Public Review for
Who squats IPv4 Addresses?

Loqman Salamatian, Todd Arnold, Italo Cunha, Jiangchen Zhu, Yunfan Zhang, Ethan Katz-Bassett, Matt Calder

This paper analyzes the phenomenon of squatted IP space: IPv4 addresses that operators use although they have not been allocated to them. This is possible because larger IPv4 blocks exist that have been allocated to organizations which never announced them in the global routing system. The authors draw on a very large data set of traceroutes and develop a heuristic to identify how squat space is used, by whom, and what the implications for Internet routing and the operator communities are. The reviewers appreciated the value of the analysis and the insights into squatted IP space and agreed on the paper's potential. For a revised version, they requested clarifications concerning to which degree the validation of the results can be performed, given that no ground truth is available, and further evaluation of the contributions that are possible thanks to the use of a non-public data set. As a result, the new version presented here provides a significant contribution that should be of interest to everyone with an interest in the operation of Internet routing and larger networks.

Public review written by
Ralph Holz
University of Twente

Who Squats IPv4 Addresses?

Loqman Salamatian
Columbia University
ls3748@columbia.edu

Todd Arnold[1]
Army Cyber Institute, West Point
todd.arnold@westpoint.edu

Ítalo Cunha
Universidade Federal de Minas Gerais
cunha@dcc.ufmg.br

Jiangchen Zhu
Columbia University
jz3268@columbia.edu

Yunfan Zhang
Columbia University
yz4244@columbia.edu

Ethan Katz-Bassett
Columbia University
ethan@ee.columbia.edu

Matt Calder[2]
Columbia University, Meta
mattcalder@meta.com

ABSTRACT

To mitigate IPv4 exhaustion, IPv6 provides expanded address space, and NAT allows a single public IPv4 address to suffice for many devices assigned private IPv4 address space. Even though NAT has greatly extended the shelf-life of IPv4, some networks need more private IPv4 space than what is officially allocated by IANA due to their size and/or network management practices. Some of these networks resort to using *squat space*, a term the network operations community uses for large public IPv4 address blocks allocated to organizations but historically never announced to the Internet. While squatting of IP addresses is an open secret, it introduces ethical, legal, and technical problems. In this work we examine billions of traceroutes to identify thousands of organizations squatting. We examine how they are using it and what happened when the US Department of Defense suddenly started announcing what had traditionally been squat space. In addition to shining light on a dirty secret of operational practices, our paper shows that squatting distorts common Internet measurement methodologies, which we argue have to be re-examined to account for squat space.

CCS CONCEPTS

• **Networks** → **Network measurement**; **Network architectures**; *Network structure*; Public Internet;

KEYWORDS

Internet Measurements; Internet Topology; IPv4 Squatting; CGNATs; IPv4 Exhaustion; IPv4 Utilization

1 INTRODUCTION

Demand for IP address space has risen dramatically in the last decade, resulting in the depletion of the IPv4 address space [4, 78]. To address this limitation, IPv6 addressing allows for 340 *trillion trillion trillion* globally unique IPs. However, after 25 years of availability, adoption of IPv6 is still slow [19]. This slow adoption can be traced to the early success of Network Address Translation (NAT), which extended the IPv4 availability through translation of *private IPv4 addresses* [77, 99]. Administrators favored NAT due to its simplicity and low administrative overhead to deploy—in contrast to the large effort needed to either reconfigure networks or support two very different IP addressing schemes with IPv6. NAT is common on Customer Premises Equipment (CPE) (*e.g.,* home broadband routers), and Carrier Grade NAT (CGNAT) moves this functionality into the ISP network where each customer is part of the ISP's private network. This entrenched preference for IPv4 has enabled a unique challenge—what to do if the allocation of private IPv4 addresses is still not enough?

To bypass the private IP space limitation, some networks turn to *squat space*, an open secret in the network operator community [63–66] referring to large blocks of *public* IPv4 address space allocated to organizations but historically never announced to the Internet. Organizations other than the owners use the unannounced addresses as if they were additional private IPv4 space. Squatting organizations generally filter squat space from their external Border Gateway Protocol (BGP) announcements, so it is rarely observed on the Internet outside of the occasional route leak.

Squat space usage introduces a number of unanswered technical, ethical, and legal challenges. Should the legitimate owner begin utilizing the previously unannounced space, the use of addresses outside the legitimate network can disrupt services and complicate operations and troubleshooting for both the legitimate owner and the sqatting organization. Further, legitimate routes can be filtered by the squatting Autonomous System (AS) or have a lower preference than its internal route, making legitimate services unavailable to portions of the Internet. Even if the squatter does not want to access services at the legitimate prefix, route leaks remain an ongoing problem [89] presenting continued risk to the legitimate owner. The FBI argues that squatting is a crime, as it constitutes fraud and theft of services [25]. In many cases, using squat space owned by the US Department of Defense (DoD) or the UK Ministry of Defense (MoD) has resulted in both confusion and security concerns for administrators who (incorrectly) believe their traffic to be routing through either organization [2, 10, 22, 60, 75, 76].

Despite these pitfalls and legal considerations, squat space commonly appears in traceroutes and route leaks. However, to what extent squat space is used in practice, and by whom, is not well understood. Until recently, most squat space would fail IP-to-AS Number (ASN) mapping [57] because it was seldom (if ever) announced to the Internet, causing traditional traceroute analysis to silently discard those hops. However, on January 20, 2021 huge swaths of unused IPv4 address space, allocated to the DoD, were announced by AS8003 with DoD approval [56, 93, 94] and on September 10 2021 changed to be announced by AS749 [55] bringing

[1]The views expressed herein are those of the authors and do not reflect the position of the US Military Academy, Department of the Army, or Department of Defense.
[2]The author was employed at Microsoft at the time of submission.

this long-time, but controversial, practice of squatting under new scrutiny. These new announcements grew to cover over 5% of the total IPv4 address space and constitute a major portion of historically squatted prefixes, presenting an ideal dataset to understand the impact of legit squat space announcements on squat space users.

In this work, we conduct the first large-scale public study of Internet squat space usage. Our specific contributions are:

Systematically identify existing squat space (§2.1): We consider the IPv4 /8 prefixes which are generally unannounced in global routing tables as potential squat space candidates. We examine 20 years of global routing tables and identify 17 such /8 prefixes as candidate squat space.

Explore in-the-wild squat space utilization at scale (§2.2): Our main dataset is 11.6 billion traceroutes destined for Microsoft by 15 million global clients in 52K ASes and more than 190 countries. To the best of our knowledge, our dataset represents the largest and most diverse sample of traceroutes processed among papers characterizing Internet scale behaviors in the wild. We supplement this dataset with publicly available measurements from RIPE Atlas [80] and Archipelago (Ark) [13].

Design heuristics to attribute squat space users (§3.1): We combine our traceroute datasets (§2.2) with BGP and organizational data (§2.3) to develop heuristics for attributing squat space usage to AS(es). We evaluate the coverage of every dataset considered and find that the publicly available datasets are insufficient to achieve the comprehensive view of squat space usage that the Microsoft dataset provides (§4). We find that the publicly available datasets only uncover 4% of the organizations observed squatting in the Microsoft dataset.

Analyze how squat space is utilized (§5): We analyze to what extent squat space is being used, what it appears to be used for, who is using it, and the frequency of route leaks involving squat space. In particular, we design a new technique to distinguish customer-side from ISP-side squat space utilization for NAT (§3.2). We use our technique to corroborate Richter *et al.*'s findings that squat space is utilized for configuring CGNATs [79] in more than 13K /24 prefixes hosted in 304 organizations, and extend their results by detecting squat space configurations by CPE devices in 15K /24 prefixes in 474 organizations. We find a small number of networks leaking squat space, and the leaked announcements are short and bursty. Of the leaking organizations that are part of our dataset, we observe 66% to be using squat space internally, 20% do not observe internal use; and 14% we were unable to make a determination.

Study the effect of a legitimate announcement on squat space users (§6): We provide the first assessment of the DoD announcing assumed squat space, which reveals the number of organizations squatting these addresses. We found that more than 70% of the ASes continued to squat DoD prefixes after the announcements. We also show an increase in the address ranges' use as squat space, despite DoD's active announcements. In particular, we find that 53% of the organizations that appear to be squatting the DoD's squat space after January were not observed to be squatting earlier.

Quantify the partial reachability of the IPv4 space caused by squatting (§6.3): To understand how valid public routes compete with internal squat space usage, we launch traceroutes from within squatting networks towards squatted addresses, finding different behaviors across ASes—some route towards the DoD, whereas others route towards internal routers. This discrepancy corroborates the risk to routing of leveraging squatted IP addresses.

Identify the implications of squat space utilization for the research and operations communities (§7): The widespread use of squat space that we observe can cause confusion among researchers and operators, since most topology measurements and analysis are not squat space aware. Topology mapping efforts tend to assume an IP address is used on a single interface in a single location by a single network. We demonstrate the use of squat space violates this assumption in a way that results in topological distortions. Squat space usage can also cause operational headaches when the true owner begins using the addresses, forcing squatting organizations to either renumber their networks or jump through hoops to isolate the different uses without partitioning the Internet.

2 DATA

Although squat space usage is acknowledged [44, 63–66], there has been no systematic method to search for it on the Internet. By understanding the underlying motivation for squat space, its desirable properties, and historic reports of it on forums and social media [22, 75, 102], we identified which large IPv4 address blocks should potentially be considered for analysis. We explored years of global routing table snapshots (§2.3) to determine which address blocks could be available for squatting, and determine criteria for address block suitability to be considered squat space (§2.1).

Squat space is rarely announced to the Internet, but addresses do show up in traceroutes. In order to identify and categorize squatting usage, we collected and analyzed billions of traceroutes from Microsoft, Ark, and RIPE Atlas (§2.2). The large volume of traceroutes allowed us to gather sufficient data to conservatively attribute squat space usage to individual organizations (§3.1) and determine how the attributed organizations are using squat space (§3.2).

2.1 Candidate Squat Space

RFC6752 defines squat space as "the practice of an ISP using address space . . . that has been officially allocated by a [Regional Internet Registry (RIR)] to another provider, but that provider is not currently using or advertising within the Internet [44]." Drawing from operational experience within the networking community, squat space is commonly restricted to unannounced /8 prefixes that have remained absent from global routing tables over time. Based on this observation, we focus our assessment of squat space to unannounced (RFC definition) /8 prefixes (per operators). We chose to look for /8 prefixes, but that decision limits our ability to detect improper use of other IP address space (*e.g.,* /16 prefixes). Hence, our work defines a lower bound of squatting as it occurs. In addition, we only investigate the allocated and historically unannounced addresses, since squatting has been defined as use of such addresses [44]; our study does not investigate other repurposed prefixes, such as those announced by a legitimate owner or reserved address blocks such as multicast addresses on 224.0.0.0/4, resulting in possibly missing other forms of IPv4 address abuse.

In our analysis we consider potential squat space to be the /8 prefixes with less than 20% of the address space announced on

/8 Prefix	Organization	Alloc.	% Ann.
6.0.0.0/8	Army Information Systems Center	1994-02	5%
7.0.0.0/8	DoD Network Information Center	1995-04	0 %
9.0.0.0/8	IBM	1992-08	0%
11.0.0.0/8	DoD Intel Information Systems	1993-05	≈ 0%
16.0.0.0/8	Hewlett Packard	1989-05	1.5 %
19.0.0.0/8	Ford Motor Company	1995-05	≈ 0%
21.0.0.0/8	DDN-RVN	1991-07	0%
22.0.0.0/8	Defense Information Systems Agency	1993-05	≈ 0%
25.0.0.0/8	UK Ministry of Defence	2005-08	≈ 0 %
26.0.0.0/8	Defense Information Systems Agency	1995-05	≈ 0 %
28.0.0.0/8	DSI-North	1992-07	0 %
29.0.0.0/8	Defense Information Systems Agency	1991-01	0%
30.0.0.0/8	Defense Information Systems Agency	1991-01	0%
33.0.0.0/8	DLA Systems Automation Center	1991-01	≈ 0%
43.0.0.0/8	WIDE	1985-01	43%
48.0.0.0/8	Prudential Securities Inc.	1995-05	≈ 0%
56.0.0.0/8	US Postal Service	1994-06	1%

Table 1: Chosen /8 prefixes resulting from our candidate selection process. The majority of the /8 prefixes were allocated in the early days of the Internet between 1985 and 1995 when the scarcity of IPv4 addresses was not a concern. Less than 2% of the address space within the /8 prefixes were announced, on average, from September 2020 to January 2021 as illustrated in Fig. 1. The only exceptions are 43.0.0.0/8 with around 43% of its space announced and 6.0.0.0/8 with 5%.

the global routing tables prior to January 2021. We selected these ranges based on customer complaints [2, 22, 60, 75, 76, 102] and operator discussions [7, 63–66].

Table 1 shows the official owners of the 17 /8 prefixes we select for analysis in our study. The prefixes listed in Table 1 were historically used either for research (*e.g.,* 43.0.0.0/8 [61]) or internal purposes (*e.g.,* DoD or IBM). Despite making up 7.7% of all public IPv4 address space, Figure 1 shows that only a tiny fraction of the prefixes were announced to the Internet over the past 20 years; the exceptions are 33.0.0.0/8 and 16.0.0.0/8 that were fully announced until 2012 and 2019. We note that 43.0.0.0/8 satisfies these criteria for all of 2021 but did not between September–December 2020; addresses were reassigned to APNIC in early 2020 [61] and bought by several major cloud providers [95] (*e.g.,* Alibaba bought space). While the 43.0.0.0/8 prefix originally saw an increase in the proportion of address space announced after the buy-out, the percentage of announcements has stabilized around 10% as shown in Figure 1. We therefore decide to consider 43.0.0.0/8 as part of the squat space despite the surge in announcements at the end of 2020.

Squat space announcements are sparse and stable. In Figure 1 we quantify the fraction of the announced address space within the /8 prefixes in Table 1. The set of announced addresses remains stable, with less than 6% of previously unannounced addresses becoming announced from one month to the next in RIPE RIS or RouteViews dumps. The 43.0.0.0/8 prefix observed the highest increase in the number of announced addresses as APNIC gradually allocated the /8 prefix to companies in the region (mainly Alibaba). The other prefixes with the largest increases are 9.0.0.0/8, 16.0.0.0/8, 22.0.0.0.0/8 and 56.0.0.0/8, three of which are owned by organizations not under the aegis of the DoD (Table 1). Their churn is often due to subprefix announcements changing from one month to the next. We observe

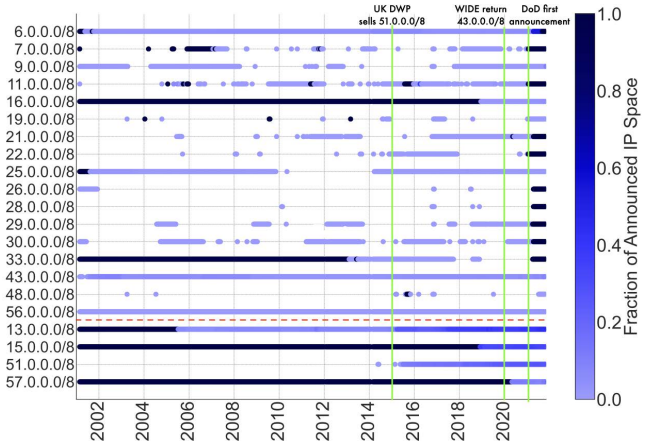

Figure 1: Historical proportion of announced subnets as seen from RIPE Routing History [82] for /8 prefixes we consider squat space (above the dashed line) and four select /8 prefixes with a historically low announcement rate (see §4.2).

a negligible fraction of the squat space being newly announced every new month, except for the AS8003 announcements at the beginning of 2021 (§6) and 43.0.0.0/8 being allocated by APNIC. We conclude that the /8 prefixes identified as squat space are not actively used in the public Internet.

No other /8 prefix should be considered squat space. To illustrate the insensitivity of our choice of threshold to define squat space, we also plot the four least announced /8 prefixes outside our defined squat space—13.0.0.0/8, 15.0.0.0/8, 51.0.0.0/8 and 57.0.0.0/8—in Figure 1. Although only 40% of 13.0.0.0/8, 15.0.0.0/8, 57.0.0.0/8 are announced at the time of our measurements, the entire /8 prefixes have historically been announced and would not be satisfactory squat space for operators. The only additional /8 prefix we could have considered is 51.0.0.0/8 which was formerly under the aegis of the United Kingdom Department of Work and Pensions but has been since slowly released to "reduce the deficit and help with the public finances" [8]. As this process happened years ago and the fraction of observable prefixes in the /8 prefix steadily increased, we made the decision not to consider 51.0.0.0/8.

2.2 Traceroute Datasets

Microsoft. Our first traceroute dataset contains paths from Microsoft clients to Microsoft's CDN collected with Odin [17]. For each month between 2020/09 and 2021/05 (except 2020/10 due to data loss), we analyze one-week snapshots totaling 11.6 billion traceroutes by 15 million clients in 52K networks across 194 countries. We limit our analysis to traceroutes containing IP addresses in the potential squat space from Table 1. For 2020 data, we work with a sampled dataset that includes only one traceroute per source IP address each month. We only use 2020 data in Section 6 to enable longitudinal results. Our 2021 data consists of all traceroutes that include a squat address, which we use in Section 4 and Section 5.

RIPE Atlas. We examine three weeks of RIPE Atlas public traceroutes [80]; one week from September 2020, one from January 2021, and one from March 2021. Each week includes more than 1.3 billion traceroutes from more than 10K probes hosted in around 5K

networks in 171 countries. We also issue our own RIPE Atlas measurements to estimate the impact of legitimate announcements of the former squat space on the internal routing of squatters (§6).

Ark. While our primary results rely on Microsoft and RIPE Atlas data sets, we also use Ark traceroutes from August 2020 to both quantify the Microsoft traceroute dataset's coverage and analyze the impact our IP-to-AS mapping (App. C). We analyze more than 390 million traceroutes to all routed /24 prefixes from August 2020 [13].

Speedchecker. We used Speedchecker [62] to assess the impact of the announcements of the DoD's squat space on the routing of squatters (§6). Speedchecker was a measurement network where users could download an app onto their mobile phone to test their Internet connections' bandwidth and voluntarily opt in to host measurements. Previous studies used Speedchecker to examine performance and inter-network connectivity and showed that it had better coverage of user networks than RIPE Atlas [5, 21, 29]. Unfortunately, Speedchecker no longer supports API access for issuing measurements as of June 30, 2021.

2.3 BGP Datasets

BGP dumps. We gather BGP announcements and routing table dumps from more than 925 routers in 379 ASes hosted in 24 facilities made available by RIPE RIS [81] and Routeviews [69]. In this paper, we collect BGP announcements observed on 2021/01/19 and 2021/01/21 (*i.e.,* before and after AS8003 started announcing DoD's prefixes on 2021-01-20) and routing table dumps for every Tuesday from August 2020 to May 2021. We also gather a week of routing tables at the beginning of every month from August 2020 to May 2021 to define the set of potential squat space (§2.1).

RIPE Routing History. We collect the RIPE Routing History Dataset [82] for the prefixes of interest (§2.1, Table 1) and their subnets. This dataset compiles all the different ASes that have announced any routable public IP prefixes forwarded by any RIPE RIS peer. The data is aggregated in 13-day time bins, which is the smallest time unit considered in the Routing History Dataset.

AS and organization metadata. To understand the size of networks we observe squatting, we rank organizations in three ways: by their (i) customer cone [14], (ii) number of users [3], and (iii) amount of assigned IPv4 address space [28], aggregated across all their ASes. To accomplish this, we augment our BGP information with AS metadata from various sources. We crawl whois registries from September 2020 and May 2021, and associate each AS with its country of registration, name, and administrative organization (§3.1). We use APNIC's per-AS users estimates [3] and customer cone data from CAIDA's AS-relationships dataset [36, 51]. Our BGP dumps tell us the number of /24 prefixes announced by each AS and we use the type of the AS as self-reported in PeeringDB [72].

3 METHODOLOGY

We look for squat space IP addresses (§2.1) showing up in our various traceroute datasets (§2.2). For each instance, we try to identify which ISP is using the squat address (§3.1) and what it is using it for (§3.2).

3.1 Squat Space Usage Attribution

As squatting raises legal and ethical concerns, attributing squat space utilization to an organization requires careful consideration. Due to the potential sensitivities, we design robust techniques to attribute squat space utilization for our set of traceroutes resulting in *strong attributions*. To gauge the impact of our cautious approach, we also contrast the difference in attribution if we use a less stringent method to obtain *weak attributions*.

Strong Attribution. To discover which networks use squat space, we search for instances where the same organization "surrounds" a squat space IP address in a traceroute according to the following:

(1) **Discard hops in Microsoft.**[3] To avoid leaking potential proprietary information, we discard all hops after the first IP address inside Microsoft's network.
(2) **Set source /24 prefix.** We insert the client IP address observed by Odin [17], RIPE Atlas [80], or Ark [13] as hop zero, allowing us to identify from which organization the traceroute originated. In the case of a NAT, the IP address observed by Odin is the public IP address, and the traceroute actually starts "behind" that IP address, which we use to identify NAT usage (§3.2).
(3) **Map traceroute IP addresses to organizations.** We map IP addresses to ASes using standard techniques [6, 35]. We map IXP fabric addresses to the corresponding member AS using data from PCH [37], PeeringDB [72], and Euro-IX [41], and map the remaining public, non-squat addresses, to the AS originating the longest covering prefix in RIPE RIS BGP snapshots [81]. Appendix C compares this approach to using *bdrmapIT* and finds that they yield identical results in more than 99% of cases. The limited differences in results are not worth the added complexity of deploying *bdrmapIT* within Microsoft's data processing pipeline. We map the resulting ASes to organizations by first applying CAIDA's AS-to-Org dataset [15] and using WHOIS as a fallback.
(4) **Squatter identification.** If a traceroute includes a contiguous sequence of hops with squat space IP addresses surrounded by public IP addresses mapped to the same organization, we attribute the squat usage to that organization. Traceroutes to unassigned addresses sometimes induce routing loops between neighboring ASes, which could cause false attribution in the RIPE Atlas and Ark datasets. However, the Microsoft traceroutes target responsive Microsoft hosts and so are not impacted.
(5) **Squatter disambiguation.** In cases where a sequence of hops with squat addresses is surrounded by two different organizations, we attribute the use of the squat space only if one of the organizations has been identified as squatting the same prefix for a different traceroute in step (4).
(6) **Removing legitimate owners.** When the squatters are the organization or sibling organizations that were legitimately allocated the prefix, we do not consider them as squatters. To do so, we manually pick all the ASes associated to the organizations in Table 1.

[3] Only applicable to Microsoft's dataset

Weak Attribution. While our technique to attribute squat space usage requires the IP addresses surrounding the observed squat space to belong to the same organization, we are also interested in ambiguous cases in order to identify additional potential squatters. Our strong methodology for positively identifying squatting organizations builds from restrictive conditions, but by relaxing some of our heuristics we can identify an additional pool of potential squatters. We design a technique to attribute squat usage when the squatted IP address is ambiguously located between two organizations with no other attributed squat space utilization. To delineate the two organizations that are adjacent to an ambiguously located squatted IP, we consider the first public non-squatted IP between the squatted IP address.

We count the number of hops per organization where each appears to be adjacent to a squat IP address with inconclusive attributions—collectively referred to as *ambiguously mapped organization* and *ambiguously mapped hops*. A less conservative approach would consist of blaming the smallest set of *ambiguously mapped organizations* such that all *ambiguously mapped hops* can be attributed. This problem can be formulated as an instance of *hitting set*. In this formalization, we define a set for each ambiguously mapped organization, where the set consists of all the ambiguously mapped hops that are adjacent to a hop of the organization. The problem then consists of choosing the smallest number of subsets to hit each ambiguously mapped hop. While this problem is known to be NP-hard, a greedy approach provides a sensible approximation to the optimal solution [26].

We use this greedy algorithm to identify a collection of organizations that we name the *weakly attributed* squatters. However, the technique suffers from two drawbacks that reduce our confidence in the attribution process compared to the strong attribution method. First, weak attribution makes the assumption that measurements traversing *ambiguously mapped organizations* are equally distributed such that the entities that are often located on *ambiguously mapped traceroutes* are more likely to be squatting. There is no way to prove that this assumption holds. So, weak attribution may skew the attribution process toward a few large transit networks that appear on many traceroutes, or towards organizations that source many traceroutes. Second, some of the private, unmapped, and squat IP addresses separating the two mapped IP addresses might belong to ASes that do not appear explicitly in our traceroutes because, for example, their routers are configured to respond to traceroutes with private and squat IP addresses. Therefore, the weak attribution process could incorrectly identify an organization, missing an invisible one. We discuss the gains in terms of organizations added through this technique in Section 5, but the rest of our method relies on the more trustworthy strong attribution.

3.2 Inferring How Squat Space is Being Used

Even when we observe squat space within an ISP, as was already shown in Richter *et al.* [79], it does not necessarily indicate that the ISP uses squat space. An ISP customer (*e.g.,* home or business) using the ISP's public address space may configure their network (and especially NAT) using squat space instead of private address space. We now discuss our approach to distinguish systematic ISP use of squat space from customer use. Our key idea is to correlate observations of the ISP's infrastructure from multiple sources within the same prefix. In this section and throughout the rest of the paper, we adopt the same naming convention to describe the different NAT types as in Lutu *et al.* [54].

CGNAT vs NAT. Many Internet users are hosted behind a NAT. When there exists a single layer of translation, *i.e.,* a single middlebox maps from a device's private IPv4 address to a public IPv4 address, we obtain what is called a *NAT44* [100]. When a customer's networking equipment (CPE) performs NAT once, and an ISP's device performs NAT again (CGNAT), we obtain nested NATs (Fig. 2, NAT444 Architecture). The CPE's NAT translates the source's private or squat address into another private or squat IP address, and the CGNAT finally translates the source address into the public IP address. This scenario is abbreviated as *NAT444* [99]. In this work, we use the discovery of two nested NATs as an indication that the carrier deploys a CGNAT on top of the customer NAT (Fig. 2, NAT444 architecture). This is important as we are focusing on uncovering whether the customer or carrier is responsible for the squat space usage behind a NAT gateway.

This problem is not trivial as a wide diversity of network configurations exist that would exhibit the same results when derived from a traceroute measurement. Consider the case of a WiFi router performing NAT. A traceroute originating from a device behind the router could result in a single unrouted IP address followed by either (i) public IP addresses (suggesting NAT44) or (ii) other unrouted IP addresses (pointing towards NAT444). However, naively labeling case (ii) as NAT444 may be incorrect; an enterprise network could configure its intranet's routers using unrouted IP addresses and deploy a single NAT44 at the border. In this case, traceroutes from sources within the enterprise would observe multiple private IP address ranges before reaching the public Internet and a naive approach could incorrectly infer two levels of NAT44, *i.e.,* NAT444. Another example is a scenario where a first NAT is performed by a WiFi router within the enterprise, and a second NAT is performed at the egress point toward the ISP, resulting in the source address being translated by the enterprise twice. In this scenario, we do have a NAT444, but it is managed by the enterprise; *i.e.,* the ISP does not perform NAT which we can attribute the squatting to.

The identification of the squat space user requires differentiating between squat space for NAT44 configuration and for NAT444 deployed at an ISP's clients or by the ISP itself. Furthermore, in our dataset, n end-users behind a single-hop NAT issuing one traceroute each will share a single source address and appear identical to one user issuing n traceroutes. To infer the existence of NAT444 and who is responsible for deploying it, we craft a technique that takes support from observations made by analyzing many traceroutes:

- The majority of devices behind a NAT are likely configured with unrouted addresses (squat space or private space), as the use of any routed IP address would render legitimate destinations inaccessible.
- Some deployments directly connect end devices to the middleboxes performing the NAT (*e.g.,* a home Wi-Fi router), but large entities such as companies might rely on traversing multiple devices before reaching the NAT.

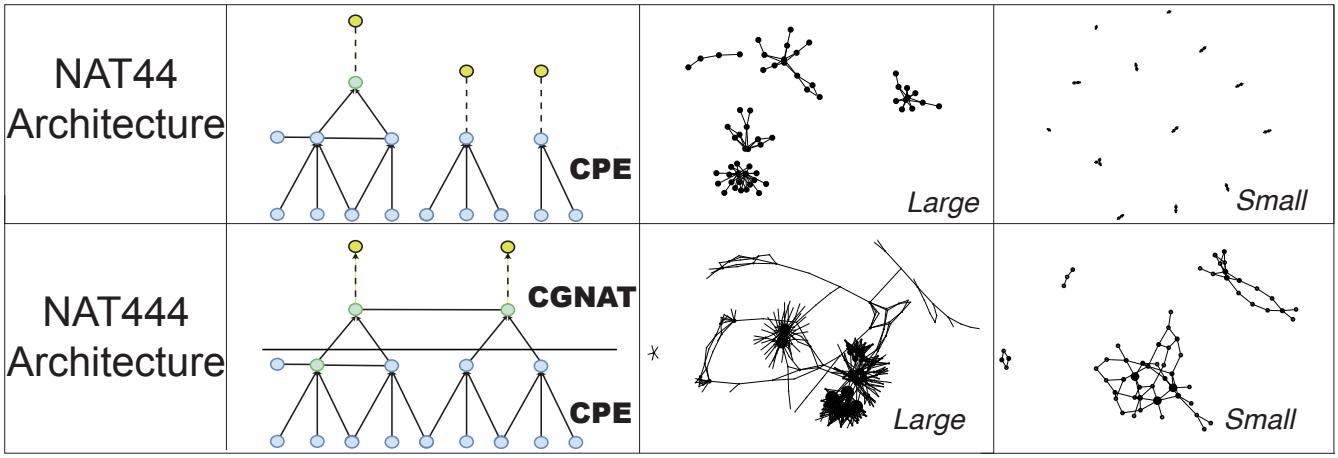

Figure 2: Diagram summarizing the two representations of the topologies that we aim to dissociate. On the left, we have a schematic representation of the traceroutes. Each node models an IP address, and each directed edge indicates the existence of a traceroute where two nodes are adjacent. Blue and green indicate different prefixes of private IP addresses, yellow represents public IP addresses, and the dashed line indicates the last link that would be considered with our technique. On the right representation, the scales of the nodes are based on degree. We note that the graph representation of the CGNAT exhibits (i) more depth to the first hop behind the CPE, (ii) fewer disconnected components, (iii) fewer sinks. We leverage the number of nodes and edges to categorize the observed magnitude of the NAT.

- NAT44 and NAT444 deployed by a single enterprise have simpler intradomain topologies when compared to NAT444 configured by ISPs, which may span a large geographical area and host multiple large enterprises.

The above observations imply that crowdsourcing measurements from sources within a /24 prefix provides us information on the topology behind the NAT and helps differentiate cases of NAT444 in enterprises from CGNAT deployed at ISPs. We use the /24 prefix granularity for CGNAT detection similarly to all previous works we are aware of [48, 79]. Furthermore, the Microsoft dataset removes the last octet of the source IP address. We discuss the trade-off of utilizing the /24 granularity in more detail in Section 3.2 (vi).

Internal connectivity graph. To capture the topology behind the NATs, we build an *internal connectivity graph*, denoted G, from traceroutes originating within a /24 prefix. Additionally, it is the smallest prefix block that is typically announced in the global routing table [54]. We discard all the traceroutes where the first public IP address maps to a different AS than the source AS. For such cases, the utilization of the squat space addresses becomes ambiguous as they could have been used to configure border routers, or worse, routers in a completely "invisible" AS with only private and squat responses to ICMP probes. The internal connectivity graph, G, has one vertex for each unrouted IP address observed in traceroutes from the /24 prefix. We consider hops from the source up to the first non-squat public IP address, *i.e.,* the graph obtained by only considering private, shared,[4] and squat IP addresses. Directed edges connect vertices that appear in consecutive hops. We now describe the steps of our approach:

[4]The 100.64.0.0/10 prefix is defined as "shared" and is intended "for ISP CGNAT deployments and NAT devices that can handle the same addresses occurring both on inbound and outbound interfaces" [44].

(i) *Calculating important graph statistics:* We compute the following statistics for each G: the number of vertices, the number of sources (*i.e.,* vertices with no incoming edge), the number of sinks (*i.e.,* vertices with no outgoing edge), the number of edges, the number of connected components, the number of edges in the minimum spanning tree of the largest component, the maximum path length across all source-sink pairs, as well as the number of squat /8 prefixes and private subnets observed by the traceroutes. We do not leverage any information relating to RTTs as they will be used to evaluate the performance of the technique in §3.2.

(ii) *From graph statistics to configurations:* Our intuition to dissociate NAT44 from NAT444 can be summarized as follows: NAT444 will result in fewer disconnected components, longer paths (sources and sinks might be in diverse locations), more nodes and edges (traceroutes are likely to traverse more routers within the ISP), and larger minimum spanning tree (*i.e.,* rich internal topology). NAT44 graphs likely comprise several disconnected components (one component per customer in the /24) with short paths between end-user devices and the NAT.

(iii) *Clustering to detect configurations:* We apply a k-means clustering on the normalized feature statistics where k is selected by minimizing the gap statistic [92]. We translate different clusters into three main categories: (a) Unknown, (b) NAT44 (c) NAT444. We allow for four more specific sub-categories of *small NAT44*, *large NAT44, small NAT444* and *large NAT444* to better discern vast deployments from smaller ones. The distinction matters as large NAT44 and small NAT444 can result in similar observed topologies.

(iv) *Leveraging centroids to define configurations:* We discuss the different clusters that the k-means algorithm has recovered and how we interpret the centroids as concrete NAT deployments.

(1) Unknown: When the centroid of the cluster possess few nodes or edges ($\leq$ 10)
(2) Small NAT44: When the centroid of the cluster results in a large number of disconnected components ($\geq$ 40), a small median hop distance from a source to its sink ($\approx$ 2) and a large number of sinks ($\approx$ 30), we conclude that multiple *small* NAT44s are in the /24 prefix.
(3) Large NAT44: When the centroid of the cluster results in a similar number of connected components ($\approx$ 40) and many sinks ($\approx$ 40) but with more nodes (as compared to NAT44) and a higher median distance between a source and a sink ($\geq$ 8), we interpret the cluster as containing /24 prefixes that each have a few large NATs.
(4) Small NAT444: When the centroid of the cluster returns fewer connected components ($\leq$ 20) and fewer sinks ($\leq$ 20), a large number of nodes and edges ($\geq$ 100) and other statistics at similar value as the large NAT44, we flag the cluster as a small NAT444.
(5) Large NAT444: When the centroid of the cluster returns long paths from each device to the first public IP address ($\geq$ 10), more edges and nodes ($\geq$ 250), more sinks than the Small NAT444 ($\geq$ 30) and several different unrouted prefixes ($\approx$ 3), we infer that the /24 prefix is behind a NAT444.

(v) *Interpreting the graphs:* In Figure 2, we draw a few examples of internal connectivity graphs and summarize how we leveraged the centroid to transform graph statistics to specific NAT deployments. The examples are selected because they are configurations with similar statistics as four centroids obtained by our k-mean clusters. In small NAT44 (upper right), we notice many disconnected components, each with few nodes, indicating simple topologies with devices performing NAT close to end-user devices. The small NAT444 (lower right) and large NAT44 architecture (upper middle) are similar to each other and hard to distinguish, so we rely on the number of private and unrouted IP /8 prefixes to distinguish between them. Finally, the large NAT444 (lower middle) includes a few large connected components with large spanning trees. The number of vertices and edges coarsely approximates the number of routers and the topology complexity behind the NATs. We label small NAT44 and large NAT44 as CPE configurations and small NAT444 and large NAT444 as CGNATs.

(vi) *Inference granularity:* We build graphs and make inferences at the finest granularity possible with our dataset: /24 prefixes. We use /24 prefixes because our Microsoft dataset does not include the last octet of source IP addresses to protect client privacy.

NAT and CGNAT deployments across prefixes less specific than a /24 prefix (*e.g.,* a /23 prefix) would be partitioned into multiple smaller graphs. Given our inference model is at the /24 prefix granularity, inferences for each /24 prefix within the less specific prefix should correctly reflect the address space use (see §4.2). A large NAT or CGNAT deployed on a prefix less specific than a /24 prefix would simply result in multiple inferences at the /24 prefix granularity.

NAT and CGNAT deployments in prefixes more specific than a /24 prefix (*e.g.,* /25 or /26) will be aggregated to the /24 level. If all the /24's subprefixes have the same use (*e.g.,* NAT or CGNAT), then we expect the resulting graph for the whole /24 prefix will have similar properties to that of a single deployment across the /24 prefix, and expect our inferences to correctly reflect address space's use. However, if a NAT or CGNAT deployment uses a more specific prefix and the subprefixes within the /24 prefix have different uses, then the resulting graph for the whole /24 prefix will have mixed properties, possibly leading to an incorrect inference. In particular, if the other subprefixes do not deploy private or squat addresses, then the resulting graph would be smaller than expected, which could lead to underestimating the scale of NAT and CGNAT deployments.

(vii) *Impact of intradomain topology.* Our inferences are sensitive to the observed intradomain topology at the origin ISP and its customers. In particular, our technique employs traceroute measurements and targets networks using private, shared or squat IP addresses. In particular, traceroutes are unable to observe a network's intradomain topology when only layer-2 devices are used between CPEs and the CGNAT device. Our graph-based inferences do not apply to these networks. Similarly, our technique assumes that routers carrying the traffic from the devices performing the client NAT translation to the router performing the CGNAT translation are only configured with private, shared, and squat IP addresses. When the routers steering traffic from the first NAT device to the second NAT device in NAT444 are configured with public IP addresses, our technique may incorrectly attribute them as internal routers as defined in Section 3.2.

Specific network topologies that do not align with the general NAT and CGNAT deployments may also lead to incorrect inferences from our graph-based approach. For example, our heuristics may fail to identify smaller CGNAT deployments where all sources traverse the same route through the organization or may misidentify a large geographically-dispersed wide-area NAT deployment as CGNAT if the intradomain topology becomes complex enough.

Despite these limitations, we believe our techniques will accurately identify large CGNAT deployments. While we do not possess a ground truth dataset, we evaluate our inferences using observations and insights from previous research (§4.2).

Routers and middleboxes. We infer that a sequence of squat IP addresses in a traceroute are assigned to *routers* or *middleboxes* if two public IP addresses surround the sequence, ignoring the source address. We further split these inferences into (i) *internal* if the surrounding public IP addresses belong to a single organization, (ii) *border* if the surrounding public IP addresses belong to different organizations, or (iii) *unmapped* when at least one of the surrounding public IP addresses cannot be mapped to an organization.

Unclassified traceroutes. The remaining traceroutes, *i.e.,* those with squat addresses after a public IP addresses but surrounded by at least one unresponsive hop and whose next visible public IP address is not in the same organization, we label *unclassified*.

4 REPRESENTATIVENESS, COVERAGE, AND EVALUATION

4.1 Coverage

Our Microsoft dataset is necessary to achieve good coverage of squat space utilization. While we strive to use public data when possible, this section demonstrates the need to include the Microsoft dataset to obtain sufficient coverage and a rich understanding of squat space usage.

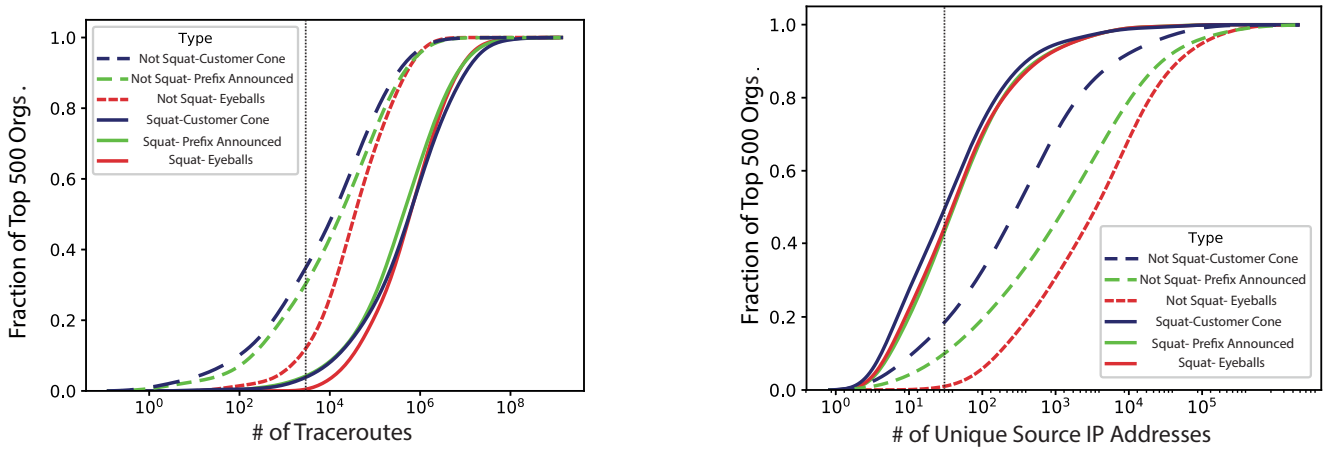

(a) Traceroutes **(b) Unique Source IP**

Figure 3: CDF of the number of distinct (a) traceroutes and (b) public source IP addresses for the top 500 organizations for all three rankings (customer cone size in blue, number of prefix announced in green and number of eyeballs in red) split by whether we identified use of squat space (dashed line for squatters). Organizations, where we do not observe squatting, possess fewer traceroutes from less diverse sources, suggesting the need to define a threshold on the number of traceroutes and unique source IP addresses to infer the absence of squat space in an organization. We draw a dotted vertical line at our selected filtering threshold as discussed in Section 4.1.

		Microsoft	RIPE Atlas	Ark
		General Coverage		
Pre-Filter	Total Source ASes (% of all ASes)	52,670 (74.3%)	6,862 (9.7%)	177 (0.2%)
	% of User Population	99.9 %	81%	6.4%
Post-Filter	Total Source ASes (% of all ASes)	10,283 (14.4%)	27 (0.1%)	0
	% of User Population	96.1%	5.4%	0
		Squatting Subset		
Traceroutes crossing Squat Space		87M	1,577K	146K
Source /24 prefixes		43K	8,867	13
Source ASes		3,952	2,381	13
/24 prefixes inside squat space		8,138	670	355
Traceroutes with strong attributions		69M	561K	16K
Identified squatting organizations		2434	108	57

Table 2: Coverage of the Microsoft, RIPE Atlas and Ark datasets before (top) and after (bottom) removing the ASes with fewer than 500 unique source IP addresses or 5000 traceroutes (§4.1 and Fig. 4). Applying the filter returns a subset of organizations where our datasets contain sufficient diversity to infer squat space utilization. We describe the exact process to generate the filter in Section 4.1. The population estimates are based on the APNIC per-AS dataset and total number of ASes is from CAIDA's January 2021 AS dataset. To achieve comprehensive coverage, the Microsoft dataset is necessary.

General coverage: In Table 2, we detail each traceroute dataset's coverage statistics. We use APNIC's population dataset [3], which uses ad-based measurements to estimate Internet user populations at the AS-level granularity. According to the APNIC dataset, RIPE Atlas and Ark, combined, have nodes in 6,891 *source ASes* hosting ≈82% of *Internet users*, whereas our Microsoft traceroutes come from 52,670 source ASes hosting ≈99.9% of users. This provides the coverage necessary for a global assessment of squat space usage.

Observing the squat space: Table 2 contrasts the view of squat space across the three different traceroute datasets. Whereas the table presents statistics for each dataset on its own, our discussion compares the coverage of the Microsoft dataset to the combined coverage of the Ark and RIPE Atlas datasets, to emphasize the scale and necessity of the Microsoft dataset compared even to the combination of large, widely used public datasets. We find that the Microsoft dataset contains 51× more traceroutes crossing squat space from vantage points spread across 4.9× more source /24 prefixes in more source ASes than the other two datasets combined. The Microsoft traceroutes traverse 23× more unique /24 prefixes inside squat space. The dataset also has a higher ratio of traceroutes with squat space attributable to an organization (traceroutes with attributions, §3.1), which results in 20× more identified squatting organizations. The Microsoft dataset can identify 70% of the squatting organizations identified by the RIPE Atlas and Ark datasets (not shown in the table). On the other hand, the RIPE Atlas and Ark datasets can only recover 4% of the squatting organizations observed in the Microsoft dataset. Because Ark and Atlas traceroutes to unassigned addresses may occasionally induce false attribution for those datasets (which does not happen for the Microsoft dataset (§3.1)), the results represent an upper bound on Atlas and Ark attribution coverage and hence a lower bound on the already substantial incremental benefit of using Microsoft data.

How many traceroutes are required to uncover squatting? While detecting a squatted address in an organization indicates that either the organization or one of its customers is using squat space, our dataset may not observe squat space in an organization because the organization does not use squat space (*i.e.,* true negative) or because we have insufficient traceroutes and miss its use of squat space addresses (*i.e.,* false negative). Figure 3 shows the CDF of the number of traceroutes (Fig. 3a) and unique source IP addresses (Fig. 3b) from the top 500 organizations, for the three rankings (§2.3), divided by whether we identified the organization

as using squat space. The top 500 organizations for each ranking cover at least 47% of users, 45% of IP addresses announced, and their customer cones cover 43% of the Internet's ASes; results for all organizations are qualitatively similar, with less pronounced differences between squatting and non-squatting organizations. For example, Figure 3a shows that we have more than 5000 traceroutes from 98% of the top organizations with the most users where we identify squatting (solid red curve, Squat-Eyeballs). Figure 3 shows that we have more traceroutes from squatting organizations (Fig. 3a) even though their traceroutes are issued from a smaller number of source IP addresses, which indicates that addresses in squatting organizations tend to issue more traceroutes than in non-squatting organizations, which could potentially indicate, for example, use of NAT or CGNAT.

We design a bootstrapping strategy to infer when there are enough measurements and enough source IP addresses from an organization to decide whether the organization does not use squat space with strong confidence. Specifically, we fix Z to be a binary property that reflects whether an organization uses squat space, X as the number of measurements, and Y as the number of sources. Furthermore, we designate $\hat{X}$ as the number of traceroutes with squat space addresses and $\hat{Y}$ as the sources whose traceroutes also observe addresses in at least one squatted /8 prefix, respectively. We are interested in finding thresholds m and n such that the likelihood of missing squatting after collecting m traceroutes from n sources, *i.e.,* $P_{\text{miss}} = P(\hat{X} = 0 \wedge \hat{Y} = 0 | Z = 1, X \geq m, Y \geq n)$, is small. We estimate P_{miss} by bootstrapping on m and n (*i.e.,* choosing n random sources and m random traceroutes from these sources) across all organizations. We iterate over a wide range m, n and $(m \wedge n)$ to estimate P_{miss}. Making $P_{\text{miss}} < 0.01$, we identify that $m = 5000$ traceroutes and $n = 50$ source addresses are good candidates as illustrated in Figure 4. Hence, in the rest of this paper, we only consider organizations with at least 5000 traceroutes and 50 different vantage points in our dataset, and deem all other organizations to have insufficient measurements to assess whether they are squatting. Although such requirements limit coverage, as is illustrated in *# source ASes after Filter* row in Table 2, this step boosts confidence in our inferences and the representativeness of our observations. Furthermore, the filter drastically reduces the number of source organizations covered from 7K to 27 for RIPE Atlas and Arkipelago, highlighting the importance of Microsoft's data for our analysis. We observe that we have enough measurements from more than 95% of the top 500 organizations based on allocated IP addresses (Fig. 5) Similar observations can be made for the highest number of end-users and largest customer cone in Figure 12 in Appendix B.

4.2 Evaluation

We confirm our observations with independent reports of active squat space usage from operators and present additional evaluation to support the validity of our technique to distinguish customer-side from ISP-side squat space utilization. Since we do not possess ground truth regarding squat space utilization, and we believe survey-based validation to be unreliable in this context (§7), confirming the accuracy of any approach is difficult. We are able to share anecdotal insights on how squat space can be deployed based on conversations regarding its usage by a large cloud provider and

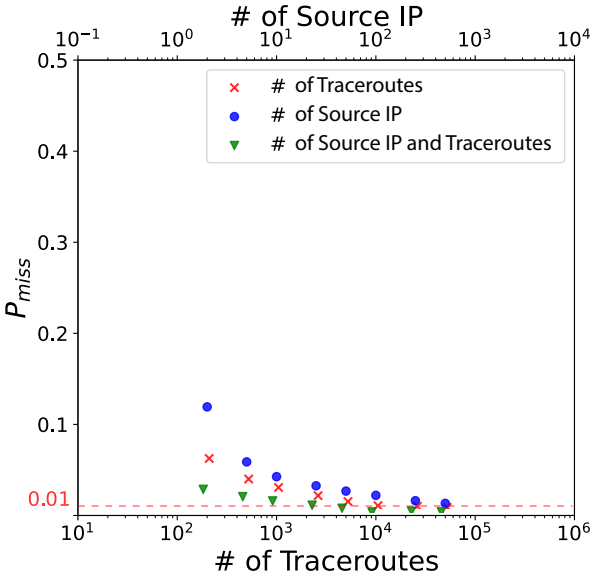

Figure 4: Computation of P_{miss} for different m (on the lower x-axis) and n (on the upper x-axis). In red (blue), we apply a single filter based on the number of traceroutes (number of sources) and no restriction on the other number. In green, we require both of the filters at the same time. We observe a decrease in the probability of missing squat space usage when we increase m and n, illustrating the need to have more than 5000 traceroutes and 50 sources to reduce P_{miss} to 0.01.

a large ISP. The exchanges provide a method for us to evaluate our attribution technique's accuracy in a small subset of cases.

We also compare our results to methodologies based on latency [54], reverse DNS (rDNS), and open port scanning. The latency technique provides circumstantial evaluation, but applies to all prefixes in our dataset. The rDNS and open port scanning techniques were used recently as ground truth to infer particular IP address usage [27, 34, 42], but apply to a restricted subset of the /24 prefixes in our dataset–only hosts with configured and informative rDNS names or that have hosts with reachable open ports. To address the unreliable nature of latency and limited coverage for the strong indicators from rDNS and open ports, we demonstrate that our technique performs well for every indicator. The strong performance across indicators independently strengthens confidence in our method's accuracy.

Operators anecdotally confirm squat space utilization. Two anecdotes from operators serve the roles of validating the existence of squat usage and some of our inferences in specific deployments. We have spoken to a very large cloud provider and received confirmation that the provider uses squat addresses to support cloud services. Interestingly, our contact could not specify when the squatting began, further confirming the practice's undocumented nature, even within large organizations. The squat space is used under a private agreement with the government to which the address space is allocated, and the addresses show up in traceroutes. This anecdote confirms our claim that squat addresses are used for other

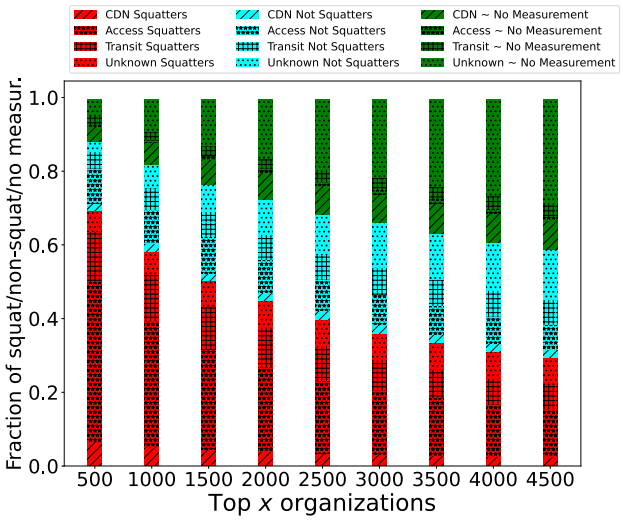

Figure 5: Proportion of squatting organizations that appear on the top x by number of IP addresses allocated. No measurement means that there are less than 5000 traceroutes and 50 sources from the network. More than 60% of the top 500 largest organizations have been observed squatting. This proportion reduces as smaller organizations are considered. We also note that we have 70% of the organizations in the top 2000 with at least 5000 measurements from 50 sources.

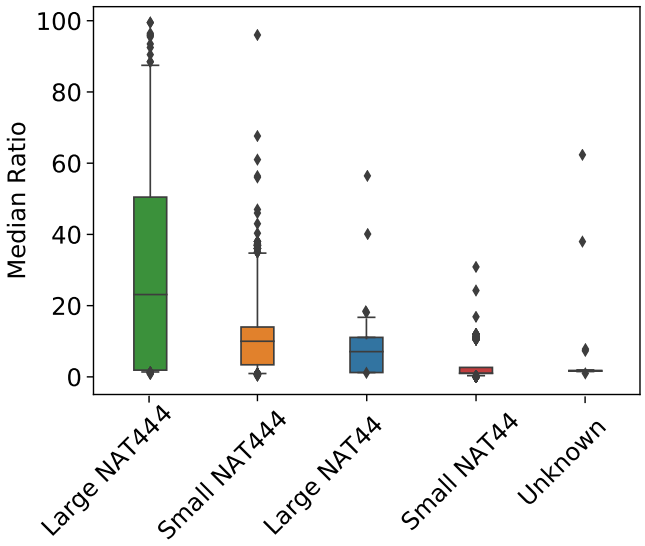

Figure 6: Boxplot of the RTT ratio between the vantage point's gateway and the devices performing the CPE NAT and the CGNAT of all the /24 prefixes within our deployment types. We observe that, on average, the prefixes we categorize as CGNAT have larger ratio of round trip time, matching the observation of earlier work [54] that CGNAT devices are further away from user devices than NAT devices.

uses in addition to customer-side NAT. Furthermore, Richter *et al.* had already observed that unrouted IP addresses appear in the traceroutes on the client-side of NATs [79]; this observation from prior work also confirms our assertion that squat space is used in the context of ISPs and their middleboxes configurations.

We also spoke with a very large broadband ISP where we saw extensive squatting, all of which our methodology attributed to customer-side use (rather than CGNAT, routers or middleboxes). Our contact said that (a) they are not aware of the ISP using any of the squat prefixes, (b) two of the examples we shared were business customers using squat space inside the business, and (c) the third example we shared was a multi-homed customer routing through a different provider, even though the source address belonged to the broadband ISP. Those anecdotes verify that our attribution approach correctly identified customer-side use in this ISP, and example (c) shows the importance of our strong attribution looking at the public addresses after the squat address to avoid "blaming" the broadband ISP.

These anecdotes establish that there are operators who agree that squatting takes place and that some configurations could trip up simple attribution approaches, but our techniques provide correct attribution in the instances we shared with the broadband ISP. Although this effort does not result in large-scale validation, it allows us to validate some of our observations and confirm the existence of the phenomenon we uncover.

We decided that the disadvantages of conducting a large-scale operator survey outweighed the advantages (discussion in §7).

Our CGNAT inferences are consistent with published techniques. To verify our CGNAT inferences, we compare our results with previous studies characterizing CGNAT and IPv4 usage.

RTT ratio: Previous work identified that CGNAT deployments have a longer propagation delay from end hosts to the public Internet than simple NAT and used this property to identify CGNAT deployments [54, 90]. Although this observation is insufficient on its own to determine whether a prefix hosts customers behind a CGNAT gateway, we assess whether our inference matches this property as a weak indication of our accuracy. In particular, we compute the ratio between the round-trip times to the first responsive hop and to the last hop preceding the first public IP address for every traceroute, and group them according to the deployment type that we have inferred in Section 3.2. Figure 6 shows a boxplot for the RTT ratios, where we observe that inferred CGNAT deployments have higher RTT ratios than CPE NATs.

Hostnames: rDNS has been used as ground truth data to identify services hosted on IPv4 addresses in prior works [34, 45]. For all source /24 prefixes where we attributed squatting to the source organization, we inspect each IP's domain name based on the publicly available reverse DNS lookup provided by Rapid7, from August 2021 [73]. By manually investigating the most commonly occurring patterns, we identify a collection of regular expressions which indicate that an IP address is administrated by a single customer (*e.g., 'cpe', 'client'*) or by the ISP (*e.g., 'cgnat', 'cgn'*). We detail the list of regular expressions in Appendix D.

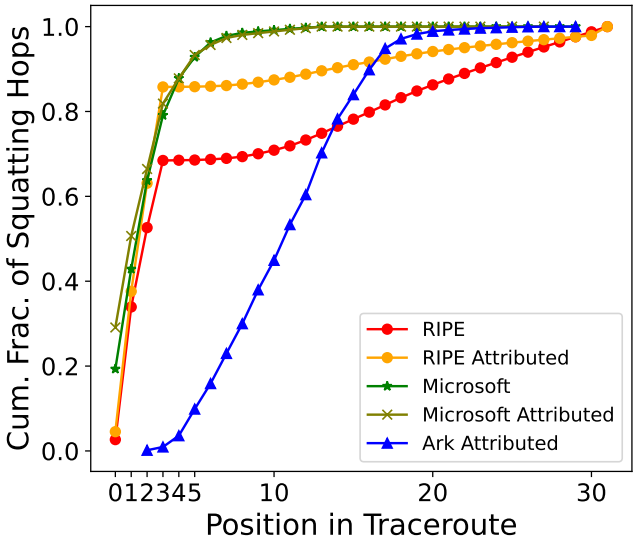

Figure 7: CDF of the position of squatting hops in traceroutes. Traceroutes from RIPE Atlas and Microsoft mostly observe squatting near sources. Ark traceroutes rarely observe squatting near the source.

In total, we are able to find 1117 different /24 prefixes for 115 different organizations with hostnames suggesting that the configuration is originating from the customer and 7 /24 prefixes and 3 organizations with a clear indication of an ISP's configuration. 0 prefix includes an address that matches both a pattern associated with customer-side use and an ISP-side configuration. Our methodology described in Section 3.2 classifies 83.5% of the prefixes that rDNS suggests have customer-side usage and 100% of the prefixes that rDNS suggests have ISP-side usage. The details of the technique and the exact matching inferences for each regular expression can be found in Appendix D.

Port and service scans: We also use port scans to verify our inferences. Because services running on their default ports are unlikely to be port-mapped, we consider hosts that run services rarely used by CGNAT (*e.g.,* RTSP, CWMP) on default ports as unlikely to be behind CGNATs. Therefore, we inspect a port scan dataset shared by Censys in August 2022 [23] to evaluate our methodology against inferences based on open ports and services. Censys infers services and manufacturers based on open port protocol headers. Of the IP addresses whose hostname matches a pattern associated with CGNAT (a total of 362K IP addresses), there are no open ports for 99.6% of the IP addresses. The total absence of open ports for the vast majority of CGNAT prefixes suggests that the vast majority of networks hosting CGNATs drop external connections. We identify more than 50 services (*e.g.,* libupnp and Home Gateway) and dozens of open ports (*e.g.,* TCP ports 7457 and 30005 as defined in TR-069 CPE WAN Management Protocol [30]) that indicate source /24 prefixes hosting single enterprises or home networks [5]. To evaluate our technique with the open port dataset, we consider three categories of prefixes: *high confidence* for non-CGNAT prefixes with more than 128 IP addresses with open ports linked to non-CGNAT devices (4 prefixes), *medium confidence* for non-CGNAT prefixes with more than 64 IP addresses with open ports (339 prefixes) and *low confidence* with more than 32 IP addresses with open ports (1032 prefixes). Our technique also performs well in the context of the Censys dataset where it correctly infers 100% of customer prefixes for high confidence non-CGNAT prefixes, 82.4% for the medium confidence ones, and 80.1% for the low confidence ones.

[5]We detail all the services and ports identified in Appendix E

5 SQUAT SPACE USAGE

Combining the Microsoft, RIPE Atlas, and Ark traceroutes (§2.2), we find a total of 89M traceroutes crossing 18,469 squat IP addresses. We identify the squatting organization in 69M traceroutes (77.5%) with strong attribution and 89M when augmented with weak attribution. We find that 65,149 (0.11%) of traceroutes require whois as a fallback and that 99.88% of traceroutes attributed with CAIDA's AS-to-Org database yield the same attribution as whois (§3.1).

We first investigate how close to the source we tend to observe squat space being used. In Figure 7, we look at the distribution of the position of all hops with squat addresses for each dataset, before and after we discard squatting hops that could not be attributed to an organization. We see that traceroutes with squat addresses in RIPE Atlas are concentrated in a few sources: 12 RIPE Atlas vantage points account for 70% of traceroutes observing squat addresses. We can attribute squatting hops close to these sources in most cases, thus the distribution of traceroutes with attributions shifts to the left. For our Microsoft traceroutes, the position of squatting hops is also biased towards sources, but the distribution is smoother than for RIPE Atlas, possibly due to the larger number of sources. We do not observe a significant difference in the position of hops that can or cannot be attributed to an organization in this dataset. Ark traceroutes, however, do not observe squatting in the origin network; squatting hops are observed at transit ASes. The positions of all squatting hops (not shown) and attributed squatting hops (shown in blue) in Ark traceroutes are quantitatively similar. This is confirmed by the fact that Ark host very few vantage points compared to Réseaux IP Européens (RIPE) and Microsoft and therefore misses a large fraction of squatting happening in the source AS. In conclusion, this tendency of squat space usage near sources means that we require a lot of sources to discover squat space in the wild.

5.1 How is Squat Space Used?

Table 3 shows the number of organizations that we attributed as utilizing squat space based on the confidence of the inference in the upper portion. The lower portion shows the number of organizations, identified to be using squat space for each purpose (§3.2) split by profile as defined in Section 5.2. An organization can show up in multiple rows if traceroutes support multiple inferences. We find that NATs are the most common use of squat space and identify CGNAT deployments in 387 organizations, suggesting that the responsibility for the squat space utilization lies on the ISP. As expected, most cases of squatting used for numbering routers and middleboxes are limited to internal devices (not on the border with other networks). We fail to infer the use of squat space for many organizations (*unclassified* row). The rightmost columns

Potential Utilization	Organizations				# of Traceroutes	Distinct Source /24s
	Access	Large Transit	Content	Unknown		
CPE NATs - Small	228	76	22	90	11M (15.9%)	8K
CPE NATs - Large	72	8	6	24	10M (14.7%)	7K
CGNATs - Small	112	30	7	76	8M (11.4%)	10K
CGNATs - Large	51	12	7	19	10M (14.4%)	3K
Internal routers and middleboxes	102	48	4	3	1.8M (2.6%)	4579
Border routers and middleboxes	2	1	0	1	42 (≤ 0.01%)	9
Unmapped routers and middleboxes	41	30	18	1	66K (≤ 0.01%)	1208
Unclassified	77	256	24	1546	27M (40.1%)	45K
Strong Attributions	593	440	80	1688	69M (77.5%)	79K
Strong + Weak Attributions	634	570	112	1832	88M (100%)	121K

Table 3: Summary of all our inferences for squat space use (§5.2). On the lower part of the table, we quantify the number of strong and weak attributions as described in Section 3.1. We split organizations with sufficient measurements by profile. On the two rightmost columns we show the number of traceroutes and distinct source IP addresses involved in the inferences. Most of our attributions relate to NAT and CGNAT configuration as was observed in [79].

show the fraction of traceroutes and /24 source prefixes in each row to capture the prevalence of squat space usage.

While our inferences are conservative and provide anecdotal indication of how squat space is used, our results provide conservative lower bounds. Additional verification using active probing (*e.g.,* IP aliasing [33, 43, 58, 87] or device fingerprinting [67, 97]) could further inform our inferences, but is challenging because it may raise complaints and because the use of squat space prevents most vantage points from even reaching the devices. We leave the pursuit of such techniques as future work.

Most of the traceroutes with squatting hops are classified (59.9%). We decided to use conservative conditions when identifying use to avoid false positives. Many (1,198 out of 3,011) squatting organizations have less than five traceroutes with squatting hops, so the limited coverage for those ASes is likely the reason for being unclassified. For the remaining organizations, we see multiple cases where the squat address is the first non-private-address hop in the traceroute, but we fail to find sufficient source addresses within the /24 prefix exhibiting identical behavior. Thus, by our criteria, they cannot be classified.

5.2 Where are the Squatters?

Geography of the squatters. We identify 2,394 squatters from 148 countries as illustrated in Figure 8. The US has the highest number of squatters per country (380 organizations), and the region with the highest number of squatters (805 organizations) is the APNIC region. APNIC was the first registry to exhaust its general-use pool of IPv4 addresses in April 2011 and has the most growth of Internet users since that period, as mentioned in the International Telecommunication Union annual reports [40]. There exists an imbalance in allocations between RIR and Internet user populations: currently there exists 0.3 IP addresses per Internet user in APNIC's region, 2.5 in ARIN's, and 1.2 in RIPE's according to BGP Potaroo [38]. This imbalance is further exacerbated when we consider the total population in the regions.

Profiling the squatters. We classify organizations into four categories by applying the following rules in order:

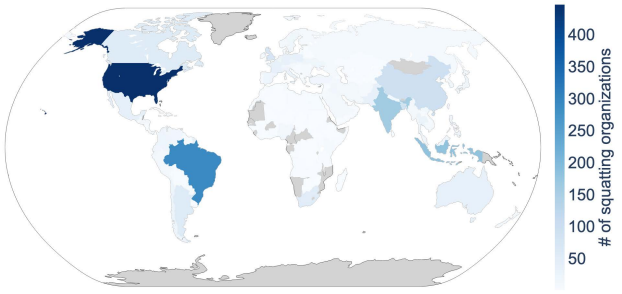

Figure 8: Chloropleth map highlighting the country of registration of the identified squatting organizations. More than 40% of the squatting organizations are registered in South-Eastern Asia, (ii) Brazilian organizations account for more than 75% of the squatters in LACNIC, and (iii) the number of squatters registered in the USA is inflated by many Content Delivery Networks (CDNs) registered there and a number of large US enterprises using squat space.

(1) Content are organizations that serve most of the content accessed by end-users [84]. They build distributed infrastructure and peer aggressively with access networks to support growing demand from an increasing number of eyeballs [85, 103, 105]. We identify CDNs by querying PeeringDB for ASes who self declare as "Content," and extend that to the managing organizations.

(2) Access organizations are ISPs for end-users. We classify an organization as *access* if its ASes have more than 10K users according to APNIC estimates of user populations [3]. Access organizations connect subscribers and their devices to the rest of the Internet and therefore assign most of their allocated IP space to end-users.

(3) Large Transit organizations are the historical highways of the Internet and stitch the many pieces of the Internet together. We classify as *transit* organizations whose ASes' customer cones include more than 100 distinct ASes. Most of their allocated IP addresses are used to configure routers and middleboxes or are delegated to customers.
(4) We classify all remaining organizations as Unknown. These organizations are not known to host content, have few users, and provide transit to less than 100 ASes in the customer cone; their AS features do not hint at any specific utilization of the squat address space.

Our combined traceroute datasets cover 251 CDN, 1,432 Access, 1,888 Transit, and 39,436 Unknown organizations.

How frequently are the squat organizations squatting? In Figure 9, we plot the frequency of traceroutes and sources observing squatting behavior in organizations that we inferred to squat. We note that the frequency of a source observing squat address is unimodal for the traceroutes (in red), with more than 55% of the organizations observing a low frequency of squat, hinting toward a punctual usage of squat space or utilization constricted to few prefixes, and 10% of the organizations with more than 50% of the traceroutes crossing squat space. For the sources, we notice an interesting jump at $x = 1$ indicating that all of the vantage points in the dataset observe addresses in at least one squat space /8 prefix, hinting toward the deployment of squat addresses very close to the clients. While we agree that deducing the operation of the squat space for an organization with few traceroutes hitting squat space is hard, we take the step to still consider them in their usage characterization in Table 3.

We also analyze the churn of the organizations that we identify to be squatting (Fig. 10). For this experiment, we focus only on three months: March, April, and May 2021 and discuss the evolution of organizations that are known to be squatting. Out of the 1800 organizations (compared to around 2486 on the whole dataset) that appear to be squatting in at least one in those months, only 950 organizations appear in all 3 months. This further illustrates the need to both define a statistical tool to ensure representativeness and the necessity for a large dataset to uncover many squatting scenarios. We also notice that 789 out of the 850 (92.8%) organizations that do not appear every month had all of their source addresses categorized as CPE NATs.

What are the most squatted /8 prefixes? In Table 4 we quantify how many organizations we identify squatting each /8 prefix, and the fraction of traceroutes traversing squat space containing hops in that /8 prefix. We observe that the /8 prefixes most frequently used as squat space belong to the DoD (*e.g.,* 7.0.0.0/8, 11.0.0.0/8, and 30.0.0.0/8), and have been recently announced (§6). The next two /8 prefixes, 16.0.0.0/8 and 19.0.0.0/8, still remain mostly unannounced. Looking at the fraction of traceroutes containing hops in each /8 prefix shows that the number of traceroutes is not strongly correlated with the number of squatting organizations (*e.g.,* 48.0.0.0/8). Possible explanations for this include (i) organizations that have few sources and few traceroutes that observe squatting, and (ii) large organizations and enterprise networks using squat space and contributing a significant number of traceroutes.

/8 Prefix	# of Org	Ratio in Traceroutes
11.0.0.0/8	1405	17.7 %
30.0.0.0/8	676	31.4 %
7.0.0.0/8	317	2.8 %
19.0.0.0/8	300	0.3 %
16.0.0.0/8	282	0.4 %
22.0.0.0/8	272	3.8 %
6.0.0.0/8	261	0.3 %
9.0.0.0/8	207	3.4 %
21.0.0.0/8	198	0.1 %
25.0.0.0/8	135	0.1 %
33.0.0.0/8	134	0.4 %
26.0.0.0/8	104	0.08 %
43.0.0.0/8	93	6.3 %
29.0.0.0/8	92	0.4 %
28.0.0.0/8	79	1.7 %
56.0.0.0/8	46	4.7 %
48.0.0.0/8	25	25.7 %

Table 4: Most squatted /8 prefixes observed (i) per organization and (ii) in total.

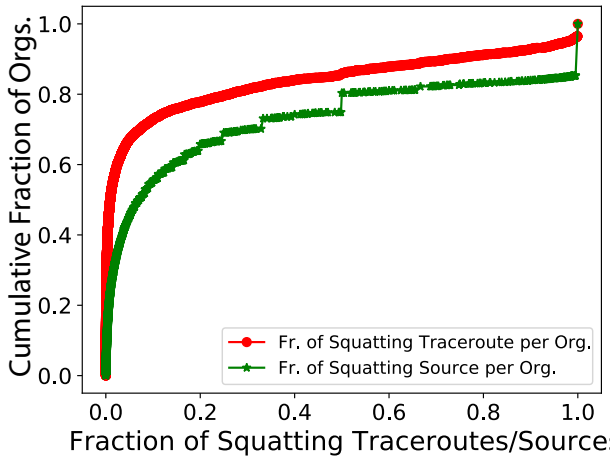

Figure 9: CDF of the frequency of traceroutes (in green) or sources (in red) that have observed squat for every organization.

We also quantify the number of different /8 prefixes squatted by organizations to understand how the squatting is taking form in the wild. While we expected that most organizations would not require squatting more than one /8 prefix, we infer that 32% of squatting organizations use multiple /8 prefixes. More surprisingly, 10% of the squatting organizations use more than five /8 prefixes, hinting toward more complex and involved squat configurations or multiple customers squatting different /8 prefixes.

5.3 Leaking Squat Space

In this section, we look into illegitimate BGP announcements of squat space. When a squatting organization receives an illegitimate announcement of a squatted prefix, the announcement will

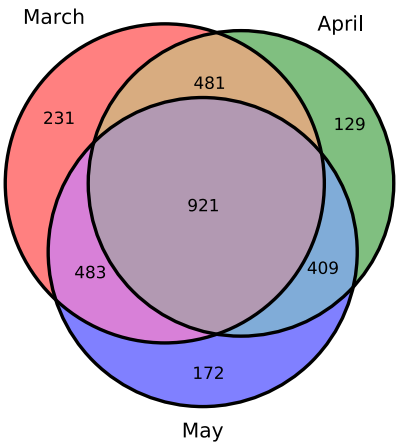

Figure 10: Churn of squatters over a 3 month period. The intersection depicts organizations that appear to be squatting in multiple months.

propagate within the squatter unless filtered, inducing the risk that packets destined to devices assigned squat addresses (*e.g.,* traffic to a router) are not delivered correctly.

Squat space route leaks. Illegitimate announcements of squat space are common and widespread. Using RIPE's historical API [82], we were able to recover all the announcements for prefixes within the squat space. We consider an announcement of DoD prefixes illegitimate when its AS-path does not traverse DoD-related entities (*i.e.,* Global Resource Systems (GRS-DoD) (AS8003), the Navy Network Information Center (NNIC) (AS 788), or the Defense Information Systems Agency (DISA) (AS 721)), and consider announcements of other squat space illegitimate when they do not originate from the corresponding assignee in ARIN's registry.

We identify 16K illegitimate announcements and 7K legitimate announcements originated by 582 ASes from January 2001 to August 2021. Figure 11 (left) plots the number of distinct ASes originating each squat space prefix each year. Figure 11 (right) shows each prefix's median visibility from BGP collector peers, *i.e.,* the fraction of peers that export a route toward the squat prefix to a RouteViews or RIPE RIS collector. While the expected behavior is that most ASes would filter announcements for squat space prefixes as recommended by best practices [99], Figure 11 shows a different reality, with many announcements propagating to BGP collectors.

We observe a gradual increase in route leaks for squat prefixes over the years, in terms of both number and visibility of announcements. This indicates that squatting may have become more common due to increasing demand for IPv4 address space. We also observe that the visibility of illegitimate DoD space announcements increased in 2021, even after GRS-DoD started legitimately announcing DoD space. We verify that the number of organizations leaking DoD space has not increased in 2021 compared to previous years. These observations indicate that transit ASes may have removed filters to allow the legitimate GRS-DoD announcements to propagate with the consequence being that leaks propagate further.

6 CASE STUDY: DOD IPV4 SPACE

Conventional wisdom predicts squat space as a serious threat to Internet functionality [7, 99]. Using squat space is risky because if the legitimate owner were ever to begin announcements, communication between the squatting organization's devices could be jeopardized and legitimate destinations within the newly announced squat space could be unreachable by the users of the squatting organization. In this section, we explore a recent case of reclaimed squat space of unprecedented size and examine the impact on the Internet when widely-used squat space is suddenly announced. This occurred in January 2021 when AS8003, authorized by the US DoD, began announcing millions of DoD owned IP addresses which had never been legitimately announced [94].

6.1 Historical Setting

After the Internet's creation in 1983, the DoD was awarded a considerable number of large IP blocks (a total of 14 /8 prefixes totaling more than 5% of the total IPv4 space, along with many smaller blocks, *e.g.,*/16 prefixes), since the Internet was created under the auspices of the Defense Advanced Research Projects Agency (DARPA) funding and project initiatives [18]. The DoD, however, only announced a fraction of its total IPv4 space to the Internet (≈20% of their allocated IPv4 address space, as shown by the light blue dots in Figure 1). The vast majority of DoD's IPv4 address space remained unannounced to the public Internet; a portion of the unannounced address space is used internally by the DoD, while some remains completely unused even within the DoD [70].

6.2 An Abrupt Change

Starting on the 20th of January 2021, AS8003 started announcing subnets in 7.0.0.0/8, 11.0.0.0/8, and 22.0.0.0/8. AS8003 increasingly announced subprefixes and eventually coalesced some into whole /8 prefixes. AS8003 extended those announcements to 29.0.0.0/8 and 33.0.0.0/8 on the 4th of March; and to 6.0.0.0/8, 26.0.0.0/8, 28.0.0.0/8, and 30.0.0.0/8 on the 15th of April. By June 2021, virtually all DoD squat space prefixes had some percentage of their IP space announced to the global routing table, with some reaching 100% (black dots in Fig. 1).

6.3 Evaluating Impact

To understand the implications of announcing a large portion of traditional squat space, we examined three timescales. On a short timescale, we asked whether the announcements appear to be filtered by transit networks in the Internet. On a medium timescale, we verified whether companies that we identified to be squatting DoD space in the past had stopped squatting or moved to non-DoD prefixes. Finally, at a longer timescale, we analyzed how the existence of an external announcement affected the reachability of DoD space from within the squatting organizations.

Very few ASes filtered out the sudden DoD announcements. Given the potential harm from misappropriating allocated addresses for squat space, we evaluate whether network operators had squat space prefixes in BGP prefix filters to prevent announcement propagation. AS8003 announced its prefixes to Hurricane Electric (AS6939), and those announced prefixes became visible from more than 80% of RouteViews and RIPE RIS peers within 30 minutes. However, it

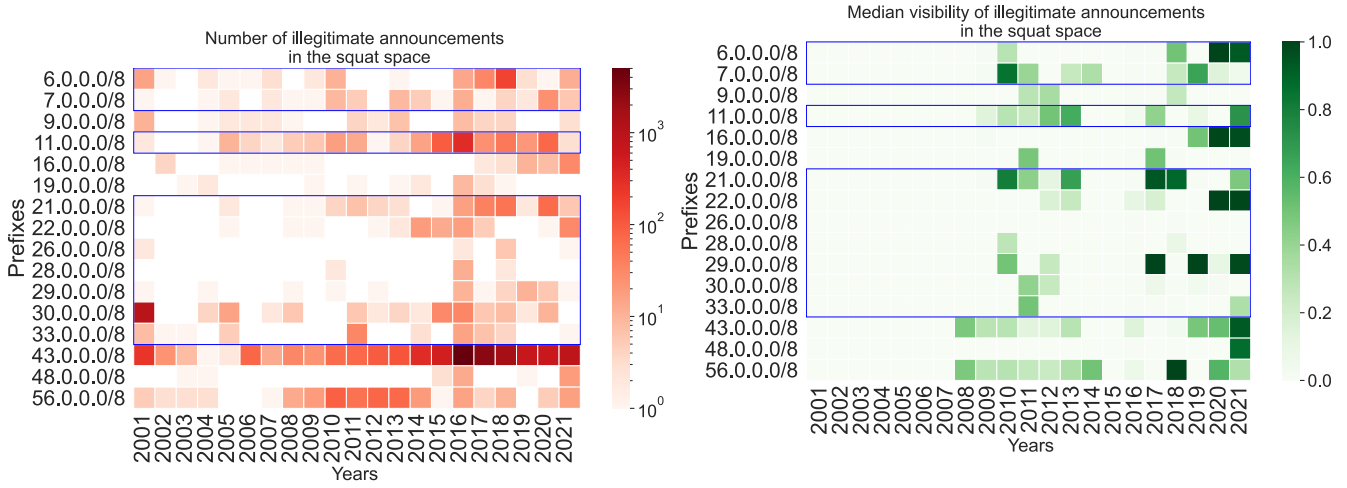

Figure 11: The number of distinct ASes originating announcements of squat space (left) and the average number of RIPE RIS BGP collector peers exporting an illegitimate route to squat space prefixes (right). 16.0.0.0/8 and 43.0.0.0/8 have assigned subnets that were historically announced with a limited visibility. We observe increased visibility of illegitimate announcements of the DoD prefixes in 2021 after AS8003 started legitimately announcing them. We frame the DoD prefixes in blue.

took more than 6 hours to get visibility at 95% of the BGP collector peers. Although propagation may have required coordination between operators (*e.g.,* emails and phone calls), the quick propagation indicates that few operators had squat space filters.

While Figure 11 shows that previous illegitimate announcements for DoD's IP space rarely propagated to more than 10% of the BGP collector peers, the Internet accepted AS8003's announcements extremely quickly. A significant difference is that AS8003's legitimate announcements were made through Hurricane Electric and propagated top-down in the Internet hierarchy. Another factor is that Hurricane Electric, being a large network provider, would be well positioned to contact other key network operators to clear any existing prefix filter(s). The rapid propagation exhibits how announcements made by big actors quickly propagate on the Internet.

Evolution of squatting organizations. Once AS8003 announcements started, organizations that squatted DoD address space would be disconnected from legitimate destinations [99]. In this section we evaluate how these organizations responded to DoD's announcements. In particular, we quantify the number of organizations squatting DoD address space before and after January 2021.

Interestingly, we find that 53% of the organizations that appear to be squatting the DoD's squat space after January were not observed to be squatting earlier. Of those 53%, we had more than 5000 traceroutes and 50 source IPs in 86% in December 2020, *i.e.,* they were in the set of organizations with sufficient measurements to infer squatting. Possibly related to this increase, we also find that the number of ASes illegitimately announcing DoD address space has surprisingly increased after January.

We also look into whether squatters identified before January 2021 continued to squat the DoD prefixes after they were announced. We identify 1,254 ASes squatting at least one DoD prefix in January 2021. We then study the traceroutes from February and March and find that many (873 out of 1,254) ASes continued to squat DoD prefixes. In summary, contrary to our expectation, the new announcements of the squat space did not result in most operators moving away from DoD address space.

How the legitimate announcements modified reachability. We identified only a partial reduction in the number of squatting ASes after the announcement, despite it being known that squat space usage could adversely affect both the squatting organization and those attempting to communicate with the legitimate address space owner [99]. Due to continued squat space usage's potential interference with routing, we sought to determine the extent of the impact of the continued use of squat space by the 873 ASes identified in the previous paragraph.

To understand the announcement's impact on squatting organizations' routing, we examine the ASes that (i) were squatting DoD prefixes in our May measurements and (ii) were hosting at least one RIPE Atlas or Speedchecker Vantage Point (VP), resulting in 497 ASes hosting a total of 5,403 VPs. We schedule traceroute measurements for two days in June 2021 from up to 10 RIPE Atlas probes and Speedchecker VPs in each organization toward every squatted IP address within the organization identified by our traceroutes.

We then bucket the traceroute results into two behaviors: (a) the route exits the squatting organization's network towards the legitimate IP address (*external routing*), or (b) the route remains inside the squatting organization's network (*internal routing*). In case a traceroute is completely unresponsive, we run an additional traceroute toward a globally reachable responsive address that we control to ensure the VP is functioning correctly. We ignore VPs if the traceroute to our reachable address is also completely unresponsive, otherwise we consider that the organization employs internal routing. One possible alternative explanation for a completely unresponsive traceroute is that the squatting organization filters legitimate announcements and does not provide a route for the squatted space to end users.

Both behaviors can be problematic. External routing can lead to inaccessibility of the internal resources configured to use squat space, while internal routing implies that part of the Internet is unreachable by the organization's users.

Out of the 497 ASes hosting VPs, we find 307 employ external routing, 57 employ internal routing, 46 exhibit both behaviors depending on the probe and the destination IP address, and 87 we discard because their measurements were unresponsive toward the globally reachable IP address. Manual analysis of a subset of the 46 ASes exhibiting both behaviors confirm that probes are hosted in different geographic locations, hinting that the squat space is used only in a region of the AS's network.

We map traceroutes that leave the source organization (external routing) to ASes and find that more than 99% of the last responsive hops are managed by Hurricane Electric. rDNS entries indicate traceroutes terminate in Ashburn, Virginia, suggesting that there is a single source for all DoD announcements and that it is possibly in the capital region of the United States. We also observe nine organizations that have internally responsive squatted IP addresses. All are organizations that previously identified as using squat space for configuring routers and middleboxes. Because the newly announced DoD prefixes currently do not appear to be used to host any content, it is hard to predict the reclaimed squat space's usage.

Summary. Our case study indicates that squat space can be legitimately announced and become globally reachable within hours. We also find that organizations squatting address space that becomes legitimately announced may not make a concerted effort to stop squatting that space (at least not in the short term). Finally, we observe that squatting organizations handle routing towards legitimate announcements of the squatted space in different ways.

7 DISCUSSION AND IMPLICATIONS

We uncovered the use of squat addresses in more than 60% of the 500 largest networks in terms of IPv4 addresses allocated (Fig. 5) and 70% of the 500 largest eyeball networks (Fig. 12), and this result represents a lower bound based on what is observable using our measurement datasets. Squatting is not restricted to the peripheral parts of the networks and is widespread enough to require a change in how we map and consider IPv4 address from the squat space.

Topological analyses of Internet measurements are not squat-aware but should be. This section discusses how common Internet measurement practices for interpreting traceroutes rely on an implicit assumption that any non-private IP address observed along a route is a unicast address assigned to a single interface on a single device in a single location in a single network that is allocated and likely announces the address, and how these practices can distort Internet topologies and analyses when applied to squat space addresses, which are reused by many networks in many locations.

AS topology distortion. CAIDA's Internet Topology Data Kit (ITDK) dataset [16] uses *bdrmapIT* [59] for IP-to-AS mappings in traceroutes, which tends to map squat space addresses based on registry assignments (App. C). We analyzed the ITDK dataset from August 2020 and found that it infers 78 inter-AS links from these squat space mappings that would not otherwise exist in the topology. These links are almost certainly incorrect, as ITDK infers the links as interdomain connections to ASes that are assigned but not announcing the squat space.

Conventional IP-to-AS mapping relies on mapping addresses to prefixes observed in BGP tables, then assigning the address to the origin AS of the announcement. This approach is also sensitive to squat space because addresses continue to be used as squat space even after there is a legitimate announcement (§6). We uncovered cases in which organizations continued squatting even months after a legitimate announcement. This observation argues for the need for a historical squat space analysis. Researchers should not map squat space addresses to ASes directly. Instead, they should use the attribution approach in our paper. We publish a script to generate an up-to-date list of squat space prefixes [83].

Geolocation distortion. Geolocation databases have a single location for a squat space address, even though the address may be reused in multiple locations globally, leading to distorted topologies when mapping from IP addresses to geographic locations. For example, the MaxMind Lite database is ubiquitous in Internet research (*e.g.,* [39, 47, 74]), and it maps DoD squat space to the US and maps 98% of the MoD squat space to the UK, even though both are used in other locations by other organizations. We illustrate an example of such mapping for a measurement originating from China:

Traceroute IP	43.254.105.0	30.255.53.5	30.255.21.9	43.254.147.9
Hop #	1	2	3	4
ASN	AS138421	AS8003	AS8003	AS138421
AS-Name	CU-CN-AS	GRS-DOD	GRS-DOD	CU-CN-AS
Geolocation	CN	US	US	CN
RTT	6 ms	7 ms	8 ms	15 ms

Speed of light constraints (≈100 km/ms) suffice to show the geolocations of the 30/8 IP addresses are incorrect, but other distortions can be more subtle and go undetected even by those careful to apply speed of light checks.

While MaxMind is known to include errors and hence should be used with care, especially for router and infrastructure addresses, a more insidious distortion occurs using rDNS, which is treated as a more reliable source of information about a router. rDNS can be used to infer properties of an IP address including its geolocation, router type, or interconnection information [27, 52, 53]. Squat space addresses can have rDNS entries, even though they are used by multiple organizations. For example, many rDNS entries have been added for 43.0.0.0/8, which we observe squatted 462,235 times in August 2021. While these entries may encode useful information about one instance of an IP address's use, they can cause distortions when assumed to apply to any observation of the address. For example, the address 43.132.247.32 resolves to the hostname `ptr-hk-43-132-247-32.qcloudmail.com`, hinting toward an IP address hosted in Hong Kong, but we came across 3 instances of this specific IP address being squatted by other organizations.

Squatting is problematic even when approved. Even in cases where legal or ethical concerns are addressed, squat space can be an administrative headache. In conversations with network operators, we learned of cases where squat space can be used with permission from the owning organization with no financial arrangement. For example, we learned that a large provider has an agreement with the UK MoD to use addresses from 25.0.0.0/8. When the MoD decides

to announce or delegate a new portion of 25.0.0.0/8 itself, it sends a 1–2 month warning to the squatting organizations, who then must scramble to renumber any parts of their network using those sub-prefixes. Failure to do so will render the squatter unable to access legitimate Internet destinations and may violate and hence invalidate the usage agreement. Examining RIPE Route History, at least 11 new /24 prefixes within 25.0.0.0/8 have been announced by UK Government affiliated ASes (206747 and 199055 [96]) since 2014, with each new announcement presumably resulting in operational scrambling by squatters. There may be additional, non-Internet facing cases that are not visible via public BGP feeds.

IPv4 squatting raises new legal questions. To the best of our knowledge, squatting of IPv4 addresses has never led to publicly documented prosecution. Following our request to investigate the legal implications of squatting, the Harvard Cyberlaw Clinic produced a 15 page memorandum analyzing the legal theories applicable to squatting [12]. In summary, the memo concludes that pursuing legal recourse against a squatting organization would prove challenging. There is neither a set of legal procedures for prosecuting squatting nor a clear established legal theory that fits IPv4 address use or IPv4 squatting. Civil suits would be difficult given a lack of clear harm in cases of squatting (in the absence of hijacking).

There are potential avenues for resolution. Other technical properties disputes, such as domain squatting, are commonly arbitrated by the WIPO Alternative Dispute Resolution[101], an assembly of more than 50 experts that specializes in resolving claims related to intellectual and technical property dispute. It appears that WIPO's arbitration process is often sufficient in this context, since different parties reach an agreement before they escalate their complaints to the penal system. Despite the fact that IPv4 address and domain squatting are conceptually very different, we posit that a similar arbitration process could appear if legitimate owners of the squat space were to start seeking actions against squatters.

Recently, efforts originating from RIPE and ARIN working groups resulted in initiatives to formally define BGP hijacking as a RIR policy violation [31, 32]. Those proposals have been since abandoned, which illustrates the tension in clearly defining hijacks and indicting potential abusers of network resources. The parallel between BGP hijacks and IPv4 squatting has its limits; BGP hijacks have worse effects on the Internet as they affect the owner of the address space directly, as well as third parties trying to communicate with the legitimate origin, while IPv4 squatting's effects are restricted within the vicinity of the network that is using squat space.

It is not clear how policy makers will address IPv4 squatting in the future. Still, in conversations, researchers on cyber-policy suggested to us that if policy makers were to become aware of the scope of this phenomenon, they might work on extending policies to simplify squatting indictment. We hope that our work will start this discussion by shedding light on how prevalent squatting is.

Surveys for squat space validation is perilous:

In general, survey feedback from operators can serve two roles:

(1) Discussing squat usage in general, not tied to our specific results.
(2) Direct validation of our inferences in specific networks.

There exists two key challenges that render a survey unlikely to provide insightful feedback regarding both points. Because of the incriminating nature of squatting socially and legally (§1,§7), a public survey may not solicit honest responses. Prior studies leveraging surveys investigated benign network properties such as the existence of CGNAT or ISP middleboxes [79, 86]. The operations community has vehemently rebuked squat space as a practice [7, 63–65] and so fear of fall out is reasonable. Even if conducted anonymously, we could not guarantee it remains so under subpoena. Furthermore, the people we know how to contact and built trust relationship throughout the years are external-facing network operations teams in charge of peering (NANOG, etc), whereas there is no reason for CGNAT teams to be external-facing.

We also do not think that an anonymous survey provides better evaluation than the techniques described in Section 4.2. Even in the best case scenario with a survey that described per-network or per-inference behavior, surveys could have dishonest or incorrect responses. We do not believe that we possess the right expertise to undertake the task of conducting and interpreting a survey on a topic with high response bias [11] where the answers might not result in building ground-truth. Bradburn *et al.* showed that perceived normative threat influences responses to questions in surveys [9], leading to potentially portraying a distorted view of squatting. It would also be extremely difficult to conduct a per-network behavior survey anonymously. We could not tie an ISP's response to a particular inference we made without deanonymizing, and so we would not be able to understand the reason for any false inferences, refine our technique, or try to generalize results to understand when our inferences do not work well.

For all those reasons, rather than conducting a survey, we decided to share anecdotes from conversations with operators (§4.2) that serve both roles (1) and also (2) in a small-scale way.

8 RELATED WORK

Squat Space. While very little academic work has been devoted to IP squatting, it has been the focus of multiple blog posts from operators, Internet organizations, and concerned Internet users. In 2007, it seems that Cisco's IP Journal discussed squat space usage in a blog post[6] discussed in an ICANN report [98]. From there, many blog posts have investigated odd behaviors end-users noticed from traceroute [22, 102] leading to conspiracy theories regarding the existence of backdoors that would be accessible by the DoD or the MoD [75, 104]. ARIN wrote a blog post to educate ISPs and other network operators on the risk and the unsustainable nature of squatting addresses [7]. More recently, the announcements of the historically unannounced DoD IPv4 space on the last evening of Donald Trump's presidency brought squatting behavior into the spotlight as a systemic problem [55, 56].

IPv4 Scarcity. With registries essentially out of IPv4 address space, IPv4 addresses are now a valuable resource. Existing work has examined this fact under the lens of address scarcity [78], IPv4 transfer markets [49], sharing mechanisms such as NAT [88], IPv6 transition [46], and utilization of IPv4 [20]. Our work complements those earlier efforts, but looks into the dubious practice of IP squatting, which has not been previously analyzed in detail. Recently, a RIPE Labs article investigated the utilization of "Future Use" 240.0.0.0/4 addresses as private addresses in Amazon Web Services [50]. Our

[6]We were not able to access the webpage in question.

methodologies both rely on traceroute, but our work differs by the wide extent of Microsoft's coverage compared to RIPE Atlas (§4) and our focus on the portion of the IPv4 space that is allocated to legitimate organizations and has started being reclaimed recently.

Security Concerns over IP Hijacks. There exists a rich literature that has analyzed BGP hijacks and the security challenges involved with IP spoofing and hijacking [24, 68, 91]. Our work differs because it investigates a different phenomenon than prefix hijacking that has a less visible effect on the Internet topology. Furthermore, squatting can happen simultaneously across many networks, and so requires more distributed vantage points to uncover, while most of the previous literature is focusing on propagated hijacks which are observable from BGP feeds from openly accessible BGP monitors.

CGNAT. There have been some works on identifying CGNAT usage [48, 54, 79]. Richter *et al.* design a technique leveraging the BitTorrent Distributed Hash Table to discover internal leakage identifications as a NAT deployment indicator [79]. To differentiate between NAT and CGNAT deployments, they build a threshold based on the amount of internal communication between addresses using private addresses. They also observe that squat space appears to be used for CGNAT deployment by mobile operators. Livadariu *et al.* provide two heuristics to detect CGNAT [48]: the first is based on the volume of requests MLab receives from a /24 prefix and the second leverages the volume of requests from a /24 prefix observed by the UCSD Network Telescope [1]. NAT Revelio employs a set of tests to identify CGNAT [54]. We do not replicate their techniques to compare their inferences against ours as our core traceroute dataset, the Microsoft one, is fixed and does not allow to run extra measurements Furthermore, UPnP measurements like those used by NAT Revelio only make sense in the context of home and some enterprise networks, and do not apply generally across all networks, as covered in our dataset. Our inferences consider large CGNAT deployments including multiple routers or middleboxes between end-user devices and the public Internet regardless of the type of the network. Our use of RTTs as an evaluation step in our analysis is based on NAT Revelio's finding that CGNAT incurs longer RTTs.

9 CONCLUSION

We presented the first large scale analysis of Internet squat space usage, using what is likely the largest traceroute dataset ever publicly analyzed in terms of vantage point distribution and coverage. Despite ethical and technical pitfalls, and the maturity of IPv6, squat space usage remains commonplace on the Internet, as we observed squat addresses used by more than 60% of the 500 largest networks in terms of IPv4 addresses allocated (Fig. 5) and 70% of the 500 largest eyeball networks (Fig. 12). In analyzing an unprecedented squat space "recovery" operation of nine /8 prefixes by the US DoD in January 2021, we found that the new BGP announcements were widely accepted, despite being well-known squat space, and that many operators continued to utilize the squat space long after the legitimate announcements were made.

While squat space attribution within a network does not imply a systematic deployment within an ISP, as customers have the autonomy to configure their (CG)NATs with squat space without the ISP's knowledge, it highlights a pervasive issue that impacts end-users, network operators, and researchers by distorting network topology mapping and geolocation. To eliminate inaccuracies introduced by squat space, we advocate for adoption of "squat space awareness" as standard practice in Internet measurement analysis.

10 ACKNOWLEDGMENTS

The authors would like to acknowledge Kendra Albert and Alexandra Bruer at the Harvard Cyberlaw Clinic for researching and writing an excellent and informative memo on the legal implications of IPv4 squatting, as well as Roxana Radu in Blatvanik School of Government at the University of Oxford for identifying potential avenues for resolving squatting complaints in the future. We are also thankful to the anonymous reviewers for their constructive feedback and Ralph Holz, our assigned editor, for guiding us through the submission process. In addition, we would like to thank the anonymous network operators for discussing some of our inferences. This paper was partially funded by NSF grant CNS-2148275, FAPEMIG, and CNPq.

REFERENCES

[1] UCSD Network Telescope Traffic Samples. https://catalog.caida.org/dataset/telescope_backscatter. Accessed: 2023-2-24.
[2] abovetopsecret.com. U.S Military Spying on Canadian? https://www.abovetopsecret.com/forum/thread761697/pg1
[3] APNIC. Visible ASNs: Customer Populations (Est.). https://stats.labs.apnic.net/aspop/.
[4] ARIN. ARIN IPv4 Free Pool Reaches Zero. https://www.arin.net/vault/announcements/2015/20150924.html.
[5] T. Arnold, E. Gürmeriçliler, G. Essig, A. Gupta, M. Calder, V. Giotsas, and E. Katz-Bassett. (How Much) Does a Private WAN Improve Cloud Performance? In *Proc. IEEE INFOCOM*, 2020.
[6] T. Arnold, J. He, W. Jiang, M. Calder, I. Cunha, V. Giotsas, and E. Katz-Bassett. Cloud Provider Connectivity in the Flat Internet. In *Proc. ACM IMC*, 2020.
[7] C. Aronson. To Squat or not to Squat? - Team ARIN. https://teamarin.net/2015/11/23/to-squat-or-not-to-squat/
[8] H. Beeman. Freeing Up Unused IP Addresses - an Update - Government Technology. https://governmenttechnology.blog.gov.uk/2015/09/02/freeing-up-unused-ip-addresses-an-update/
[9] N. M. Bradburn, S. Sudman, E. Blair, and C. Stocking. Question Threat and Response Bias. *Public Opinion Quarterly* 1978.
[10] broadbandforum.co. Reliance Jio 4G: Strange Public IPv4 Address Assigned Behind NAT. https://broadbandforum.co/threads/reliance-jio-4g-strange-public-ipv4-address-assigned-behind-nat.190267/
[11] D. E. Broadbent. Word-Frequency Effect and Response Bias. *Psychological review* 1967.
[12] A. Bruer and K. Albert. Legal Analysis of IPv4 Squatting. Harvard Law School's Cyberlaw Clinic at Harvard's Berkman Klein Center for Internet & Society, https://blogs.harvard.edu/cyberlawclinic/files/2023/02/IPv4-Squatting-Legal-Analysis-FINAL.pdf,~2023,.
[13] CAIDA. IPv4 Prefix-Probing Traceroute Dataset, 2020/08. https://www.caida.org/catalog/datasets/ipv4_prefix_probing_dataset/.
[14] CAIDA. The CAIDA UCSD AS Classification Dataset, 2020–2021. https://www.caida.org/catalog/datasets/as-classification.
[15] CAIDA. The CAIDA UCSD AS to Organization Mapping Dataset, 2020/01. https://www.caida.org/data/as_organizations.xml.
[16] CAIDA. Macroscopic Internet Topology Data Kit (ITDK), 2020/08. https://www.caida.org/catalog/datasets/internet-topology-data-kit/
[17] M. Calder, R. Gao, M. Schröder, R. Stewart, J. Padhye, R. Mahajan, G. Ananthanarayanan, and E. Katz-Bassett. Odin: Microsoft's Scalable Fault-Tolerant CDN Measurement System. In *Proc. USENIX NSDI*, 2018.
[18] D. Clark. The Design Philosophy of the DARPA Internet Protocols. In *SIGCOMM Comput. Commun. Rev.*, 1988.
[19] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey. Measuring IPv6 Adoption. In *Proc. ACM SIGCOMM*, 2014.
[20] A. Dainotti, K. Benson, A. King, B. Huffaker, E. Glatz, X. Dimitropoulos, P. Richter, A. Finamore, and A. C. Snoeren. Lost in Space: Improving Inference of IPv4 Address Space Utilization. *IEEE Journal on Selected Areas in Commun.* 2016.
[21] T. K. Dang, N. Mohan, L. Corneo, A. Zavodovski, J. Ott, and J. Kangasharju. Cloudy with a Chance of Short RTTs: Analyzing Cloud Connectivity in the Internet. In *Proc. ACM IMC*, 2021.

[22] DSL Reports. Why is My First Hop to a DoD Assigned IP Address? - Rogers | DSLreports Forums. https://www.dslreports.com/forum/r25679029-Why-is-my-first-hop-to-a-DoD-assigned-IP-address

[23] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman. A Search Engine Backed by Internet-Wide Scanning. In *Proc. ACM CCS*, 2015.

[24] T. Ehrenkranz and J. Li. On the State of IP Spoofing Defense. In *Proc. ACM Trans. on Internet Tech.*, 2009.

[25] C. Elverson. Your AS is Mine: BGP/IP Hijacking, the ICBM of the Cyber World. NANOG 77. https://pc.nanog.org/static/published/meetings/NANOG77/2108/20191028_Elverson_Your_As_Is_v1.pdf

[26] B. Escoffier, J. Monnot, and V. T. Paschos. Weighted Coloring: Further Complexity and Approximability Results. *Inform. Process. Lett.* 2006.

[27] A. D. Ferguson, J. Place, and R. Fonseca. Growth Analysis of a Large ISP. In *Proc. ACM IMC*, 2013.

[28] A. R. for Internet Numbers. Using Whois. https://www.arin.net/resources/registry/whois/

[29] A. Formoso, J. Chavula, A. Phokeer, A. Sathiaseelan, and G. Tyson. Deep Diving into Africa's Inter-Country Latencies. In *Proc. IEEE INFOCOM*, 2018.

[30] T. B. Forum. CPE WAN Management Protocol - Broadband Forum. https://www.broadband-forum.org/download/TR-069_Amendment-6.pdf

[31] C. Friacas and J. Palet Martinez. Resource Hijacking is a RIPE Policy Violation. https://www.ripe.net/participate/policies/proposals/2019-03/draft

[32] C. Friaças and J. P. Martinez. Arin-Prop-266: BGP Hijacking is an ARIN Policy Violation. https://www.arin.net/participate/policy/proposals/2019/ARIN_prop_266_v3/

[33] S. Garcia-Jimenez, E. Magana, D. Morato, and M. Izal. Pamplona-traceroute: Topology Discovery and Alias Resolution to Build Router Level Internet Maps. In *Global Information Infrastructure Symposium*, 2013.

[34] P. Gigis, M. Calder, L. Manassakis, G. Nomikos, V. Kotronis, X. Dimitropoulos, E. Katz-Bassett, and G. Smaragdakis. Seven Years in the Life of Hypergiants' Off-nets. In *Proc. ACM SIGCOMM*, 2021.

[35] V. Giotsas, T. Koch, E. Fazzion, I. Cunha, M. Calder, H. V. Madhyastha, and E. Katz-Bassett. Reduce, Reuse, Recycle: Repurposing Existing Measurements to Identify Stale Traceroutes. In *Proc. ACM IMC*, 2020.

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

[37] P. C. House. Internet Exchange Directory. https://www.pch.net/ixp/dir.

[38] G. Huston. Resource reports. https://resources.potaroo.net/iso3166/v4cc.html

[39] C. Iordanou, G. Smaragdakis, I. Poese, and N. Laoutaris. Tracing Cross Border Web Tracking. In *Proc. ACM IMC*, 2018.

[40] ITU. ICT Facts and Figures. In *ITU Telecom World*, 2017.

[41] E. IX. IXP Database (IXPDB). https://ixpdb.euro-ix.net/en/.

[42] T. Javed, M. Haseeb, M. Abdullah, and M. Javed. Using Application Layer Banner Data to Automatically Identify IoT Devices. In *SIGCOMM Comput. Commun. Rev.*, 2020.

[43] K. Keys, Y. Hyun, M. Luckie, and kc claffy. Internet-Scale IPv4 Alias Resolution with MIDAR. In *IEEE/ACM Trans. Netw.*, 2013.

[44] A. Kirkham. RFC 6752: Issues with Private IP Addressing in the Internet. https://tools.ietf.org/html/rfc6752.

[45] Y. Lee and N. Spring. Identifying and Analyzing Broadband Internet Reverse DNS Names. In *Proc. ACM CoNEXT*, 2017.

[46] S. L. Levin and S. Schmidt. IPv4 to IPv6: Challenges, Solutions, and Lessons. *Telecommunications Policy* 2014.

[47] Z. Li, D. Levin, N. Spring, and B. Bhattacharjee. Internet Anycast: Performance, Problems, & Potential. In *Proc. ACM SIGCOMM*, 2018.

[48] I. Livadariu, K. Benson, A. Elmokashfi, A. Dhamdhere, and A. Dainotti. Inferring Carrier-Grade NAT Deployment in the Wild. In *Proc. IEEE INFOCOM*, 2018.

[49] I. Livadariu, A. Elmokashfi, A. Dhamdhere, and kc claffy. A First Look at IPv4 Transfer Markets. In *Proc. ACM CoNEXT*, 2013.

[50] Q. Lone. 240/4 as Seen by RIPE Atlas. https://labs.ripe.net/author/qasim-lone/2404-as-seen-by-ripe-atlas/

[51] M. Luckie, B. Huffaker, kc claffy, A. Dhamdhere, and V. Giotsas. AS Relationships, Customer Cones, and Validation. In *Proc. ACM IMC*, 2013.

[52] M. Luckie, B. Huffaker, A. Marder, Z. Bischof, M. Fletcher, and kc claffy. Learning to Extract Geographic Information from Internet Router Hostnames. In *Proc. ACM CoNEXT*, 2021.

[53] M. Luckie, A. Marder, M. Fletcher, B. Huffaker, and kc claffy. Learning to Extract and Use ASNs in Hostnames. In *Proc. ACM IMC*, 2020.

[54] A. Lutu, M. Bagnulo, A. Dhamdhere, and kc claffy. NAT Revelio: Detecting NAT444 in the ISP. In *Proc. PAM*, 2016.

[55] D. Madory. Wait, Did AS8003 Just Disappear? https://www.kentik.com/blog/wait-did-as8003-just-disappear/

[56] D. Madory. The Mystery of AS8003. https://www.kentik.com/blog/the-mystery-of-as8003/

[57] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz. Towards an Accurate AS-Level Traceroute Tool. In *Proc. ACM SIGCOMM*, 2003.

[58] P. Marchetta, V. Persico, and A. Pescapè. Pythia: Yet Another Active Probing Technique for Alias Resolution. In *Proc. ACM CoNEXT*, 2013.

[59] A. Marder, M. Luckie, A. Dhamdhere, B. Huffaker, kc claffy, and J. M. Smith. Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In *Proc. ACM IMC*, 2018.

[60] Microsoft. Unusual Account Activity. https://answers.microsoft.com/en-us/outlook_com/forum/all/unusual-account-activity/a9b09af4-7ed4-477b-92e8-cf2e7673a858

[61] J. Murai. Announcement regarding IPv4 Address Block 43/8. https://blog.apnic.net/2020/03/25/announcement-regarding-ipv4-address-block-43-8/

[62] N.A. Speedchecker Ltd. Available at: http://www.speedchecker.com.

[63] NANOG. Rogers Canada Using 7.0.0.0/8 for Internal Address Space. https://mailman.nanog.org/pipermail/nanog/2011-May/036483.html

[64] NANOG. Another LTE Network Turns Up as IPv4-Only Squat Space + NAT. https://mailman.nanog.org/pipermail/nanog/2012-July/050386.html

[65] NANOG. Bogon List Update for Prefix for 5.1.0.0/19. https://mailman.nanog.org/pipermail/nanog/2012-May/048510.html

[66] NANOG. DoD IP Space. https://mailman.nanog.org/pipermail/nanog/2019-November/104001.html

[67] nmap. Nmap: the Network Mapper. https://nmap.org.

[68] O. Nordström and C. Dovrolis. Beware of BGP attacks. *SIGCOMM Comput. Commun. Rev.* 2004.

[69] U. of Oregon. Route Views Archive Project. http://routeviews.org.

[70] G. A. Office. Internet Protocol Version 6: DoD Needs to Improve Transition Planning. https://www.gao.gov/assets/gao-20-402.pdf.

[71] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In *Proc. ACM IMC*, 2016.

[72] PeeringDB. PeeringDB. http://www.peeringdb.com.

[73] Rapid7. Project Sonar, Reverse DNS. https://www.rapid7.com/research/project-sonar/

[74] J. Reardon, Á. Feal, P. Wijesekera, A. E. B. On, N. Vallina-Rodriguez, and S. Egelman. 50 Ways to Leak Your Data: An Exploration of Apps' Circumvention of the Android Permissions System. In *Proc. USENIX Security*, 2019.

[75] Reddit. If You Have an Android, You Can Find Out Where Your Information Is Being Routed. My Canadian Cell Info Is Going To UK Ministry of Defence... https://www.reddit.com/r/conspiracy/comments/1sql0i/if_you_have_an_android_you_can_find_out_where/

[76] Reddit. Reddit: Can Anyone Help me Shed Light on this? Apparently MOD are Accessing my Email Account. https://www.reddit.com/r/unitedkingdom/comments/3q4emc/can_anyone_help_me_shed_light_on_this_apparently/

[77] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. RFC 1918: Address Allocation for Private Internets. https://datatracker.ietf.org/doc/html/rfc1918

[78] P. Richter, M. Allman, R. Bush, and V. Paxson. A Primer on IPv4 Scarcity. *SIGCOMM Comput. Commun. Rev.* 2015.

[79] P. Richter, F. Wohlfart, N. Vallina-Rodriguez, M. Allman, R. Bush, A. Feldmann, C. Kreibich, N. Weaver, and V. Paxson. A Multi-Perspective Analysis of Carrier-Grade NAT Deployment. In *Proc. ACM IMC*, 2016.

[80] RIPE NCC. RIPE Atlas. https://atlas.ripe.net/.

[81] RIPE NCC. RIPE Routing Information Service (RIS). https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris.

[82] RIPE NCC. RIPEstat Routing History. https://stat.ripe.net/widget/routing-history.

[83] L. Salamatian. GitHub Code Repository – IPv4Squatting. https://github.com/Burdantes/IPv4Squatting.

[84] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy. An Analysis of Internet Content Delivery Systems. In *Proc. USENIX OSDI*, 2002.

[85] B. Schlinker, I. Cunha, Y.-C. Chiu, S. Sundaresan, and E. Katz-Bassett. Internet Performance from Facebook's Edge. In *Proc. ACM IMC*, 2019.

[86] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making Middleboxes Someone Else's Problem: Network Processing as a Cloud Service. In *Proc. ACM SIGCOMM*, 2012.

[87] J. Sherry, E. Katz-Bassett, M. Pimenova, H. Madhyastha, A. Krishnamurthy, and T. Anderson. Resolving IP Aliases with Prespecified Timestamps. In *Proc. ACM IMC*, 2010.

[88] N. Škoberne, O. Maennel, I. Phillips, R. Bush, J. Zorz, and M. Ciglaric. IPv4 Address Sharing Mechanism Classification and Tradeoff Analysis. *IEEE/ACM Trans. Netw.* 2013.

[89] T. Strickx. How Verizon and a BGP Optimizer Knocked Large Parts of the Internet Offline Today. In *The Cloudflare Blog*, 2019.

[90] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband Internet Performance: A View from the Gateway. In *Proc. ACM SIGCOMM*, 2011.

[91] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In *Proc. ACM IMC*, 2019.

[92] R. Tibshirani, G. Walther, and T. Hastie. Estimating the Number of Clusters in a Data Set via the Gap Statistic. In *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 2001.

[93] C. Timberg. A secretive Pentagon program that started on Trump's last day in office just ended. The mystery has not. https://www.washingtonpost.com/technology/2021/09/10/pentagon-internet-protocol-addresses-trump/
[94] C. Timberg and P. Sonne. Minutes before Trump left office, millions of the Pentagon's dormant IP addresses sprang to life. https://www.washingtonpost.com/technology/2021/04/24/pentagon-internet-address-mystery/
[95] A. Toonk. IPv4 for Sale - WIDE and APNIC selling 43.0.0.0/8. https://toonk.io/ipv4-sale-wide-and-apnic-selling-43-8/index.html
[96] ukcloud. Public Sector Cloud Offerings. https://ukcloud.com/
[97] Y. Vanaubel, J.-J. Pansiot, P. Mérindol, and B. Donnet. Network Fingerprinting: TTL-based Router Signatures. In *Proc. ACM IMC*, 2013.
[98] L. Vegoda. IPv4 Squatting: Awareness Raising and Research. https://www.icann.org/en/blogs/details/ipv4-squatting-awareness-raising-and-research-23-3-2010-en
[99] J. Weil, V. Kuarsingh, C. Donley, C. Liljenstolpe, and M. Azinger. RFC 6598: IANA-Reserved IPv4 Prefix for Shared Address Space. https://datatracker.ietf.org/doc/html/rfc6598
[100] D. Wing and M. IETF78. Nat Tutorial.
[101] WIPO. Wipo Alternative Dispute Resolution (ADR) for Intellectual Property Offices and Courts. https://www.wipo.int/amc/en/center/specific-sectors/ipoffices/
[102] Wireshark. T-Mobile : Clever or Insane. https://blog.wireshark.org/2010/04/t-mobile-clever-or-insane/
[103] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger. Leveraging Interconnections for Performance: The Serving Infrastructure of a Large CDN. In *Proc. ACM SIGCOMM*, 2018.
[104] Writer, Steve. Report Accuses BT of Supplying Backdoors for GCHQ and NSA. https://www.csoonline.com/article/2134231/report-accuses-bt-of-supplying-backdoors-for-gchq-and-nsa.html
[105] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat. Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering. In *Proc. ACM SIGCOMM*, 2017.

A HOW TO REPLICATE THE RESULTS OF THIS PAPER

In this section, we describe how the different results in this paper can be replicated. We share our Github repository analysis code[7] [83]. The Microsoft traceroutes are considered sensitive and cannot be made public. All the other data used in this paper can be accessed publicly (*e.g.,* RIPE Atlas traceroutes) or by contacting us.

A.1 Generating Squat Space

We generate a list of unrouted IPv4 prefixes (§2.1) every month. We provide a script that downloads one snapshot per week in a given month from each RIPE RIS collector to find the unrouted prefixes. We detail its exact utilization in the GitHub repository.

A.2 Gathering Public Datasets

Route History: We construct a script to crawl RIPE Route History information. We share this code in our repository. This data is used to generate Figures 1 and 11.

Arkipelago: We download Arkipelago traceroutes used to generate the August 2020 Internet Topology Data Kit (ITDK). We create AS-level mappings from these measurements performing IP-to-AS (§3.1). We use the AS-level mapping in our case study about bdrmapIT in Section 7 and Appendix C.

RIPE Atlas: We gathered RIPE Atlas measurements using the BigQuery interface and write a query to identify traceroutes crossing squat space. Below we provide the example query for March 2021:

```
  WITH
    traceroutes AS (
  SELECT
    msm_id,
    prb_id,
    src_addr,
    dst_addr,
    hops,
    start_time
  FROM
    `ripencc-atlas.measurements.traceroute` AS t
  WHERE
    DATE(start_time) BETWEEN "2021-03-08"
    AND "2021-03-15")
SELECT
  msm_id,
  prb_id,
  src_addr,
  dst_addr,
  hops,
  h.hop_addr,
  start_time
FROM
  traceroutes,
  UNNEST(hops) AS h
WHERE
  h.hop_addr IN (
  SELECT
    string_field_1
  FROM
    `squatspace.march-2021`)
  AND (src_addr,
    dst_addr,
    start_time) IN (
  SELECT
    (src_addr,
      dst_addr,
      MAX(start_time))
  FROM
    `ripencc-atlas.measurements.traceroute`
  WHERE
    DATE(start_time) BETWEEN "2021-03-08"
    AND "2021-03-15"
  GROUP BY
    src_addr,
    dst_addr)
```

We create a parser to convert all traceroutes, regardless of their measurement platform, to the following format:

```
src_ip|src_asn|internal_info|ip_path
   |pos_squat_address|rtt_path|timestamp
```

A.3 Attribution

We built a script that takes all the traceroutes and attributes to each squat address the following metadata:

```
trace_id|src_asn|squat_ip|squat_index|squat_asn
   |squat_org|squat_rir|src_ip|squat_index_updated
   |date|data_source|legitimate|attribution
   |squat_prefix
```

Its exact utilization is explained in the repository [83]. We use this output to draw Figures 5 and 6, and to compute results in Section 5.

A.4 AS-level meta information

We provide a processing pipeline to generate the AS-level meta information dataset combining information from CAIDA's AS Relationships dataset, PeeringDB, and APNIC AS Population estimates. This dataset is used to profile organization using squat space for Figures 3 and 9, and to build Table 2. To identify country of registration for Figure 8, we use the country informed in the whoIS database.

A.5 Statistics

Filter Selection and Coverage: We compute the total number of measurements and number of source IP addresses per source AS during the period of measurements. For Microsoft, we cannot share that dataset as it would result in leaking sensitive information, but we share the counts for the other datasets. For RIPE Atlas, we create a query on Google Big Query to count the number of measurements for each AS during the period of measurements. We use this data to generate Figures 3 and 5.

Churn of squatters: We provide the script to generate Figure 10 by looking at the churn of squatting entities across three months.

[7] https://github.com/Burdantes/IPv4Squatting

A.6 Squat Space Leakage

We gather BGP routing table every week using BGPstream [71] and use RIPE Routing History to collect historical announcements of the squat space. The repository includes the scripts to gather both datasets, which are used to generate Figure 11.

B TOP x ORGANIZATIONS WITH RESPECT TO EYEBALLS AND CUSTOMER CONE

To determine whether our coverage was sufficient, we analyzed the number of organizations in the top x for which we had sufficient measurements to determine whether or not they were squatting (Section 4.1). We considered the top networks according to three separate metrics: total IPv4 address space allocation (Figure 5), customer cone [14] (Figure 12a), and number of users, as estimated by APNIC's AS Population dataset (Figure 12b).

Similar to Figure 5, Figures 12a and 12b show we obtained a large volume of measurements from a wide range of networks and network types. For example, we cover more than 90% of the top 50 networks ranked by the number of users, 80% for IPv4 space and 70% for customer cone. This lower coverage is not surprising since we determine representativeness to require at least 5000 measurements and 50 VPs passing through the AS, and our analysis considers primarily our Microsoft dataset (§2.2), which has more measurements from ASes with large user populations. On the other hand, networks with large customer cones are often transit providers with few human users, so they are less likely to originate many measurements. Exceptions to this are networks such as Comcast (AS7922), which functions as a hybrid transit/access network.

C DECISION TO NOT USE *BDRMAPIT*

We evaluated whether *bdrmapIT* [59] can improve the attribution of squatting to organizations. First, around 75% of the squat IP addresses in the Microsoft traceroute dataset appear before the first hop with a public address. These attributions are unlikely to be changed by *bdrmapIT*. For squatting hops surrounded by responsive hops appearing in the middle of the traceroute, we compare the identified squatting organization when using classic IP-to-AS mapping and when using *bdrmapIT*. We focus only on the Ark's datasets, since *bdrmapIT* was originally designed with the intention to work on Ark's data. Our first finding is that *bdrmapIT* maps squat space IP addresses to organizations using registry information. This is problematic in practice, as a squat space address can be reused by multiple organizations, but *bdrmapIT* will consistently map it to the same organization in all traceroutes. This observation argues for AS mapping techniques that are aware of squat space, which we expand upon in Section 7. To understand whether we should use *bdrmapIT* for non-squat space addresses, we use *bdrmapIT* on hops using public addresses only, then apply our squat attribution algorithm on the resulting AS-level path. Of the 146,123 traceroutes with squat space, using *bdrmapIT* results in attribution of the squat address for 10,943 traceroutes, and using our IP-to-organization mapping attributes squat space in 10,433 traceroutes. Of 10,013 traceroutes they both attribute, the two approaches agree on 9862 traceroutes (98.5%). In 4.0% (420) of the traceroutes, IP-to-AS mapping attributes a squatting organization while *bdrmapIT* is not able to, and 8.4% (920) reciprocally. The limited gain obtained by *bdrmapIT* does not justify the added complexity of deploying it within Microsoft's data processing pipeline.

Keyword	Number of IPs	/24s	Orgs	Matching Inferences
NAT-related keywords				
'kabel'	2K	15	1	100%
'comcastbusiness'	1K	8	1	100%
'business'	4K	24	7	88%
'broad'	41K	187	20	82%
'fibre'	2K	13	5	85%
'cpe'	47K	260	17	85%
'user'	29K	110	10	84%
'cable'	39K	186	10	80%
'client'	14K	74	13	78%
'customer'	16K	84	16	76%
Merged	220K	1117	115	83.5%
CGNAT-related keywords				
'cgn'	1K	6	2	100%
'cgnat'	1K	1	1	100%
Merged	1K	7	3	100%

Table 5: Evaluation using inferred prefix use from reverse DNS. We see that our methodology agrees with every prefix with CGNAT hints and performs with above 80% performance for most prefixes with customer control.

D EVALUATION USING REVERSE DNS

We introduce in Table 5 the breakdown of our inferences for a list of 12 strings that we were qualitatively able to associate with NATs. We consider the first 10 strings to refer to user-facing NAT deployments and the last two to refer to CGNAT deployments. We only consider /24 prefixes with consistent naming, *i.e.,* prefixes where at least 60% of the rDNS entries include the keyword. We note that our inferences closely match the use inferred from hostnames: the 'matching inferences' column shows the fraction of prefixes for which we (correctly) inferred a customer NAT deployment in the top 10 rows, and the fraction of prefixes for which we (correctly) inferred a CGNAT deployment in the last 2 rows. Furthermore, we see that some naming conventions apply across many organizations, indicating that our technique generalizes and strengthens our belief that we can differentiate cases of operator- and customer-configured squat addresses.

E PORT ANALYSIS

In this appendix, we detail the methodology to select the ⟨open port, application⟩ tuples we use to identify prefixes that are *not* deploying a CGNAT. We consider the port and service scanning results conducted by Censys [23], which scans the entire routable IPv4 address space and performs protocol probing on more than 4000 popular ports. The port scan that we used for the port and service analysis was conducted between 2022/07/15 and 2022/08/15. If a host was scanned by Censys for multiple times during this period, we only consider the last scan result of that host in our analysis.

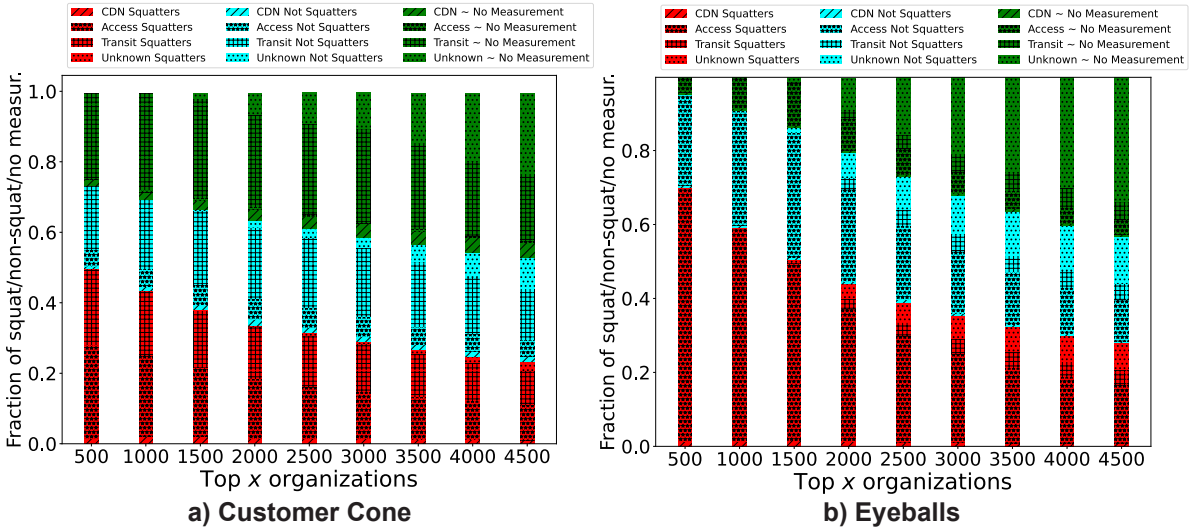

Figure 12: Proportion of squatting organizations that appear on the top x by (a) the size of the customer cone and (b) the number of eyeballs.

We selected 19 ⟨open port, application⟩ tuples that we expect will not be forwarded downstream by a CGNAT. We infer that IP addresses can receive connections on these ⟨open port, application⟩ tuples are likely not behind a CGNAT. We used a restrictive methodology when selecting the tuples to avoid false positives. (i) We only consider ⟨open port, application⟩ tuples where the application is running on its default port, which makes configuring static port-forwarding on the CGNAT not viable. (ii) We intentionally did not consider any ⟨open port, application⟩ tuples that could potentially be used to manage a CGNAT, such as 22/SSH, 80/HTTP, and 443/HTTPS. The full list of ⟨open port, application⟩ tuples that we selected, along with the number of unique IPv4 addresses accepting connections and their covering /24 prefixes, can be found in Table 6.

Fig. 13 shows the number of prefixes (y axis) with a given number of IP addresses with open ports (x axis). The blue bars count the IP addresses reachable on any open port, and the green bars count the IP addresses reachable on at least one of our chosen ⟨open port, application⟩ tuples. The black, red, and yellow lines denote the thresholds for prefixes inferred to *not* deploy CGNATs with low-, medium-, and high-confidence (§4.2), respectively.

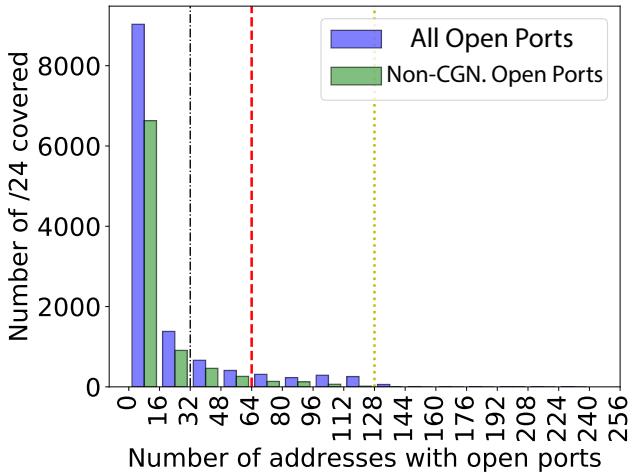

Figure 13: Number of prefixes with a given number of IPs reachable on any port (blue bars) and non-CGNAT ports (green bars, Table 6). Prefixes to the right of the black, red, and yellow lines are inferred to *not* deploy a CGNAT with low-, medium-, and high-confidence, respectively.

Application	Open Ports	# Unique IPs	# of /24s covered
HTTP/CWMP	7547	28239	2816
	58000	13458	589
	30005	6674	754
RTSP	554	11721	2381
FTP	21	7637	2154
PPTP	1723	6910	2131
DNS	53	16054	1963
TELNET	23	6529	1963
RDP	3389	116K	2304
NTP	123	5566	2071
SIP	5060	3865	805
SMB	445	2631	946
MYSQL	3306	1706	1046
OPENVPN	1194	1380	766
	443	80	30
SMTP	465	216	167
	587	220	165
POP3	110	219	181
IMAP	143	214	179
Merged	—	111616	8610

Table 6: List of ⟨open port, application⟩ tuples that we expect to be unreachable behind a CGNAT, along with the number of unique IPv4 address and IPv4 prefixes associated with each tuple.