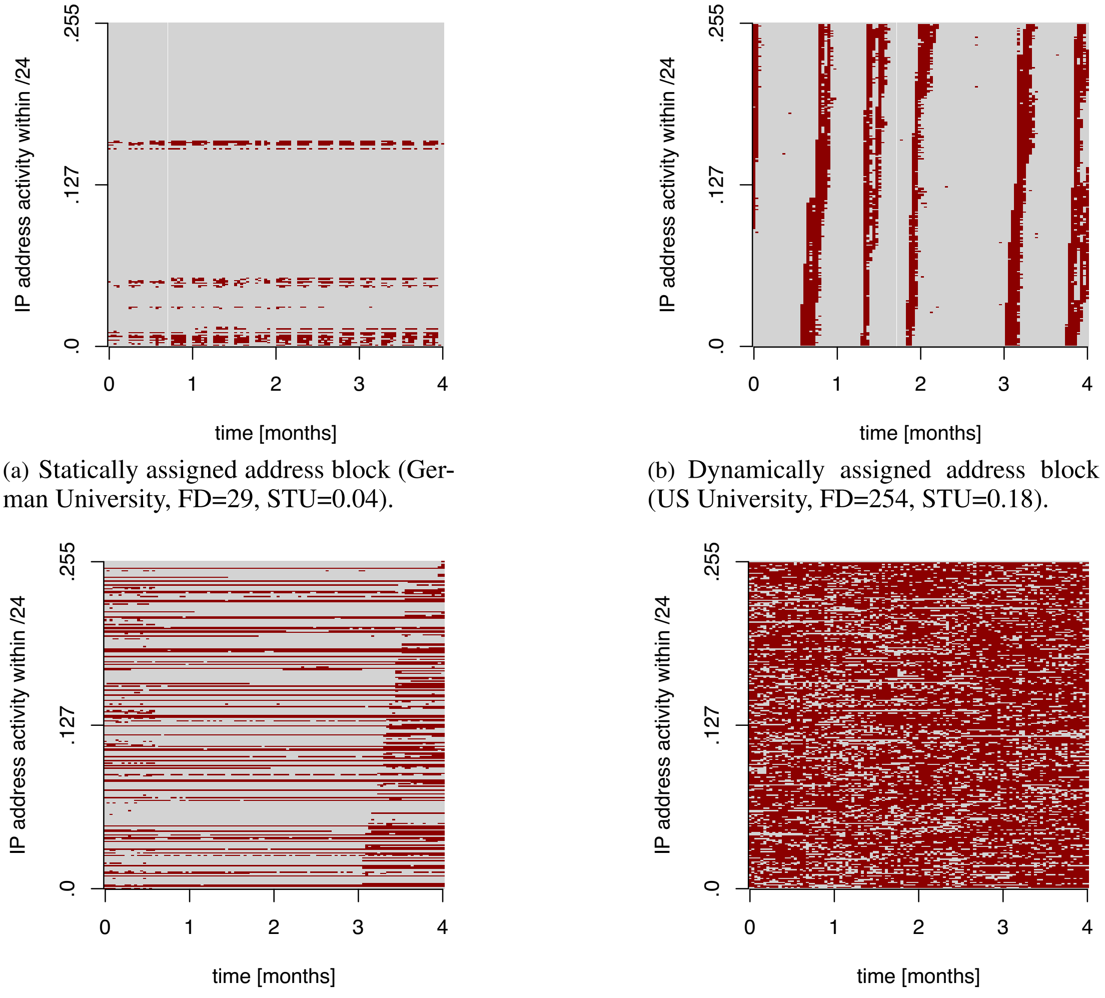

(a) Statically assigned address block (German University, FD=29, STU=0.04).

(b) Dynamically assigned address block (US University, FD=254, STU=0.18).

(c) Dynamically assigned address block with residential users (US ISP, FD=175, STU=0.26).

(d) Dynamically assigned address block with residential users (German ISP, FD=254, STU=0.75)

Figure 3.10: Regular activity patterns: Interplay between address assignment practice and user-behavior, annotated with filling degree (FD) and spatio-temporal utilization (STU) values.

address pools. While Figure 3.10(b) shows a round-robin IP address assignment in an underutilized pool, Figure 3.10(c) shows dynamic addressing with a very long lease time (i.e., the duration for which a specific subscriber holds an IP address), with some IP addresses having almost continuous activity and others having infrequent activity. Figure 3.10(d) shows another mode of dynamic addressing, wherein the ISP sets the lease time to a maximum of 24 hours, thus causing hosts to be frequently reassigned a different IP address. We refer to the activity patterns in Figure 3.10 as *in situ* activity, as they result from address assignment practice and its interplay with end-user behavior in one administratively configured *situation;* that is, we have no evidence that the situation, nor the activity pattern, changed due to network reconfiguration. An important observation is that *in situ* activity in address blocks varies significantly amongst those that have different address assignment configurations.

Changed patterns: Modification of assignment practice. As shown in Figure 3.11, we also observe activity patterns that are temporally or spatially inconsistent. This is some evidence that the patterns'

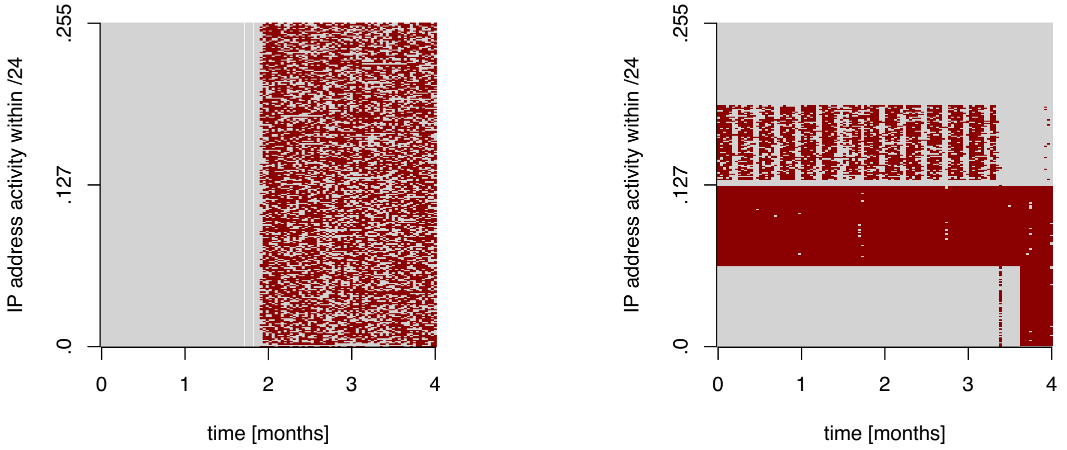

(a) German University, FD=256, STU=0.32.

(b) German University, FD=187, STU=0.38.

Figure 3.11: Modified assignment practice.

dynamics are not the result of constant address assignment policy, but, rather, are the result of address (a) reallocation, (b) assignment reconfiguration, and/or (c) repurposing.

We next study address activity pattern at large scale. In particular, we are first interested in detecting which portions of the address space show a consistent address assignment pattern as opposed to blocks that show major changes in their activity pattern. We then dive into the former, activity patterns that are the result of address assignment practice in conjunction with end-user behavior. Here, we put a particular emphasis on the resulting *utilization* of address blocks.

3.5.1 Block Activity Metrics

In order to comprehensively characterize IP address activity, it is imperative to use metrics that capture the activity spatially, i.e., over the IP address space of an address block, and temporally, i.e., across time. To capture address activity patterns, next, we introduce two metrics:

IP address filling degree (FD): this metric captures the number of active IPs within an address block within a window of time. There is not a single address block size that is ideal, but we chose a /24 Classless inter-domain Routing (CIDR) prefix, i.e., the smallest distinct, globally-routed entity. This is a compromise, since we recognize that both smaller prefixes are sometimes more appropriate, as in Figure 3.11(b), and that larger prefixes sometimes exhibit uniform patterns of activity, e.g., Figure 3.10(b). Values of this metric range from 1 to 256. We will later see that this metric is particularly helpful in dissecting static from dynamic addressing mechanisms.

Spatio-temporal utilization (STU): this metric captures the aggregate activity of active IPs over time. We define utilization as the fraction: spatio-temporal activity divided by the maximum spatio-temporal activity, for a given block and window of observation (time). Relying on our four months worth (112 days) of daily activity data, the spatio-temporal activity can range from 1, where one single IP address was active for one day, up to $112 \times 256 = 28672$, where all addresses in a block were active every day, which would be the maximum spatio-temporal activity. STU is this value normalized as a fraction with range 0 to 1.

We annotated Figures 3.10 and 3.11 with their respective values for filling degree (FD) and spatio-temporal utilization (STU). In these examples, FD varies from values as low as 29 to as high as 256. The STU varies from 0.04 up to 0.75.

3.5.2 Detecting Change

As a first-order partitioning of the active IPv4 address space, we are interested in identifying address blocks with a significant change in address assignment practice during our observation interval. Per Section 3.4.2, we know that some portion of address churn on longer timescales affects larger address ranges ("bulky events") than do short-term changes. To quantify changes in address assignment, we rely on our spatio-temporal utilization metric. In particular, Figure 3.12 shows the maximum change in spatio-temporal utilization on a month-to-month basis for each active /24 block. Here, we observe that the majority (90%) of the /24 blocks cluster around the origin, i.e., they do not show a major change in their utilization. Another 10% of the active address blocks, on the other hand, are located more closely to the tails of the CDF, these are blocks for which we observe significant changes in address activity.

To dissect address blocks into *major change* and *minor change* blocks, we set a threshold at $X = \pm 0.25$. We decided to use this threshold, as it retains cases of heavy *in situ* change, e.g., Figure 3.10(b), but excludes of major configuration change, e.g., Figure 3.11. Based on this threshold, we find that as many as 9.8% of the active /24 blocks show major change in their address activity within our four months period, while 90.2% of the blocks show no more than minor change. Thus, we separate blocks that likely underwent reallocation or change in address assignment practice (*major change*, Figure 3.11) from those that did not (Figure 3.10).[9]

3.5.3 Static and Dynamic Addressing

Having culled out those blocks with major changes, we next focus on the activity characteristics of steady address blocks. Since we have observed that the address assignment policy greatly influences its activity patterns, we would like to identify specific assignment practices. We pay particular attention to *utilization* characteristics associated with these practices. We argue that an address block's utilization is determined by (a) its address assignment policy and (b) the behavior of its users and their hosts.

Static vs. dynamic addressing: As a first cut, we are interested in how *static* and *dynamic* addressing mechanisms compare when it comes to address space utilization. In the static case, the ISP assigns a fixed IP address for each device/subscriber. Dynamic addressing, on the other hand, automatically assigns IP addresses from predefined ranges. In order to apply our metrics, we wanted an initial set of blocks that are known to be likely statically or dynamically assigned. To this end, we used PTR (reverse DNS) records and tagged /24 blocks containing addresses with consistent names that suggest static (keyword `static`) as well as dynamic (keyword `dynamic, pool`) assignment, a well-known methodology [196,227,280]. In total, we find 456K dynamic /24 address blocks and 262K static address blocks. We then compare their activity based on our dataset. Figure 3.13 shows a CDF of the filling degree (active IPs per /24) for the two subsets of static or dynamic /24s as well as for the entirety of our dataset. Comparing the curves for dynamically and statically assigned address blocks, we see a stark difference: While 75% of static /24s show a filling degree lower than 64 IPs, more than 80% of the dynamic /24s show a very high filling degree, i.e., higher than 250 IP addresses. When comparing these observations to our entire dataset, we observe that about 50% of the entire visible address space

[9] We acknowledge that some changes in address assignment might result in only minor STU change and that others might result in larger STU change. We chose a threshold based on anecdotal examination of activity patterns.

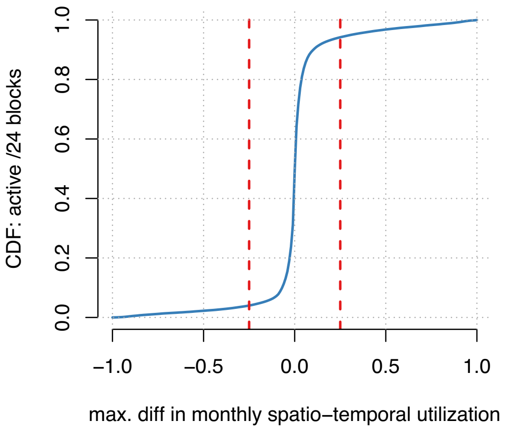

Figure 3.12: Maximum monthly change in spatio-temporal utilization per /24 block. We select 90% of the blocks in the stable region.

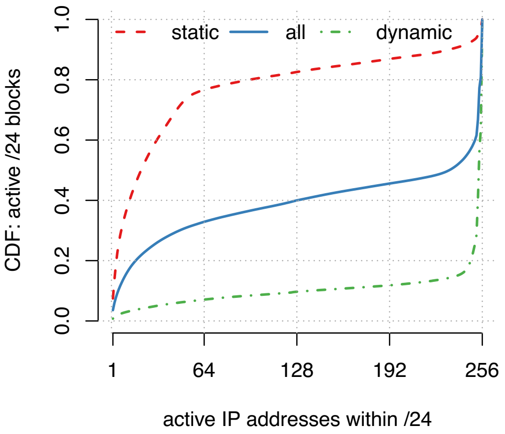

Figure 3.13: Filling degree of active /24 blocks, where we dissect some identifiable blocks to be static or dynamic using reverse DNS.

shows a very high filling degree (higher than 250). Another 30%, by contrast, show filling degrees lower than 64. If our DNS-derived samples are representative, most sparsely populated /24 blocks are statically assigned and most dynamic pools cycle, i.e., have every address assigned at least once, during our observation window of 4 months, resulting in a high filling degree. However, about 20% of the active /24s that remain have varying filling degrees. These are either statically-assigned blocks with higher utilization or dynamically-assigned blocks with quite little utilization, e.g., those with long lease times as in Figure 3.10(c)).

Dynamic address pools: We find that dynamically-assigned /24 prefixes generally show a very high filling degree, with more than 250 active IP addresses in more than 80% of the cases. Activity patterns of dynamic address pools heavily depend on the configured assignment policy, i.e., the *pool size* in relation to the number of connecting devices. Figures 3.10(b) and 3.10(d) both show dynamic addressing patterns, however we see that their utilization is very different. To shed more light into such dynamic pools, we make use of our second metric, the *spatio-temporal utilization*. Focusing on those 1.2 million /24 blocks that have a very high filling degree (larger than 250, and hence likely dynamically assigned), Figure 3.14, shows their spatio-temporal utilization as a percentage of their maximum possible utilization.[10] Here, we see that most of these address blocks have high utilization, with most blocks at more than 80%. In fact, we even see some 60K /24 blocks with 100% spatio-temporal utilization. This extraordinary utilization hints that they might contain shared proxy or gateway addresses; we will revisit these in Section 3.6. We also see more than 450K /24 prefixes with a utilization lower than 60% and 200K /24s with a utilization even lower than 20%.

3.5.4 Potential Utilization

Figure 3.13 makes it clear that the spatio-temporal utilization of address blocks differs dramatically. We find that *static* vs *dynamic* addressing mechanisms play an important first-order role, and now present some estimates on an address block's maximum potential spatio-temporal utilization. We acknowledge that the activity seen trough the CDN logs present us with a *lower bound* both in terms of active IP addresses as well as with regards to their spatio-temporal utilization. We constrain this exercise to only

[10] Figure 3.14 looks similar when only considering dynamic address blocks as identified using reverse DNS PTR records.

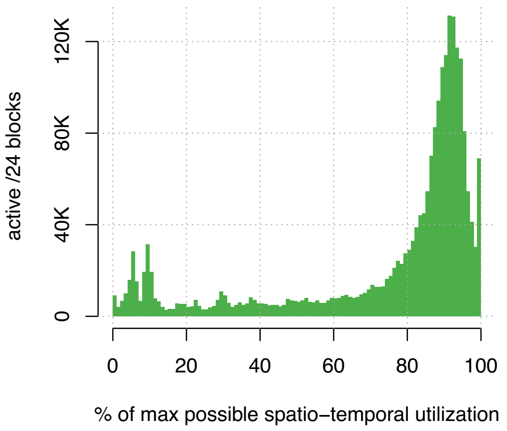

Figure 3.14: Spatio-temporal utilization within /24 blocks with >250 active IP addresses.

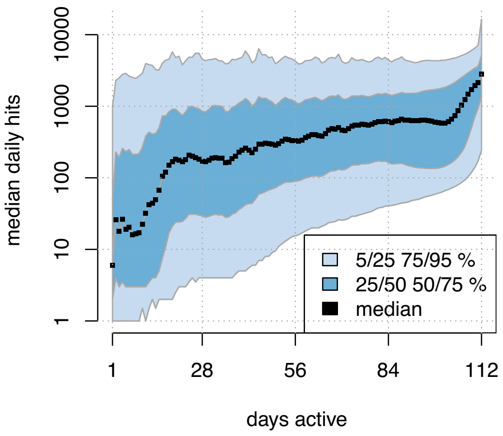

Figure 3.15: Median daily hits per IP address binned by activity (days). The y-axis is log-scaled.

those blocks known to be active, i.e., those that are known to be allocated, globally-routed, and in operation. We argue that increasing utilization in these blocks is—*in some instances*—a mere configuration issue. Sometimes this means switching from static to dynamic assignment, but other times it means only reconfiguring an existing dynamic pool.

Specifically, we find that more than 30% of the active IP address blocks, more than 1.5M /24 blocks, have a filling degree lower than 64 active IP addresses. Our DNS PTR-based tagging method suggest that static address assignment practices are the main driver for low spatio-temporal utilization of IP address space. We acknowledge that hosts that do never contact the CDN might also affect the filling degree of certain address blocks. On the other hand, for the 50% of the active /24 address that appear to be dynamically managed, we find that the majority have high spatio-temporal utilization, i.e., more than 80%. However, we also find that about one third of dynamic blocks show low spatio-temporal utilization; Figure 3.10(b) is a striking such example. We argue that—as these address blocks are already dynamically assigned—reducing their pool sizes could instantly free significant portions of address space.

3.6 Traffic and Devices

Up to this point, we studied the activity of an IPv4 address with respect to time and to neighboring addresses, e.g., in a /24 prefix. We've seen a variety of address activity patterns and associated addressing mechanisms. Next, we take another dimension into account: traffic. In particular, we'd like to answer these questions: *(i)* How does address activity correlate with traffic? *(ii)* Do we see a long-term trend with respect to the fraction of traffic associated with the heavy-hitter addresses? *(iii)* How does traffic contribution relate to the number of connected end hosts? Afterward, in Section 3.7, we will combine traffic metrics and host estimates with the activity measurements of Section 3.5 to obtain a comprehensive perspective of the active IPv4 address space.

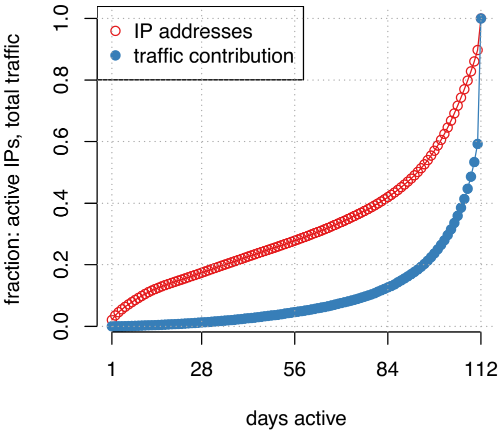

Figure 3.16: Cumulative fraction of active IP addresses in each bin, cumulative traffic contribution per bin.

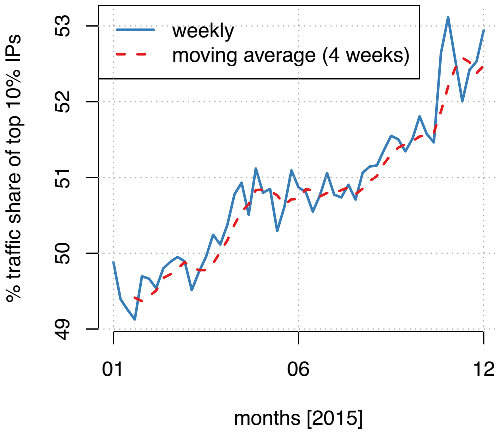

Figure 3.17: Relative share of total traffic of top 10% IPs. The y-axis starts at 49%.

3.6.1 Activity vs. Traffic

Firstly, we are interested in how the binary notion of activity of IP addresses is related to the volume of traffic that the CDN delivers to them. For this, we rely on our dataset that captures the number of daily HTTP requests as issued by each individual IP address (as described in Section 3.2). We group all IP addresses that were active during our 4-month (112 days) period into 112 bins, corresponding to the number of days each individual IP address was active. Figure 3.15 shows the median *daily* hits that were issued by the total count of IP addresses in each bin, where we only consider days where an IP address issued at least one hit. We also show the 5, 25, 75 and 95 percentiles for each bin (the y-axis is log-scaled). Note the strong correlation between temporal activity of IP addresses and their daily traffic contribution. While addresses that were only active for a few days issue only a median of fewer than 100 requests per day, the traffic contribution is much higher for addresses that were active on more days. Indeed, we see that the traffic contribution significantly increases for IP addresses that were active almost every day ($\geq$ 110 days), and those addresses that were active every day show an even higher median daily traffic contribution. This observation becomes clearer when looking at Figure 3.16, where we plot cumulative fractions of the total number of IP addresses falling into each bin (red) and their contribution to the CDN client's total traffic (blue). While only less than 10% of IPv4 addresses were active every single day, these addresses account for more than 40% of the CDN's total traffic. The combination of continuous daily activity over the course of four consecutive months as well as the significantly larger contribution in overall traffic suggests that those 10% of the active IPv4 addresses include gateways, e.g., NAT routers and web proxies, aggregating the traffic of multiple users as well as WWW client bots (e.g., employed by search engines or content aggregators).

3.6.2 Traffic Consolidation

Given that we have reached a stage in which the number of active IPv4 addresses has stagnated, we were curious whether there is an observable trend over 2015 of increasing traffic concentration in the heavy-hitter addresses. To visualize this, we show in Figure 3.17 the traffic share of the 10% of addresses with the greatest traffic. (Note that the y-axis starts at 49%.) Here, we use our weekly dataset to show how this trend has been developing over the entirety of the year 2015. Figure 3.17 indeed shows a clear trend of traffic consolidation. While in January 2015, those IPv4 addresses received a share between 49% and

50%, we see that their traffic share steadily increased over the course of the year. As of December 2015, the top 10% of the active IPv4 addresses consume an additional 3% of the total traffic that the CDN serves, which we believe is a notable increase over one year. Given the stagnating count of active IPv4 addresses, we expect IPv4 traffic consolidation to continue except, e.g., when and where alleviated by IPv6.

3.6.3 Estimating Relative Host Counts

Having understood that the characteristics of activity of an IP address vary dramatically, both regarding its utilization as well as volume of traffic, we are next interested in how many hosts reside in a given address block. With the increasing prevalence of address sharing mechanisms (e.g., Carrier-Grade NAT [236]), active IP addresses are no longer an accurate metric to quantify the number of hosts in a given address block, e.g., to reason about Internet penetration, activity of individual users, or address activity in general.

While we do not have data available that provides us with a definitive number of connected hosts per IP address, we will estimate by *HTTP User-Agent strings*, as a proxy. Whenever a Web object is requested from a server, the respective client application identifies itself by providing a User-Agent string within the HTTP request header. We extended the CDN data-collection platform to store a random sample of HTTP User-Agent strings of connecting hosts. Due to the high volume of this data, we only store the User-Agent field for 1 out of 4K HTTP requests, and we restrict this analysis to the last month of our observation period.

In the canonical case, the User-Agent identifies the browser version, OS version as well as the screen resolution. However, in more recent times, primarily driven by smartphone applications, which typically identify themselves and their version number with an individual User-Agent string, we see much higher diversity in these strings [281]. HTTP User-Agent strings have been used in the past to quantify host populations behind NAT devices in residential networks [190]. Here, we use them only as a *relative measure of host counts* per address block, i.e., we do not claim to be able to exactly quantify host populations. This is mainly because (a) the coarse-grained sampling of this dataset and (b) the fact that some single devices introduce multiple User-Agent strings (e.g., a smartphone running many applications) while, simultaneously, those and other devices running the same applications might share an IP address which, thus, will consolidate *unique* User-Agent strings (on a client address); the former can result in overestimation and the latter in underestimation of the host population. We point out that some hosts might also use arbitrary strings as User-Agents, which do not necessarily reflect a specific application, or strings with unique transaction identifiers per request, which can result in an overestimation of the number of connected hosts in some instances. In future work, we plan to further assess the validity of the collected User-Agent strings.

Figure 3.18 shows, for each active /24 block, the number of User-Agent samples (x-axis) versus the number of *unique* User-Agent Strings (y-axis). Thus, the x-axis value is an estimate of *traffic volume* (based on sampled requests/hits) issued by hosts in a block and the y-axis value is a relative measure of the *number of hosts* residing in a block. Overall, we see a strong correlation between traffic and hosts. Upon closer look, we can dissect the area in the plot in three groups: The first (and largest) group of /24 blocks ranges from the center of the figure to the lower left. Indeed, here we find the bulk of address blocks, e.g., from residential ISPs. Then, we have blocks that are shifted more towards the right, but show a low number of unique User-Agent strings (bottom right in the figure). By further investigation, we found that these blocks are mainly related to automated activity, e.g., crawling bots, which issue a large number of requests, but do so with one (or very few) User-Agent string(s). More interestingly, we

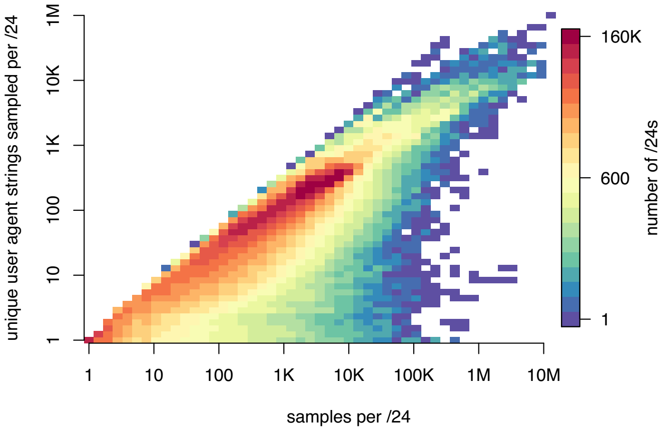

Figure 3.18: Diversity of User-Agent strings per /24 block.

see a third region, in the top right, of a huge number of requests, and a very high diversity of User-Agent strings. A closer inspection of these blocks reveals that it is precisely those blocks that correspond to gateways, aggregating the traffic of thousands of end-users. We manually inspected the top 5K blocks in the top-right region of the plot. Using WHOIS information, we find that more than half of these blocks belong to ISPs located in Asia and that the majority is in use by cellular operators.

3.7 Deriving Demographics

In this section we combine our activity metrics (spatio-temporal, traffic, relative host counts) to provide a comprehensive perspective of the active IPv4 address space. Our three features are fundamentally different in nature, which manifests itself also in different scaling of our derived values per address block. Hence, to project our features onto a unified scale, we first need to *normalize* our measures of traffic and the relative host count. Our measure of spatio-temporal utilization is already normalized to a range $(0, 1]$. We normalize the traffic contribution as well as the relative host count, by using a log-transform of the value per /24 block and divide it by the maximum log-transformed value of all active /24 blocks. Having these three normalized values per /24 block in hand, we next bin the resulting values into 10 intervals of a length of 0.1. This results in a 3-dimensional array with 1000 entries. We now assign each /24 block to one of these bins within our matrix.[11]

3.7.1 Internet-wide Demographics

Figure 3.19 shows a 3D-visualization of our feature matrix, where we indicate the number of /24 address blocks falling into each bin by scaling the size of the respective sphere. We can make several observations from this plot: *(i)* We see a strong division of address blocks along the spatio-temporal utilization axis. While one set of blocks is clustered towards values with a very small spatio-temporal utilization (less than 0.2), another set is clustered towards very high spatio-temporal utilization. Recalling Section 3.5, this can mainly be attributed to varying addressing mechanisms. *(ii)* Taking the traffic contribution into account, we see that densely utilized address blocks typically have a higher traffic volume. However, this observation is not always true; we also see significant portions of the address

[11] For traffic contribution, the median of the 1st/5th/10th bin corresponds to $4/1.5M/44B$ monthly hits; For relative host density, the median of the 1st/5th/10th bin corresponds to $2/2K/500K$ unique sampled User-Agents strings.

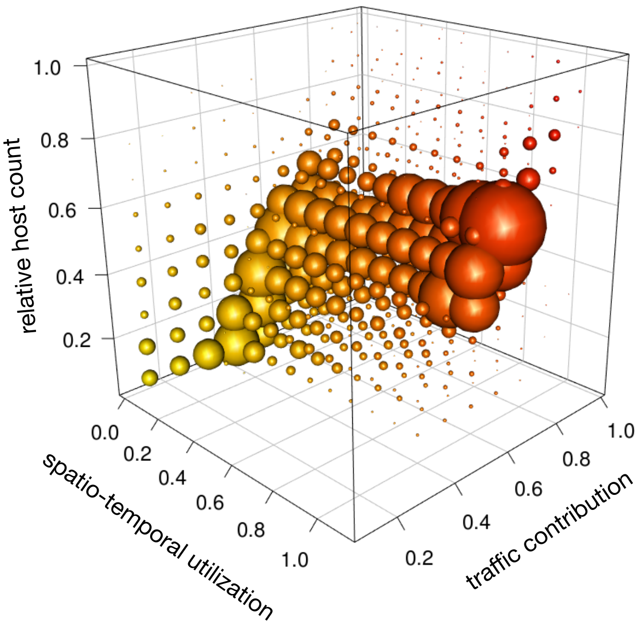

Figure 3.19: Characterization of the active IPv4 address space: Spatio-temporal activity, traffic contribution, relative host count per /24.

space with high traffic volume in sparsely-populated areas. *(iii)* When relating these two features to our host count measure, we again see a higher host count for highly-utilized and traffic-intensive blocks. In particular, we see only a very tiny portion of /24 blocks that fall into the highest bin for the host count metric. These blocks typically also show a maximum spatio-temporal utilization and maximum traffic contribution (small spheres at the top-right). It is important to notice that blocks contained in these small spheres are responsible for a significant share of the CDN's overall traffic.

3.7.2 Regional Characteristics

Lastly, we dissect the address space by regional registries. Recall that the address space is subject to management from 5 different organizations (RIRs, Section 3.3). Each RIR applies different management policies and the current state of address exhaustion also varies per RIR. We hence believe that this grouping can assist in understanding the current status of the address space in each of these regions and support policy decisions when it comes to managing the last remaining blocks and re-allocations of address blocks already in use. Figure 3.20 shows an address space categorization for the five RIRs. Here, we plot the spatio-temporal utilization and traffic contribution on the x and y axes, and indicate the relative host counts by the color scale (gray: low relative host count, red: high relative host count). Again, we adjust the size of the circles to reflect the number of /24s falling into each bin.

We can see that about half of the active address space within the ARIN region clusters towards the left, i.e., shows low utilization, low traffic contribution. However, we note that there are some heavily active address blocks also in this region (small red dots at x = 0.2 / y = 0.8,0.9). We see that the other regions have more of their address space being highly-utilized, which is especially true for LACNIC and AFRINIC. A possible explanation for this behavior is that LACNIC and AFRINIC were incorporated much later than the other RIRs and had address conservation as a primary goal from the very beginning [232]. Noticeably for the APNIC and AFRINIC regions, we see a significant chunk of /24 blocks

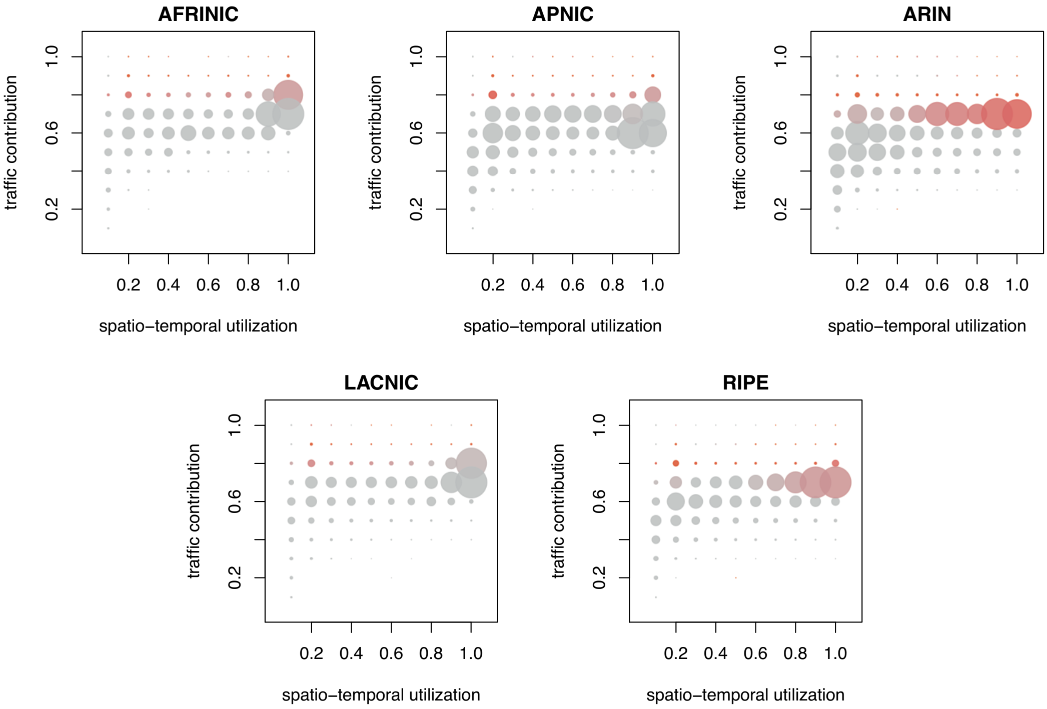

Figure 3.20: Breakdown of IP address space characterization per RIR. Color encodes the relative host count.

towards the top-right of the figures (x = 1.0, y = 0.8), which also show a very high relative host count. This hints towards increased proxying/gateway deployments which is more pronounced in these regions when compared to, e.g., ARIN.

3.8 Implications

Our findings have a number of direct implications both for researchers, network operators, and for Internet governance.

Implications to measurement practice: We count 1.2 billion active, globally-unique IPv4 addresses, more than has been reported previously, except by statistical estimation [283], boding well for future use of such statistical models and techniques driven by sampled observation. Our address count analysis implies that remote active measurements are insufficient for census or complete survey of the Internet, particularly at IP address-level granularity. Also, our passive measurements have shown extensive churn in IPv4 addresses on all timescales, which implies that any census needs to be qualified by the observation frequency and period.

Implications to Internet Governance: The 1.2 billion active addresses we count represent 42.8% of the possible unicast addresses that we see advertised in the global routing table. If we restrict our implications to the 6.5 million /24 prefixes in which we observed active WWW clients addresses (Table 5.3), i.e., exclude blocks that may be dedicated to network infrastructure and services, we see that roughly 450 million address may have been unused. If some large subset actually are unused, today, one could imagine reallocating them for use in IPv6 transition mechanisms that require IPv4 addresses, e.g., NAT64 and DNS64 [63, 64], or as a commodity whose supply might last years in a marketplace, based on past rates of growth in IPv4 address use (Figure 3.3).

IPv4 address markets are an operational reality, governed by the respective RIR policies [232]. A pertinent implication of our work for these markets is that our metrics, combined with the appropriate vantage points, are able to determine spatio-temporal utilization of network blocks. This can aid RIRs in determining the current state of address utilization in their respective regions, in determining if a transfer conforms with their transfer policy (e.g., four of five RIRs require market transfer recipients to justify need for address space) as well as in identifying likely candidate buyers and sellers of addresses.

Implications to network management: It is feasible for networks to monitor their traffic (e.g., at border routers) and employ our metrics and perform our analysis on a continual basis. Measuring spatio-temporal utilization would enable an operator to more efficiently manage the IPv4 addresses they assign, especially in networks such as those discussed in Section 3.6. Networks that make gains in efficiency by discovering unnecessary address blocks may decide to become sellers in the IPv4 transfer marketplace. More generally, we believe that our measurements can serve as input for fruitful discussions on address assignment practices and their eventual effect on address space utilization.

Implications to network security: Our observations of many, disparate rates of change in the assignment of IP addresses to users has consequences for maintaining host-based access controls and host reputations. A host's IP address is often associated with a reputation subsequently used for network abuse mitigation, e.g., in the form of access control lists and application rate-limits that specifically use those IP network blocks or addresses as identifiers with which some level of trust is (or is not) associated. Unfortunately, in this way, addresses and the network blocks become encumbered by their prior uses and the behavior of users within. This happens when reputation information is stale. The implication of our work here is that it can *inform* host-based access control and host reputation, e.g., by determining the spatial and temporal bounds beyond which an IP addresses reputation should no longer be respected. Further, our change detection method (Section 3.5.2) could be used to trigger expiration of host reputation when networks are renumbered or repurposed.

Implications to content delivery: Details about user activity at the address level are valuable in CDN operation. A key responsibility of CDNs is to map users to the appropriate server(s) based on criteria including performance and cost [210]. Details about active IP addresses and network blocks are increasingly important when the CDN uses end-user mapping [87], where client addresses are mapped to the appropriate server.

3.9 Chapter Summary

In this chapter, we studied global IPv4 address activity. We first present an overview of the visibility from various vantage points into address activity. While we find that no single vantage point can ultimately capture address activity in its entirety, our CDN logs give us both detailed and broad insights into global address activity of Web clients. In total, we capture the activity of 1.2 billion IPv4 addresses across the year of 2015, the highest number of active IPv4 addresses ever measured. In our historical analysis of IPv4 address counts, we make an interesting observation: After many years of constant growth, the number of IPv4 addresses that contact the CDN have stagnated since 2014. Thus, the exhaustion of the available IPv4 address space is not only visible when studying address allocations (Chapter 2), but is now manifested directly in address activity. Hence, we have now entered an era in which simple address counts do not capture the increasingly complex situation of usage of the IPv4 address space. In turn, we develop techniques and analyses that allow us to study structural and qualitative aspects and dynamics of the post-exhaustion IPv4 address space.

We find that the population of active IPv4 addresses shows substantial dynamics, both on short as well as on long timescales. Over the course of one year, almost 25% of the active IPv4 address population changed. We find that address churn is ubiquitous: All networks contribute to this churn, to a varying degree. Most dynamics, including long-term changes, are entirely hidden from the global routing table and thus reflect changes within autonomous systems. We then study address activity in detail and reveal a diverse set of activity patterns of individual addresses and address blocks. We attribute this churn to *operational changes* in address assignment practices within networks as well as to *regular* address activity. In order to quantify address activity, and to measure the utilization potential of the available IPv4 address space, we develop metrics that allow us to capture address activity in the stagnant IPv4 Internet: spatio-temporal aspects of address activity, address-associated traffic volume, and relative host counts.

Our analysis reveals significant potential in possibly unused as well as underutilized address blocks. Our findings bode well for address reassignment practices to increase the utilization of the available IPv4 address space, our first mitigation approach to IPv4 address exhaustion (as introduced in Section 1.3). Networks that make gains in efficiency by discovering unnecessary address blocks may decide to become sellers in the IPv4 transfer marketplace. While address restructurings within networks are up to the respective network operators, address markets are guarded by policies. Our metrics and findings have the potential to aid policymakers when regulating such markets. In fact, we find that address utilization vastly differs across different administrative regions of the address space.

Our measurements of address activity show that address exhaustion affects the dynamics of address activity. Besides the stagnation of active IPv4 address counts, we observe an increasing concentration of traffic on fewer, heavily active IPv4 addresses, suggesting increasing deployment of gateways to multiplex more users behind fewer public addresses. Augmenting our findings with our metric of relative host counts, we find highly uneven utilization, with some portions of the address space exhibiting enormous activity, where single IPv4 addresses represent potentially thousands of individual end users. Our metrics can capture such activity and our findings can also help to adjust address reputation systems, which now need to deal with this situation, with IPv4 address sharing across space and time. The findings in this chapter motivate us to study address sharing mechanisms in detail in the next chapter.

4

Carrier-Grade NAT to the Rescue

In the previous chapter, we studied IPv4 address activity on a broad scale. Our finding of increasing traffic concentrated on fewer, highly-active IPv4 addresses suggests increasing gateway deployments in end-user networks. This observation motivates us to study one particular mitigation strategy in detail: Carrier-Grade NAT. To accomplish this, we next develop techniques and analyses to identify and pinpoint Carrier-Grade NAT deployment in the Internet. Our techniques allow us to quantify the prevalence of CGN deployment at scale, providing empirical data on which technology different ISPs in different parts of the world use to mitigate their IPv4 scarcity issues. They also allow us to study dominant properties of the identified CGNs, which reveals both problems that networks face when rolling out CGN, as well as the impact that this technology has on the end users residing in that ISP.

Today, Network Address Translation (NAT) [112] is ubiquitous at the edge of home networks to meet both the ISPs' desire to conserve IP addresses and the users' requirement of connecting a multitude of devices. IP address scarcity has long moved beyond home networks and onto the global stage [232]. Nowadays, large ISPs are confronted by address shortages, and hence turn to a well-worn coping technique: NAT. Instead of aggregating small populations of tightly-knit users and devices within one residence under a single IP address, *Carrier-Grade NATs* (CGNs) apply NAT to many independent and disparate endpoints spanning physical locations. On one level we can view CGNs as representing a second instantiation of a well-known technique for combating address shortages. While tempting, conflating CGNs with small edge-based NATs represents a false equivalence, for two reasons: (i) by operating at large scales, CGNs face issues not present in residential settings, which have received more examination, and (ii) CGNs generally represent a second level of address translation—i.e., CGNs operate *in addition* to existing edge-based NAT—and therefore compound some of the issues that address translation raises.

While we know anecdotally that ISPs deploy CGN, we are not aware of quantitative studies of the prevalence and operation of CGNs in the wild. In this chapter, we take a first step toward developing an empirical understanding of these increasingly crucial pieces of Internet infrastructure. We make four high-level contributions:

Operator Perspectives on CGNs: We begin by presenting a survey of operators in Section 4.1. We distributed a questionnaire on pertinent mailing lists, seeking to shed light on operators' motivations and experiences with CGN operation in the wild. We received illuminating input from 75 operators. Our survey reveals widespread adoption of CGN technology—with over half of the responding operators having deployed CGNs or planning to in the near future—despite the resulting operational difficulties.

Measurement Methodology: One of the key characteristics of CGNs is their *transparent* operation from the perspective of endpoints. While transparency has its benefits (e.g., clients require no setup process to use a CGN), it complicates detection and measurement of CGNs. Multiple levels of address translation increase the difficulty further as each step overwrites any evidence a previous NAT left in the traffic. Therefore, the sender of a packet cannot tell if or how many times the source address will be translated on the path towards a destination, and the recipient cannot know the original source of the packet. To address these difficulties we introduce two methods in Section 4.3 for exploring CGNs. First, we observe that some nodes in the BitTorrent DHT mistake addresses internal to a CGN for external addresses and therefore propagate ("leak") these to other nodes. Therefore, we are able to derive a broad understanding of the deployment of CGNs by probing the DHT. Our second set of methods relies on extensions to the Netalyzr measurement platform [170], which allow us to study the presence and detailed properties of CGNs based on locally available addressing information, repeated connectivity tests, as well as a new method that leverages the stateful nature of NATs and uses TTL-limited probes to force retention of state in some hops while allowing it to expire in others.

Studying Global CGN Presence: IPv4 address scarcity manifests differently for different networks in different parts of the world [232]. Our CGN detection methods give us a broad and unprecedented view into the global deployment of CGNs, which we present in Section 4.4. Our vantage points cover more than 60% of the Internet's "Eyeball ASes" that connect end users to the Internet. We find the CGN penetration rate to be 17–18% of all Eyeball ASes. Moreover, we find that CGN deployment is ubiquitous in cellular networks with more than 90% of all cellular ASes deploying CGNs. We also find a direct relationship between regions with higher perceived IPv4 address scarcity and CGN deployment.

Understanding CGN Behavior: CGNs tackle a massive resource distribution problem, whereby scarce public IPv4 addresses are multiplexed using a relatively small set of internal IPv4 addresses and a limited port space across thousands of end hosts. CGNs can be configured in a multitude of ways, with currently little known about CGN configurations, dimensioning, and behavior in the wild. Hence, in Section 4.5 we make our final contribution: a deep dive into the properties of deployed CGNs. We analyze the internal address ranges used by CGNs, which reveals that some ISPs even face scarcity of internally used ("private") address space. We also find CGN placement is diverse, ranging from 1–12 hops from the user. We find that the methods CGNs use to distribute available public IP addresses and port numbers to their subscribers vary dramatically. We then assess how CGNs restrict user connectivity and compare our insights about CGNs to the properties of commonly deployed CPE (customer premises equipment) NATs.

Finally, we note that while our methods and data provide an unprecedented view into the use and properties of CGNs in the wild, we only partially illuminate the CGN landscape. Each of our measurement approaches has limitations that somewhat restrict their scope. For instance, since mobile devices rarely use BitTorrent, our DHT crawl does not shed significant light on the use of CGNs within mobile ISPs. Our study constitutes an initial view into the deployment of CGNs with much future work to be done to better understand the impact of these critical components of the modern Internet.

4.1 An Operator's Perspective

To gain a better understanding of the real-world challenges that IPv4 address scarcity poses and how ISPs are coping, in late 2015 we circulated a survey on a dozen of network operator mailing lists and

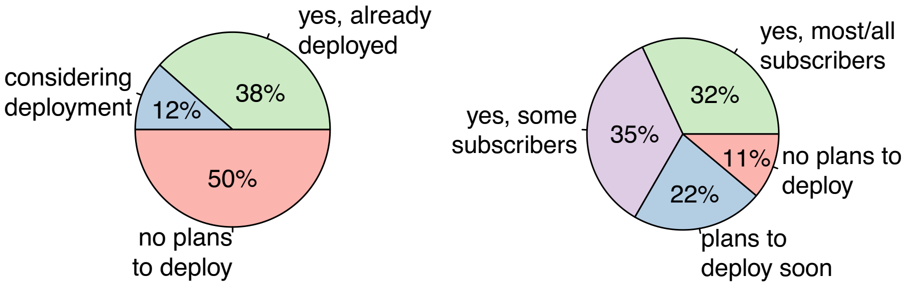

(a) IPv4 Carrier-Grade NAT.

(b) IPv6.

Figure 4.1: ISP survey: Status of Carrier-Grade NAT deployment and IPv6 deployment.

eventually collected responses from 75 ISPs located all over the world. These ISPs run the gamut in terms of size and type, including cellular and residential ones. While we do not claim the respondents form a statistically unbiased sample, we note that we received answers from operators in all regions of the world, spanning the whole spectrum of ISPs (cellular, residential) ranging from small rural ISPs in Africa up to Fortune 50 companies, connecting millions of subscribers to the Internet. Thus, we do believe that the approaches and concerns raised by these ISPs deserve our attention. Next, we summarize the survey responses.

IPv4 Address Space Scarcity: More than 40% of the responding ISPs indicate that they directly face IPv4 address scarcity issues. Some ISPs report a subscriber-to-IPv4 address ratio as high as 20:1. However, others point out that while their subscriber-to-address ratio is 1:1, internal subnetting and fragmentation make address space management cumbersome, especially when attempting to accommodate new customers. Another 10% of the respondents indicate that while they do not yet face scarcity, they believe it is looming in the near future. The ISPs not facing IPv4 address scarcity are mainly ones that received significant blocks of address space many years ago, as well as ISPs in the African region.[1] Interestingly, three ISPs also indicated that they face scarcity of *internal* IPv4 address space. These networks leverage CGN but also need internal address space for their internal management.

IPv4 Address Space Markets: Three of the responding ISPs report that they have bought IPv4 addresses, while another 15 ISPs indicate that they have considered procuring additional addresses. However, ISPs indicate concern regarding buying address space, including price of available address blocks (named by 60%), fear of obtaining "polluted" address blocks with a bad reputation from previous use (44%) and uncertainty regarding the ownership of blocks (42%).

CGN Deployment vs. IPv6 Deployment: Figure 4.1 shows the respondents' approach to CGN and IPv6. Almost 40% of the ISPs indicate they deploy IPv4 CGNs, with another 12% considering CGN in the near-term. Typically, ISPs note incremental CGN deployments, either targeting new customers or shifting specific subsets of subscribers into CGN deployment. That is, *most CGN deployments are partial*. Next we find that 32% of the ISPs indicate IPv6 deployment to most or all of their subscribers, while another 35% have partial IPv6 deployments for some subscribers. The dominant transition mechanism noted is dual stack. Some ISPs also provide customers with an internal (CGN) IPv4 address and a publicly reachable IPv6 address. This arrangement will likely gain popularity in the near future as IPv4 connectivity will remain necessary until full IPv6 deployment.

CGN Concerns: Participating ISPs also had the option to inform us about possible concerns when operating CGNs. The responding ISPs raised several concerns regarding the setup and the operation

[1] Africa is the only region in which the IPv4 address pool is not yet depleted.

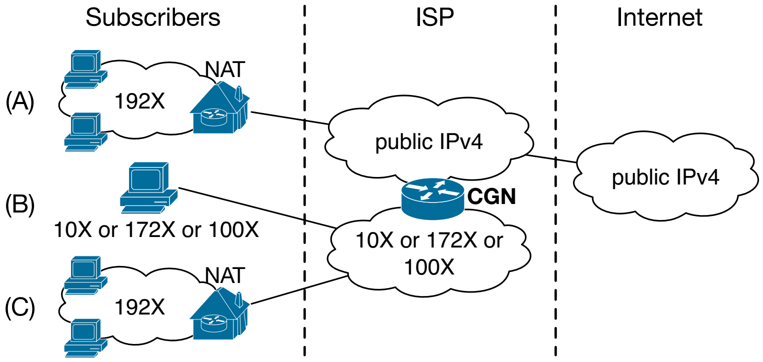

Figure 4.2: NAT scenarios. A resides behind a single in-home NAT, B behind a single carrier-grade NAT, and C behind both an in-home and carrier-grade NAT (NAT444).

of CGNs. A primary concern is that some applications (e.g., online gaming) do not work seamlessly with their CGN setups, causing subscriber complaints that remain difficult for the ISPs to resolve at the best of times. Additional concerns relate to traceability of users behind CGNs. Losing the ability to directly identify users can raise two kinds of problems. First, ISPs may be legally required to be able to map flows to subscribers. Second, IP addresses accrue reputations as they get used—e.g., as sources of spam—and therefore by sharing IP addresses among users the reputation is also shared and can cause problems for some users.

In addition, operators voiced concerns about a lack of well-developed best practices for configuring and dimensioning CGNs, rendering operating these devices cumbersome. In particular, operators need to resort to experimentation on aspects such as the distribution of external IP addresses and port ranges to customers, and whether to use distributed or centralized CGN infrastructure. Respondents named the port space as well as the amount of state CGNs need to maintain as primary challenges when configuring CGNs. Accordingly, ISPs report widely varying dimensioning of their CGNs in practice, ranging from static 1:1 NAT per customer—to prepare for the future—to limits of 512 sessions per customer due to heavy NATing.

4.2 Background

As we sketch above, the lack of ready access to new IPv4 addresses is leading ISPs to alternate technologies to accommodate their addressing needs. One such approach is to leverage Carrier-Grade NATs (CGN). When an ISP uses CGN, it provides subscribers with internal IP addresses and then applies address translation to their traffic. CGNs often introduce multiple layers of address translation since subscribers often run NAT devices on their own edge networks (e.g., as built into most CPE devices in users' homes). We refer to the case of subscribers whose packets are translated once before they reach the public Internet as *NAT44* and to the case where packets are translated twice as *NAT444* [261, 276].

Figure 4.2 illustrates various addressing structures in common use on the Internet. In each of the scenarios the ISP has a pool of public IPv4 addresses that are used differently by various subsets of its customers. The ISP gives each subscriber in group *A* a single public IP address. The subscriber in turn runs a subscriber-side NAT44 device to share this IP address among all the devices on the internal network. This is typical for many residential subscribers. Subscriber *B* receives an internal IP address from the ISP which a NAT translates into a public IP before packets reach the wide-area network. This case of a carrier-side NAT44 device is common within cellular networks. Finally, subscriber *C*'s network is

Range	Shorthand	RFC	Comments
192.168.0.0/16	192X	1918	Commonly used in CPE
172.16.0.0/12	172X	1918	
10.0.0.0/8	10X	1918	
100.64.0.0/10	100X	6598	for CGN deployments

Table 4.1: Address space reserved for internal use.

identical to subscriber A's in that a local NAT is used to facilitate connectivity for a multitude of internal devices. However, in this case, instead of providing a single external IP address the ISP provides the subscriber with a single internal IP address, which in turn it translates with a CGN before traffic reaches the wide-area network. This is a case of NAT444, or two layers of address translation. An ISP that runs a CGN does not necessarily NAT all of its subscribers. Many ISPs only NAT new subscribers and some even have various classes of subscribers and allow customers to choose their type of connectivity, which may come at different prices (some ISPs charge their customers for a public IP address, e.g., [206]).

On the basis of the terminology used in the IETF, we now define several NAT-related terms we use throughout the remainder of the paper.

Address Types: We distinguish IP addresses both in terms of their location relative to a NAT, as well as in terms of their numeric value. We refer to an address on the edge-facing, client-local side of a NAT as *internal* vs. *external* when nearer to the network core. An address is *reserved* if it resides in prefixes (as set forth e.g., by RFC 1918 [230]) that should not get announced to the global routing table, and *routable* otherwise. Table 4.1 lists those address ranges reserved for internal use by the IETF. [2]

NAT Mappings: NATs keep state that maps each internal IP address and port number tuple to an external IP address and port number tuple. Unless manually configured, NATs create mappings on-demand once a local host behind the NAT (i.e., with an internal IP address) sends a packet from its IP_{int}:$port_{int}$ endpoint to a remote IP_{dst}:$port_{dst}$. The NAT then records an IP_{ext}:$port_{ext}$ tuple, translates the packet and sends it to the destination host. When the external host replies to IP_{ext}:$port_{ext}$, the NAT finds the corresponding entry in its mapping table, translates the destination address to IP_{int}:$port_{int}$ and forwards the packet internally.

Mapping Types: NAT behavior differs in the reuse of existing mappings and in the filtering rules for the usage of established mappings. A *symmetric NAT* creates different mappings for subsequent packets sent from the same IP_{int}:$port_{int}$ endpoint to different IP_{dst}:$port_{dst}$ endpoints. This behavior significantly impedes NAT traversal and makes symmetric NATs the most restrictive type of NAT. Other types of NAT reuse existing mappings regardless of their IP_{dst}:$port_{dst}$. They differ in their filtering policy, here listed in decreasing order of restrictiveness: *port-address restricted NATs* only allow incoming packets from the very IP_{dst}:$port_{dst}$ that was initially contacted from the host inside the NAT, *address restricted NATs* require a matching IP_{dst}, but allow packets from varying port numbers, while *full cone NATs* allow incoming packets from any external host once a mapping is created. This makes full cone NATs the most permissive type of NAT [252].[3]

Mapping Timeouts: As with any stateful middlebox, NATs must manage their internal state and therefore track active flows. The NAT must release mappings that are no longer needed. NATs generally use both TCP state tracking and timeouts to prune unnecessary NAT mappings. Recommended minimum timeouts are 120 seconds for UDP [62] and 2 hours for TCP [132].

[2] Technically some reserved addresses are in fact routable; we focus here on their intended use.

[3] This terminology allows arranging NATs according to their restrictiveness and improves readability, therefore we use it despite being discouraged by the IETF [62].

Port Allocation: NATs differ in their selection of an external *port*$_{ext}$ number for a new session. NATs implementing *port preservation* attempt to retain the original source port as the external port (i.e., *port*$_{int}$ =*port*$_{ext}$), unless there is a collision and an alternate port must be chosen. Other NATs—especially large NATs—assign ranges of the external port space to each internal host and then assign external ports on-demand from this pool in sequential or random order [62].

IP Pooling: Large NATs typically use multiple external IP addresses, called *NAT pooling*. Upon connecting, a subscriber typically gets allocated a public IP address out of the pool. NATs employing *paired pooling* always use the same *IP*$_{ext}$ for a given *IP*$_{int}$. Otherwise, a NAT is said to use *arbitrary pooling*. In our methodology, the presence of NAT pools will play an important role when it comes to dissecting home NAT deployments from CGN deployments.

Hairpinning: Consider the communication between two hosts—A and B—behind the same NAT. When A sends a packet to B it will use B's *IP*$_{ext}$:*port*$_{ext}$. When the NAT receives this packet it can detect that the destination of the packet is in fact itself and therefore direct the packet to B's *IP*$_{int}$:*port*$_{int}$. This behavior is called *hairpinning* [62, 132]. If the NAT leaves the source *IP*$_{int}$:*port*$_{int}$ in place when forwarding the packet, then the hosts can discover their internal IP addresses when communicating behind the same NAT.

Related Work

IETF RFCs contain most of the available literature about CGNs. In particular, RFC 6888 specifies basic requirements for CGNs [215], whereas RFC 6544 [250] and RFC 5128 [264] describe two popular mechanisms for NAT traversal: ICE and UDP/TCP hole punching, respectively. As a result of NAT's added complexity, RFCs also describe how CGNs affect application-level functionality [109, 264].

Several academic studies have tried to identify NAT deployment in home networks using UPnP queries [106, 186] or IP ID header fields [68], by passively observing IP TTLs and HTTP User-Agent strings [190], and by applying NAT detection to NetFlow traces [171]. Müller et al. conducted an active, topology-based traversal of cascaded large-scale NATs [198]. One NAT test presented in our work is an augmented version of their methodology. Ford et al. studied the effectivity of different NAT punching techniques in NAT-ed networks [119]. The studies conducted by Wang et al. [274] performed a comprehensive active measurement campaign to understand middleboxes present in cellular networks. In contrast to Netalyzr their tool relied on rooted handsets to modify packets at the IP and TCP layers. Donley et al. [71] studied the impact of CGN deployment on Web browsing performance in one ISP. Ohara et al. [211] simulated how CGNs can impact on TCP connection establishment in mobile networks. Finally, Skoberne et al. presented a theoretical taxonomy of NAT deployments and discuss their possible impact on network performance [261]. Richter et al. [235] measured an increasing concentration of traffic on fewer IPv4 addresses during 2015, hinting at an increasing use of CGN deployment in the Internet.

Little is known about actual CGN hardware deployed in the wild and their consequences for the different stakeholders. We cannot readily identify how NAT vendors implement their equipment and how ISPs take advantage of them. To partly overcome this limitation, we rely on vendor manuals and network operator tutorials to obtain deeper insights into practical considerations of CGN deployment [18, 19, 89, 276].

	Peers	Unique IPs	ASes
Queried	21.5M	15.5M	18.8K
Learned	192.0M	62.1M	26.7K

Table 4.2: BitTorrent DHT data: *Queried*: Peers that were issued and replied to *find_nodes* requests. *Learned*: All peer information we gathered.

4.3 Detecting CGN at Scale

Our first set of methodologies aim to investigate the breadth of CGN deployment in the Internet. In general terms our CGN detection mechanisms leverage both internal and external observations of IP addresses associated with a given host to detect discrepancies and therefore presence of address translation. We use two techniques to obtain internal observations: implicit and explicit. Our implicit observations come via standard BitTorrent clients leaking internal address information, while our explicit observations come from users running active measurements on our behalf the Netalyzr tool. We stress that we strive for conservativeness in our CGN detection methods. That is, we would rather provide a sound lower bound on CGN presence than using a more speculative approach that identifies more CGNs of questionable validity.

4.3.1 Detecting CGNs via BitTorrent

The BitTorrent Distributed Hash Table (DHT) [17] represents a distributed data structure that links hosts looking for specific content with hosts that have that content without using centralized infrastructure. The nodes that make up the DHT form a connected graph so that search queries for specific content are propagated to a node with the given information. Each node is identified by a 160 bit *nodeid* which is randomly chosen by the node itself (and is unique with high probability). To form the graph, DHT nodes both maintain a list of DHT peers and provide an interface for other nodes to query this list. Further, the nodes on this list must be periodically validated with *bt_ping* messages to ensure reachability. This in turn means that the contact information a node A has for node B—in the form of an IP address and port number—represents B's location from A's perspective. We observe that the nodes represent vantage points that we do not control but can none-the-less probe to learn about host-to-host connectivity. We find that this connectivity is sometimes represented by internal IP addresses. That is, the path between two hosts does not traverse the publicly routed Internet, but takes place completely within a private network (e.g., within an ISP). Additionally, these hosts are clearly also able to communicate outside of this private network and therefore are behind some form of NAT. We developed a crawler to collect connectivity information from the BitTorrent DHT and then leverage that data to form an understanding of CGN deployments.

Crawling the BitTorrent DHT: We developed a crawler that starts with DHT nodes learned from the BitTorrent bootstrap servers and issues a series of *find_nodes* requests to DHT nodes with a random query. The response to *find_nodes* is a list of up to eight "close" peers where closeness is calculated using the XOR distance between the query and each *nodeid* in the node's list of peers [191]. We issue five queries, which provides connectivity information—*nodeid*, IP address and port number—for roughly 40 nodes. We then in turn query the newly learned peers in the same fashion. Our crawler also participates in the DHT and therefore accepts incoming requests from nodes that have learned about our crawler through the source information in our requests.

	Internal Peers		Leaking Peers		
Range	**Total**	**Unique IPs**	**Total**	**Unique IPs**	**ASes**
192X	565.9K	11.2K	186.8K	162.2K	4.1K
172X	336.6K	85.0K	52.9K	33.9K	1.0K
10X	1.3M	328.5K	283.9K	194.4K	2.2K
100X	1.5M	251.5K	192.0K	165.8K	723

Table 4.3: Peers reported via reserved IP addresses (left) and the corresponding peers that leaked them (right).

As we note above, in some instances peers reply to *find_nodes* with information about nodes that include reserved IP addresses (Table 4.1), indicating the probed peer can reach the reported peer without crossing the publicly routed Internet. We refer to this behavior as *internal address leakage*. When we learn an internal address for a given *nodeid* we refer to this node as an *internal peer*. When our crawler finds a node leaking internal peers we issue an additional ten *find_nodes* queries in the hopes of finding additional internal peers. We continue issuing *find_nodes* queries in batches of ten for as long as we continue to harvest internal peers.

Note that within BitTorrent the *nodeid* is the sole identity notion for a given peer. However, as peers can have multiple endpoints (internal, external), as well as multiple IP addresses/ports due to dynamic IP address assignment or due to BitTorrent clients modifying the local port number, we identify a unique peer by the full tuple of *(IPaddress:port, nodeid)*. As a positive side effect, this also eliminates possible biases due to DHT poisoning [273], where peers announce themselves with a foreign *nodeid*.

BitTorrent Dataset: The dataset we use in the remainder of this paper comes from a one-week crawl starting on March 3, 2016.[4] Table 4.2 summarizes the dataset. We probed more than 21M peers across nearly 19K ASes. These probed DHT nodes in turn revealed contact information for 192M peers across more than 26K ASes. Of these 192M peers, 107.7M peers and 36.7M unique IP addresses responded to *bt_ping* probes. Table 4.3 shows an overview of the leaked contact information, where we break private peers down based to the internal address space range. Among the peers crawled, we find more than 700K peers leaking contact information for more than 3.7M internal peers (i.e., peers with IP addresses in a reserved range) across more than 5K ASes.

We observe that both the number of BitTorrent speakers as well as the extent of leakage, is highly uneven when comparing ASes, and even within the address space of specific ASes. To illustrate, Figure 4.3 compares the address space advertised in the global routing table for 3 selected ASes (x-axis) against the number of unique IP addresses with active BitTorrent peers per address block (y-axis). We normalize the address space advertised per AS by deaggregating all their prefixes to /24 granularity and arranging them in numerically increasing order. Hence, the bin width does not reflect a fixed portion of address space, but a fraction of the total address spaced announced by the respective AS. The figure shows the number of unique BitTorrent speaker IP addresses (orange bars), as well as the number of unique internal IP addresses leaked by BitTorrent peers in the respective region. Comcast has an average of only 5 active peers per /24 and shows very little leakage of internal peers. AS8402, on the other hand, has a larger BitTorrent population and we observe hotspots of internal peer leakage, hinting towards partial CGN deployment. Extreme cases, such as AS45815, show very high leakage from peers throughout their address space.

Identifying CGNs: Our dataset clearly shows the presence of NATs via leaked internal peers. Next we seek to establish the degree to which these NATs are network-level CGNs as opposed to simple

[4] We have additional crawls from late 2015 and early 2016 that show consistent results to those we present in this paper.

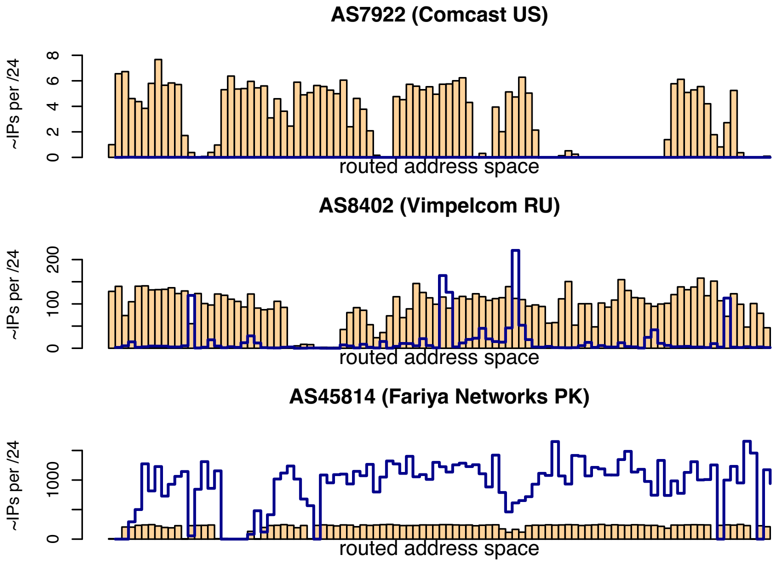

Figure 4.3: Active BitTorrent peer IPs (orange bars) vs. internal IP addresses leaked (blue lines), averaged per /24, for selected ASes.

home NAT deployments. First, to detect any type of NAT using the BitTorrent dataset there must be multiple BitTorrent clients that directly communicate within some internal network. Next, to determine the presence of a large network-level CGN we require *NAT pooling* behavior (Section 4.2). In other words, within a single AS we require (*i*) multiple peers with different external IP addresses to leak internal peers and (*ii*) intersections in the internal peers leaked across multiple external IP addresses. Moreover, we require the internal peers within a cluster to reside within the same internal address range (e.g., 10/8).

To detect this behavior on a per-AS level we next form a graph for each AS where each peer is a vertex and each edge between a public peer A and an internal peer B indicates that A leaks contact information for B. Note, when constructing graphs we only consider internal peers which were leaked exclusively by peers residing in a single AS. This excludes leaking relationships caused by VPN tunnels. Figure 4.4 shows a small subset of the graphs for two ASes as an illustration. Figure 4.4(a) shows there is only isolated leaking within AS 7922 (Comcast). We find more than 1K peers leaking internal addresses within AS 7922. However, we also find that each leaked internal peer is leaked by exactly one external peer. In contrast, Figure 4.4(b) shows strong clusters within AS 12874 (FastWEB) consisting of multiple peers behind different external IP addresses that leak a large number of internal peers, which form large intersections.[5] This shows that our clustering methodology is effective in separating home NAT deployments from network-level CGNs.

We next construct a graph for each AS in our dataset and determine the largest connected cluster for each AS. Figure 4.5 shows our clustering results. Here, we plot a point for each AS and position it according to the size of the largest cluster we found in this AS (if any). In particular, the x-axis shows the number of unique public IP addresses contained in the largest cluster and the y-axis shows the number of unique internal IP addresses contained in the largest cluster. We only find a small number of ASes that contain large clusters in the 192X space (top left figure). We find ASes with large clusters to be more prevalent in the other, larger, internal address ranges. This supports our hypothesis that 192X address space is primarily used in small home NAT environments. While in principle a cluster with at least two different

[5] We manually confirmed CGN presence in AS 12874 and also verified the discovery of internal peers (via *NAT Hairpinning*, Section 4.2) and leakage of internal peers in this AS by running a regular BitTorrent client on a host behind CGN in this AS.

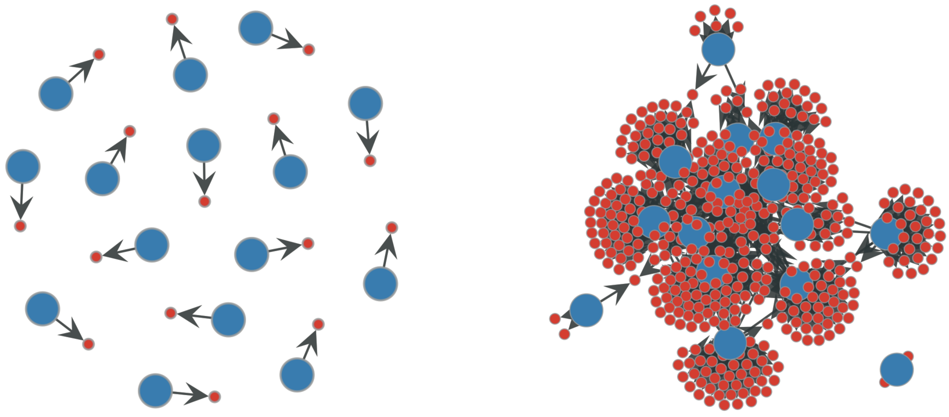

(a) Isolated leaking relationships (AS7922, Comcast, 192X internal space).

(b) Clustered leaking relationships (AS12874, FastWEB, 100X internal space).

Figure 4.4: Peer leakage in non-CGN vs. CGN ASes. Large blue vertices are BitTorrent peers leaking peers with internal IP addresses (small red vertices).

external IP addresses is indicative of *NAT pooling*, we only determine CGN presence for an AS when the largest connected cluster contains at least five public IP addresses and five private IP addresses. This is to address possible misclassifications arising from dynamic addressing, e.g., a home network with internal NAT deployment that changes its public IP address. We annotated Figure 4.5 with our detection boundary. While we show network-wide results in Section 4.4, we note that this methodology shows CGN usage in roughly 10% of the probed ASes for which our crawler queries at least 200 peers.

DHT Data Calibration: Our BitTorrent-based CGN detection relies on three key properties of the DHT peers: *(i)* BitTorrent peers behind the same NAT can learn internal endpoints of other peers, *(ii)* peers export internal endpoints via the DHT, and *(iii)* peers only propagate contact information for peers that have been validated via direct interaction. We verified *(i)* and *(ii)* by running two popular and unmodified BitTorrent clients (uTorrent on Windows and Transmission on Linux) and measuring the control traffic they exchange as part of their regular operation. We confirmed that these peers learned their internal endpoints when located behind NATs that allow multicast communication as well as behind NATs that have *Hairpinning* enabled. Further, these peers forward packets with internal source IP addresses (see Section 4.2). We also validated the latter within an ISP that deploys CGN. Therefore, we conclude that BitTorrent clients can—if the circumstances allow it—learn their internal endpoints and propagate that information via the DHT when requested.

Finally, we assume hosts follow the BitTorrent DHT specification [17] and only propagate reachability information for peers they learn after reachability has been directly validated by the host itself. Otherwise, hosts would propagate potentially dubious reachability information and likely we would detect CGN presence in practically any AS that hosts enough peers. To validate our assumption, we setup a common BitTorrent client (uTorrent on Windows) with a *nodeid* of ID_{us} and let it interact with the DHT. At the same time on a different host we crawled the DHT requesting ID_{us}. We queried 100K peers and were given contact information for ID_{us} by $1,387$ peers. We found that only 18 of these DHT peers (1.3%) did *not* validate the reachability of ID_{us} before propagating the information. This shows that our assumption that DHT peers follow the specification and properly validate reachability before propagating contact information is sound.

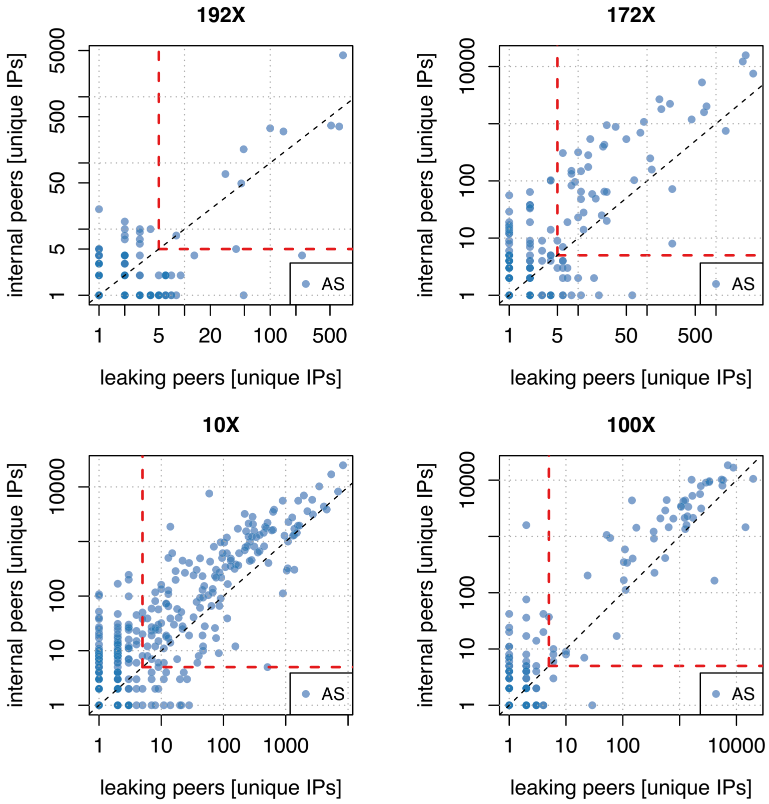

Figure 4.5: Size of the largest connected cluster of leaking and internal BitTorrent peers per AS. The x-axis shows the number of public IP addresses, the y-axis the number of internal IP addresses contained in the largest cluster.

4.3.2 Detecting CGNs via Netalyzr

To complement our observations from crawling the BitTorrent DHT, we leverage ICSI's Netalyzr network troubleshooting service [170]. While the BitTorrent DHT provides a useful set of specific information from end hosts, Netalyzr allows us to define explicit tests we wish to run from end hosts. These tests interact with a suite of custom-built test and measurement servers and return the results to our data collection server. We developed a set of tests aimed at illuminating NAT behavior and deployed these in 2014. While Netalyzr provides the potential to gather much richer information than we find in the BitTorrent DHT, we are at the mercy of individual users to access Netalyzr and run the tests. Users run Netalyzr via one of three supported clients: a Java applet for Web browsers, a command-line client, or an Android client available in the Google Play store [9].

In the context of understanding CGNs, Netalyzr offers two advantages over our BitTorrent crawl. First, since BitTorrent is not heavily used on mobile devices the Android version of Netalyzr extends our view into this important network type.[6] Second, Netalyzr allows us to directly obtain the IP addresses used by the host, including (i) the local IP address of the device that executes Netalyzr, IP_{dev}, (ii) the external IP address of the CPE router device as learned via UPnP (where available), IP$_{cpe}$, and (iii) the public IP address as seen by the Netalyzr server, IP$_{pub}$. We categorize IP$_{dev}$ and IP$_{cpe}$ in four categories: (i) private address from one of the reserved blocks for this purpose, (ii) *unrouted* for addresses that are nominally public, but do not appear in the routing table, (iii) *routed match* for case where the address is routable, appears in the routing table and matches IP$_{pub}$ (i.e., the non-NAT case) and (iv) *routed*

[6] Note, while mobile devices can join wifi networks we scope our measurements to those on cellular data networks.

	Cellular	**Non-cellular**	
	IP$_{dev}$	**IP**$_{dev}$	**IP**$_{cpe}$
Address Space	N=8.6K	N=567.5K	N=229.8K
192X	0.2%	92.4%	8.9%
172X	2.5%	1.1%	0.8%
10X	58.7%	6.2%	4.8%
100X	17.3%	0.0%	1.9%
unrouted	12.5%	0.0%	0.0%
routed match	5.7%	0.0%	83.0%
routed mismatch	3.0%	0.3%	0.5%

Table 4.4: Address ranges seen for the device IP address and for the router's external IP address.

mismatch for the case where the address is routable, appears in the routing table but does not match IP$_{pub}$.

Cellular Networks: Detecting the use of CGNs in cellular networks is straightforward in Netalyzr because there are no devices between the mobile device and the ISP and therefore the classification of the ISP-assigned IP$_{dev}$ directly indicates the presence of a CGN.[7] While straightforward, we require five observations from an AS before we include it in our study to ensure our conclusions are sound and are not the result of some unexpected behavior. The second column of Table 4.4 shows the breakdown of IP$_{dev}$ for all our cellular Netalyzr sessions. We find 94% of the sessions—i.e., all cases except the *routed match* case—show a translated address. A first view of CGN deployment in cellular networks on a per-AS basis shows 63.8% exclusively assign internal IP addresses to mobile devices. Similarly, we find 6.0% of ASes exclusively assign public IP addresses to devices and show no signs of address translation. Meanwhile, 30.3% assign a mixture of internal IP address and public IP addresses to devices. We note that another 5.0% of ASes assign internal addresses from public blocks that are actually routed, but still perform address translation!

Non-Cellular Networks: We next shift to Netalyzr's detection of CGNs in non-cellular networks. In these networks, IP$_{dev}$ is often assigned by a device in the device's network and therefore when CGNs are present there are multiple address translations happening in the path (NAT444). This in turn makes detection difficult. First, we winnow our analysis to ASes that have at least ten Netalyzr sessions in our dataset.[8] The third column of Table 4.4 shows that IP$_{dev}$ is nearly always a private address, as expected. In addition to IP$_{dev}$, Netalyzr uses UPnP [73] to attempt to determine IP$_{cpe}$ for the first hop CPE device. The fourth column of Table 4.4 shows the breakdown of IP$_{cpe}$ for the 40% of cases where UPnP provides the address. In 83% of the cases, IP$_{cpe}$ is a public IP address from the ISP, hence no CGN is present. The remaining 17% of the cases clearly point to multiple NATs. However, whether these are ISP-based CGNs or multiple small-scale NATs in the edge networks is not clear. Therefore, we add two steps to disambiguate the situation.

First, we observe that CPE routers often make assignments from the 192X block (Table 4.4, column 3), whereas the CGNs we find via BitTorrent and in the cellular environment more often make assignments outside the 192X block (Figure 4.5 and Table 4.4, column 2). Therefore, we use Netalyzr's list of IP$_{dev}$ assignments to determine the top ten /24 blocks from which CPE devices make assignments (covering 95% of assignments). We then conclude that any IP$_{cpe}$ that falls within one of these blocks was likely

[7] Exceptions could be caused by users manipulating their network access with VPN tunnels or by users who run their own cellular access point (e.g., femtocells). Netalyzr's Android client collects enough data to allow us to prune such cases from our analysis.

[8] Note, we require more observations in the non-cellular case (ten) compared to the cellular case (five) because the situation is not as straightforward due to the presence of in-path network equipment in the edge network. This makes the breadth of behavior we observe larger and in turn we need more observations to draw sound conclusions.

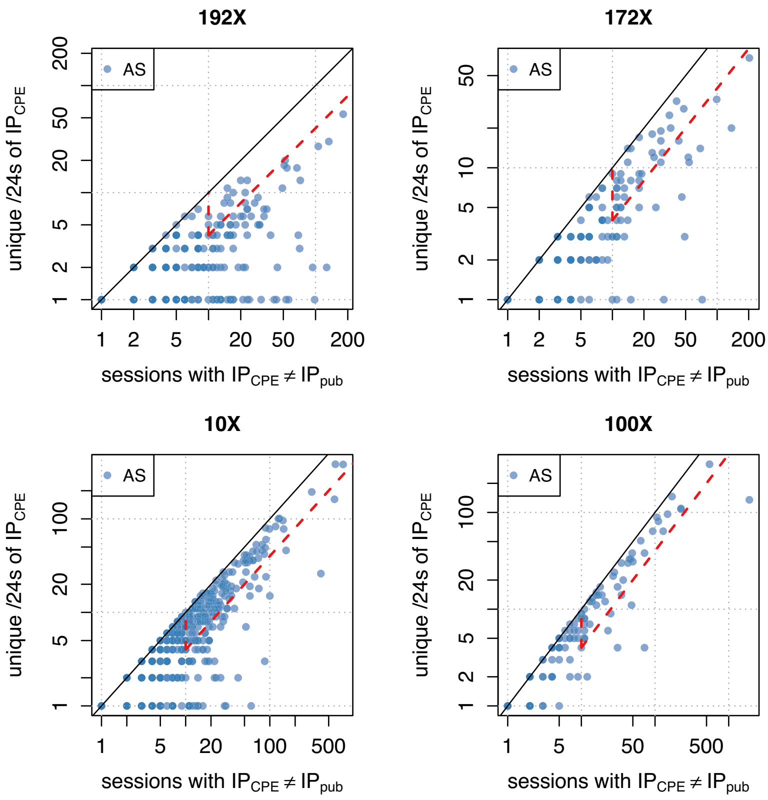

Figure 4.6: Netalyzr CGN candidate ASes: Sessions where IP$_{cpe}$ does not match IP$_{pub}$ (x-axis) vs. unique /24s of IP$_{cpe}$ addresses (y-axis).

assigned by another local CPE device and not a CGN. Applying this filter removes more than half the ambiguous situations and leaves us with 7.9% of Netalyzr's sessions that may be CGNs.

As a second step, we observe that due to their scale, CGNs necessarily must more broadly assign addresses than would be necessary in a small-scale edge network. Therefore, to conclude a CGN is present we require IP$_{cpe}$ diversity within an AS. Specifically, an AS must have $N \geq 10$ Netalyzr sessions that may be behind a CGN. We expect that as the number of Netalyzr sessions increases, our observations of address space diversity will, as well. Therefore, those sessions must span at least $0.4 \times N$ internal /24 address blocks are deemed to indicate a CGN is deployed.[9] Figure 4.6 shows a point for each AS in our dataset, with the x-axis showing the number of ambiguous multiple NATing situations we observe and the y-axis showing the number of /24 address blocks we observe within the AS. The dashed line represents our CGN detection cutoff point. Similar to our observations in our BitTorrent dataset, the 192X address space is sparsely used for CGNs, while more CGNs are present in the other reserved address blocks. Overall, our method detects CGN presence in almost 15% of the covered ASes.

Our CGN detection is no doubt imperfect. However, we note that our heuristics start with cases where our data *conclusively* indicates multiple address translators are present. Further, manual validation against our survey results, ISP' websites and threads on operator mailing lists lends confidence to our conclusions. Finally, as we note above, our methods for labeling CGNs are conservative. For instance, there are points to the right of dashed line in Figure 4.6 that likely represent undetected CGNs. These

[9] We note that we do not expect address diversity to infinitely scale with the number of observations. However, given our data this heuristic works well. Furthermore, adding additional complexity to the methodology without grounding in empirical observation is not useful.

	routed ASes (N=52K)		**eyeball ASes, PBL (N=2.9K)**		**eyeball ASes, APNIC (N=3.1K)**	
	covered	CGN-positive	covered	CGN-positive	covered	CGN-positive
BitTorrent	2,724 (5.2%)	254 (9.40%)	1,673 (57.7%)	180 (10.8%)	1,824 (59.6%)	204 (11.2%)
Netalyzr non-cellular	1,367 (2.6%)	195 (14.3%)	866 (29.8%)	151 (17.4%)	929 (30.4%)	174 (18.7%)
BitTorrent ∪ Netalyzr	3,166 (6.0%)	421 (13.3%)	1,791 (61.7%)	306 (17.1%)	1,946 (63.6%)	350 (18.0%)
Netalyzr cellular	218 (0.4%)	205 (94.0%)	175 (6.0%)	162 (92.6%)	171 (5.6%)	161 (94.2%)

Table 4.5: Coverage and detection rates of our methods as fraction of all routed ASes as well as Eyeball ASes, primarily connecting end users, as derived from PBL and APNIC.

points represent many Netalyzr sessions that show much address diversity—but not enough to meet our threshold. Our validations and conservative cutoffs leave us confident in the determinations we make, at the likely expense of not identifying all CGN deployments.

4.4 A Network-wide View of CGN Deployment

We now summarize our measurements of global CGN deployment based on the methodologies we develop in Section 4.3. Table 4.5 reports our results in terms of ASes where we detect at least partial CGN deployment. We report our results within the context of three different AS populations in the three big columns in the table. The second big column of the table considers the entire population of roughly 52K routed ASes. Meanwhile, the third and fourth columns represent the results in the context of so-called "eyeball" ASes (ASes that connect end users to the Internet). The third big column considers the population of ASes that the Spamhaus Policy Block List [263] identifies as including the equivalent of at least 2,048 addresses in "end user" blocks. The last big column considers the population of ASes to be those in the APNIC Labs AS Population list [45] that have at least 1,000 samples. Our datasets cover 6.0% of the ASes in the Internet, but over 60% of the eyeball ASes. Given that our methodologies rely on user-driven tools (Netalyzr and BitTorrent clients) it is unsurprising that we cover an order of magnitude more eyeball ASes.

In terms of CGN deployments, we find that 13.3% of all non-cellular ASes use CGNs. However, the penetration jumps to 17–18% when considering only non-cellular eyeball ASes. In cellular networks the use is over 92% in all cases. These results show that CGNs are a reality for many Internet users. We also note that while we are able to cover roughly twice as many ASes with our BitTorrent dataset, the Netalyzr measurements find CGNs in higher proportions. This is expected and underscores important aspects of each methodology. While we are able to opportunistically leverage the information from the BitTorrent DHT, we are unable to direct or control the measurements. So, while BitTorrent has a large footprint the data is noisy. On the other hand, Netalyzr must coax people to explicitly run the tool and therefore the population is not large. However, once run we directly control the measurements and can gather more data directly (e.g., via probing UPnP). Finally, we note that the table shows that Netalyzr often does not add significantly to the coverage, but does add significantly to the CGN deployment results. Therefore, the BitTorrent detection should be viewed as a lower bound on CGN penetration.

Finally, we return to the impetus of NAT in the first place: address scarcity. Figure 4.7 shows our results partitioned by RIR regions. The left-hand plot shows that the percentage of covered eyeball ASes within each region does not show a significant regional bias.[10] The middle plot in the figure shows the percentage of the eyeball ASes we find to deploy CGNs. Here we observe that APNIC and RIPE show more than twice the CGN penetration of the other regions. These are also the two regions that ran

[10] We use the PBL eyeball AS list for this plot.

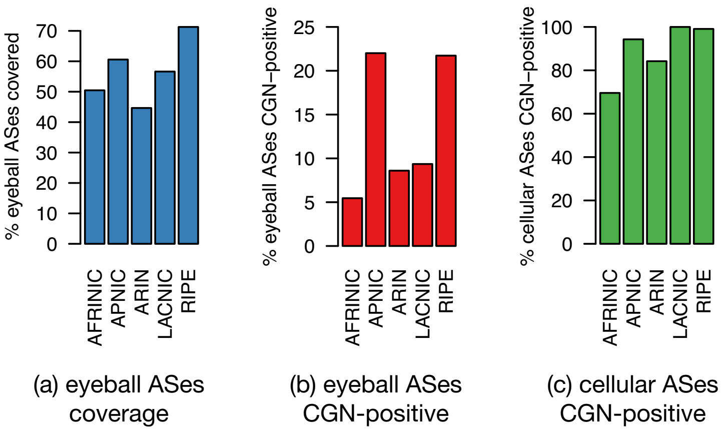

(a) eyeball ASes coverage

(b) eyeball ASes CGN-positive

(c) cellular ASes CGN-positive

Figure 4.7: Eyeball AS coverage and CGNs per region.

out of IPv4 addresses first. Meanwhile, we find the lowest CGN penetration in AFRINIC, which is the only region that has not yet exhausted its supply of IPv4 addresses. The last plot in the figure shows the CGN penetration in cellular networks by region. AFRINIC is again an outlier in this plot with “only” two-thirds of the ASes leveraging CGNs.

4.5 Drilling into CGN Properties

Having a broad perspective of CGN deployment in today’s Internet in hand, we next drill into the properties of the detected CGNs. NATs can be configured in a multitude of ways and as our survey results indicate, configuring a NAT at carrier-scale presents a massive resource distribution problem, including (*i*) public IP addresses, (*ii*) private IP addresses and (*iii*) ephemeral port numbers. The CGN creates state in the form of NAT mappings with finite lifetimes (timeouts) to associate these resources depending on the NAT mapping type. A CGN’s configuration directly affects (*i*) the degree of resource sharing, i.e., how many subscribers can reside behind a given set of public IP addresses, as well as (*ii*) the number of simultaneous flows available to individual subscribers.

In this section, we study the configuration of our identified CGNs. In particular, we study (*i*) which address ranges ISPs use internally, (*ii*) how CGNs assign IP addresses and ports to their subscribers, (*iii*) topological properties of CGNs (i.e., the location of the NAT), and (*iv*) the kind of NAT mappings deployed CGNs commonly employ. Where appropriate, we contrast findings for CGNs with our findings for commonly deployed CPE devices.

4.5.1 Internal Address Space Usage

Our two probing methods enable us to evaluate properties of the address space behind detected CGNs. Figure 4.8(a) shows per AS the internal address space ranges used within non-cellular as well as cellular CGNs. Overall, we observe that naturally the largest private range (10X) is the most commonly used space for CGNs, followed by the 100X block newly allocated specifically for CGN deployments [275]. We also observe CGNs deploying the smaller 172X and 192X address spaces. Interestingly, roughly 20% of the ASes use multiple ranges of reserved address space in their CGN deployment. We speculate that the size of individual blocks does not suffice or, more likely, that such private address space is

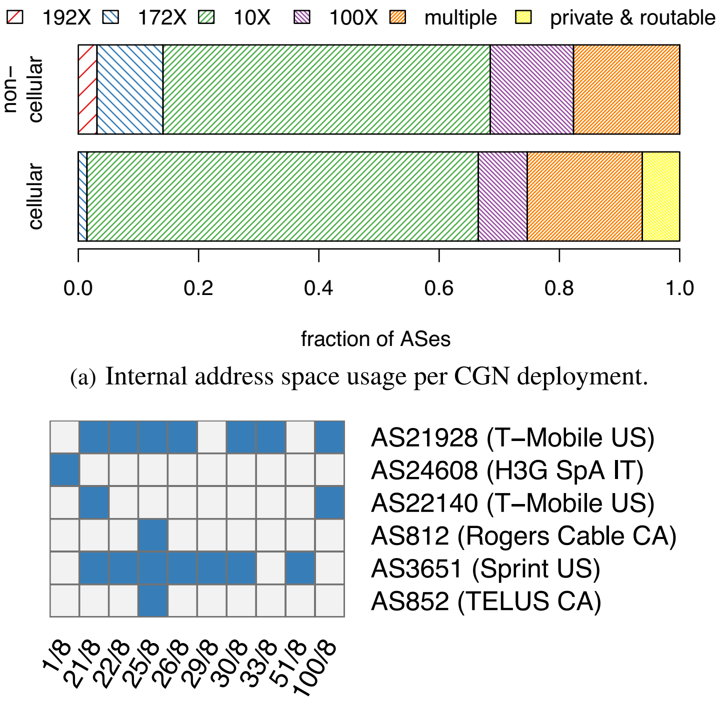

(a) Internal address space usage per CGN deployment.

	1/8	21/8	22/8	25/8	26/8	29/8	30/8	33/8	51/8	100/8
AS21928 (T–Mobile US)		■	■	■	■		■	■		■
AS24608 (H3G SpA IT)	■									
AS22140 (T–Mobile US)		■								■
AS812 (Rogers Cable CA)				■						
AS3651 (Sprint US)		■	■	■	■	■	■		■	
AS852 (TELUS CA)				■						

(b) ASes using routable space as internal space.

Figure 4.8: Internal address space in CGN deployments.

already in use in other parts of the organization. Some cellular ISPs in fact use routable address space for their internal CGN deployments. In Figure 4.8(b) we show which routable address blocks make up the most prominent cases we detected. While most of the routable address space used is not routed in practice (such as the 25.0.0.0/8 block, allocated to the UK Ministry of Defense), some ISPs use address space within their internal deployment (e.g., 1.0.0.0/8) that is publicly routed by other ASes. We contacted a representative of one of these major ISPs who confirmed that their internal deployment of routable address space results from scarcity of internal address space. Thus, some ISPs evidently *experience a shortage of internal address space* and adopt drastic measures at the expense of potential security and connectivity problems once public and internal addressing collides. Moreover, this address range proliferation renders troubleshooting CGNs even more cumbersome.

4.5.2 Port and IP Address Allocation

Next, we study how CGNs allocate ports and IP addresses to their subscribers. We start with the former. NAT port allocation may adopt the following strategies [62]: (i) *port preservation*, where the NAT attempts to maintain the local port of the flows; (ii) *sequential use*, where the NAT allocates ports in a sequential order; (iii) *random use*, where the NAT allocates ports without a clearly identifiable pattern.[11]

[11] We allow some leeway in determining port behavior. For example, we identify port preservation if at least 20% of ports remain preserved, and we declare *sequential use* if every two subsequent connections exhibit a numerical port difference smaller than 50. This accounts for situations in which NATs can not allocate the original or subsequent sequential port because of already existing mappings.

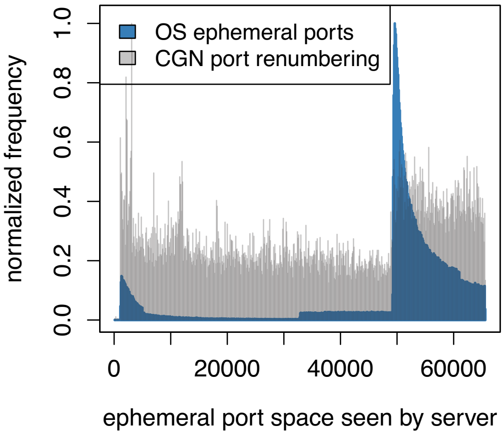

Figure 4.9: Ephemeral port space seen by our server from non-CGN vs. CGN connections without port preservation.

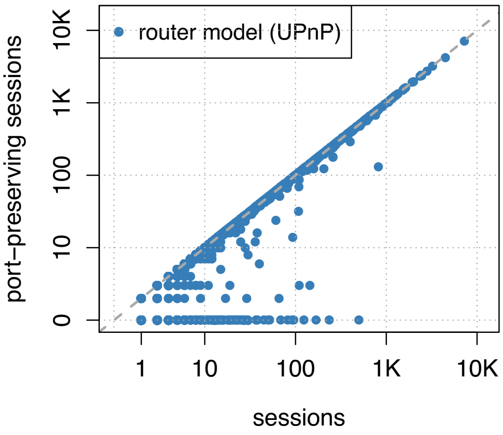

Figure 4.10: Port preservation behavior of CPE equipment. 92% of UPnP sessions are from devices behind port-preserving CPEs.

Measuring port translation: During one execution of Netalyzr (a "session"), its client opens 10 sequential TCP flows to an echo server listening on a high port number unlikely to be proxied. These TCP flows enable us to reason about the port allocation strategy implemented by the CGN, by comparing the local ephemeral port number, as chosen by the device, with the source port as seen by our server. Figure 4.9 shows the distribution of source port numbers as observed by our server. We show two histograms, one for port-preserved sessions and one for sessions exhibiting port translation. While operating systems employ ephemeral port ranges [95], CGNs translating port numbers utilize the entire port space. This observation could prove useful for server-side attempts (e.g., by content providers) to identify whether a client-side IP address belongs to a CGN.

Port translation of CPE routers: In non-cellular networks, where users' packets typically get NATed by a CPE, our measurements might be affected by CPEs employing port translation. To assess the impact of CPEs on port translation, Figure 4.10 shows for each CPE model (inferred using UPnP) the number of *non-CGN* sessions where our server saw the same ports as chosen by the device. We observe that in more than 92% of non-CGN sessions the CPE did *not* alter the source port numbers. Hence, while some CPE do translate ports, their effect on our analysis remains small.

Network-wide port allocation strategies: Figure 4.11 shows the distribution of port allocation strategies for each CGN-positive AS. We sort ASes with a "pure" allocation strategy (left part of the plot) and move ASes with mixed allocation strategies to the right. We observe a uniform port allocation strategy for about a third of the non-cellular ASes and for about 50% of the cellular ones. For the rest, *CGN behavior is heterogeneous.* We can attribute this to distributed CGN deployments, where users of the same ISP reside behind multiple CGNs, and the fact that NAT devices do not necessarily behave consistently, changing their behavior under load and over time [188]. Table 4.6 summarizes the dominating port allocation strategy per AS.

Chunk-based port allocation: In addition to the classification in our three categories of port allocation strategies, we also identify CGNs with random chunk-based allocation, where each subscriber receives a fixed port block [89]. Given sufficient data, we can infer the size of such port "chunks". Figure 4.12 shows an example of chunk-based port management: AS12978 allocates 4K ports per subscriber. For each recorded session in this AS, our server observes source ports translated in no particular order (neither preserved nor sequential) and that all ports of a given session fall within a well-defined range.

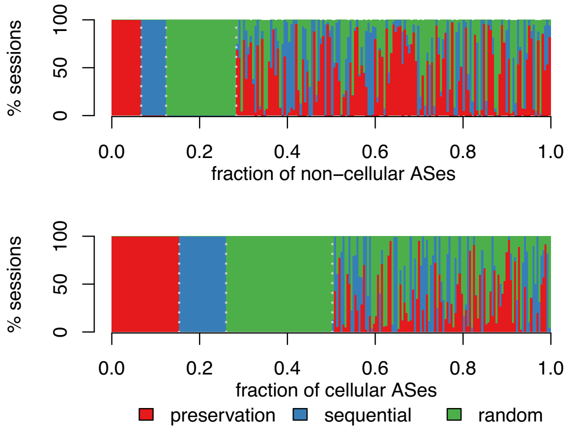

Figure 4.11: Distribution of observed port allocation strategies per CGN AS.

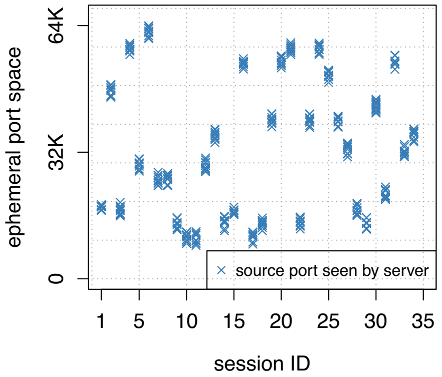

Figure 4.12: Observed ports per session, chunk-based random allocation strategy (4K ports per subscriber, AS12978).

Port allocation strategy		Non-cellular	Cellular
Port-preservation		41.2%	27.9%
Sequential		22.2%	26.0%
Random		35.6%	44.7%
Random (with chunk allocation)		9 (4.6%)	8 (3.7%)
	$\text{CS} \leq 1\text{K}$	4	2
Chunk size (CS)	$1\text{K} < \text{CS} \leq 4\text{K}$	2	3
	$4\text{K} < \text{CS} \leq 16\text{K}$	3	3

Table 4.6: Port allocation strategies observed for CGN ASes. For ASes implementing chunk-based random port allocation we estimate the per-subscriber chunk size.

In order to detect chunk-based allocation for all ASes, we require (*i*) at least 20 sessions with random port translation and (*ii*) all sessions exhibiting port numbers within a range smaller than 16K ports. Using this approach we identified 17 ASes using chunk-based allocations, shown in Table 4.6. While a minority of the identified CGNs, it allows us to reason directly about the dimensioning of the CGN for example in terms of the number of subscribers sharing a given IP address. We find 6 ASes in which subscribers only receive a port chunk smaller than 1K; for 3 of them, the chunk size falls to 512 ports—a scarily small number given that loading a single Web page can result in many dozens of TCP connections to fetch its various objects [76], resulting in a sizeable overall number of concurrent connections in residential networks [34]. The size of the port chunks then also directly translates into the maximum number of users per public IP address: we find 64 subscribers per IP address in the case of a 1K port chunk.

NAT pooling behavior: For the majority of the CGN-positive ASes, we observe *paired pooling behavior*, i.e., subsequent TCP sessions are bound to the same public IP address (recall Section 4.2), with varying port allocation patterns. However, we find that 21% of the CGNs also employ *arbitrary pooling behavior*, i.e., Netalyzr reported multiple global IP addresses during the duration of the test for more than 60% of the sessions. This list includes major ISPs spanning all geographic regions. IETF guidelines [62] discourage this behavior due to its complicating effect on applications (particularly SIP and RTP), which use multiple ports on the same end host but do not negotiate IP addresses individually [62, 161].

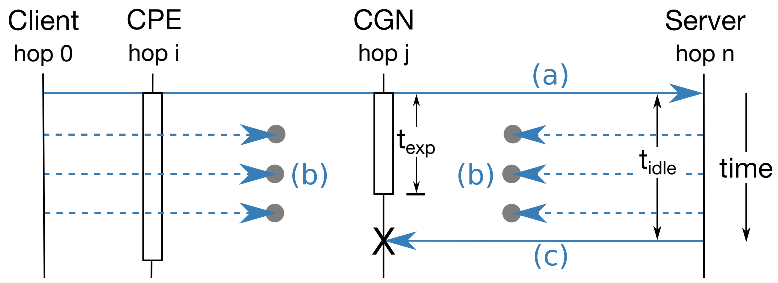

Figure 4.13: Reachability experiment scenario, consisting of initialization packet (a), keepalive packets (b), and probe packet (c). This example uses the parameters $i \leq ttl_c < j$ and $ttl_s < n - j$. If $t_{exp} < t_{idle}$ the mapping in the CGN (hop j) expires before the server replies with a probe packet.

	CGN detected	No CGN detected
IP address mismatch	67.6%	30.9%
IP address match	0.5%	0.9%

Table 4.7: Detection rate of TTL-driven NAT enumeration.

4.5.3 CGN-specific Measurements

To extract additional details about the CGNs under study, we extended the Netalyzr test suite with two tests. In this section we describe the tests' operation; the corresponding findings follow subsequently.

TTL-driven NAT enumeration: An extension of previous work [198], this method identifies the precise on-path location and the mapping timeouts of cascaded NATs. To do so, it leverages the stateful nature of NATs and their need to remove the state of idle UDP flows from the translation table (recall Section 4.2). During the test we repeatedly perform a reachability experiment that selectively detects stateful middleboxes in a chosen subset of the path, as depicted in Figure 4.13. As annotated, the test consists of three stages: *(a)* the client creates a UDP flow to our server, *(b)* both endpoints transmit TTL-limited "keepalive" probes for an idling period t_{idle} in order to keep the flow's state alive up to but not at the hop under test, *(c)* the server checks whether it can still reach the client. If not, we conclude that the hop under test is a NAT that has removed the flow's state. Our test enumerates the path between client and server by systematically performing iterations of this reachability experiment using different parameters for the client-side initial TTL ttl_c, server-side initial TTL ttl_s, and elapsed idle time t_{idle}.

We acknowledge three limitations of this approach. First, as a crowd-sourced test relying on user involvement, we need to limit the idle period of the test. We test idle times up to 200 seconds, the maximum possible value without prolonging the overall runtime of a Netalyzr test session. Hence, NATs with a mapping timeout larger than 200 seconds go unnoticed, leading to an underestimation of the actual number of NATs: in 30.9% of the tests (see Table 4.7) we do not find an expired mapping, while a mismatch between the client's local and server-perceived IP addresses (Section 4.3.2) evidently indicates NAT deployment. In the following, we only consider cases in which we could successfully observe an expired mapping. Second, based on the results of the reachability experiment we cannot distinguish between NATs and other stateful middleboxes such as stateful firewalls. However, we find stateful middleboxes without address translation in only less than 0.5% of our tests (see Table 4.7). We exclude these cases in the following analysis. Third, for a reliable expiration of the keepalive packets, the technique requires stable path lengths. Due to the large number of reachability experiments per test session (~60), we can detect and filter results with unstable paths.

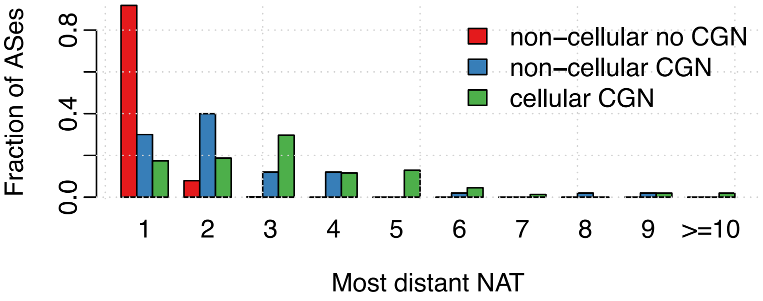

Figure 4.14: Maximum NAT distance from the subscriber.

STUN test: To study the mapping types of CGNs, we implemented a STUN [251] test in the Netalyzr test suite in October 2015. STUN determines the mapping type implemented by on-path NATs. STUN sends probe packets to a public STUN server (which answers certain probe packets from a different port and/or IP address) and waits for the respective replies.[12]

From the TTL-driven NAT enumeration test (deployed in September 2014) we have collected more than 38K sessions, whereas the STUN test (deployed in October 2015) produced 23K sessions. To be able to contrast sessions from within CGN-positive networks against CGN-negative ones, we augment the results from both tests with the results from our CGN detection tests (Section 4.3). We further apply filtering rules to the results of both tests to ensure that we have collected at least three sessions from a particular network (combination of AS number and CGN classification type, e.g. “cellular CGN”). After applying the filters, this leaves us with 18K sessions from the NAT enumeration test running via both non-cellular (70%) and cellular networks (30%). The results cover 608 ASes, whereof 43% (259 ASes) deploy CGN. For the STUN test we count 20K sessions from non-cellular (87%) and cellular networks (13%). The STUN results span 720 ASes including 170 CGN-positive ASes (24%).

4.5.4 Topological Properties of CGNs

Figure 4.14 shows the distribution of the number of hops between the client and the most distant NAT detected, grouped per AS and its respective CGN deployment status. We detected NATs as far away as 18 hops from the client. As expected, most of the NATs in CGN-negative ASes (92%) sit just one hop away from the client, i.e., they are typically located right on the CPE router. Compared to that, most CGNs are located two to five hops away from the client (64% of non-cellular and 73% of cellular ASes). In non-cellular ASes the CGN distance mostly ranges from two hops up to six hops. In the case of cellular ASes, however, the CGN distance ranges from one hop to two hops and up to 12 hops away from the client. In fact, we find that for 10% of the cellular ASes, the CGN is located six or more hops away from the client. A large number of hops between client and CGN hints at a centralized CGN infrastructure with large aggregation points, which has the potential of affecting the accuracy of IP geolocation databases when locating the external IP address of clients behind CGN.

4.5.5 Flow-Mapping Properties of CGNs

The type of NAT mapping (recall Section 4.2) as well as its state-keeping timeout directly affect the reachability of a host located behind a NAT, and thus has a profound effect on applications that rely on peer-to-peer connectivity [97, 119] or long-lived sparse flows [274].

[12] For more details on the operation of STUN, we refer to [251].

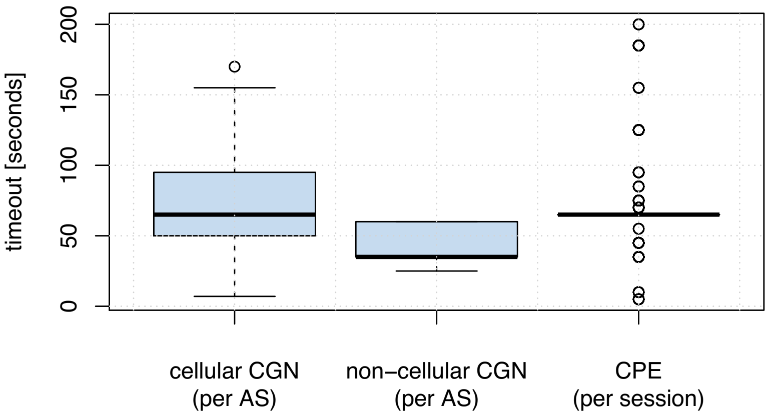

Figure 4.15: UDP mapping timeouts of CPEs and CGNs.

Mapping Timeouts: Figure 4.15 shows the UDP mapping timeouts for the detected CGNs, both in the cellular as well as in the non-cellular case. Here, we aggregate our CGN-positive sessions on a per-AS level. An AS is represented by its most frequent timeout value (mode). We also report timeout values that we detected for CPE devices (shown in the right boxplot), where we show a boxplot of all recorded sessions. In NAT444 scenarios (non-cellular CGN) we need to make sure to report the timeout of CGNs rather than the CPE NATs. Therefore, to reason about CGN mapping timeouts, we only consider sessions that were detected as CGN (Section 4.3) and where our TTL-driven NAT enumeration detected the NAT at a distance of three or more hops away from the client. We observe that 74% of detected NATs expire idle UDP state after 1 minute or less, but we find values ranging from 10s to 200s.[13] CGNs in cellular networks exhibit a larger median mapping timeout (65s) compared to non-cellular networks (35s). For CPE NATs we predominantly measured a timeout of 65s. We find higher variability and a lower median of timeout values for non-cellular CGNs when compared to CPE NATs. Low CGN timeout values might in turn negatively affect the longevity of sparse UDP flows that are also exposed to CPE NATs. While we find lower timeout values for CGNs compared to CPEs, we acknowledge that this property does not necessarily hold true for CGNs in general, as our test can not detect timeout values larger than 200 seconds.

Mapping types: Figure 4.16 shows our STUN results. We order the observed mapping types from most restrictive (*symmetric NAT*) to most permissive (*full cone NAT*). In Figure 4.16(a) we show the NAT mapping type as observed for CPE routers, while the bars in Figure 4.16(b) indicate the most permissive type of NAT mapping for our CGN-positive ASes. Recall that when multiple NAT devices reside on the path, STUN reports the most restrictive behavior of them, which also determines eventual NAT traversal. Hence, we argue that the most permissive STUN type provides a good approximation for the CGN behavior, because there cannot be a STUN result less restrictive than the CGN. We observe that, while exhibiting some diversity, less than 2% of the tests showed CPE NATs with very restrictive symmetric NATs. In contrast to CPE NATs, we observe 11% of non-cellular CGN ASes whose most permissive mapping type is symmetric. Among these networks we find many popular large European ISPs. For cellular networks we observe a bimodal outcome, with a large fraction of both restrictive (40% symmetric) and permissive (20% full cone) NAT types. We see large operators on both ends of the spectrum, with major cellular networks in the US deploying CGNs with symmetric mapping types.

[13] Note that our timeout detection mechanism uses a 10 second probing interval. Hence, reported values can differ up to 10 seconds from the actual NAT timeout.

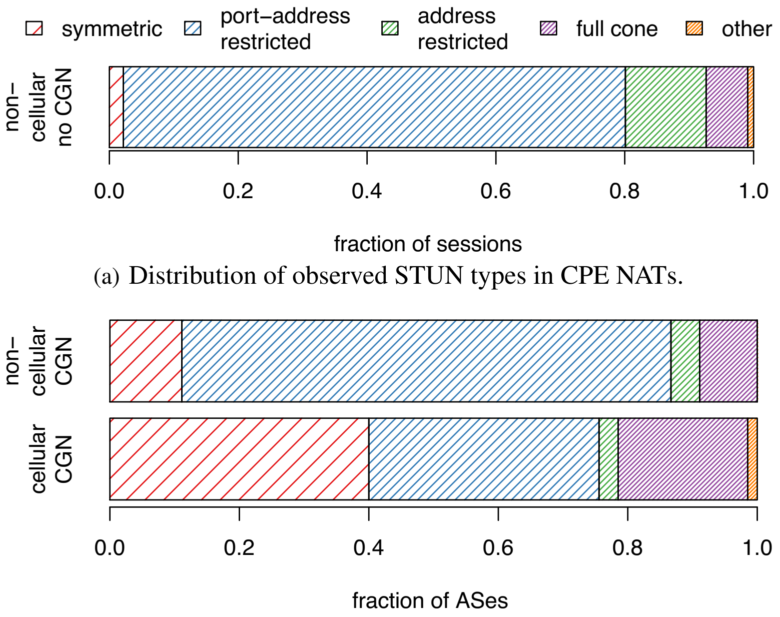

(a) Distribution of observed STUN types in CPE NATs.

(b) Most permissive STUN type per AS (only CGN sessions).

Figure 4.16: STUN results per AS.

Thus, we often measure stricter NAT mapping policies for CGN-positive sessions when compared to common home CPE devices. We conclude that a large fraction of ISPs deploy CGNs that use symmetric flow mappings, which limits the customers' ability to establish direct connections. For this reason, the IETF lists an endpoint-independent mapping (which symmetric NATs violate) as their first requirement for CGNs [62, 132].

4.6 Implications

Our analysis shows that ISPs widely deploy CGN. We find that more than 17% of eyeball ASes and more than 90% of cellular ASes rely on CGNs (Section 4.4), with particularly high deployment rates in Asia and Europe—regions in which IPv4 address scarcity cropped up first, as the respective registries ran out of readily available IPv4 addresses in 2011 and 2012. Thus, adopting CGN presents a viable alternative to buying IPv4 address space from brokers. CGNs actively extend the lifetime of IPv4 and hence also fuel the demand of the growing market for IPv4 address space [232], which in turn affects market prices and possibly hampers the adoption IPv6.

CGNs directly affect "how much Internet" a subscriber receives, by (*i*) limiting available ephemeral port space, (*ii*) restricting the directionality of connections, and (*iii*) limiting connection lifetimes due to finite state-keeping budgets. Studying our identified CGN deployments, we find a wide spectrum of configurations and degrees of address sharing (Section 4.5). On the limiting end of the spectrum, we find ISPs allocating as little as 512 ephemeral ports per subscriber (Section 4.5.2), multiplexing up to 128 subscribers per public IP address. Comparing NAT flow mapping types and timeouts of CGNs to commonly deployed CPE hardware, we find that in many instances CGNs use more restrictive flow mapping types when compared to their home counterparts (Section 4.5.5). This rules out peer-to-peer connectivity, complicating modern protocols such as WebRTC [140] that now need to rely on rendezvous servers.

We argue that the lack of *guidelines* and *regulations* for CGN deployment compounds the situation. While the IETF publishes best practices for general NAT behavior [62, 132, 214] as well as basic requirements for CGN deployments [215] (which, incidentally, many of our identified CGNs violate), dimensioning NATs at carrier-scale in a way that minimizes collateral damage remains a black art. Our finding that some large ISPs find the need to employ publicly routable (indeed, sometimes routed) address space for internal CGN deployment (Section 4.5.1) underlines the graveness of the situation. While it remains out of scope for this work to precisely measure the effect of CGNs on end-users' applications, we believe that our observations can serve as input for establishing such guidelines. Our findings should also interest regulators, who in some countries already impose acceptable service requirements on Internet performance (e.g., the FCC's measurements of advertised vs. achieved throughput [116]). We argue that the presence and service levels of CGNs should be readily identifiable in ISPs' offerings. Unfortunately, we find that most ISPs do not cover CGN deployment in their terms of service. Lastly, our findings document further erosion of the meaningfulness of IP address reputation, address-based blacklisting, IP-to-user attribution, and geolocating end users (Sections 4.5.2 and 4.5.4), which become all but infeasible in the presence of CGN.

4.7 Chapter Summary

In this chapter, we developed and applied techniques to detect CGN presence in the Internet and to extract dominant characteristics of the identified CGN instances. Our methods, based on harvesting internal IP addresses from the BitTorrent DHT and on extensions to the Netalyzr active measurement framework, prove effective in uncovering CGN deployments: We cover more than 60% of ISPs that connect end users to the Internet, and we detect and study more than 500 instances in ISPs across the globe.

CGNs have so far received comparably little attention, not just in terms of empirical research, but also in terms of developing best practices and regulations. One possible reason for the lack of empirical research might be that CGNs are often considered as an intermediary solution to temporarily mitigate IPv4 address scarcity (e.g., [159]), until full IPv6 deployment. Our data, however, shows deployment in a multitude of networks. In particular, we find that virtually all (more than 94%) of the tested cellular networks deploy CGN and more than 17% of the tested non-cellular end-user ISPs deploy CGN. Thus, CGN deployment has now moved onto the global stage, and CGNs are a reality for a large number of end users. We find that deployment rates are particularly pronounced in Europe and Asia Pacific, the two regions that ran out of IPv4 address space first. These observations underline that deploying CGN is a popular and widely used way to mitigate IPv4 scarcity. Given broad deployment of CGNs in networks across the globe, it is highly likely that these middleboxes are here to stay for the foreseeable future.

CGNs tackle a massive resource distribution problem and can be configured and dimensioned in a multitude of different ways. CGNs directly impact and limit the available connectivity of end users behind the CGN and change the activity characteristics of the gateway IP addresses fundamentally. Our analysis reveals a striking variability in the dimensioning, configuration, and placement of CGNs. We find that the degree of resource sharing varies substantially across the identified CGNs, with some ISPs dedicating just as few as 512 ports per subscriber. Moreover, CGNs restrict the connectivity available to end users to varying degrees, due to the use of differently strict NAT mapping types and timeout values. Mappings of public IP addresses to internal IP addresses as well as port translation behavior of the identified CGNs varies greatly, creating a scenario which renders the activity of gateway IP addresses unpredictable, and makes attribution of end users to IP addresses all but infeasible. We also uncover another problem that some ISPs face when deploying CGN: Shortage of *internal* address space. As a result, large ISPs opt to assign publicly routable address space to hosts behind the CGN. This practice has the potential to result in severe and hard-to-track connectivity problems, if these addresses are both in public as well as in internal use. This unintended use of routable address space also raises questions regarding ownership of IP addresses and address pollution issues.

Widespread CGN deployment has direct implications for application developers, network operators, content providers, law enforcement agencies, and regulators. While the IETF publishes best practices for general NAT behavior [62, 132, 214] as well as basic requirements for CGN deployments [215] (which many of our identified CGNs violate, e.g., by introducing short NAT timeouts), there exists no common ground on what degree of resource sharing is appropriate. This limited availability of guidelines and the absence of regulation thereof has resulted in a highly heterogeneous landscape regarding the configuration and scale of CGN deployment. Our findings in this chapter can serve as a basis for developing best practices, guidelines, and possibly regulation of CGN deployment. More detailed best practices could aid network operators to configure these devices with minimal harm for the affected end users and the Internet as a whole.

5 The Shift Towards IPv6

In the previous chapters, we focused exclusively on the IPv4 address space and evaluated mitigation strategies to increase its utilization. The long-term solution to the IPv4 address exhaustion problem will be the transition over to IPv6, vastly expanding the available address space. In this chapter, we study several connectivity- and traffic-related aspects of the IPv4 Internet, and its IPv6 counterpart, to shed light on the current state of IPv6 adoption and to illuminate several challenges in the process of transitioning to IPv6. Given the centrality of the Internet Protocol in the protocol stack, changing IP requires fundamental changes in hardware and software. In contrast to previously discussed approaches, such as the Carrier-Grade NAT approach, the transition over to IPv6 is a truly global effort. IPv6 will only solve address scarcity issues once broadly adopted by networks. Since its standardization in 1995, there have been many initiatives to promote IPv6 adoption and deployment [16]. Broad IPv6 adoption was intended to happen long before IPv4 address scarcity commenced [66, 94]. Yet, despite all these efforts, the transition to IPv6 has been slow and challenging in production environments [21, 88].

Presenting comprehensive statistics on IPv6 adoption and identifying the connected challenges is difficult for several reasons. True IPv6 adoption requires connectivity within the Inter Domain Internet, within ISP networks and content providers, within local networks, and at the respective network boundaries. It also relies on rendezvous mechanisms (most predominantly DNS) to support IPv6, to eventually allow exchange of data over IPv6. However, IPv6 adoption does not end within the network: IPv6 must be supported by operating systems and applications as well, on both ends of a connection. The complexity of this interplay makes it challenging to find the right metrics to assess and track IPv6 adoption. It is even harder to determine and pinpoint where precisely in this puzzle challenges and barriers for IPv6 adoption exist. In this chapter, we strive to inform the ongoing IPv6 transition, by first analyzing several IPv6 adoption metrics from different vantage points in the Internet, and then pinpointing some of the barriers and challenges that the Internet community faces on its way to broad IPv6 deployment.

The main contributions of this chapter are as follows:

- We develop and apply techniques to reconstruct the control plane over IPv4 and IPv6 at two Internet Exchange Points (IXPs), connecting hundreds of networks. We find that IPv6 Interdomain connectivity still lags behind IPv4 connectivity, albeit with a trend towards increasing IPv6 connectivity in a more recent snapshot. Taking traffic into account, however, we find that IPv6 peering links are less likely to carry traffic than their IPv4 counterparts. We then find that

IPv6 traffic is heavily concentrated on very few peering links. Our findings caution against taking inter-domain peering links at face value to track IPv6 adoption.

- We develop a methodology to classify the application-layer protocol of the exchanged traffic at our IXP. Traffic classification presents a challenge here, since our available dataset consists of packet samples. Applying our method to the exchanged traffic, we find that while the overall application mix seems consistent with widely reported statistics, each peering link carries an individual and different set of applications. Our findings illuminate the heterogeneity of the carried traffic across peering links. Individual applications need to be adapted in order to support and allow data exchange over IPv6 instead of IPv4 [61,259,282]. Given the heterogeneity of the application mix, we can expect a disparate potential for IPv6 adoption and traffic for individual peering links, depending on the involved networks and the applications they run or carry.
- We then study the interplay of connectivity and traffic over IPv4 and IPv6 in a residential network. This vantage point gives us the ability to precisely discern what portions of the traffic *are* and *could be* exchanged over IPv4/IPv6. Using a custom method to tag subscribers and their respective traffic flows to be either IPv4 or IPv6 capable, we illuminate a detailed picture of *barriers* that prevent traffic from being carried over IPv6 as well as the potential for traffic that could immediately be carried over IPv6, once service providers offer IPv6.

The remainder of this chapter is structured as follows: We review related work concerning IPv6 adoption in Section 5.1. We study connectivity and traffic over IPv4 and IPv6 in Section 5.2. We then present a method and analysis to reveal the application mix over individual peering links in Section 5.3 and then present our analysis of IPv4 and IPv6 interaction in a dual-stack ISP in Section 5.4 and summarize the findings and implications of this chapter in Section 5.5.

5.1 Related Work

The ability to track IPv6 adoption heavily depends on availability of relevant datasets [90]. Some works have reported the IPv6 traffic share at multiple vantage points in the Internet. In 2008, most IPv6 traffic at a tier-1 ISP in the US was DNS and ICMP [165]. While initiatives such as the "World IPv6 day" in 2011 ignited the increase of IPv6 traffic at various vantage points [255], by 2013 the share of IPv6 traffic at European IXPs or at 260 network providers was still below 1% [96,233]. Nonetheless, every year IPv6 traffic experiences a many-fold increase [96]. This development has encouraged studies on dual-stack networking performance [65,88,204,217], active measurements of the Internet's IPv6 infrastructure [70,185] and analyses of the AS-level topology [105,127]. Moreover, a large body of literature has focused on measuring IPv6 adoption among ISPs and service providers [91,96,105,127,164,165]. Several large content providers publish statistics on IPv6 usage of their services, e.g., Google [130] and Akamai [30]. Some works seek to understand the root causes that slow down IPv6 adoption and find a slower pace of adoption at the edge compared to core networks [105], or poor IPv6 quality in the early days of this transition [203]. As of today, the IPv6 control and data planes are—when applicable—*almost* on par with IPv4 [181], while both control planes show signs of convergence with respect to AS relationships [127]. In parallel to the research community, standardization bodies have invested decades to address IPv6-related aspects. Relevant to our work are fallback mechanisms for dual-stack applications [277] (*happy eyeballs*) and their implementations (see e.g., [20,143,144,256]).

We complement this body of work with measurement studies in the inter-domain Internet, as observed at IXPs as well as a measurement study within a dual-stack ISP.

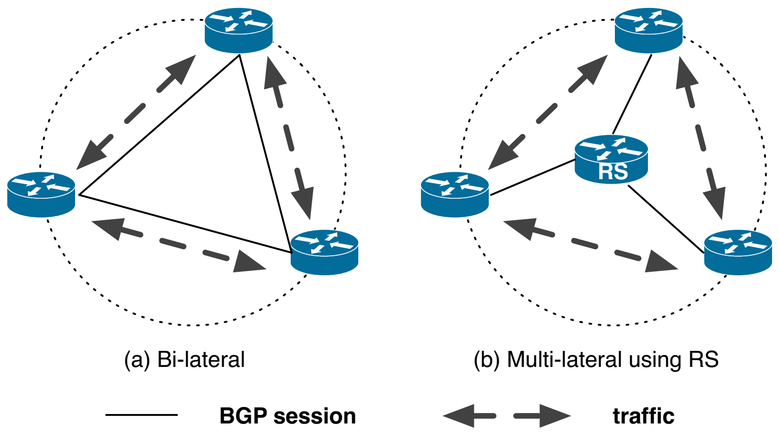

(a) Bi-lateral

(b) Multi-lateral using RS

Figure 5.1: IXP peering options.

5.2 Inter-Domain: Connectivity and Traffic at IXPs

In order to exchange data, networks, more specifically Autonomous Systems, need to establish peering links over which they can then exchange data. Peerings can either be established using direct interconnections (also known as *private peering*), or using IXPs (also known as *public peering*). In the case of private peering, the interconnecting networks establish a physical point-to-point connection between their border routers. This typically happens within co-location facilities such as Equinix [113]. Another way for networks to interconnect is by leveraging IXPs. In this case, networks establish a physical connection to a layer-2 switching fabric, which is operated by the IXP. Having layer 2 connectivity to all other participating networks (called members in the following), networks can then establish BGP sessions either directly with each other or via a Route Server (RS) [234] at such IXPs.

There are currently some 350+ Internet eXchange Points (IXP) worldwide, and some of the largest and most successful IXPs have more than 600-700 members and carry as much traffic as some of the global Tier-1 ISPs [86]. With membership growth rates of 10-20% per year [86] and annual traffic growth rates of 50-100%, these IXPs have emerged as key entities in the Internet infrastructure where a vast majority of today's peering connections are established [29,129,279]. This growing importance of IXPs for the Internet peering ecosystem and the IXPs' increasing popularity with the full spectrum of Internet players have come in full view with recent studies such as [29, 86, 129, 184]

In the following, we will study the peerings at two such IXPs. Here, we pay particular attention to differences in IPv4 and IPv6 connectivity and the respective traffic patterns.

5.2.1 Peering Options at IXPs

The typical way to establish connectivity between two ASes is to establish a direct BGP session between two of their respective border routers. Initially, if two IXP member ASes wanted to exchange traffic via the IXP's switching fabric, they had to establish a *bi-lateral (BL)* BGP peering session at the IXP. However, as IXPs grew in size, to be able to exchange traffic with most or all of the other member ASes at an IXP and hence reap the benefits of its own membership, a member's border router had to maintain more and more individual BGP sessions. This started to create administrative overhead, operational burden, and the potential of pushing some router hardware to its limit.

To simplify routing for its members, IXPs introduced RSes and offered them as a free value-added service to their members. In short, an IXP RS is a process that collects routing information from the RS's peers or participants (i.e., IXP members that connect to the RS), executes its own BGP decision process, and re-advertises the resulting information (i.e., best route selection) to all of the RS's peer routers. Figure 5.1 shows the flow of control plane information (BGP sessions) and data packets (data plane) for both traditional bi-lateral peering as well as peering via the RS at an IXP. The latter is referred to as *multi-lateral (ML)* peering because it typically involves more than two BGP partners.

A member AS connects to the RS via a single BGP session to set up BGP peering with all other IXP members that peer with the RS.[1] Clearly, this lowers the maintenance overhead, in particular for small ASes. Note, however, that using the RS (i.e., ML peering) does not preclude BL peering by one and the same member AS. In particular, larger ASes can take advantage of the RS while still having the option to establish BL peerings with selectively-chosen IXP members. For example, if a large member AS finds the capabilities of the RS to be insufficient for its needs (e.g., with respect to traffic engineering or filtering) or prefers to have more control over the peerings with its most important peers, it can use BL peerings with the latter and ML peerings with those members that peer with the IXP's RS.

5.2.2 IXP Route Servers: Design

In the following, we describe a typical Route Server configuration, based on the BIRD routing daemon [12]. The software was developed by CZ.NIC Labs and has been actively supported and widely used in the IXP community. This configuration that has been abstracted from the Euro-IX RS example [5] and is the basis of the one in operational use by one of the IXPs with which we have an ongoing collaboration.

Like all routing daemons, BIRD maintains a Routing Information Base (RIB) which contains all BGP paths that it receives from its peers – the Master RIB. However, when using BIRD as RS, it can be configured to (i) maintain peer-specific RIBs and (ii) use them instead of the Master RIB for peer-specific BGP best path selection (see Figure 5.2). When configured this way, each member AS that peers with the RS maintains a BGP session with the RS, which results in a peer-specific RIB. When IXP member AS X advertises a prefix to the RS, it is first put into the AS X-specific RIB. Next, if this prefix passes the AS X-specific import filter, it is added to the RS' Master RIB. If the export filter of AS X allows it, then this prefix will also be added to each AS Y-specific RIB, where AS Y is any other IXP member that connects to the RS. Then the RS performs a peer-specific best path selection and exports the prefix by re-advertising it to each AS Y.

IXPs typically apply import filters to ensure that each member AS only advertises routes that it should advertise. To derive import filters, the IXPs usually rely on route registries such as IRR [7]. This policy limits the likelihood of unintended prefix hijacking and/or advertisements of bogon prefixes including private address space. With respect to export filters, they are typically triggered by the IXP members themselves to restrict the set of other IXP member ASes that receive their routes. The commonly used vehicle for achieving this objective is to tag route advertisements to the RS with RS-specific BGP community values [128]. These values are set on a per route basis and restrict to which members the route can be propagated. Thus, by using export filters, peers of the RS can express policies.

[1] For redundancy purposes, IXPs typically operate two RSes and encourage their members to connect to both RSes.

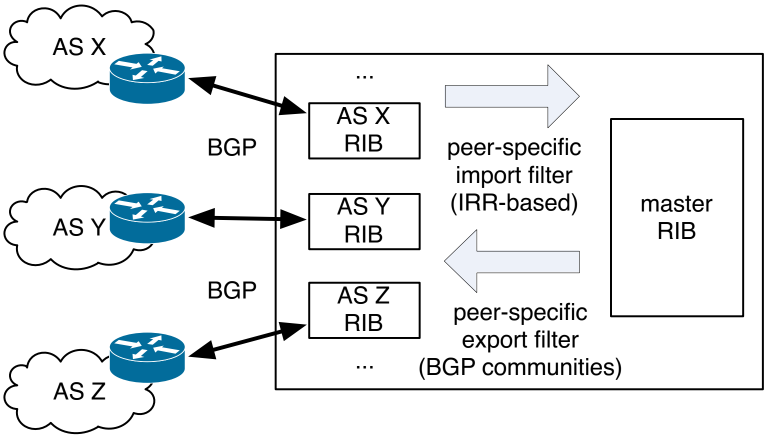
AS X
AS Y
AS Z
BGP
BGP
...
AS X RIB
AS Y RIB
AS Z RIB
...
peer-specific import filter (IRR-based)
peer-specific export filter (BGP communities)
master RIB

Figure 5.2: BIRD route server: Example setup with peer-specific RIBs and import/export filtering.

5.2.3 IXP Datasets

In the following, we introduce our two IXPs and the datasets that are available to us to study connectivity and traffic exchanged at each IXP.

Large IXP (L-IXP): This IXP is one of the largest IXPs in Europe and worldwide. It operates a layer-2 switching fabric that is distributed over a number of colocations/datacenters within a metropolitan area. For this IXP, we cover a 4-week period in August/September 2013 (denoted as *2013* in the following), and a 4-week period in January 2017 (denoted as *2017* in the following). In *2013*, this IXP had 496 members and a peak traffic volume of 3 Tbps. In *2017*, this IXP had 718 members and a peak traffic volume of 5 Tbps.

Medium IXP (M-IXP): This medium-sized IXP operates a layer-2 switching fabric, and is present in several locations. For this IXP, we cover a 4-week period in December 2013 (denoted as *2013*) in the following. As of late 2013, this IXP had 101 members and its peak traffic exceeded 250 Gbps.

Route Server Snapshots

For both IXPs, we have access to the data from their BIRD route server deployment. The unique advantage of having access to these IXP-provided control plane measurements is that they are rich enough to accurately and completely reconstruct the "ground truth" in terms of connectivity at these IXPs that has been established with the help of the RS, i.e., the IXPs' multi-lateral peering fabrics.

Traffic

For each of our IXPs, we have access to data plane measurements in the form of traffic that is routinely collected from the IXPs' public switching infrastructures. More precisely, for each IXP, the available datasets consist of massive amounts of sFlow records [258], sampled from their public switching infrastructure. The measured sFlow records contain Ethernet frame samples that have been collected using random sampling (1 out of 16K). sFlow captures the first 128 bytes of each sampled frame. Thus, they contain full Ethernet, network- and transport-layer headers as well as some bytes of payload for each sampled packet. For further details about relevant aspects of these collected sFlow records (e.g., absence of sampling bias, removal of irrelevant traffic), see [29]. For the *2017* snapshot of the L-IXP, we have

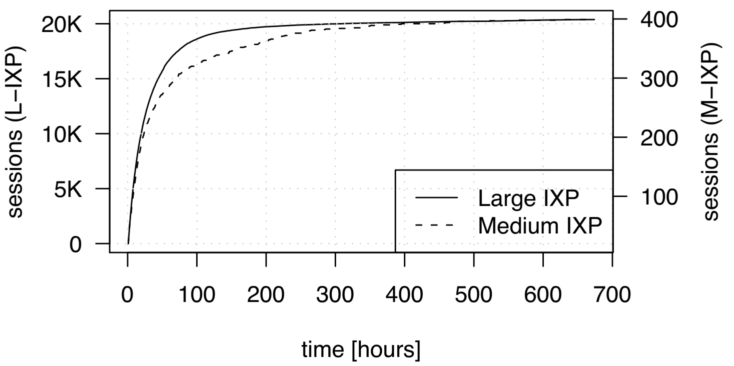

Figure 5.3: Inferred bi-lateral BGP sessions over time.

access to IPFIX records, collected from their public switching infrastructure using random sampling (1 out of 10K). The IPFIX records give us access to Ethernet, network- and transport-layer addresses and flags, but do not contain payload information.

We rely on 4 continuous weeks of collected sFlow/IPFIX for each IXP and snapshot. Having access to these IXP-provided data plane measurements makes it possible to examine the connectivity that has been established without the use of the IXPs' RSes (i.e., bi-lateral peerings). However, more importantly, these measurements provide valuable but hard-to-obtain information about how the two parties of an IXP peering use that link.

5.2.4 Peering Links at IXPs

Inferring Peering Links at IXPs

Relying on our IXP-provided measurements, we show in this section how we can get close to recovering the actual peering fabrics at the IXPs. To determine if IXP members AS X and AS Y are using a ML peering at the IXP, we rely on the IXP-provided RS data. More specifically, for the L-IXP, we first check if AS X and AS Y peer with the RS. If so, we next check in the peer-specific RIB of AS Y for a prefix with AS X as next hop. If we find such a prefix, we say that AS X uses a ML peering with AS Y. If we also find AS Y in the peer-specific RIB of AS X as next hop, we say that the ML peering between AS X and AS Y is *symmetric or bi-directional*; otherwise, we say that the ML peering between AS X and AS Y is *asymmetric*. Given that the RS at the M-IXP only uses a Master RIB but no peer-specific RIBs, we re-implement the per-peer export policies based upon the Master RIB entries to determine peerings via the RS. More specifically, if there is a route for a prefix in the Master RIB with AS X as next hop, we postulate a ML peering with all member ASes that peer with the RS, including AS Y, unless the community values associated with the route explicitly filter the route via the peer-specific export filter to AS Y.

To determine if IXP members AS X and AS Y are using a BL peering at the IXP, we rely on the IXP-provided traffic measurements. In particular, to conclude that AS X and AS Y established a BL peering at the IXP, we require that there are sFlow records in the IXP-provided traffic data that show that BGP data was exchanged between the routers of AS X and AS Y over the IXP's public switching infrastructure.[2] We cannot however differentiate between asymmetric and symmetric BL peerings with these data plane measurements.

[2] The routers' IP addresses have to be within the publicly known subnets of the respective IXP.

		M-IXP (2013)	**L-IXP (2013)**	**L-IXP (2017)**
	member ASes	101	496	718
IPv4	ML symmetric	3,140	65,599	143,057
	ML asymmetric	594	14,153	25,432
	BL only	61	5,705	8,421
	BL and ML	399	14,673	17,013
	total peerings	**3,795 (75%)**	**85,457 (70%)**	**176,910 (69%)**
IPv6	ML symmetric	1,173	34,596	75,246
	ML asymmetric	434	5,086	23,870
	BL only	75	3,727	5,267
	BL and ML	223	4,256	8,979
	total peerings	**1,682 (33%)**	**53,409 (35%)**	**104,383 (41%)**

Table 5.1: Inferred multi-lateral and bi-lateral peering links at our IXPs.

Note that our methodologies yield a lower bound for BL peerings and an upper bound for ML peerings, but there is evidence that these bounds are in general very tight. For example, with respect to BL peerings, our method is not significantly biased by the sFlow sampling rate because the numbers are very stable once we use data from more than two weeks. Indeed, Figure 5.3 shows that for the L-IXP, the additional BL peerings seen in the third (fourth) week are less than 1% (0.5%). As far as ML peerings are concerned, our method does not account for the fact that some RS peers might reject the advertisements of the RS, which can result in some over-counting by our method. At the same time, we find pairs of member ASes that use the provided layer 2 connectivity both for ML as well as for BL peering.

IPv4 and IPv6 Peering Links at IXPs

Our best efforts to reconstruct the actual ML and BL peering fabrics of our IXPs is summarized in Table 5.1. We further break down (where possible) each of the ML and BL peering fabrics into links that are used for either IPv4 or IPv6 and in a symmetric or asymmetric manner. For each IXP, we also tally the total number of peerings along with the peering degree (percentage of established peering links compared to the number of possible peering links).

We observe a dense peering mesh at our IXPs, with ML peerings outnumbering BL peerings by a ratio of 4:1 and 8:1 at the L-IXP and the M-IXP respectively. Thus, connectivity at these IXPs is clearly driven by their RSes and the resulting ML peerings. We see an increase in IPv6 connectivity when comparing 2013 and 2017 snapshots of our L-IXP. However, the degree of connectivity (i.e., the number of established peering links divided by the total possible peering links) for IPv6 still ranges at only 41% at our L-IXP in 2017, compared to 69% for IPv4. While the degree of connectivity went up from 35% to 41% over a four year time range, many peering links are still IPv4-only. This increase in IPv6 connectivity at our IXP is clearly dominated by ML peerings and suggests that RSes play a vital role when it comes to establishing the IPv6 control plane and has the potential to provide immediate IPv6 connectivity to a multitude of networks. Our connectivity findings for IPv6 agree with earlier studies that show that the IPv6 BGP control plane, while still lacking behind IPv4, is emerging [105].

5.2.5 Traffic on Peering Links

Next, we are interested in how many of the established peerings reported in Table 5.1 are actually "used"; that is, see traffic. Furthermore, we are interested in the actual distribution of traffic when taking the type of the peering link into account. This will aid to put our connectivity-based findings into proper perspective.

Identifying Traffic-Carrying Peering Links

To identify a traffic-carrying peering between AS X and AS Y, we look for sFlow/IPFIX records that (i) contain MAC addresses which belong to AS X and AS Y, respectively, and (ii) have IP addresses that are not part of the IP address space assigned to the IXP. Thus, we only count the exchange of non-local IP traffic, which allows us to clearly separate control traffic (i.e., BGP sessions) and actual data traffic, thus we can distinguish between BL peerings with and without traffic. Once we identified a traffic-carrying peering link, we assign it to be either BL or ML, depending on our earlier introduced inference. For a small portion of the traffic (less than 0.5% for both IXPs) we did not find a corresponding BL or ML peering link. We discard this traffic from our analysis.[3]

In this context, for IXP member ASes that peer with other member ASes at the IXP both bi-laterally and multi-laterally, we are faced with the problem of determining whether the observed traffic between two such ASes is traversing the BL or ML peering link between them; that is, identifying the traffic-carrying peering(s). Taking a pragmatic approach, when two IXP member ASes peer with one another at the IXP both bi-laterally and multi-laterally, we tag the BL peering between them as the traffic-carrying peering and associate any observed traffic with it. Intuitively, our argument for this approach is based on the observation that compared to ML peering, establishing a BL peering requires work (e.g., manually setting up BGP sessions) and is an indication of joint incentives and needs between the involved parties. On the other hand, peering multi-laterally at the IXP (i.e., using the RS) is designed to be easy and informal, making it in general possible to exchange traffic with all the RS's peers from the get-go. To provide empirical support for our argument, we manually searched for Looking Glasses (LGes) that query the routing tables of member routers that peer both bi-laterally and multi-laterally with other members. We found six such LGes with sufficient capabilities to reason about the best path selected. In all cases, advertisements via BL sessions were selected as best path over advertisements from the RS.[4]

Traffic-Carrying Peering Links: IPv4 and IPv6

Table 5.2 summarizes the results of our analysis of the traffic-carrying links. When compared to Table 5.1, the first column in Table 5.2 shows that most peering links (i.e., more than 80% at both IXPs) are actually "used" in the sense of the binary attribute "carry traffic/no traffic". Moreover, we note that the ratio of traffic-carrying peerings is largest for BL peerings, followed by symmetric ML peering, followed by asymmetric ML peering.

Moving beyond this binary classification of peering links, this usage picture sharpens when examining the second column of Table 5.2. This column shows for each IXP the number of peerings responsible for 99.9% of the IXP's total traffic in terms of bytes, for IPv4 and IPv6 individually. Hence, all peerings that

[3] Possible explanations for this traffic are either non-detected BGP sessions or peerings using protocols other than BGP (e.g., static routing).

[4] Selection of BL over ML was typically done by setting the local preference to a higher value for routes received via BL sessions. However, we point out that is not necessarily true for all peerings.

IPv4	**M-IXP (2013)**		**L-IXP (2013)**		**L-IXP (2017)**	
	all	99.9p	all	99.9p	all	99.9p
% BL	93.5	47.7	92.4	55.6	87.2	53.5
% ML sym.	83.7	24.0	85.9	31.3	84.1	22.9
% ML asym.	38.5	7.89	23.8	5.43	21.0	5.8
links total	2,968	918	67,915	28,849	135,993	44,590
traffic contribution	99.5%		99.4%		97.5%	

IPv6	**M-IXP (2013)**		**L-IXP (2013)**		**L-IXP (2017)**	
	all	99.9p	all	99.9p	all	99.9p
% BL	74.9	7.17	76.2	4.92	87.9	14.6
% ML sym.	52.2	0.48	54.0	0.52	77.8	1.2
% ML asym.	25.3	0.07	30.4	0.04	62.6	0.7
links total	819	24	24,159	556	79,335	3,065
traffic contribution	0.50%		0.63%		2.46%	

Table 5.2: Percentage of links that carry traffic (all traffic vs. top 99.9% of all traffic), their corresponding type and contribution to overall IXP traffic.

collectively see less than 0.1% of the overall traffic are discarded. When imposing such thresholds to eliminate peerings that carry only comparably little amounts of traffic, we observe a drastic reduction of the number of active peerings. Indeed, the main take-away from this thresholding exercise is that it puts the connectivity-related findings reported in the previous section into proper perspective. Specifically, it demonstrates that while RSes increase connectivity and are responsible for the larger part of peerings, the majority of those ML peerings typically does not carry much traffic. At the same time, the smaller number of BL peerings that are established at IXPs carry in general the bulk of the traffic.

Focusing on the differences between IPv4 and IPv6, we notice that only a tiny fraction of the large number of IPv6 peerings that have been established at our IXPs carry any significant traffic volumes. In fact, about 98% of all IPv6 peering links at both IXPs in 2013 carried no significant traffic volumes (to meet our threshold). Comparing our findings from 2013 with 2017 for the L-IXP, we see that the fraction of low-traffic links has slightly decreased from 98% in 2013 down to 96% in 2017. In absolute numbers, that means that, as of 2017, we see 3,065 IPv6 peerings with significant traffic versus 44,590 IPv4 peerings, a ratio of roughly 7%. Hence, the existence of a peering link between two ASes does not reveal much about whether such peerings are actually used. Since a majority of inter-domain peering links at IXPs are established via the Route Server, and not via bilateral agreements, the importance of an inter-AS link (and the corresponding use as metric e.g., to track IPv6 adoption) is questionable. This cautions against relying purely on BGP control plane date to meaningfully track IPv6 adoption.

Traffic Concentration on Peering Links for IPv4 and IPv6

Next, we move a step further and assess the traffic concentration across peering links, where we pay particular attention to comparing the properties of IPv4 and IPv6 traffic-carrying peering links. As per Table 5.2, the majority of peering links at our IXPs do no carry significant traffic volumes. We observe that the fraction of peering links that carry only comparably little volumes is much higher for IPv6, when compared to IPv6. Figure 5.4 shows a CCDF of the traffic contribution of individual peering links, where we treat IPv4 and IPv6 peerings independently.

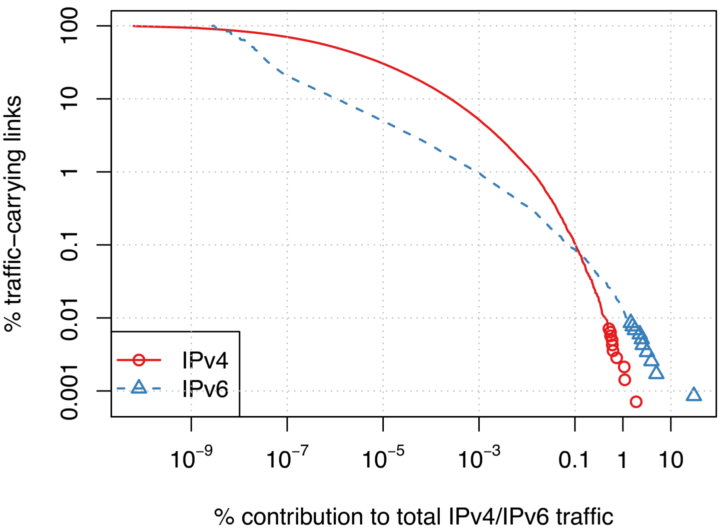

Figure 5.4: CCDF: Traffic contribution of IPv4 and IPv6 peering links, relative to the total IPv4 and IPv6 traffic at this IXP.

Indeed, we observe that IPv6 traffic is much more concentrated across a very small number of peering links. This becomes more clear when looking at the tail of the distribution: The single IPv6 peering that carries most IPv6 traffic at this IXP contributes more than 30% to the overall IPv6 traffic! In contrast, the top traffic-carrying IPv4 peering link only contributes 1.8% to the overall IPv4 traffic.

5.2.6 Summary

In this section, we studied IPv4 and IPv6 connectivity and traffic in the Inter-domain Internet, leveraging data from two IXPs. We find that as of 2017, the total of IPv6 peerings at our IXPs is still less than 60% of the number of IPv4 peering links. However, we find that IPv6 connectivity improved over the course of the last 4 years. IXP Route Servers make it easy for networks to immediately establish hundreds of peerings with other IXP members, requiring just one single BGP session with the Route Server. Hence, Route Servers can support IPv6 adoption, providing immediate connectivity. However, the ease with which multi-lateral peerings can be established also cautions against taking inter-domain peering links at face value to measure IPv6 adoption.

When taking traffic into account, we find that most IPv6 peering links only carry marginal traffic volumes. As of 2017, only less than 4% of the established IPv6 peerings carry the bulk of the exchanged traffic, 99.9%. In IPv4, more than 32% of the established peerings carry the top 99.9% of traffic. Thus, IPv6 traffic is significantly more concentrated on few peering links, when contrasted with its IPv4 counterpart. In fact, more than 30% of all IPv6 traffic (2.5% of the overall traffic) is carried over one single IPv6 peering link. Thus, while connectivity between IPv6 networks is there, few of them exchange significant traffic volumes over IPv6.

5.3 Inter-Domain: Application Mix Heterogeneity

So far, we studied the control and data planes for IPv4 and IPv6 at IXPs. We find that IPv6 peering links at our IXPs carry much less traffic, when compared to their IPv4 counterparts. To illuminate some of the possible reasons that refrain traffic to be carried over IPv6 instead of IPv4, a deeper understanding

Name	Timerange	Sampling	Packets	Bytes	IPv4 / IPv6	TCP / UDP
09-2013	2013-09-02 to 2013-09-08	1/16K	9.3B	5.9TB	99.36 / 0.63	83.7 / 16.3
12-2012	2012-12-01 to 2012-12-07	1/16K	8.5B	5.5TB	99.64 / 0.36	83.1 / 16.9
06-2012	2012-06-04 to 2012-06-10	1/16K	7.3B	4.6TB	99.80 / 0.20	80.7 / 19.3
11-2011	2011-11-28 to 2011-12-04	1/16K	6.4B	4.2TB	99.93 / 0.07	79.8 / 20.2
04-2011	2011-04-25 to 2011-05-01	1/16K	5.3B	3.5TB	99.94 / 0.06	79.2 / 20.3

Table 5.3: Overview of dataset characteristics. The number of packets/bytes refer to the number of packets collected i.e., after sampling.

of the various exchanged traffic components is necessary. The choice, or even the ability, to use IPv6 as preferred protocol depends on various factors. Not only the network itself, but also the end hosts and the applications running on end hosts need to support, and choose, IPv6 instead of IPv4 [61, 259, 282].

However, due to the heterogeneity of the Internet and its complex topology and global scope, there are no simple answers to questions like “What are the most popular applications in today’s Internet?” or “What is the application mix in today’s Internet?” In fact, as more and more networks consider factors such as cost, performance, security, ease-of-use, and flexibility when deciding about which kind of traffic to send over which type of peering links, the application mix can be expected to differ from link to link.

We are interested in how representative commonly-reported aggregate statistics concerning the Internet’s application mix are in view of the network’s enormous heterogeneity. To this end, we first develop a new methodology to classify traffic from packet-sampled traffic traces. Packet sampling is a widely employed technique when monitoring high-bandwidth infrastructures and is commonly used by large ISPs and IXPs. We then rely on traffic traces collected at our large IXP (L-IXP) and apply our traffic classification methodology to infer the application mix on the thousands of public peering links at this IXP. Our results show that the heterogeneity of the Internet extends directly to the application mix of its traffic, and we illustrate the observed heterogeneity by providing insight into how and why the application mix can differ from interconnection to interconnection and among different types of networks.

5.3.1 Dataset Characteristics

For our classification, we rely on packet-sampled traffic traces captured from the public switching fabric of the L-IXP, as introduced in Section 5.2.3. We use five snapshots (selected from a period that spans 2.5 years), each covering a full week (168 consecutive hours). Table 5.3 lists the pertinent properties of these traces. Unless mentioned otherwise, we use the week 09-2013 as default snapshot. Since our newly obtained IPFIX snapshots from this IXP do not contain payload information (Section 5.2.3), we here rely solely rely on our sFlow snapshots.

Recall that sFlow captures the first 128 bytes for each sampled Ethernet frame. Thus, each packet includes the full link layer (Ethernet), network layer (IP), and transport layer (TCP/UDP) protocol headers as well as a limited number of payload bytes. In the most common case, where the IPv4 and TCP protocols are used, this leaves 74 bytes worth of payload information (if TCP option fields are set, the available payload is further reduced by a few bytes). In the following, we consider only IPv4 traffic, as the fraction of IPv6 is still below 1% in the considered snapshots. We note, however, that this methodology is also applicable for IPv6 traffic.

The sampled nature of our datasets poses significant challenges when trying to apply traditional traffic-classification approaches (see Section 5.3.2 for details). To assess the impact of sampling on the visi-

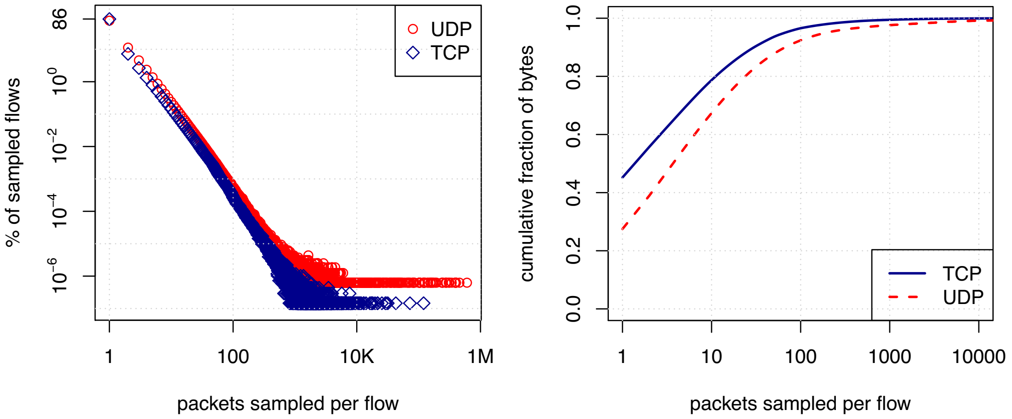

(a) Samples per flow (1200s timeout).

(b) Fraction of total bytes related to number of packets sampled per flow.

Figure 5.5: IXP data sampling characteristics relevant for traffic classification.

bility of "full" flows, we aggregate the packets sampled at our IXP using the typical 5-tuple aggregation consisting of source and destination IP addresses, source and destination port numbers, and the transport protocol. Figure 5.5(a) shows the number of packets that are sampled for each flow, using a 1200s timeout. It shows that we see only a single packet for some 86% of the sampled TCP flows (76% for sampled UDP flows). We also observe flows for which we sample several hundreds of thousands of packets over the course of one week. Surprisingly, UDP flows dominate the heavy-hitter flows and closer inspection reveals that most of the large UDP flows are related to recursive DNS interactions between name servers. Accordingly, Figure 5.5(b) shows the cumulative total number of bytes related to flows for which we sample less or equal than x packets. It shows that in the case of TCP, more than 45% of the bytes are sampled from flows for which we sample only a single packet (27.5% for UDP). Since we only observe packets, we cannot rely on any per-flow properties nor can we expect to sample packets at any specific position of a flow e.g., the first packet(s). Moreover, we cannot expect to have any visibility into the bidirectional nature of any of the flows–all that sampling gives us is a "random set of packets."

5.3.2 Classification Approach

Related Work

Application classification has attracted the attention of researchers for many years and has resulted in a large number of different methods and studies. However, the characteristics of our datasets (i.e., sampling, no bidirectional visibility) pose new challenges for application classification. In particular, since most of the existing classification approaches require information that is not available in our datasets (e.g., unsampled packet traces, flow statistics), these methods are not directly applicable in our context.

Before presenting our new application characterization method, we first provide a condensed taxonomy of existing classification approaches. To this end, we follow closely the description presented in [169] and focus on those aspects of the different approaches that prevent them from being directly applicable to the types of datasets we are considering. For a more detailed discussion of the various existing application classification approaches, we refer to extensive surveys such as [81, 101, 169, 202, 272].

Port-based approach: Many applications typically run on fixed port numbers which can be leveraged to classify packets to their corresponding applications. The drawbacks of port-based classification are that *(i)* applications can rely on random port numbers (e.g., as Peer-to-Peer (P2P) applications) and *(ii)* applications might use well-known port numbers to obfuscate traffic (e.g., see [193]). On the positive side, port-based classification has been shown to be still effective [189], is robust to sampling and can be applied to our dataset in a straight-forward manner. Note that port-based classification was already performed for the sFlow data captured at this IXP in [29].

Payload-based approach: Also referred to as Deep Packet Inspection (DPI), payload-based classification produces very accurate results by relying on application-specific signatures (i.e., known byte patterns of known protocols). Application signatures are typically based on protocol handshakes and can often be assembled using only the first few payload-carrying packets that are exchanged between the communicating hosts (i.e., an `HTTP GET` request followed by an `HTTP/1.{0,1}` reply). The payload-based approach is often used to establish ground truth for the application mix of traffic traces (see e.g., [84] for a comparative study). While we have access to the initial bytes of the payload of each sampled packet, we do not necessarily sample the first packet(s) of flows that contain application signatures. In addition, we cannot inspect bidirectional payload patterns of flows using our datasets.

Flow features-based approach: By utilizing flow properties (e.g., the total number of packets, average packet size), several approaches focus on classifying flows as belonging to specific applications without inspecting the payload of packets. Since we do not have per-flow information, these approaches are not applicable to our datasets.

Host behavior-based approach: This class of approaches classifies traffic by profiling the detailed network interaction of hosts (e.g., which destinations are contacted on which ports [163] or the network-wide interactions of hosts [151]). The various approaches in this class have been shown the be particularly effective for characterizing P2P applications [162]. While we are not able to perform fine-grained profiling of hosts due to the sampled nature of our data, we do make use of properties inferred from the *social* behavior of hosts to uncover instances of Peer-to-Peer traffic.

Building Blocks

The foundation of our classification approach outlined below is the ability to attribute *some* of the sampled packets to their respective applications by mainly using payload signatures and partly relying on port numbers. In particular, we rely on signatures which we derived from the *L7-filter* [2] and the *libprotoident* library [33] for well-known protocols such as HTTP, SMTP, POP3, IMAP, NNTP and SSH. We also make use of application signatures derived from protocol specifications [3, 14] for BitTorrent. We also used available signatures to detect other P2P protocols (e.g., eDonkey) but their contributions in terms of classifying packets were insignificant. We verified all application signatures using manually generated traffic traces. For SSL-based protocols (we focus on HTTPS, NNTPS, POP3S, and IMAPS), we use signatures indicating an SSL handshake and consider SSL handshake packets on the well-known port number of the respective application (e.g., 443 for HTTPS) as belonging to that application.

To ensure the accuracy of our application signatures (i.e., keeping the false positives low by limiting the number of signatures), we restrict our set of application signatures and port numbers and only consider applications that *(i)* generate significant traffic and *(ii)* are reliably detectable using application signatures and, if needed, port numbers. For example, we do not try to classify Skype traffic because its detection remains unreliable unless one uses specialized approaches [72].

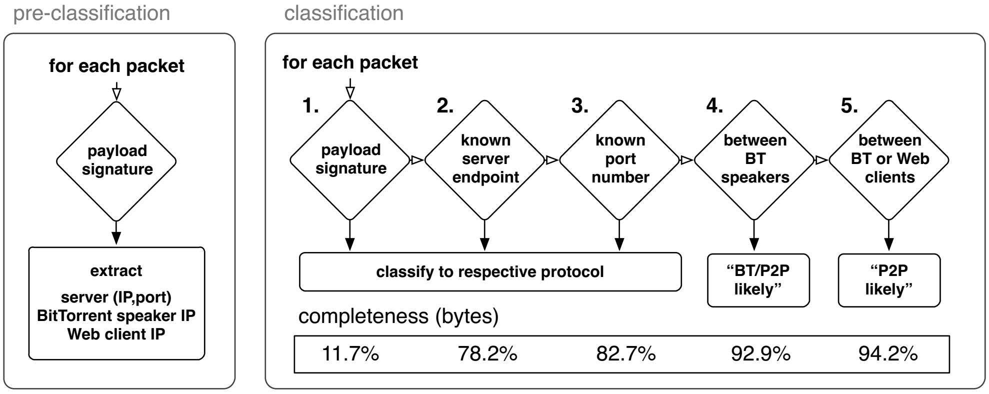

Figure 5.6: Classification Pipeline annotated with the cumulative bytes classified in each step.

Classification Pipeline

Figure 5.6 illustrates our classification pipeline. In particular, our classification approach requires that the given traffic trace be processed twice, first in a *pre-classification* phase and then in a *classification* phase. The purpose of the first phase is to derive *state*, which will then be leveraged in the *classification* phase to attribute packets to their respective endpoints, revealing the corresponding application.

I. Pre-classification phase

The goal of the pre-classification step is to extract server *endpoints* and IP addresses of clients, which will be used as state in the subsequent classification phase. In this phase, we rely solely on payload-based classification using our validated signatures (as well as SSL signatures on well-known ports). For each packet that belongs to a client-server application, we save the server endpoint, i.e., its (IP, port) tuple. To identify the server-side of a packet, we rely on directed signatures (e.g., HTTP request vs. HTTP reply). For packets matching a BitTorrent signature, we save the SRC and DST IPs but not the port numbers. Since most BitTorrent traffic that matches our signatures is UDP-based which, due to its connectionless nature, is more susceptible to spoofing as well as other phenomena such as BitTorrent DHT poisoning for control traffic (e.g., [273]), we only count an IP address as BitTorrent speaker if we sample at least 2 packets that originate from/are sent to that IP address matching our signatures. Additionally, we save IP addresses of HTTP clients. In this pre-classification, we identify more than 2.7M HTTP server endpoints (1.43M unique IP addresses), and 210K HTTPS endpoints. On the client side, we identify 37.7M HTTP client IPs as well as 38.9M BitTorrent speakers, where the overlap between HTTP client IPs and BitTorrent speakers is 12.4M IP addresses.

II. Classification phase

We next process that same trace again and ensure that each packet proceeds through the classification pipeline shown in Figure 5.6. Once a packet can be attributed to an application, no further processing will be done for that packet.

Step 1: Payload signature matching. We match our previously extracted application signatures on each packet. Just by matching application signatures, we are able to classify 11.7% of the bytes exchanged at our IXP. This unexpected high number (recall that application signatures typically occur only in the first packets of a flow) is mainly the result of a proliferation of UDP-based BitTorrent data transfers, i.e., μTP [14]. μTP is a transport protocol based on UDP and includes its own header in every single packet. Thus, its classification is robust to sampling – in stark contrast to TCP traffic. The proliferation of μTP has also been reported in earlier studies [117, 178] as well as the rise of UDP-based applications using own headers in every packet [118]. In total, 11.3% of the packets matched a signature, of which 84.5% matched the BitTorrent UDP signature, another 11.7% matched an HTTP signature, 2.4% an SSL handshake on port 443, 0.94% a BitTorrent TCP signature, and 0.46% other signatures.

Step 2: Server endpoint matching. If a packet does not contain a valid application signature, we then check if the source or the destination (IP, port) tuple of the respective packet is a known server endpoint, as identified in our pre-classification step. If so, we classify the packet as belonging to the specific application. In this step, we classify 66.5% of bytes! This result highlights the efficiency of using a stateful application characterization approach. While we cannot sample application signatures on a per-flow basis, aggregating the information on a per (IP, port) endpoint basis largely overcomes the challenge posed by packet sampling. At the same time, we achieve a high confidence by relying on strong payload-based classification. This method works particularly well for popular client-server based applications, most prominently HTTP, where a large number of connections is destined to a comparably small number of server endpoints. To assess the impact of possibly stale endpoints (e.g., hosts that do not run the classified application on their server endpoint after some time), we repeated the classification by only using server endpoints that were identified within a time frame of 24 (12) hours, which reduced our completeness by only 1% (2%) of the bytes.

Step 3: Port-based classification. We next use a short list of 15 known port numbers (mapping to 13 applications) to classify respective packets as belonging to the corresponding application. In this step, we classify another 4.5% of all bytes. The largest contributor to this third step is RTMP [135] (1.7%), for which no reliable signature is available. Interestingly, a significant fraction of traffic on port 1935 (RTMP) is HTTP traffic (and was thus already classified in the previous step), likely RTMP-inside-HTTP [15]. Generally, we note that port-based classification can still be used reliably (but is not necessarily complete) when used in a conservative fashion, confirming prior studies [189]. For example, we observe that only less than 0.3% of the TCP traffic on port 80 did not match an endpoint which was detected using HTTP signatures (in the pre-classification). However, we find that more than 10% of the total HTTP traffic is not seen on port 80, and the most popular encountered non-standard ports are 8080 (3.8% of HTTP traffic), 1935 (2.9% of HTTP traffic) and 8000 (0.6% of HTTP traffic).

Step 4: Packet exchanged between BitTorrent speakers. In this step, we consider packets that were not classified in a prior step and classify them as "BT/P2P likely" if they are exchanged between two previously identified BitTorrent speakers (N=38.9M). This step enables us to classify an additional 10.2% of the IXP's traffic. Depending on the individual client's configuration and capabilities, BitTorrent relies on TCP and UDP as transport protocol for data exchange as well as for exchanging control messages (e.g., DHT queries). While we are able to classify the bulk of BitTorrent UDP traffic (recall that we classified more than 11% of the traffic just using signatures), we are not able to classify the bulk of TCP traffic exchanged between BitTorrent speakers. In this step we account for this portion of the traffic. To provide further empirical support for this approach, we inspected partly sampled TCP messages of the peer-wire protocol [3] which corresponds to the transfer of *chunks*. By extrapolating the number of *piece* messages of the BitTorrent peer-wire protocol and multiplying it with the observed chunk size (16K in 99% of all cases), we can estimate that the pure content volume (excluding headers and control traffic) exchanged via BitTorrent TCP peer-wire connections is around 8%. Thus, it follows that the majority of the traffic classified in this step, i.e., traffic exchanged between identified BT speakers, is indeed very likely BitTorrent traffic. To acknowledge the lowered confidence and the possibility of other protocols contributing to this class, we classify these packets as "BT/P2P likely".

Step 5: Packet exchanged between Web clients or BitTorrent speakers. As a tie-breaking criteria, we classify all packets that are exchanged between either Web clients or BT speakers (N=64.2M) as "P2P

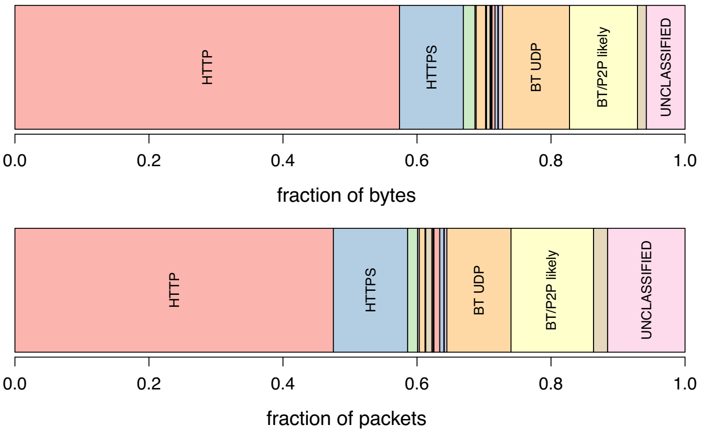

Protocol	% bytes	% pkts
HTTP	57.39	47.52
HTTPS	9.53	11.08
RTMP	1.72	1.48
NNTP	1.41	0.87
NNTPS	0.63	0.38
SMTP	0.53	0.91
DNS	0.45	0.87
SSH	0.42	0.61
other known	0.68	0.74
BT UDP	10.00	9.57
BT/P2P likely	10.14	12.31
P2P likely	1.32	2.10
unclassified	5.78	11.56

Figure 5.7: Application mix (September 2013) for packets and bytes.

likely”. We hence extend the set of hosts that are likely end users to not only consist of BT speakers, but also IP addresses that fetch content from HTTP(S) servers. The underlying assumption here is that end hosts might run other P2P applications, but not BitTorrent. Traffic exchanged between such end hosts was consequently not classified as P2P traffic in the previous step. We only classify another 1.3% of the IXP’s total traffic by using this heuristic. This small number suggests that most P2P likely traffic is indeed exchanged between BitTorrent speakers and was already classified in the previous step.

Using this classification approach, we are able to attribute 82.7% of the IXP’s overall traffic directly to its corresponding application (Steps 1-3). More than 78% of the traffic can be classified either directly using payload signatures or by matching the packet to server endpoints identified using payload signatures – we only fall back to port-based classification for 4.5% of the traffic. Another 11.5% of the traffic is classified as “BT/P2P likely” using our heuristics based on the social behavior of hosts.

5.3.3 Aggregate Application Mix at an IXP

In this section, we discuss properties of the observed application mix. Figure 5.7 shows the result of our classification method when applied to the IXP’s traffic, both in terms of packets and bytes (flow statistics are not obtainable from our packet-sampled traces). We observe that HTTP(S) clearly dominates the application mix with a share of more than 65% of the bytes. While the increasing dominance of HTTP for a multitude of applications has been reported in prior studies (e.g., [219]), the other significant share of traffic is composed of the BitTorrent UDP and BT/P2P likely class, accounting for some 20% of the exchanged bytes. Other protocols such as email, newsgroups, RTMP etc. account for roughly 6% of the bytes exchanged at the IXP.

Figure 5.8(a) shows a timeseries of the contributions of the various applications for the 09-2013 trace. While we see that HTTP(S) always dominates (its share never drops below 55%), we observe a typical diurnal pattern indicating more pronounced HTTP(S) usage in the busy hour in the late afternoon. The share of BitTorrent/P2P peaks in the off-hours. Interestingly, we observe a second peak of BT/P2P activity each day, which is likely due to BitTorrent users in various time zones. Also the protocols in the “other known” category dominate in the off-hours. NNTP(S) is the largest contributor to this category and is reportedly used for file-sharing [189].

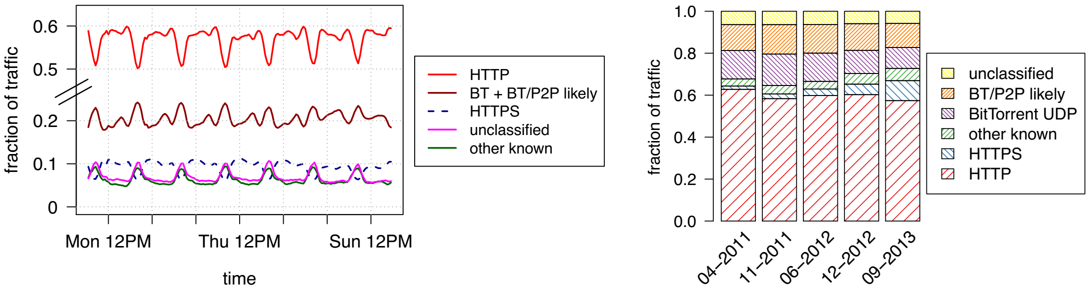

(a) Weekly timeseries of the application mix.

(b) Evolution of the application mix over the last 2.5 years.

Figure 5.8: Application mix over time.

Study	Network Type	Method	Year	HTTP(S)	Bytes other known	BT/P2P	unclassified
[173]	5 large ISPs (peerings, Global)	payload-based	2009	52.1%	24%	18.3%	5.5%
[173]	110 Networks (peerings, Global)	port-based	2009	52%	10%	1%	37%
[189]	Large ISP (access, Europe)	payload-based	2009	57.6%	23.5%	13.5%	10.6%
[125]	Large ISP (backbone, US)	payload-based	2010	60%	28%	12%	N/A
[96]	260 Networks (peerings, Global)	port-based	2013	69.2%	4%	<7%	20%
[10]	Various (N/A, North America)	payload-based [11]	2014	≈ 70%	N/A	6%	N/A
[10]	Various (N/A, Europe)	payload-based [11]	2014	≈ 65%	N/A	15%	N/A
[10]	Various (N/A, Asia-Pacific)	payload-based [11]	2014	≈ 60%	N/A	30%	N/A
[10]	Various (N/A, Latin America)	payload-based [11]	2014	≈ 65%	N/A	9.4%	N/A

Table 5.4: Reported application mix in other studies (fixed, IPv4).

Next, we use five snapshots to infer the application mix as observed at this IXP during the last 2.5 years. The results for the exchanged bytes are shown in Figure 5.8(b). We observe that while the IXP's aggregate application mix is relatively stable, there is a significant increase of HTTPS traffic during these 2.5 years, from 1.9% in April 2011 to 11.1% in September 2013. Note that while in the snapshots from November 2011 to December 2012, both the share of HTTPS and HTTP traffic increased, there is a simultaneous decrease in HTTP and steep increase in HTTPS in 2013, suggesting significant switchover from HTTP to HTTPS in 2013.

5.3.4 The Application Mix: A Moving Target

The Aggregate View

The Internet's application mix has been the topic of numerous past studies by networking researchers and commercial companies alike. In the following, we report how the observed application mix at our IXP compares to other recent studies that not only relied on traffic data from different vantage points (and hence different types of peering links) but also used different application classification methods. Recall that in this study, we are only considering traffic that traverses the IXP's public peering links and have no visibility into the traffic that is sent over the private peering links established at this IXP. Table 5.4 lists some of the pertinent prior studies and provides information about the reported application mix, the type of traffic data used, and (where available) the classification method used.[5] A cursory comparison of the results of these studies with our findings suggests that the application mix of the

[5] Note that the applications belonging to the "other known" traffic class vary across studies.

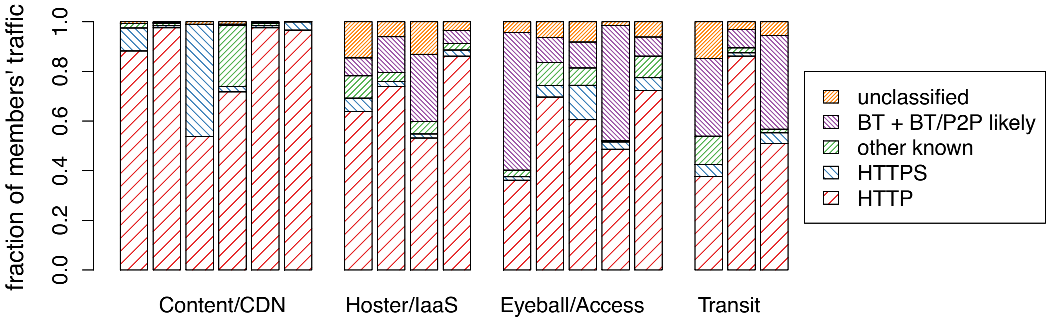

Figure 5.9: Application mix of the top 15 traffic-contributing member ASes grouped by business type, and the three most traffic-contributing transit ASes.

Internet is rather homogeneous. That is, HTTP(S) dominates with a share of roughly 60%, no matter where in the network and with what methodology the application mix was measured. Other protocols such as BitTorrent or P2P seem to vary by region from around 10% to 30%, but these variations could also be in part due to varying classification approaches.

Beyond the Aggregate Application Mix

Next, we take a closer look at the apparent homogeneous nature of the Internet's application mix and examine in detail the application mix of the traffic that traverses the peering links of specific networks.

Figure 5.9 shows the application mix for each of the top-15 traffic-contributing member ASes of our IXP and top-3 traffic-carrying transit providers that are also IXP members. The type of the top-15 traffic-contributing IXP members is either *Content/CDN*, *Hoster/IaaS* or *Eyeball/Access*, and together they are responsible for 59% of the all the traffic (in bytes) seen at this IXP.[6] We see that for all networks of type *Content/CDN* HTTP(S) traffic clearly dominates, with shares close to 100%. While most of these networks still rely mainly on HTTP, we notice one prominent network (third bar from the left) that has almost a 50/50 ratio of HTTP and HTTPS traffic. This example suggests that the earlier reported growth in HTTPS is mainly driven by some big content providers switching over to HTTPS. Overall, for networks of type *Content/CDN* we observe little or no application-mix heterogeneity on their individual links. Networks of the type *Hoster/IaaS* show a more diverse profile when it comes to their application mix. While HTTP still dominates, we see surprisingly no significant amount of HTTPS traffic. At the same time, these networks also exchange other types of traffic of various protocols as well as significant shares of BitTorrent traffic and unclassified traffic. Note that BitTorrent is also increasingly used to deliver video content or software [107]. In short, the diverse application mix contributed by Web hosters reflects the fact that they offer infrastructure services to a wide variety of companies and individuals, which in turn make different use of the provided resources. The results for *Eyeball/Access* networks show that the application mix of networks connecting end-users to the IXP also varies significantly. While for some of them, HTTP(S) (along with small fractions of other traffic such as email, RTMP, news) clearly dominates, we also see eyeball networks with more than 50% of BitTorrent traffic—the two networks with significant BitTorrent contributions are serving eastern European countries, while the other three networks are serving users in central Europe. This observation suggests that the differences in BitTorrent usage also reflect geography (i.e., varying application popularity). The application mix

[6] We determined the types of the IXP members using manual classification.

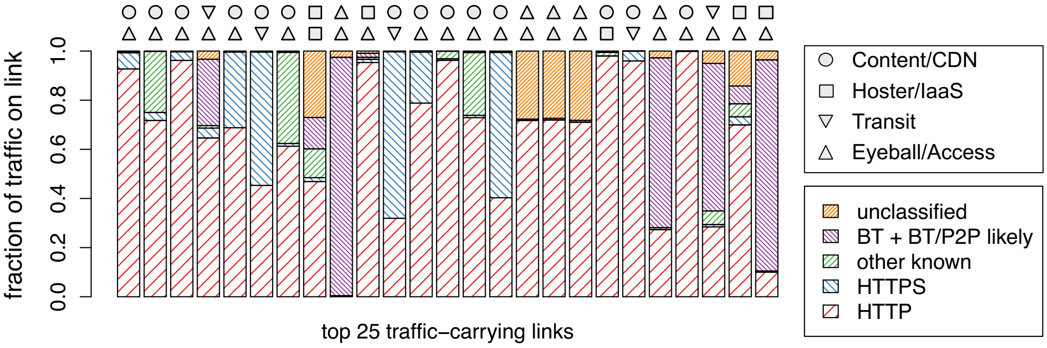

Figure 5.10: Application mix of the top 25 traffic-carrying links.

seen for *Transit* networks is in general quite diverse, as they typically carry traffic from a wide range of different networks.

The picture of the Internet's application mix sharpens even more when we look at the application mix seen on individual peering links. Figure 5.10 shows the application mix for the top-25 traffic-carrying bidirectional links at our IXP. The figure also includes the business types of the networks on either side of these peerings. Based on this set of links, which see significant traffic, we observe a variety of different application mixes. While all Content-to-Eyeball links carry exclusively HTTP(S) and few other known applications, BitTorrent is the clear winner on two links between Eyeball networks. Thus, when taking into account the business types of two networks associated with a peering link, we notice a strong dependency on the resulting applications mix. The few links that show a more heterogeneous application mix are usually transit links or, interestingly, links involving Hosters and IaaS providers. When looking at the top-25 unidirectional links (not shown), we see a similar pattern, where for Content-to-Eyeball links the resulting application homogeneity (i.e., HTTP) is even more dominant.

5.3.5 Summary

We develop a traffic classification methodology that is by and large able to overcome the challenges posed by packet-sampled traffic through the use of a *two-pass* classification approach based on endpoint-aggregation. Using our new methodology we can attribute more than 78% of the bytes exchanged over the public switching infrastructure of a large IXP to their respective applications by relying on strong payload-based classification.

We attribute another 11.5% when including a heuristic based on communication patterns and classify an additional 4.5% using port-based classification. In the process, we observe that the aggregate application mix as seen at our IXP is largely consistent with that reported in other recent studies. However, when dissecting the traffic and examining the application mix of Internet traffic that traverses individual public peering links, we show that the application mix becomes heterogeneous but is strongly influenced by the business types of the networks on either side of a peering link.

Our measurements of application mix heterogeneity across peering links highlight that shifting traffic from IPv4 over to IPv6 will not just be a matter of providing connectivity. Individual applications need to be adapted in order to support and allow data exchange over IPv6, instead of IPv4 [61, 259, 282]. Given the heterogeneity of the application mix, we can expect a disparate potential for IPv6 adoption and traffic for individual peering links, depending on the involved networks and the applications they run or carry.

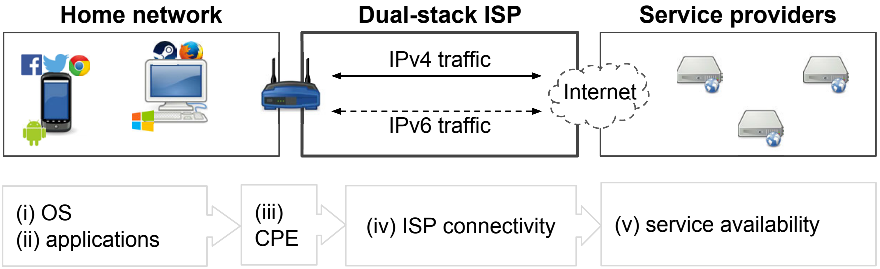

Figure 5.11: IPv6 traffic in dual-stack networks. Barriers are present at home networks (operating systems, applications and CPEs), ISPs (offered DSL connectivity), and at service providers.

5.4 Intra-Domain: Barriers and Intent for IPv6 Traffic

We have found a disparity between IPv6 connectivity and traffic, and we have seen that the application mix exchanged on individual peering links is highly heterogeneous. We next seek to understand in detail the interplay between available connectivity and the actual exchange of traffic over either IPv4 or IPv6. To accomplish this, we study this problem from the perspective of 12.9K subscribers of a dual-stack ISP. This vantage point gives us a unique opportunity to analyze both the connectivity that is available to the subscribers and the service providers, and the actual traffic exchange over either IPv4 and IPv6. We focus here on protocols that use DNS as a rendezvous mechanism, since DNS queries and answers allow us to infer the presence, or the request for, IPv6 capability as well as the subsequent data exchange.

To exchange data over IPv6, all components on the path from a source to a destination need to fully support IPv6 (see Figure 5.11). This includes *(i)* end-user devices and operating systems supporting IPv6, *(ii)* applications making proper use of the available connectivity options (see [267]), *(iii)* customer premises hardware (CPEs) supporting and providing IPv6 to the home network [47, 260], *(iv)* the ISP assigning IPv6 addresses to the subscribers CPEs [102], and finally *(v)* content providers enabling their services over IPv6 [192]. Moreover, even if all of the above conditions apply, i.e., all components *support* IPv6, a second dimension of the problem is whether IPv6 will be preferred over IPv4, as modern applications employ a technique named "*happy eyeballs*" to *choose* between IPv4 and IPv6 according to the current network conditions [277].

A client that uses happy eyeballs potentially initiates two TCP connections: One to the IPv4 endpoint of the requested service, another one to the IPv6 endpoint. The client will then choose the protocol that completed the TCP handshake first, yet typically giving preference to IPv6 by delaying the IPv4 connection initiation (e.g., Firefox and Chrome only initiate an IPv4 connection if the IPv6 connection attempt did not complete after 300ms [277]).

Our main findings can be summarized as follows:

(i) Even though this ISP supports IPv6 connectivity, a large number of subscribers can not *use* IPv6. We find that in some cases the ISP does not provide IPv6 connectivity to its subscribers. More often, however, the CPE router limits IPv6 connectivity.

(ii) Consequently, IPv6-ready services exchange a significant amount of traffic over IPv4. IPv4-only speaking devices and fallback mechanisms further increase the share of IPv4 traffic for these services. On the other end of the spectrum, we observe a strong *intent* for IPv6 traffic, i.e., clients request a significant share of content to be carried over IPv6. However, many service providers only provide IPv4 and hence cannot fulfill these requests over IPv6.

(iii) Due to dual-stack applications' preference for IPv6, dual-stack networks could face a rapid and substantial increase of the IPv6 traffic share if even just a few major service providers enable IPv6 for high-traffic domains.

The rest of this section is organized as follows: We describe our methodology in Section 5.4.1 and introduce our dataset in Section 5.4.2. Section 5.4.3 presents our findings on the interplay between connectivity and traffic. We discuss limitations of our approach and vantage point in Section 5.4.4.

5.4.1 Methodology

The focus of our study is the traffic at a residential broadband network of a dual-stack ISP. As shown in Figure 5.11, IPv4 and IPv6 traffic coexist at such a vantage point. Whether IPv4 or IPv6 is used depends on a large variety of factors, as mentioned above.Hence, a dual-stack ISP presents a unique opportunity to study the interactions of this ecosystem and its influence on the share of IPv6 traffic. To this end, we first need to discover the connectivity options of the two engaged parties, i.e., the subscribers (the client side) and the service providers (the server side). With this information in hand we can proceed to study which traffic is exchanged over which protocol, and why.

Measuring IPv6 Connectivity

Connectivity of subscribers ("client side"). Broadband network providers typically rely on Remote Authentication Dial-In User Service (RADIUS [237]) to assign IP addresses to subscribers. With this protocol, CPEs obtain IP addresses, usually a single IPv4 address that multiplexes devices (NAT). This protocol specification also supports the delegation of IPv6 addresses to subscribers [22, 102, 254]. If the CPE receives an IPv6 prefix assignment, we say that the subscriber obtains IPv6 connectivity from the ISP. Traffic statistics later tell us whether the subscriber's devices make actual use of this assigned IPv6 prefix.

Since not all devices within home networks support IPv6, the raw traffic statistics are necessary but not sufficient to infer if a device within a subscriber's premise can use IPv6. We use `AAAA` DNS requests as an indicator for the presence of IPv6-speaking devices. Most dual-stack applications follow the *happy-eyeballs* proposed standard (see [277]), and issue `A` as well as `AAAA` DNS requests. If the requested service is available over IPv6, the device attempts to connect simultaneously to two addresses contained in the DNS resource records (`RRs`); one being IPv6 and the other IPv4. An application that adheres to the example implementation then establishes two TCP connections and uses the one that completed the handshake faster. Some implementations introduce a preference towards IPv6. For example, Apple devices issue an IPv6 connection immediately after a successful `AAAA` request if the `A` response did not arrive already, or if historical RTT data suggests a difference > 25 ms [256]. Given that most DNS clients issue `AAAA` requests first [195], some dual-stack devices do not always attempt a connection over both IPv4 and IPv6 although they issue requests for both `RRs`.

One important fact regarding IPv6-speaking devices is that many resolver libraries still issue `AAAA` requests, even in the absence of global IPv6 connectivity [4]. Thus, we can use this information (i.e., `AAAA` requests from subscribers without IPv6 connectivity) to further identify CPEs that offer link-local IPv6 connectivity even if the ISP does not provide IPv6 connectivity to them.

Connectivity of services ("server side"). In this work we use the term "service" to refer to content and functionality that is available on the Internet via a *Fully-Qualified Domain Name* (FQDN). For example, at `www.google.com` we can find a search service as well as plain content. If the network infrastructure that hosts a service supports IPv6, a service provider willing to make its services available over IPv6 just needs to update the corresponding DNS AAAA and possibly PTR resource records (RRs) [192]. Henceforth, we can analyze DNS traffic to infer if a service is IPv6-ready by looking for non-empty AAAA responses in our traces. However, as we may not be able to observe all AAAA RRs (e.g., if the clients are not IPv6 enabled), we complement passive data with active measurements, i.e., we actively request A and AAAA records for FQDNs found in our trace.[7]

From IPv6 Connectivity to IPv6 Usage

Now that we are aware of the *connectivity* options of subscribers and services (IPv4 and/or IPv6), we proceed to study the exchanged traffic. To accomplish this, we first need to annotate each flow in our trace with the respective subscriber and service.

Matching flows to names. One of the building blocks for our methodology is the ability to associate the DNS requests issued by an IP address to the network flows it generates, i.e., reproduce the mapping between FQDNs and server IPs for each subscriber. This problem has been already explored (see, e.g., [69, 194, 216]), and we extend it to include the connectivity information. We note that in the case of dual-stack networks, the IP addresses of the flows and those of the DNS traffic are not necessarily the same. Therefore, we cannot directly use the source IP of a DNS request as a rendezvous. Instead, we keep track of the IPv4 and IPv6 addresses assigned to each subscriber. In addition, we need to update this mapping according to the TTL values of the DNS response RRs. We are aware that related studies have reported violations of the TTL field by clients [82, 194]. For example, Callahan et al. [82] observe that 13% of the TCP connections use expired records and attribute it to security features present in modern Web browsers. In this work, we opt for a conservative approach and strictly use the TTL expiration values. In addition, we do not consider negatively cached responses, e.g., a service without a AAAA RR. Our rationale is that although negative answers should, in principle, be cached according to the SOA record [36], some resolvers do not respect this [177]. At times, we will consequently not observe a AAAA request for services without AAAA RR and may mis-attribute it to a device that does not support IPv6.

Annotating flows. We next annotate each flow with the following information: *(i)* whether the ISP has delegated an IPv6 prefix to the subscriber's CPE, *(ii)* the FQDN associated with the flow, where possible, and *(iii)* if the subscriber issued an A and/or a AAAA DNS request. After collecting the trace we extend this annotation with the following information: *(iv)* if the subscriber makes use of its assigned IPv6 prefix at all, and with *(v)* the connectivity options for the FQDN i.e., whether the service is available over IPv4 and/or IPv6.

5.4.2 Dataset

The dataset used throughout this study covers all IP traffic generated by 12.9K DSL subscribers of a residential broadband network during a period of 45 hours in winter 2015–2016. We implemented a custom tool built on top of the *libtrace* library [32] to produce two streams of data from raw network data. The first stream consists of packet summaries, including packet size, SRC and DST IP addresses, and port numbers. For TCP packets, we also save TCP flags, SEQ, and ACK numbers. The second stream consists of full-sized packets of DNS traffic (UDP port 53). We then process our packet summaries to obtain flow-level statistics. Namely, we aggregate the packet summaries into the 5-tuple and expire inactive flows after 3,600s. For TCP flows we also compute the time difference between the SYN packet

[7] We conducted these additional measurements shortly after the data collection.

Trace	#bytes	#flows
TCP$_{v4}$	80.5%	53.1%
TCP$_{v6}$	10.7%	4.7%
UDP$_{v4}$	7.4%	18.2%
UDP$_{v6}$	1.1%	21.7%
total	64.5T	356.2M

Table 5.5: Total traffic over IPv4/IPv6 and TCP/UDP.

Service Side	**Subscriber Side**			**total**
	IPv4-only	*IPv6-inactive*	*IPv6-active*	
IPv4-only	5.4%	20.1%	22.4%	47.9%
IPv6-ready	3.2%	9.2%	15.4%	27.8%
IPv6-only	0.0%	0.0%	< 0.1%	< 0.1%
Unknown	3.4%	8.8%	12.1%	24.2%
total	11.9%	38.1%	49.8%	100%

Table 5.6: Traffic contribution (sum of IPv4 and IPv6) partitioned by the state of IPv4/IPv6 connectivity of subscribers and service providers.

and the `SYN ACK` packet to estimate TCP handshake times.[8] Given the location of our monitor within the aggregation network, these "handshakes" only capture the wide-area delays (backbone RTTs) and do not include delays introduced by the access- and home network (see [189] for details on the technique). Finally, we remark that the dataset was collected, processed, and analyzed at an isolated and secured segment infrastructure of the ISP. The toolset operates in an automated fashion and anonymizes line IDs and addresses before writing the annotated flows to the disk. Table 5.5 summarizes the dataset collected for this study.

DNS transactions. We processed 141.9M DNS transactions, where we denote a transaction as an `A` or a `AAAA` request with a valid DNS response. 69.6% of these entries are of type `A` and 30.4% of type `AAAA`. Out of these DNS transactions, 0.6% and 36.0% of the `A` and—respectively—`AAAA` requests could not be resolved (NXDOMAIN). The high ratio of unresolved `AAAA` requests is the result of content that is indeed requested for IPv6, but still not accessible over IPv6 (see §5.1). 39% of the `A` requests were sent over IPv6, and 28% of the `AAAA` requests over IPv4. Additionally, we actively queried `A` and `AAAA` records for all FQDNS found in the trace. In total, we successfully queried 1.34M FQDNs of which 1.17M had only an `A` record, 169K had both `A` and `AAAA` records, and 474 only had a `AAAA` record, but no `A` record.

Flow-level statistics. Table 5.5 shows a breakdown of the contribution of TCP and UDP traffic, dissected by IP version. Unsurprisingly, TCP$_{v4}$ dominates in terms of traffic volume. However, the share of IP$_{v6}$ is substantial (11.9%) especially when compared to older measurement studies at other vantage points [96,255]. Web traffic sums up to 86.6% of the trace volume (13.5% over IPv6).[9] We find that QUIC contributes 2.8% of the overall trace volume (39.5% over IPv6). Considering the relative UDP contributions over IPv4 and IPv6, we see that the share of UDP$_{v6}$ flows is well above the UDP$_{v4}$ share. A closer look reveals that this bias is introduced by DNS traffic: DNS accounts for 71.0% of all UDP flows and 75.3% of DNS flows are sent over IPv6.

Classification coverage. We are able to associate up to 76.1% of the traffic to services using the flow-classification approach described in §5.4.1. While our coverage statistics are consistent with the base results reported in [194], we remark that ours are lower than related methods because our method *i)* does not use a warm-up period to account for already cached DNS `RR`s, *ii)* relies on each subscriber's own DNS traffic, and *iii)* adheres to the TTL values included in DNS responses.

[8] We exclude flows with retransmissions of packets with the `SYN` flag set.

[9] TCP traffic on ports 80 and 8080 (HTTP), 443 (HTTPS), and UDP traffic on port 443 (QUIC).

5.4.3 A Dual-stack ISP Perspective on IPv6 traffic

The Subscriber Side

We find three classes of DSLs among the 12.9K subscriber lines of this vantage point: *i) IPv4-only*: lines that do not get IPv6 connectivity from the ISP (17.3%), *ii) IPv6-inactive*: lines provisioned with IPv6 connectivity but no IPv6 traffic (29.9%), and *iii) IPv6-active:* lines with IPv6 connectivity as well as IPv6 traffic (52.9%).

***IPv4-only* subscribers.** This set of lines corresponds to subscribers for which the ISP has still not activated IPv6 connectivity (e.g., old contracts). They contribute 12.0% to the overall trace volume. 26.6% of their traffic is exchanged with services that are available over IPv6. We notice that some devices issue `AAAA` DNS requests, most likely because some CPEs create a link-local IPv6 network. In fact, for 11.6% of the traffic related to IPv6 services we observe a `AAAA` request. This first observation is relevant for *IPv6-adoption* studies, as it indicates that in some cases DNS traffic may not well reflect the actual connectivity. This shows that many devices are already prepared to use IPv6 connectivity, waiting for the ISP to take proper action.

***IPv6-inactive* subscribers.** For 36.1% of the DSLs we do not observe any IPv6 traffic, even though the ISP assigned IPv6 prefixes to the CPEs. One explanation is that the CPE has not been configured to enable IPv6 on the home network (see e.g., [110, 133, 268]). Thus, the ISP provides IPv6 connectivity, but the end-devices only have internal IPv4 addresses (e.g., RFC1918), assigned from the CPE. Consequently, we find that only 1.7% of the traffic from these subscribers can be associated with a `AAAA` request, likely because most devices suppress `AAAA` requests in the absence of a link-local IPv6 address. Other, less likely, explanations are that none of the devices present at premises during the trace collection support IPv6 (e.g., Windows XP), or the subscribers do not contact services available over IPv6. The latter is unlikely, as 24.1% of the traffic in this subscriber class is exchanged with IPv6-ready services.

***IPv6-active* subscribers.** Subscribers in this category actively use the provided IPv6 connectivity. The share of IPv6 traffic out of their total traffic for these subscribers is almost twice as high (21.5%) when compared to the overall trace (11.9%). When only considering traffic exchanged between IPv6-active subscribers and services that are indeed available over IPv6, the ratio is even higher (69.6%). Yet, that leaves us with 30% of the traffic exchanged between two IPv6-enabled hosts being carried over IPv4. This can be caused either by end-user devices not requesting content over IPv6 (no `AAAA RR`) or end-user devices choosing IPv4 over IPv6 because of their happy eyeball implementation. Indeed, when only considering traffic for which the client requested both IPv4 and IPv6 (`A` and `AAAA`), the share of IPv6 in this category raises up to 85.1%. This is an important observation for service providers and operators, as it implies that enabling IPv6 can increase the share of IPv6 traffic from/in dual-stack networks rapidly.

The Service Provider Side

We next shift our focus from subscribers to services (FQDNs). Similar to the previous section, we define three categories. We say that a service is *IPv4-only* if it only has a valid non-empty `A RR`. *IPv6-only* services are those which only have a valid non-empty `AAAA RR`. A service that is *IPv6-ready* has valid and non-empty `A` and `AAAA RR`s. We report in Table 5.6 how these three categories of services contribute to the total traffic and intersect them with the three subscriber categories.

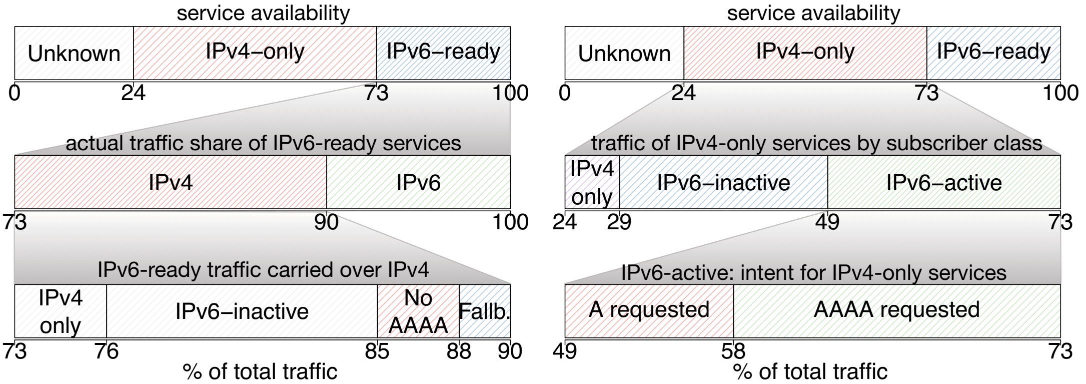

(a) **IPv6 barriers.** Top: service availability. Center: IP version that carries *IPv6-ready* content. Bottom: Reason why traffic is carried over IPv4 instead of IPv6.

(b) **IPv6 intent.** Top: service availability. Center: Breakdown of *IPv4-only* traffic by subscribers' type. Bottom: traffic from *IPv6-active* subscribers to *IPv4-only* services.

Figure 5.12: Barriers and intent for IPv6 traffic in a dual-stack ISP.

***IPv4-only* services (only A RR).** As expected, this set of services dominates the share of traffic (47.9%). However, for 36.2% of this traffic we observe a preceding AAAA request from the subscriber requesting the content, which implies that this traffic has the potential to be served over IPv6 if the corresponding service providers enable IPv6.

***IPv6-only* services (only AAAA RR).** We find around 500 services that *appear to be* available only over IPv6, accounting for less than 0.1% of the traffic. Manual inspection reveals that most of them are mere connectivity checkers. Some service providers add strings to hostnames, which may appear as an IPv6-only service (e.g., both *host.domain.org* and *hostv6.domain.org* have a AAAA RR, but only the former has an A RR).

***IPv6-ready* services (A and AAAA RRs).** These services generate a significant amount of traffic (27.8%). However, as many subscribers from this dual-stack network cannot use IPv6, the actual share of IPv6 traffic within this class of services is only 38.6%.

IP traffic: Barriers and Intent for IPv6

As shown in Table 5.6, the upper bounds for IPv6 traffic share when looking at services and subscribers independently is roughly 2 and respectively 4 times the actual IPv6 traffic share. At the same time, not all traffic in the cross-product of *IPv6-active* subscribers and *IPv6-ready* services is carried over IPv6. We next proceed to study the root causes that lead to this lower-than-possible IPv6 share. To this end, we use the term *IPv6 barriers* to reason about traffic to and from IPv6-ready services, which is carried over IPv4 instead of IPv6. Correspondingly, we use the term *IPv6 intent* to reason about traffic to and from IPv4-only services, of which some portion could be carried over IPv6, as requested by the subscribers.

IPv6 barriers. Figure 5.12(a) illustrates why traffic related to *IPv6-ready* services is exchanged over IPv4. On the top of the figure we show a bar summarizing all traffic in the trace according to the service availability. As previously stated, 27.8% of the traffic relates to services available over IPv6. Nevertheless, the majority of it (61.4%) is actually exchanged over IPv4 (see middle bar). In the bottom bar we illustrate why data is exchanged over IPv4 instead of IPv6. Most of this traffic (70.5%) is carried over IPv4 because the subscribers do not use IPv6 connectivity at all (*IPv4-only* and *IPv6-inactive*). We

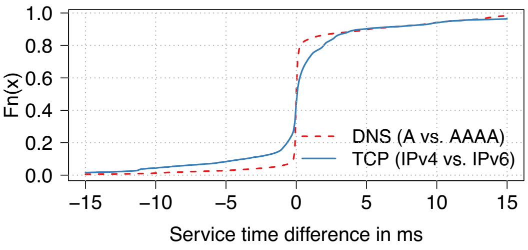

Figure 5.13: ECDF: Differences between IPv6 and IPv4 TCP handshake and DNS resolution times per hostname. Positive values indicate longer transactions for IPv6 and `AAAA` RRs.

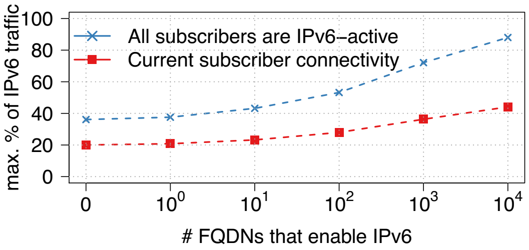

Figure 5.14: Estimation of the maximum *possible* share of IPv6 traffic when IPv4-only FQDNs enable IPv6. We sort FQDNs by their contribution in terms of bytes.

make two observations for the remainder of this traffic (which is generated by *IPv6-active* subscribers). The majority of it has no associated `AAAA` request, which can primarily be attributed to end-devices that do not support IPv6: they do not issue `AAAA` requests. For another 40% of the IPv4 traffic from *IPv6-active* subscribers to *IPv6-ready* services we observe a `AAAA` request. These are likely flows generated by devices that fall back to IPv4 as a result of the *happy-eyeballs* algorithm.

IPv6 intent. Figure 5.12(b) illustrates what fraction of the traffic of *IPv4-only* services (top bar) could be carried over IPv6. While the bar in the middle depicts how much of this traffic they exchange with each subscriber category, the bottom bar shows the traffic characteristics for the *IPv6-active* subscribers. In particular, we observe that end-user devices in the *IPv6-active* group issue `AAAA` requests for 62.5% of this traffic. Thus, there is a strong intent for IPv6 traffic that cannot yet be satisfied by the service side. In fact, our measurement likely even underestimates this value because we do not take into account negatively-cached `AAAA` RRs (see §5.4.1).

Happy eyeballs. Given that part of the traffic carried over IPv4, which could be carried over IPv6, can be attributed to (un-)happy eyeballs, we now study two metrics concerning dual-stack applications and devices, i.e., the RTT estimates and the DNS resolution times (see [256]). Our RTT estimate corresponds to the backbone RTTs (§5.3.1). For the DNS resolution time (`A` vs. `AAAA`), we only consider transactions with non-empty responses and for which we find just one request and one response in the same UDP flow. We aggregate these per hostname and compute the median only for those hostnames with at least 10 samples. Generally, dual-stack services offer similar conditions, i.e., around 80% of the values are within a range of 10 ms. Under such conditions, happy-eyeball implementations likely select IPv6, as indicated by our earlier results. This observation is important for service providers transitioning to IPv6, as it implies that after enabling IPv6 they can expect a significant increase of IPv6 traffic if they already exchange high volumes of data with dual-stack consumer networks. We note that the final *choice* of connectivity is subject to how different implementations adapt to network conditions [20, 143, 144], i.e., by delaying the initiation of the IPv4 TCP connection by different thresholds (e.g., Chrome and Firefox delay the potential initiation of the IPv4 connection by 300ms) [277].

Case Studies

We next describe two case studies: a large search provider and a large CDN. Our case studies illustrate two opposite facets of the transition to IPv6. These providers together contribute to 35.7% of the overall and 73.1% of the IPv6 traffic. They both operate various AS numbers as well as caches inside ISPs. To identify their traffic, we rely on the origin ASN as derived from the IP addresses in the flows. To identify traffic from caches, we obtain a list of the Fully Qualified Domain Names (FQDNs) associated with IP addresses managed by these ASNs, derived from DNS PTR records.

A large search provider. Our first case study is a service provider that actively supports and promotes IPv6. 37.6% of its traffic is IPv6, and it alone contributes 69.9% of all IPv6 traffic in the trace. After annotating 91.8% of the traffic with FQDNs, we corroborate that almost all content—not all traffic relates to search services—requested by users at this vantage point is available over IPv6 (98.7%). *IPv4-only* and *IPv6-inactive* subscribers generate 74.1% of the IPv4 traffic while the share of IPv6 traffic for the *IPv6-active* subscribers is 70.5%. This observation suggests that for this provider the connectivity of the subscribers is the main obstacle for the increase in IPv6 traffic.

A large CDN. We are able to annotate 84.7% of the CDN traffic with FQDNs. Only 2.5% of the traffic is carried over IPv6, and only 3.3% of the CDN traffic relates to *IPv6-ready* services. This implies that here the bottleneck for IPv6 is the server side, since only 2.1% of the content requested with a AAAA is actually exchanged over IPv6.

Transition to IPv6. Service providers willing to transition to IPv6 need to update the corresponding DNS RRs. To illustrate the potential impact of this process on the share of IPv6 traffic, we next concentrate on *IPv4-only* services. We present in Figure 5.14 an upper bound for the share of IPv6 traffic when the top traffic-contributing FQDNs enable IPv6. We produce two estimates. The first one assumes that there are no changes in the subscribers connectivity. The second one assumes that all subscribers become *IPv6-active*. Note, we do not take into consideration 24.2% of the bytes in the trace as we cannot associate them with a service. Enabling IPv6 connectivity for all subscribers immediately doubles the upper bound for the IPv6 traffic share (to almost 40%). However, to reach IPv6 traffic shares close to 90%, more than 10K FQDNs need to enable IPv6 connectivity. That said, and as shown earlier in this paper, *IPv4-only* devices and *happy-eyeballs* fallbacks to IPv4 can reduce this share.

5.4.4 Discussion

We are well aware that our vantage point is not representative of the Internet as a whole. While this particular ISP promotes IPv6 connectivity, others opt to deploy Carrier Grade NATs to combat IPv4 address scarcity. Yet, we argue that our observations most likely apply—to varying degrees—to other dual-stack ISPs as well, since anecdotal reports from ISPs and operators report similar issues when it comes to shifting traffic to IPv6 (e.g., [133, 134, 268]). Hence, these observations can aid ISPs and service providers by providing guidance on how to provision for IPv6 as well as insights on traffic dynamics during the transition phase. For example, *IPv4-only* service providers could exchange up to 30% of their traffic over IPv6 if they enable IPv6. By contrast, although 53% of the IPv4 traffic to *IPv6-ready* services involves subscribers whose CPEs most likely do not provide IPv6 connectivity to their home network, *happy eyeballs* usually *chooses* IPv6 over IPv4 (85%). We posit that IPv6 traffic shares will likely be subject to sudden increases when CPE devices enable IPv6 support in the home network. Virtual CPEs [83] could make it easier for operators to transition their subscribers to IPv6 and troubleshoot IPv6-related problems. Hence, avenues for future work include a closer investigation of issues specific to devices and applications as well as a characterization of *happy-eyeballs* fallbacks to IPv4.

5.4.5 Summary

We study the interplay between connectivity and actual traffic exchange from 12.9K residential broadband users. We infer connectivity provided by the ISP to the subscribers using RADIUS data, and rely on DNS queries and replies to infer the request or non-request for IPv6 connectivity by the client as well as the availability of IPv6 connectivity by the requested service provider. Mapping DNS requests to the exchanged traffic, we gain a detailed picture of both the IPv6 capability in end hosts and service providers, and as well the resulting traffic exchange.

We reveal obstacles hampering IPv6 traffic in dual-stack ISPs, including CPE devices not supporting IPv6, applications falling back to IPv4, and a broad lack of IPv6 support among service providers. In spite of such obstacles, we report a pronounced increase, intent, and potential for growth regarding IPv6.

5.5 Chapter Summary

The Internet's transition to IPv6 presents a tremendous operational effort, since it requires far-reaching changes in the network itself as well as at the edge, i.e., changes in home devices, operating systems, and applications. At the same time, the heterogeneity of the network makes it difficult to draw an accurate picture of IPv6 adoption and to pinpoint the barriers for the ongoing IPv6 transition. In this chapter, we studied connectivity- and traffic-related aspects relevant for IPv6 adoption. Our findings have a number of implications for researchers, regulators, and network operators alike.

Our analysis of interconnectivity between ASes at our two IXPs shows pronounced IPv6 connectivity. IXP Route Servers make it easy for networks to immediately establish hundreds of peerings with other IXP members, requiring just a single BGP session with the route server. Hence, route servers can support IPv6 adoption, providing immediate connectivity. However, the ease with which multi-lateral peerings can be established also cautions against taking inter-domain peering links at face value to measure IPv6 adoption. Most of the identified IPv6 peerings carry very little traffic. Thus, while *connectivity* between networks over IPv6 exists, few of them exchange significant traffic volumes over IPv6. These findings caution both network operators and researchers tracking IPv6 adoption. Policies that encourage networks to receive IPv6 address allocations (and to use them, e.g., [249]) are not necessarily effective, since provisioning inter-domain connectivity is only the first step towards transition.

To gain a better understanding of the individual traffic components, we devise a methodology to determine the application mix in sampled traffic. Our classification of the prevalent applications of the exchanged traffic over IPv4 peering links shows pronounced heterogeneity across peering links. For traffic to be carried over IPv6, applications and operating systems need to be adapted. The large diversity in terms of traffic illuminates that shifting traffic from IPv4 over to IPv6 presents us with a much more complex problem beyond providing inter-domain connectivity. IPv6 adoption starts with the IPv6 capability of individual applications and operating systems. The shift of traffic from IPv4 and IPv6 will likely present us with an individual challenge for each individual peering link, and be dependent on the individual traffic components it carries. While we see a stark concentration of the application mix on Web traffic, we find a large number of application protocols, and it is unclear to what extent these applications are yet ready to support IPv6. For details, measurements and experiences with IPv6 compatibility of individual applications, we refer to [61, 259, 282].

Our analysis of traffic exchange over IPv4 and IPv6 in a dual-stack ISP sharpens the picture and reveals several *barriers* for IPv6 adoption. These barriers cause traffic—even if both the client and the server speak IPv6—to be carried over IPv4 instead of IPv6. In particular, we find that CPE routers are likely the cause for a significant portion of traffic that can only be carried over IPv4. An implication here is that IPv6 support (and possibly default settings) of CPE devices could almost immediately increase the share of IPv6 traffic in this ISP by a factor of two. On the other hand, we find that there is a strong *intent* for IPv6 (i.e., clients requesting content to be served over IPv6), suggesting that service providers that enable their content over IPv6 can expect immediate shifts of traffic from IPv4 over to IPv6. Making services available over IPv6 is hence not a mere additional offering, but corresponding hardware and interconnectivity must be properly provisioned prior to enabling IPv6.

6 Conclusion

The Internet is in the midst of its first fundamental disruption: The exhaustion of the IPv4 address space. We measure a widespread impact of IPv4 address exhaustion on the broader Internet. As of 2017, we observe increasing scarcity reflected in growing address markets, a stagnation of the number of active IPv4 addresses, widespread IPv4 gateway deployment, and ongoing IPv6 adoption. Networks make individual decisions on which mitigation strategy—or combination of strategies—to adopt, how to implement it, and when. The combination of individual decisions results in different ramifications for the Internet as a whole and for its stakeholders, including end users, ISPs, content providers, and governance bodies.

6.1 Summary

In the first part of this dissertation we studied the history and the interplay between management and governance decisions on the one hand, and their resulting impact on IPv4 address usage dynamics on the other. The size of the IPv4 addresses space was a design decision. When IPv4 was introduced, the early pioneers and designers of the Internet protocol could not foresee the growth and the disruptive power of the Internet that was to follow soon. Thus, an informed decision on how to dimension such a purely virtual resource space was impossible at that time. Scarcity was not seen as a looming issue, and hardware constraints, such as limited and expensive memory, led to a standardization of 32-bit long address fields in the IP header. Besides the "too few digits" problem, IPv4 addresses present an unprecedented case of a scarce virtual resource that requires truly global coordination. In the face of a rapidly growing Internet and a looming scarcity problem, processes, institutions, and policies to govern this virtual resource were established. Stricter address allocation policies have proven effective, but continued rapid growth of the Internet and the corresponding demand for IPv4 addresses lead to the exhaustion of IANA's address pool in 2011. In the subsequent years, four out of five of the Regional Internet Registries also exhausted their pools. From an allocation perspective, the IPv4 address space is now close to fully exhausted. However, the underlying issues when it comes to governing the address space prevail. As of 2017, we still lack a widely deployed technology that prevents networks from illegitimately routing and using address blocks that are not assigned to them. Moreover, the question of ownership rights for some IP address blocks is still not unanimously clarified. Despite these challenges, our measurements show that Internet governance decisions have a direct and measurable effect on IPv4 address space consumption and utilization. This observation bodes well for informed policymaking to provide IPv4 resource liquidity, and to set the right incentives further growth of the Internet. Future scenarios for address management could include a more competitive environment among the different registries, or even a re-centralization of the registries. Despite the technical and legal challenges in this tussle, we see an increasing number of listed IPv4 address transfers. Thus, IPv4 address scarcity is real and networks are willing to pay for additional IPv4 addresses. While purchasing IPv4 addresses is evidently a viable option for many network operators, the number of monthly transferred addresses is lower compared to pre-exhaustion allocation rates. This observation suggests that networks likely also use other means to mitigate their individual scarcity issues.

In the second part of this dissertation we assessed and analyzed global IPv4 address activity. Empirical measurements of address activity are imperative to understand IPv4 exhaustion, its current status and its developments, since address allocation is decoupled from actual address *use*. Our measurements show that after years of constant growth, the number of active IPv4 addresses contacting a major global Content Distribution Network (CDN) has stagnated since 2014. IPv4 address exhaustion does not only present us with a management and allocation issue any longer, but now manifests directly in actual IPv4 address activity. Stagnating address counts, but a growing IPv4 Internet, bring a new reality upon us: We have now entered an era in which the growth of activity in the IPv4 Internet is not measurable by counting active IPv4 addresses. Instead, we developed techniques and analyses to study structural properties of address activity. Our year-long observation of address activity shows client activity from 1.2 billion IPv4 addresses. We find that IPv4 address activity exhibits substantial dynamics; up to 25% of the active address pool changes over the course of one year. We are able to identify and attribute address activity patterns to network restructurings, user behaviors, and, in particular, various address assignment practices. Our measurements show that there is still significant potential for increasing the utilization of the IPv4 address space. We find that static address assignments harbor large reserves of potentially unused addresses, and that the utilization of dynamically assigned address blocks could be increased by reconfiguration of the respective address pools. Networks could consolidate their address space, accommodate additional hosts and devices, or become sellers in the IPv4 marketplace. On the other end of the spectrum, we find increasing concentration of Web traffic across fewer, heavily active IPv4 addresses, pointing towards increasing gateway—and in particular Carrier-Grade NAT—deployment. Our measurements reveal a wide spectrum of address use, from lightly used statically assigned address blocks, dynamic address blocks, to gateways accommodating potentially thousands of users. IPv4 exhaustion brought a reimagined IPv4 upon us, one that entails increased address sharing in both space and time.

In the third part of this dissertation we studied the deployment of Carrier-Grade NAT in the Internet. We develop techniques to detect the presence of CGN devices in ISPs and to distill properties of the identified CGN instances. While CGNs are often referred to as a short-term mitigation strategy for IPv4 address exhaustion, we find them widely deployed in networks across the globe. In fact, we find that CGN is the norm in cellular networks (94% of cellular ISPs deploy CGN) and is also increasingly popular in residential settings (17% of the non-cellular ISPs deploy CGN). Given widespread deployment, CGNs are likely to remain a vital part of the Internet infrastructure for the foreseeable future. Our analysis reveals a striking variability in the dimensioning, configuration, and placement of CGNs. We find that the degree of resource sharing varies vastly across the identified CGNs, with some ISPs dedicating just as little as 512 ports per subscriber. Furthermore, CGNs affect the connectivity of end users by restricting the available connectivity to varying degrees, due to the use of different NAT mapping types, and timeout values. Thus, "how much Internet" a subscriber receives behind a CGN differs substantially across different ISPs. Here, IPv4 address exhaustion directly affects end users. CGN deployment allows for virtually instant IPv4 scarcity relief for networks that connect end users, but also changes the semantics and the corresponding activity of public-facing gateway IP addresses. Mappings of public

IP addresses to internal IP addresses as well as port translation behavior of the identified CGNs varies dramatically, creating a scenario which renders the activity of gateway IP addresses unpredictable, and makes attribution of end users to IP addresses all but infeasible. We also uncover pitfalls that CGN deployments pose for network operators: Large ISPs face a shortage of internal IPv4 address space and hence opt to use routable address space behind their CGN deployment. This practice has the potential to affect the connectivity of these address blocks, if concurrently in public use. We argue that the limited availability of guidelines on how to configure and operate CGNs and the absence of regulations on how to allocate CGN resources pose major obstacles here, directly affecting the Internet's end users.

In the last part of this dissertation we studied dominant aspects of the evolving IPv6 Internet, and barriers faced when transitioning to IPv6. Our analysis of inter-domain connectivity shows that IXPs, and in particular Route Servers at IXPs, provide an easy and quick way for networks to establish IPv6 peerings. Consequently, we see IPv6 Inter-Domain connectivity increasing, yet it still lags behind its IPv4 counterpart. Contrasting our findings with traffic, however, changes the picture: Most IPv6 peering links carry insignificant traffic volumes. As a result, connectivity-based metrics to track IPv6 adoption need to be taken with caution. Our study of the application mix carried over the various peering links reveals substantial heterogeneity, suggesting that the shift over to IPv6 will likely be an individual task for each network involved and heavily dependent on the individual mix of traffic components they carry. Finally, studying the detailed interactions between IPv4 and IPv6 connectivity and traffic in a dual-stack ISP revealed that there are still barriers (e.g., lack of proper IPv6 support and configuration of CPE routers) present that prevent significant traffic portions from being carried over IPv6, even if the ISP provides IPv6 connectivity to its subscribers. Our findings calls for attention from application developers and vendors of home networking equipment. On the other hand, we see a strong *intent* for IPv6 traffic, i.e., users of this ISP request a large fraction of content to be served over IPv6. This intent can—to a large extent—not yet be satisfied by the server-side. Our measurements suggest that we can expect a prolonged period with coexistence of IPv4 and IPv6, during which we can face sudden and disparate shifts of traffic from IPv4 over to IPv6.

6.2 Future Directions

The Internet continuously evolves and the problem of IPv4 address scarcity can play out in a multitude of different ways. Our analysis shows potential for growth both in the IPv4 address space by increasing deployment of address conservation mechanisms, as well as in IPv6 adoption. Thus, **continuous reappraisal** of the presented measurements will be imperative to identify dominant future trends and their ramifications. The findings in this dissertation might also prove helpful in the context of future, and thus still unknown, instances of scarcity of virtual resources. They might also inform design choices for future instances of protocols that either come with inherent resource limitations (e.g., physical resources, spectrum frequencies), or protocols that introduce resource limitations, e.g., by defining a finite name or address space.

Currently, IPv4 is the dominant addressing protocol on the Internet, and disabling support for IPv4 is not a viable option for network operators. However, operating IPv4 and IPv6 concurrently in the long term requires substantial resources and increases the cost of operating the Internet [139]. As of today, it is still unclear when IPv6 will reach a degree of deployment that makes it the first-class citizen on the Internet, replacing IPv4 as the dominant network-layer protocol. Whenever that point is reached, the community faces another challenge: Incrementally switching off IPv4 and preparing the Internet for **IPv6-only operation**. A challenging question will be at what point network operators can, or should, eventually consider disabling support for IPv4 and what the impact of such a decision will be. Detailed measurements that document the lifetime of devices on the Internet, as well the criticality of individual devices and services will be necessary to inform such a decision.

Yet, there are many more looming challenges that we face in the short- and mid-term. These challenges are of both governmental and technical nature. The institutions governing IP address allocations are still in the process of introducing and refining policies to regulate IPv4 address transfer markets and to provide incentives for increasing IPv6 adoption. An avenue for future work includes ongoing study of **address policies** and their direct impact on IPv4 address utilization, and possibly connected incentives for IPv6 adoption. The RIRs currently apply different policies, with some requiring recipients of transfers to justify the need for additional IPv4 address space, while others removed this requirement. It remains an open question whether the market for IPv4 address space will fulfill its function to increase IPv4 utilization until broader deployment of IPv6, or whether it can result in address hoarding by wealthy institutions speculating on increasing prices. Empirical measurements of address utilization of transferred blocks could help shape such transfer policies.

The adoption of one or another technology to mitigate IPv4 scarcity by different network operators is primarily a monetary decision [139]. A detailed **analysis of the economic aspects of IPv4 scarcity and its mitigation** could both aid network operators to make informed decisions as well as foster efficient policymaking. Broader adoption of IPv6 could result in higher availability of compatible hardware, software, and skilled network administrators, leading to dropping overall costs for IPv6 adoption and ease of deployment. Increasing IPv6 adoption, in turn, could result in lower prices of IPv4 addresses, which could motivate other networks to purchase IPv4 addresses and pursue an IPv4-only approach in the mid-term. It is a challenge to understand the interplay between IPv4 address prices, the costs and benefits of deploying CGN, and the costs and benefits for individual networks of rolling out IPv6. Understanding and modeling the monetary impact and the interdependencies of these options could both allow predictions on the future development of this tussle, and support ISPs when making business-critical decisions.

The exhaustion of the IPv4 address space, increasing IPv6 adoption, and the **concurrent operation of two addressing protocols** (IPv4 and IPv6) also raises new operational challenges for availability, compatibility, and security. The increasing scarcity of IPv4 addresses and the lack of resource certification could lead to instability of the routing system and increasing occurrences of prefix hijacks. Future work could involve studying the impact of address shortage and increase of prefix hijacking events, and how to protect against them. Increasing IPv6 deployment presents us with a whole new set of technical challenges. IP address reputation systems and corresponding best practices were designed with IPv4 in mind, but IPv6 addressing allows for much more flexibility. Avenues for future work include detailed assessment of the activity of the emerging IPv6 address space as well as the development of host reputation approaches in this new address space. Current rates of IPv6 adoption suggest that we face a mid- to long-term scenario during which IPv4 and IPv6 coexist. Dual-stacked hosts will thus be reachable via two IP addresses and protocol stacks. Applications need to be able to seamlessly interact in a dual-stack environment. Firewalls and intrusion detection systems need to find ways to match IPv4-IPv6 pairings to perform consistent filtering for both protocols, and to identify possible cross-protocol attacks. The operation of the Internet with two concurrently active addressing protocols requires attention of the Internet community, including network operators, application developers, hardware vendors, regulators, and end users.

List of Figures

List of Tables

Bibliography

[1] Internet Addresses Census dataset, PREDICT ID: USC-LANDER/internet_address_census_it55w-20130723/rev3638. Traces taken 2013-07-23 to 2013-08-25. Provided by the USC/LANDER project (http://www.isi.edu/ant/lander).

[2] Application Layer Packet Classifier for Linux (L7-filter). http://l7-filter.sourceforge.net/.

[3] BitTorrent Protocol Specification v 1.0. https://wiki.theory.org/BitTorrentSpecification.

[4] Current implementation of AI_ADDRCONFIG considered harmful. https://goo.gl/prXWfz.

[5] Euro-IX Resources: Traffic, Reports, and Best Practices. https://www.euro-ix.net/resources.

[6] Internet-Wide Scan Data Repository. https://scans.io/.

[7] IRR - Internet Routing Registry. http://www.irr.net.

[8] Measurement Lab (M-Lab). https://www.measurementlab.net/.

[9] Netalyzr for Android. Google Play. https://play.google.com/store/apps/details?id=edu.berkeley.icsi.netalyzr.android.

[10] Sandvine Global Internet Phenomena, 1H 2014. https://www.sandvine.com/downloads/general/global-internet-phenomena/.

[11] Sandvine Traffic Classification. https://www.sandvine.com/technology/traffic-classification.html.

[12] The BIRD Internet Routing Daemon. http://bird.network.cz.

[13] UCSD Network Telescope. http://www.caida.org/projects/network_telescope/.

[14] uTorrent Transport Protocol Specification. http://www.bittorrent.org/beps/bep_0029.html.

[15] What is RTMTP? https://www.wowza.com/glossary/rtmpt.

[16] World IPv6 Launch. https://goo.gl/hOoXNo.

[17] A. LOEWENSTERN AND A. NORBERG. DHT Protocol (BEP-05). http://www.bittorrent.org/beps/bep_0005.html.

[18] A. TABDILI. Carrier Grade NAT: Requirements and Challenges in the Real World. http://www.menog.org/presentations/menog-10/Amir%20Tabdili%20-%20Carrier%20Grade%20NAT.pdf, 2012. MENOG 10.

[19] A10 NETWORKS. Carrier Grade NAT (CGN) Deployment Guide. https://www.a10networks.com/sites/default/files/resource-files/A10-DG-Carrier_Grade_NAT_%28CGN%29_Large_Scale_NAT_%28LSN%29.pdf.

[20] ABEN, E. Hampering Eyeballs - Observations on Two "Happy Eyeballs" Implementations. https://goo.gl/qUW6s.

[21] ABEN, E., TRENAMAN, N., KIESSLING, A., AND WILHELM, R. Lost Starts - Why Operators Switch off IPv6, 2016. NANOG 66.

[22] ABOBA, B., ZORN, G., AND MITTON, D. RADIUS and IPv6. RFC 3162, 2001.

[23] ADDREX. IPv4 Address Broker. http://www.addrex.net.

[24] ADRIAN, D., DURUMERIC, Z., SINGH, G., AND HALDERMAN, J. A. Zippier ZMap: Internet-wide Scanning at 10 Gbps. In *8th USENIX Workshop on Offensive Technologies* (2014).

[25] AFRINIC. AFRINIC Service Agreement 2013. https://www.afrinic.net/en/services/rs/rsa.

[26] AFRINIC. IPv4 Allocation Policy (AFPUB-2005-v4-001). http://www.afrinic.net/en/library/policies/126-afpub-2005-v4-001.

[27] AFRINIC. Out-Of-Region Use of AFRINIC Internet Number Resources (AFPUB-2014-GEN-002-DRAFT-01). http://afrinic.net/en/community/policy-development/policy-proposals/1157-out-of-region-use-of-afrinic-internet-number-resources.

[28] AFRINIC. Policy Development Process in the AFRINIC service region (AFPUB-2010-GEN-005). http://www.afrinic.net/en/community/policy-development/251-policy-development-process-in-the-afrinic-service-region-afpub-2010-gen-005.

[29] AGER, B., CHATZIS, N., FELDMANN, A., SARRAR, N., UHLIG, S., AND WILLINGER, W. Anatomy of a Large European IXP. In *ACM SIGCOMM* (2012).

[30] AKAMAI TECHNOLOGIES. IPv6 Adoption Visualization. https://goo.gl/QWN7u8.

[31] AKAMAI TECHNOLOGIES. State of the Internet Report. https://www.akamai.com/us/en/our-thinking/state-of-the-internet-report.

[32] ALCOCK, S., LORIER, P., AND NELSON, R. Libtrace: A Packet Capture and Analysis Library. *ACM CCR 42*, 2 (Apr 2012).

[33] ALCOCK, S., AND NELSON, R. Libprotoident: Traffic Classification Using Lightweight Packet Inspection. *TR, University of Waikato* (2012).

[34] ALCOCK, S., NELSON, R., AND MILES, D. Investigating the Impact of Service Provider NAT on Residential Broadband Users. *TR, University of Waikato* (2010).

[35] ALT, L., BEVERLY, R., AND DAINOTTI, A. Uncovering Network Tarpits with Degreaser. In *ACSAC* (2014).

[36] ANDREWS, M. Negative Caching of DNS Queries (DNS NCACHE). RFC 2308, 1998.

[37] ANTONAKAKIS, M., PERDISCI, R., DAGON, D., LEE, W., AND FEAMSTER, N. Building a Dynamic Reputation System for DNS. In *USENIX Security Symposium* (2010).

[38] APNIC. APNIC IPv4 Address Pool Reaches Final /8 (2011-04-15). http://www.apnic.net/publications/news/2011/final-8.

[39] APNIC. IPv4 Address Transfer Logs. ftp://ftp.apnic.net/public/transfers/apnic/.

[40] APNIC. IPv4 Transfers Listing Service. http://www.apnic.net/services/become-a-member/manage-your-membership/pre-approval/listing.

[41] APNIC. Policy development process (APNIC-111). http://www.apnic.net/publications/media-library/documents/policy-development/development-process.

[42] APNIC. Policy environment for Internet number resource distribution in the Asia Pacific (APNIC-125, 2011-05-09). https://www.apnic.net/policy/policy-environment/text.

[43] APNIC. prop-017: Recovery of unused address space (2004-02-26). https://www.apnic.net/policy/proposals/prop-017.

[44] APNIC. Transfer, merger, acquisition, and takeover policy (APNIC-123). http://www.apnic.net/policy/transfer-policy.

[45] APNIC LABS. Customers per AS Measurements. Description: https://labs.apnic.net/?p=526 Dataset: http://stats.labs.apnic.net/aspop.

[46] APNIC LABS. IPv6 Measurement Maps. https://https://stats.labs.apnic.net/ipv6/, 2017.

[47] ARIN. ARIN IPv6 Wiki: Broadband CPE. https://goo.gl/Wydr3Q.

[48] ARIN. ARIN Policy Development Process. https://www.arin.net/policy/pdp.html.

[49] ARIN. ARIN's Number Resource Policy Manual (NRPM) Version 2014.02. https://www.arin.net/policy/nrpm.html.

[50] ARIN. Draft Policy ARIN-2014-1: Out of Region Use. https://www.arin.net/policy/proposals/2014_1.html.

[51] ARIN. Inter-RIR and Specified Transfers of Internet Number Resources. https://www.arin.net/knowledge/statistics/transfers.html.

[52] ARIN. Inter-RIR Transfers. https://www.arin.net/resources/request/transfers_8_4.html.

[53] ARIN. Legacy Registration Service Agreement Statistics. https://www.arin.net/knowledge/statistics/legacy.html.

[54] ARIN. Legacy Registration Services Agreement. https://www.arin.net/fees/agreements/legacy.html.

[55] ARIN. Legacy Registration Services Agreement Outreach (2009-03-03). https://www.arin.net/resources/legacy/outreach.html.

[56] ARIN. Legacy Registration Services Agreement v3.0. https://www.arin.net/resources/agreements/legacy_rsa.pdf.

[57] ARIN. Registration Services Agreement v11.0. https://www.arin.net/resources/agreements/rsa.pdf.

[58] ARIN. RPKI Terms of Service Agreement. https://www.arin.net/resources/rpki/tos.pdf.

[59] ARIN. Specified Transfer Listing Service. https://www.arin.net/resources/transfer_listing/.

[60] ARIN. Transfers to Specified Recipients. https://www.arin.net/resources/request/transfers_8_3.html.

[61] ARKKO, J., AND KERANEN, A. Experiences from an IPv6-Only Network. RFC 6586 (Informational), Apr 2012.

[62] AUDET, F., AND JENNINGS, C. Network Address Translation (NAT) Behavioral Requirements for Unicast UDP. RFC 4787 (Best Current Practice), Jan 2007. Updated by RFCs 6888, 7857.

[63] BAGNULO, M., MATTHEWS, P., AND VAN BEIJNUM, I. Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers. IETF RFC 6146, April 2011.

[64] BAGNULO, M., SULLIVAN, A., MATTHEWS, P., AND VAN BEIJNUM, I. DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers. IETF RFC 6147, April 2011.

[65] BAJPAI, V., AND SCHÖNWÄLDER, J. IPv4 versus IPv6 - Who connects faster? In *IFIP Networking* (2015).

[66] BAKER, F., LI, X., BAO, C., AND YIN, K. Framework for IPv4/IPv6 Translation. RFC 6144 (Informational), Apr 2011.

[67] BEIJNUM, I. V. *BGP: Building reliable networks with the border gateway protocol.* Page 61ff. O'Reilly Media, Inc., 2002.

[68] BELLOVIN, S. M. A Technique for Counting NATted Hosts. In *IMW* (2002).

[69] BERMUDEZ, I., MELLIA, M., MUNAFÀ, M., KERALAPURA, R., AND NUCCI, A. DNS to the Rescue: Discerning Content and Services in a Tangled Web. In *ACM IMC* (2012).

[70] BEVERLY, R., LUCKIE, M., MOSLEY, L., AND CLAFFY, K. Measuring and Characterizing IPv6 Router Availability. In *PAM*. 2015.

[71] BOCCHI, E., KHATOUNI, A. S., TRAVERSO, S., FINAMORE, A., GENNARO, V. D., MELLIA, M., MUNAFO, M., AND ROSSI, D. Impact of Carrier-Grade NAT on Web Browsing. In *IWCMC* (2015).

[72] BONFIGLIO, D., MELLIA, M., MEO, M., RITACCA, N., AND ROSSI, D. Tracking Down Skype Traffic. In *IEEE INFOCOM* (2008).

[73] BOUCADAIR, M., PENNO, R., AND WING, D. Universal Plug and Play (UPnP) Internet Gateway Device - Port Control Protocol Interworking Function (IGD-PCP IWF). RFC 6970 (Proposed Standard), Jul 2013.

[74] BUSH, R. Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI). RFC 7115 (Best Current Practice), Jan. 2014.

[75] BUSH, R., AUSTEIN, R., BELLOVIN, S., AND ELKINS, M. The RPKI & Origin Validation. NANOG 52, 2001.

[76] BUTKIEWICZ, M., MADHYASTHA, H. V., AND SEKAR, V. Understanding Website Complexity: Measurements, Metrics, and Implications. In *IMC* (2011).

[77] BUTLER, K., FARLEY, T., MCDANIEL, P., AND REXFORD, J. A Survey of BGP Security Issues and Solutions. *Proceedings of the IEEE 98*, 1 (2010).

[78] C. MARSAN. Stanford move rekindles 'Net address debate (2000-01-24). Network World, Vol. 17, No. 4.

[79] CAI, X., AND HEIDEMANN, J. Understanding Block-Level Address Usage in the Visible Internet. In *ACM SIGCOMM* (2010).

[80] CAIDA. Ark Measurement Infrastructure. `http://www.caida.org/projects/ark/`.

[81] CALLADO, A., KAMIENSKI, C., SZABO, G., AND KELNER, B. G. J., FERNANDES, S., AND SADOK, D. A Survey on Internet Traffic Identification. *IEEE Comm. Surveys and Tutorials* (2009).

[82] CALLAHAN, T., ALLMAN, M., AND RABINOVICH, M. On Modern DNS Behavior and Properties. *ACM CCR 43*, 3 (2013).

[83] CANTÓ, R., LÓPEZ, R. A., FOLGUEIRA, J. L., LÓPEZ, D. R., ELIZONDO, A. J., AND GAMERO, R. Virtualization of Residential Customer Premise Equipment. Lessons Learned in Brazil vCPE Trial. *Information Technology 57*, 5 (2015).

[84] CARELA-ESPAÑOL, V., BUJLOW, T., AND BARLET-ROS, P. Is Our Ground-Truth for Traffic Classification Reliable? In *PAM* (2014).

[85] CERF, V. IAB Recommended Policy on Distributing Internet Identifier Assignment and IAB Recommended Policy Change to Internet "Connected" Status. RFC 1174 (Informational), 1990.

[86] CHATZIS, N., SMARAGDAKIS, G., FELDMANN, A., AND WILLINGER, W. There is More to IXPs than Meets the Eye. *ACM CCR 43*, 5 (2013).

[87] CHEN, F., SITARAMAN, R. K., AND TORRES, M. End-User Mapping: Next Generation Request Routing for Content Delivery. In *ACM SIGCOMM* (2015).

[88] CHO, K., LUCKIE, M., AND HUFFAKER, B. Identifying IPv6 Network Problems in the Dual-stack World. In *ACM SIGCOMM Network Troubleshooting Workshop* (2004).

[89] CISCO. NAT Administration Guide, StarOS Release 17. `http://www.cisco.com/c/dam/en/us/td/docs/wireless/asr_5000/17-0/PDF/17-NAT-Admin.pdf`.

[90] CLAFFY, K. Tracking IPv6 Evolution: Data We Have and Data We Need. *ACM CCR 41*, 3 (2011).

[91] COLITTI, L., GUNDERSON, S. H., KLINE, E., AND REFICE, T. Evaluating IPv6 Adoption in the Internet. In *PAM* (2010).

[92] COMER, D. E. *Internetworking with TCP/IP, Volume 1: Principles, Protocols, and Architectures, Fourth Edition*, 4th ed. Prentice Hall PTR, Upper Saddle River, NJ, USA, 2000.

[93] COTTON, M., AND VEGODA, L. Special Use IPv4 Addresses. RFC 5735 (Best Current Practice), Jan. 2010. Obsoleted by RFC 6890, updated by RFC 6598.

[94] CURRAN, J. An Internet Transition Plan. RFC 5211 (Informational), Jul 2008.

[95] CYMRU. Ephemeral Source Port Selection Strategies. `https://www.cymru.com/jtk/misc/ephemeralports.html`.

[96] CZYZ, J., ALLMAN, M., ZHANG, J., IEKEL-JOHNSON, S., OSTERWEIL, E., AND BAILEY, M. Measuring IPv6 Adoption. In *ACM SIGCOMM* (2014).

[97] D'ACUNTO, L., POUWELSE, J., AND SIPS, H. A measurement of NAT & Firewall Characteristics in Peer to Peer Systems. In *ASCI* (2009).

[98] DAINOTTI, A., BENSON, K., KING, A., CLAFFY, K., GLATZ, E., AND DIMITROPOULOS, X. Estimating Internet Address Space Usage Through Passive Measurements. *ACM CCR 44*, 1 (2014).

[99] DAINOTTI, A., BENSON, K., KING, A., CLAFFY, K., KALLITSIS, M., GLATZ, E., AND DIMITROPOULOS, X. Estimating Internet address space usage through passive measurements. *ACM CCR 44*, 1 (2014), 42–49.

[100] DAINOTTI, A., BENSON, K., KING, A., HUFFAKER, B., GLATZ, E., DIMITROPOULOS, X., RICHTER, P., FINAMORE, A., AND SNOEREN, A. Lost in Space: Improving Inference of IPv4 Address Space Utilization. *IEEE J. on Sel. Areas in Comm. 34*, 6 (Jun 2016), 1862–1876.

[101] DAINOTTI, A., PESCAPE, A., AND CLAFFY, K. Issues and Future Directions in Traffic Classification. *IEEE Network Magazine* (2012).

[102] DEC, W., SARIKAYA, B., ZORN, G., MILES, D., AND LOURDELET, B. RADIUS Attributes for IPv6 Access Networks. RFC 6911, 2013.

[103] DEERING, S., AND HINDEN, R. Internet Protocol, Version 6 (IPv6) Specification. RFC 2460 (Draft Standard), 1998. Updated by RFCs 5095, 5722, 5871, 6437, 6564, 6935, 6946, 7045, 7112.

[104] DENG, W., ZHU, P., LU, X., AND PLATTNER, B. On Evaluating BGP Routing Stress Attack. *Journal of communications 5*, 1 (2010), 13–22.

[105] DHAMDHERE, A., LUCKIE, M., HUFFAKER, B., CLAFFY, K., ELMOKASHFI, A., AND ABEN, E. Measuring the Deployment of IPv6: Topology, Routing and Performance. In *ACM IMC* (2012).

[106] DICIOCCIO, L., TEIXEIRA, R., MAY, M., AND KREIBICH, C. Probe and Pray: Using UPnP for Home Network Measurements. In *PAM* (2012).

[107] DIGITAL TRENDS. Is it too late for BitTorrent to shake its seedy image? `http://www.digitaltrends.com/opinion/bittorrents-image-problem/`, 2013.

[108] DONALD CODLING, FBI. Attribution: Growing Challenges For LEAs. CWAG Conference of Western Attorneys General 2013 `http://www.cwagweb.org/pdfs/2013/privacy_conference_ppts/cwag%20preso%2013.ppt`.

[109] DONLEY, C., HOWARD, L., KUARSINGH, V., BERG, J., AND DOSHI, J. Assessing the Impact of Carrier-Grade NAT on Network Applications. RFC 7021 (Informational), Sep 2013.

[110] DRAKE, K. You have IPv6. Turn it on. `https://goo.gl/maSZRM`, 2016.

[111] DURUMERIC, Z., WUSTROW, E., AND HALDERMAN, J. A. ZMap: Fast Internet-Wide Scanning and its Security Applications. In *USENIX Security Symposium* (2013).

[112] EGEVANG, K., AND FRANCIS, P. The IP Network Address Translator (NAT). RFC 1631 (Informational), May 1994. Obsoleted by RFC 3022.

[113] EQUINIX. Global Data Centers and Colocation Services. `http://www.equinix.com/locations/`.

[114] EUROPOL. Closing the Online Crime Attribution Gap: European Law Enforcement tackles Carrier-Grade NAT (CGN). `https://www.europol.europa.eu/newsroom/news/closing-online-crime-attribution-gap-european-law-enforcement-tackles-carrier-grade-nat-cgn`.

[115] FAN, X., AND HEIDEMANN, J. Selecting Representative IP Addresses for Internet Topology Studies. In *ACM IMC* (2010).

[116] FCC. Measuring Broadband America. `https://www.measuringbroadbandamerica.com/`.

[117] FINAMORE, A., MELLIA, M., MEO, M., MUNAFO, M., AND ROSSI, D. Experiences of Internet traffic monitoring with Tstat. *Network, IEEE 25*, 3 (2011), 8–14.

[118] FINAMORE, A., MELLIA, M., MEO, M., AND ROSSI, D. KISS: Stochastic Packet Inspection Classifier for UDP Traffic. *IEEE/ACM Trans. Networking* (2010).

[119] FORD, B., SRISURESH, P., AND KEGEL, D. Peer-to-Peer Communication Across Network Address Translators. In *USENIX ATC* (2005).

[120] FORD, P., REKHTER, Y., AND BRAUN, H. Improving the Routing and Addressing of IP. *IEEE Network 7*, 3 (1993).

[121] FULLER, V., LI, T., YU, J., AND VARADHAN, K. Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy. RFC 1519 (Proposed Standard), 1993. Obsoleted by RFC 4632.

[122] GARTNER. Forecast: Internet of Things - Endpoints and Associated Services, Worldwide, 2016. `https://www.gartner.com/doc/3558917/forecast-internet-things--endpoints`.

[123] GEHLEN, V., FINAMORE, A., MELLIA, M., AND MUNAFÒ, M. M. Uncovering the big players of the web. In *TMA* (Berlin, Heidelberg, 2012), TMA'12, Springer-Verlag, pp. 15–28.

[124] GEOFF HOUSTON. IP Address Report. `http://www.potaroo.net`.

[125] GERBER, A., AND DOVERSPIKE, R. Traffic Types and Growth in Backbone Networks. In *OFC/NFOEC* (2011).

[126] GERICH, E. Guidelines for Management of IP Address Space. RFC 1466 (Informational), 1993. Obsoleted by RFC 2050.

[127] GIOTSAS, V., LUCKIE, M., HUFFAKER, B., AND CLAFFY, K. IPv6 AS Relationships, Clique, and Congruence. In *PAM* (2015).

[128] GIOTSAS, V., AND ZHOU, S. Improving the Discovery of IXP Peering Links through Passive BGP Measurements. In *Glob. Internet* (2013).

[129] GIOTSAS, V., ZHOU, S., LUCKIE, M., AND KC CLAFFY. Inferring Multilateral Peering. In *CoNEXT* (2013).

[130] GOOGLE. IPv6 Statistics. https://www.google.com/intl/en/ipv6/statistics.html.

[131] GUEYE, B., ZIVIANI, A., CROVELLA, M., AND FDIDA, S. Constraint-Based Geolocation of Internet Hosts. *IEEE/ACM Trans. Networking 14*, 6 (2006), 1219–1232.

[132] GUHA, S., BISWAS, K., FORD, B., SIVAKUMAR, S., AND SRISURESH, P. NAT Behavioral Requirements for TCP. RFC 5382 (Best Current Practice), Oct 2008. Updated by RFC 7857.

[133] GYSI, M. Residential IPv6 at Swisscom, an Overview. https://goo.gl/QO2SZF, 2012.

[134] GYSI, M. Status of Swisscom's IPv6 Activities, Outlook and Opportunities. Swiss IPv6 Council IPv6 Business Conference.

[135] H. PARMAR AND M. THORNBURGH. Adobe's real time messaging protocol. http://www.adobe.com/devnet/rtmp.html, 2012.

[136] HAO, S., SYED, N. A., FEAMSTER, N., GRAY, A. G., AND KRASSER, S. Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine. In *USENIX Security Symposium* (2009).

[137] H.D. MOORE. Project Sonar. https://community.rapid7.com/community/infosec/sonar/blog.

[138] HEIDEMANN, J., PRADKIN, Y., GOVINDAN, R., PAPADOPOULOS, C., BARTLETT, G., AND BANNISTER, J. Census and Survey of the Visible Internet. In *ACM IMC* (2008).

[139] HOGG, S. Dual-Stack Will Increase Operating Expenses. http://www.networkworld.com/article/2222870/cisco-subnet/dual-stack-will-increase-operating-expenses.html, 2012.

[140] HOLMBERG, C., HAKANSSON, S., AND ERIKSSON, G. Web Real-Time Communication Use Cases and Requirements. RFC 7478 (Informational), Mar 2015.

[141] HORTON, M., AND ADAMS, R. Standard for interchange of USENET messages. RFC 1036, Dec. 1987. Obsoleted by RFCs 5536, 5537.

[142] HUBBARD, K., KOSTERS, M., CONRAD, D., KARRENBERG, D., AND POSTEL, J. Internet Registry IP Allocation Guidelines. RFC 2050 (Historic), 1996. Obsoleted by RFC 7020.

[143] HUSTON, G. Bemused Eyeballs. https://labs.apnic.net/?p=188, 2012.

[144] HUSTON, G. Revisiting Apple and IPv6. https://goo.gl/qjKdv5, 2015.

[145] HUSTON, G. In Defence of NATs. In *Glob. Internet* (2017).

[146] HYUN, Y., HUFFAKER, B., ANDERSEN, D., LUCKIE, M., AND CLAFFY, K. C. The IPv4 Routed /24 Topology Dataset, 2014. http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml.

[147] ICANN. Global Policy for the Allocation of the Remaining IPv4 Address Space (ratified in 2009). https://www.icann.org/resources/pages/remaining-ipv4-2012-02-25-en.

[148] ICANN. Multi-Stakeholder Discussion: Legacy Internet Protocol (IP) Numbers in the Current Policy Environment (2012). Audio Transcript. http://toronto45.icann.org/node/34377.

[149] ICANN BLOG. Recovering IPv4 Address Space. http://blog.icann.org/2008/02/recovering-ipv4-address-space/.

[150] IEEE. 802 LAN/MAN Standards Committee. http://www.ieee802.org/.

[151] ILIOFOTOU, M., GALLAGHER, B., ELIASSI-RAD, T., XIE, G., AND FALOUTSOS, M. Profiling-By-Association: A Resilient Traffic Profiling Solution for the Internet Backbone. In *CoNEXT* (2010).

[152] INFORMATION OF SCIENCES INSTITUTE, USC. Internet Address Survey Binary Format. http://www.isi.edu/ant/traces/topology/address_surveys/binformat_description.html, 2012.

[153] INTERNATIONAL TELECOMMUNICATION UNION. Statistics. http://www.itu.int/en/ITU-D/Statistics/Pages/default.aspx.

[154] INTERNET ACTIVITIES BOARD. Meeting Minutes. June 18-19, 1992. https://www.iab.org/documents/minutes/minutes-1992/iab-minutes-1992-06-18/.

[155] IP TRADING. IPv4 Address Broker. http://www.iptrading.com.

[156] IPV4 MARKET GROUP. IPv4 Address Broker. http://www.ipv4marketgroup.com.

[157] IPV4 MARKET GROUP BLOG. Don't Sign an RSA During Your 8.2 IPv4 Transfer. http://ipv4marketgroup.com/?s=do+not+sign+rsa.

[158] IRR. List Of Routing Registries. http://www.irr.net/docs/list.html.

[159] JIANG, S., GUO, D., AND CARPENTER, B. An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition. RFC 6264 (Informational), Jun 2011.

[160] JIN, Y., SHARAFUDDIN, E., AND ZHANG, Z. L. Identifying dynamic IP address blocks serendipitously through background scanning traffic. In *CoNEXT* (2007).

[161] JOHNSTON, A. B. *SIP: Understanding the Session Initiation Protocol*, 3rd ed. Artech House, Inc., Norwood, MA, USA, 2009.

[162] KARAGIANNIS, T., BROIDO, A., FALOUTSOS, M., AND CLAFFY, K. Transport layer identification of P2P traffic. In *ACM IMC* (2004).

[163] KARAGIANNIS, T., PAPAGIANNAKI, K., AND FALOUTSOS, M. BLINC: multilevel traffic classification in the dark. In *ACM SIGCOMM* (2005).

[164] KARIR, M., HUSTON, G., MICHAELSON, G., AND BAILEY, M. Understanding IPv6 Populations in the Wild. In *PAM* (2013).

[165] KARPILOVSKY, E., GERBER, A., PEI, D., REXFORD, J., AND SHAIKH, A. Quantifying the Extent of IPv6 Deployment. In *PAM* (2009).

[166] KATZ-BASSETT, E., JOHN, J. P., KRISHNAMURTHY, A., WETHERALL, D., ANDERSON, T., AND CHAWATHE, Y. Towards IP geolocation using delay and topology measurements. In *ACM IMC* (2006).

[167] KATZ-BASSETT, E., MADHYASTHA, H., ADHIKARI, V., SCOTT, C., SHERRY, J., VAN WESEP, P., KRISHNAMURTHY, A., AND ANDERSON, T. Reverse Traceroute. In *NSDI* (2010).

[168] KHAN, A., HYON-CHUL, K., KWON, T., AND CHOI, Y. A comparative Study on IP Prefixes and their Origin ASes in BGP and the IRR. *ACM CCR 43*, 3 (2013).

[169] KIM, H., CLAFFY, K., FOMENKOV, M., BARMAN, D., FALOUTSOS, M., AND LEE, K.-Y. Internet Traffic Classification Demystified: Myths, Caveats, and the Best Practices. In *CoNEXT* (2008).

[170] KREIBICH, C., WEAVER, N., NECHAEV, B., AND PAXSON, V. Netalyzr: Illuminating The Edge Network. In *ACM IMC* (2010).

[171] KRMICEK, V., VYKOPAL, J., AND KREJCI, R. NetFlow Based System for NAT Detection. In *ACM CoNEXT* (2009).

[172] KUROSE, J. F., AND ROSS, K. W. *Computer Networking: A Top-Down Approach (6th Edition)*, 6th ed. Pearson, 2012.

[173] LABOVITZ, C., LEKEL-JOHNSON, S., MCPHERSON, D., OBERHEIDE, J., AND JAHANIAN, F. Internet Inter-Domain Traffic. In *ACM SIGCOMM* (2010).

[174] LACNIC. Policy Development Process (version 2.0 - 01/09/2008). `http://www.lacnic.net/en/web/lacnic/proceso-de-desarrollo-de-politicas`.

[175] LACNIC. Policy Manual v2.1 (2014-03-25). `http://www.lacnic.net/web/lacnic/manual-2`.

[176] LACNIC. Registration Services Agreement. `http://lacnic.net/docs/rsa-en.pdf`.

[177] LAGERHOLM, S., AND ROSELLI, J. Negative Caching of DNS records. Tech. rep., Microsoft, 2015.

[178] LEE, C., LEE, D. K., AND MOON, S. Unmasking the Growing UDP Traffic in a Campus Network. In *PAM* (2012).

[179] LEPINSKI, M., AND KENT, S. An Infrastructure to Support Secure Internet Routing. RFC 6480 (Informational), Feb. 2012.

[180] LEPINSKI, M., AND TURNER, S. An Overview of BGPSEC. IETF Internet-Draft draft-ietf-sidr-bgpsec-overview-05 `https://tools.ietf.org/html/draft-ietf-sidr-bgpsec-overview-05`, 2014.

[181] LIVADARIU, I., ELMOKASHFI, A., AND DHAMDHERE, A. Characterizing IPv6 Control and Data Plane Stability. In *IEEE INFOCOM* (2016).

[182] LIVADARIU, I., ELMOKASHFI, A., AND DHAMDHERE, A. On IPv4 transfer markets: Analyzing reported transfers and inferring transfers in the wild. *Computer Communications 111* (2017), 105 – 119.

[183] LIVADARIU, I., ELMOKASHFI, A., DHAMDHERE, A., AND CLAFFY, K. A First Look at IPv4 Transfer Markets. In *ACM CoNEXT* (2013).

[184] LODHI, A., LARSON, N., DHAMDHERE, A., DOVROLIS, C., AND CLAFFY, K. Using PeeringDB to Understand the Peering Ecosystem. *ACM CCR 44*, 2 (2014).

[185] LUCKIE, M., BEVERLY, R., BRINKMEYER, W., AND CLAFFY, K. Speedtrap: Internet-Scale IPv6 Alias Resolution. In *ACM IMC* (2013).

[186] LUTU, A., BAGNULO, M., DHAMDHERE, A., AND CLAFFY, K. NAT Revelio: Detecting NAT444 in the ISP. In *PAM* (2016).

[187] M. PRINCE. In June number of IPv6 addresses (/64 masked) crossed IPv4 addresses connecting to @CloudFlare for the first time. Twitter `https://twitter.com/eastdakota/status/765699957449895936`.

[188] MACDONALD, D., AND LOWEKAMP, B. NAT Behavior Discovery Using Session Traversal Utilities for NAT (STUN). RFC 5780 (Experimental), May 2010.

[189] MAIER, G., FELDMANN, A., PAXSON, V., AND ALLMAN, M. On Dominant Characteristics of Residential Broadband Internet Traffic. In *ACM IMC* (2009).

[190] MAIER, G., SCHNEIDER, F., AND FELDMANN, A. NAT Usage in Residential Broadband Networks. In *PAM* (2011).

[191] MAYMOUNKOV, P., AND MAZIERES, D. Kademlia: A Peer-to-Peer Information System Based on the XOR Metric. In *Peer-to-Peer Systems*. Springer, 2002.

[192] MCCONACHIE, A. How To Make Your Website Available Over IPv6. `https://goo.gl/Vs2IuO`, 2014.

[193] MOORE, A. W., AND PAPAGIANNAKI, K. Toward the Accurate Identification of Network Applications. In *PAM* (2005).

[194] MORI, T., INOUE, T., SHIMODA, A., SATO, K., ISHIBASHI, K., AND GOTO, S. SFMap: Inferring Services over Encrypted Web Flows Using Dynamical Domain Name Graphs. In *TMA* (2015).

[195] MORISHITA, Y., AND JINMEI, T. Common Misbehavior Against DNS Queries for IPv6 Addresses. RFC 4074, 2005.

[196] MOURA, G. C. M., GANAN, C., LONE, Q., POURSAIED, P., ASGHARI, H., AND VAN EETEN, M. How Dynamic is the ISPs Address Space? Towards Internet-Wide DHCP Churn Estimation. In *Workshop on Research and Applications of Internet Measurements* (2015).

[197] MUELLER, M., KUERBIS, B., AND ASGHARI, H. Dimensioning the Elephant: An Empirical Analysis of the IPv4 Number Market. *info 15*, 6 (2013), 6–18.

[198] MÜLLER, A., WOHLFART, F., AND CARLE, G. Analysis and Topology-based Traversal of Cascaded Large Scale NATs. In *ACM HotMiddlebox* (2013).

[199] NATIONAL SCIENCE FOUNDATION. A Brief History of the NSF and the Internet. `https://www.nsf.gov/news/special_reports/cyber/internet.jsp`.

[200] NATIONAL SCIENCE FOUNDATION (NSF). Amendment 7 to Cooperative Agreement Between NSI and U.S. Government (1997-12-03). `http://archive.icann.org/en/nsi/coopagmt-amend7-03dec97.htm`.

[201] NCC, R. YouTube Hijacking: A RIPE NCC RIS case study. `http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study`.

[202] NGUYEN, T. T. T., AND ARMITAGE, G. A Survey of Techniques for Internet Traffic Classification using Machine Learning. *IEEE Comm. Surveys and Tutorials* (2008).

[203] NIKKHAH, M., AND GUÉRIN, R. Migrating the Internet to IPv6: An Exploration of the When and Why. *IEEE ToN* (2015).

[204] NIKKHAH, M., GUÉRIN, R., LEE, Y., AND WOUNDY, R. Assessing IPv6 Through Web Access a Measurement Study and Its Findings. In *ACM CoNEXT* (2011).

[205] NORDMARK, E., AND GILLIGAN, R. Basic Transition Mechanisms for IPv6 Hosts and Routers. RFC 4213 (Proposed Standard), Oct. 2005.

[206] NOW (NEW ZEALAND ISP). What if I need a public IP Address? https://support.nownz.co.nz/support/solutions/articles/5000504832-what-if-i-need-a-public-ip-address-.

[207] NRO. Extended Allocation and Assignment Reports. http://www.nro.net/statistics.

[208] NRO. Free Pool of IPv4 Address Space Depleted. http://www.nro.net/news/ipv4-free-pool-depleted.

[209] NTT AMERICA. Request for IPv4 and IPv6 Address Space v3.7. https://whois.gin.ntt.net/ipj.txt.

[210] NYGREN, E., SITARAMAN, R. K., AND SUN, J. The Akamai Network: A Platform for High-performance Internet Applications. *SIGOPS Oper. Syst. Rev. 44*, 3 (2010).

[211] OHARA, Y., NISHIZUKA, K., CHINEN, K., AKASHI, K., KOHRIN, M., MURAMOTO, E., AND MIYAKAWA, S. On the Impact of Mobile Network Delays on Connection Establishment Performance of a Carrier Grade NAT Device. In *ACM AINTEC* (2014).

[212] PACKET CLEARING HOUSE (PCH). Routing Origin Inconsistency. https://prefix.pch.net/applications/routing-origin-inconsistency/.

[213] PADMANABHAN, R., DHAMDHERE, A., ABEN, E., KC CLAFFY, AND SPRING, N. Reasons Dynamic Addresses Change. In *ACM IMC* (2016).

[214] PENNO, R., PERREAULT, S., BOUCADAIR, M., SIVAKUMAR, S., AND NAITO, K. Updates to Network Address Translation (NAT) Behavioral Requirements. RFC 7857 (Best Current Practice), Apr 2016.

[215] PERREAULT, S., YAMAGATA, I., MIYAKAWA, S., NAKAGAWA, A., AND ASHIDA, H. Common Requirements for Carrier-Grade NATs (CGNs). RFC 6888 (Best Current Practice), Apr 2013.

[216] PLONKA, D., AND BARFORD, P. Context-aware Clustering of DNS Query Traffic. In *ACM IMC* (2008).

[217] PLONKA, D., AND BARFORD, P. Assessing Performance of Internet Services on IPv6. In *IEEE ISSC* (2013).

[218] PLONKA, D., AND BERGER, A. Temporal and Spatial Classification of Active IPv6 Addresses. In *ACM IMC* (2015).

[219] POPA, L., GHODSI, A., AND STOICA, I. HTTP as the narrow waist of the future Internet. In *ACM HotNets* (2010).

[220] POSTEL, J. User Datagram Protocol. RFC 768 (INTERNET STANDARD), Aug. 1980.

[221] POSTEL, J. Assigned numbers. RFC 790 (Historic), 1981. Obsoleted by RFC 820.

[222] POSTEL, J. Internet Control Message Protocol. RFC 792 (INTERNET STANDARD), Sep 1981. Updated by RFCs 950, 4884, 6633, 6918.

[223] POSTEL, J. Internet Protocol. RFC 791 (Internet Standard), 1981. Updated by RFCs 1349, 2474, 6864.

[224] POSTEL, J. Transmission Control Protocol. RFC 793 (INTERNET STANDARD), Sept. 1981. Updated by RFCs 1122, 3168, 6093, 6528.

[225] PUJOL, E., RICHTER, P., AND FELDMANN, A. Understanding the Share of IPv6 Traffic in a Dual-stack ISP. In *PAM* (2017).

[226] QUAN, L., HEIDEMANN, J., AND PRADKIN, Y. Trinocular: Understanding Internet Reliability Through Adaptive Probing. In *ACM SIGCOMM* (2013).

[227] QUAN, L., HEIDEMANN, J., AND PRADKIN, Y. When the Internet sleeps: correlating diurnal networks with external factors. In *ACM IMC* (2014).

[228] RAMACHANDRAN, A., AND FEAMSTER, N. Understanding the Network-level Behavior of Spammers. *ACM CCR 36*, 4 (2006).

[229] REKHTER, Y., MOSKOWITZ, B., KARRENBERG, D., AND DE GROOT, G. Address Allocation for Private Internets. RFC 1597 (Informational), Mar. 1994. Obsoleted by RFC 1918.

[230] REKHTER, Y., MOSKOWITZ, B., KARRENBERG, D., DE GROOT, G. J., AND LEAR, E. Address Allocation for Private Internets. RFC 1918 (Best Current Practice), Feb 1996. Updated by RFC 6761.

[231] REYNOLDS, J., AND POSTEL, J. Assigned numbers. RFC 1060 (Historic), 1990.

[232] RICHTER, P., ALLMAN, M., BUSH, R., AND PAXSON, V. A Primer on IPv4 Scarcity. *ACM CCR 45*, 2 (2015).

[233] RICHTER, P., CHATZIS, N., SMARAGDAKIS, G., FELDMANN, A., AND WILLINGER, W. Distilling the Internet's Application Mix from Packet-Sampled Traffic. In *PAM* (2015).

[234] RICHTER, P., SMARAGDAKIS, G., FELDMANN, A., CHATZIS, N., BOETTGER, J., AND WILLINGER, W. Peering at Peerings: On the Role of IXP Route Servers. In *ACM IMC* (2014).

[235] RICHTER, P., SMARAGDAKIS, G., PLONKA, D., AND BERGER, A. Beyond Counting: New Perspectives on the Active IPv4 Address Space. In *ACM IMC* (2016).

[236] RICHTER, P., WOHLFART, F., VALLINA-RODRIGUEZ, N., ALLMAN, M., BUSH, R., FELDMANN, A., KREIBICH, C., WEAVER, N., AND PAXSON, V. A Multi-perspective Analysis of Carrier-Grade NAT Deployment. In *ACM IMC* (2016).

[237] RIGNEY, C., WILLENS, S., RUBENS, A., AND SIMPSON, W. Remote Authentication Dial In User Service (RADIUS). RFC 2865, 2000.

[238] RIPE NCC. Early Registration Transfer (ERX) Project. `http://www.ripe.net/lir-services/resource-management/erx`.

[239] RIPE NCC. Ingrid Wijte on the RIPE policy mailing list. (2012-05-23). `http://www.ripe.net/ripe/mail/archives/address-policy-wg/2012-May/006981.html`.

[240] RIPE NCC. IPv4 Transfer Listing Service. `http://www.ripe.net/lir-services/resource-management/listing`.

[241] RIPE NCC. IPv4 Transfer Statistics. `http://www.ripe.net/lir-services/resource-management/ipv4-transfers/table-of-transfers`.

[242] RIPE NCC. Policy Development Process in RIPE (ripe-614). http://www.ripe.net/ripe/docs/ripe-614.

[243] RIPE NCC. Policy for Inter-RIR Transfers of Internet Resources (ripe-644). http://www.ripe.net/ripe/docs/ripe-644.

[244] RIPE NCC. RIPE NCC Standard Service Agreement (ripe-533). http://www.ripe.net/ripe/docs/ripe-533.

[245] RIPE NCC. RPKI Statistics. http://certification-stats.ripe.net.

[246] RIPE NCC. Services to Legacy Internet Resource Holders (ripe-605). http://www.ripe.net/ripe/docs/ripe-605.

[247] RIPE NCC. Status of Legacy IPv4 Address Space (2011-09-12). https://labs.ripe.net/Members/xavier/status-of-legacy-ipv4-address-space.

[248] RIPE NCC. Post Depletion Adjustment of Procedures to Match Policy Objectives, and Clean-up of Obsolete Policy Text. http://www.ripe.net/ripe/policies/proposals/2013-03, 2013.

[249] RIPE NCC. IPv4 Address Allocation and Assignment Policies for the RIPE NCC Service Region (ripe-606). http://www.ripe.net/ripe/docs/ripe-606, 2014.

[250] ROSENBERG, J., KERANEN, A., LOWEKAMP, B. B., AND ROACH, A. B. TCP Candidates with Interactive Connectivity Establishment (ICE). RFC 6544 (Proposed Standard), Mar 2012.

[251] ROSENBERG, J., MAHY, R., MATTHEWS, P., AND WING, D. Session Traversal Utilities for NAT (STUN). RFC 5389 (Proposed Standard), Oct 2008. Updated by RFC 7350.

[252] ROSENBERG, J., WEINBERGER, J., HUITEMA, C., AND MAHY, R. STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs). RFC 3489 (Proposed Standard), Mar 2003. Obsoleted by RFC 5389.

[253] RUBI, E. M. The IPv4 Number Crisis: The Question of Property Rights in Legacy and Non-Legacy IPv4 Numbers. *American Intellectual Property Law Association (AIPLA) Quarterly Journal 39* (2011), 477.

[254] SALOWEY, J., AND DROMS, R. RADIUS Delegated-IPv6-Prefix Attribute. RFC 4818, 2007.

[255] SARRAR, N., MAIER, G., AGER, B., SOMMER, R., AND UHLIG, S. Investigating IPv6 Traffic. In *PAM* (2012).

[256] SCHINAZI, D. Apple and IPv6 - Happy Eyeballs. https://goo.gl/XBP9g4, 2015.

[257] SCHULMAN, A., AND SPRING, N. Pingin' in the Rain. In *ACM IMC* (2011).

[258] InMon – sFlow. http://sflow.org/.

[259] SHIN, M.-K., HONG, Y.-G., HAGINO, J., SAVOLA, P., AND CASTRO, E. M. Application Aspects of IPv6 Transition. RFC 4038 (Informational), Mar 2005.

[260] SINGH, H., BEEBEE, W., DONLEY, C., AND STARK, B. Basic Requirements for IPv6 Customer Edge Routers. RFC 7084, 2013.

[261] SKOBERNE, N., MAENNEL, O., PHILLIPS, I., BUSH, R., ZORZ, J., AND CIGLARIC, M. IPv4 Address Sharing Mechanism Classification and Tradeoff Analysis. *IEEE/ACM ToN 22*, 2 (2014).

[262] SOLENSKY, F. Continued Internet Growth. In *IETF 18* (1990), pp. 59–61. http://www.ietf.org/proceedings/18.pdf.

[263] SPAMHAUS. The Policy Block List. https://www.spamhaus.org/pbl/.

[264] SRISURESH, P., FORD, B., AND KEGEL, D. State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs). RFC 5128 (Informational), Mar 2008.

[265] SWITCH. Swiss Tele Communication System for Higher Education. http://www.switch.ch/.

[266] TANENBAUM, A. S., AND WETHERALL, D. J. *Computer Networks*, 5th ed. Prentice Hall Press, Upper Saddle River, NJ, USA, 2010.

[267] THALER, D., DRAVES, R., MATSUMOTO, A., AND CHOWN, T. Default Address Selection for Internet Protocol Version 6 (IPv6). RFC 6724, 2012.

[268] TIKAN, T. IPv6 Deployment in Estonia. https://goo.gl/vTQUpH, 2015.

[269] TOONK, A. How accurate are the Internet Route Registries (IRR). (2009-03-28). http://www.bgpmon.net/how-accurate-are-the-internet-route-registries-irr/.

[270] TOONK, A. Securing BGP routing with RPKI and ROA's. (2011-01-19). http://www.bgpmon.net/securing-bgp-routing-with-rpki-and-roas/.

[271] UNION, I. T. ICT Facts & Figures: The world in 2016. https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2016.pdf, May 2016.

[272] VALENTI, D., ROSSI, D., DAINOTTI, A., PESCAPÈ, A., FINAMORE, A., AND MELLIA, M. Reviewing Traffic Classification. In *TMA* (2013).

[273] WANG, L., AND KANGASHARJU, J. Real-world sybil attacks in BitTorrent mainline DHT. In *IEEE GLOBECOM* (2012).

[274] WANG, Z., QIAN, Z., XU, Q., MAO, Z. M., AND ZHANG, M. An Untold Story of Middleboxes in Cellular Networks. In *ACM SIGCOMM* (2011).

[275] WEIL, J., KUARSINGH, V., DONLEY, C., LILJENSTOLPE, C., AND AZINGER, M. IANA-Reserved IPv4 Prefix for Shared Address Space. RFC 6598 (Best Current Practice), Apr 2012.

[276] WING, D. NAT Tutorial. In *IETF 78* (2010).

[277] WING, D., AND YOURTCHENKO, A. Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555, 2012.

[278] WONG, B., STOYANOV, I., AND SIRER, E. G. Octant: A Comprehensive Framework for the Geolocalization of Internet Hosts. In *NSDI* (2007).

[279] WOODCOCK, B., AND FRIGINO, M. 2016 Survey of Internet Carrier Interconnection Agreements. Packet Clearing House. www.pch.net/resources/Papers/peering-survey/PCH-Peering-Survey-2016.pdf, November 2016.

[280] XIE, Y., YU, F., ACHAN, K., GILLUM, E., GOLDSZMIDT, M., AND WOBBER, T. How Dynamic are IP Addresses? In *ACM SIGCOMM* (2007).

[281] XU, Q., ERMAN, J., GERBER, A., MAO, Z., PANG, J., AND VENKATARAMAN, S. Identifying Diverse Usage Behaviors of Smartphone Apps. In *ACM IMC* (2011).

[282] YORK, D. *Migrating Applications to IPv6: Make Sure IPv6 Doesn't Break Your Applications*. O'Reilly Media, Inc., 2011.

[283] ZANDER, S., ANDREW, L., AND ARMITAGE, G. Capturing Ghosts: Predicting the Used IPv4 Space by Inferring Unobserved Addresses. In *ACM IMC* (2014).