

Inferring Carrier-Grade NAT Deployment in the Wild

Ioana Livadariu*, Karyn Benson†, Ahmed Elmokashfi*, Amogh Dhamdhere† and Alberto Dainotti†

*Simula Research Laboratory †CAIDA/UCSD

Email: {ioana,ahmed}@simula.no, {karyn,amogh,alberto}@caida.org

Abstract—Given the increasing scarcity of IPv4 addresses, network operators are resorting to measures to expand their address pool or prolong the life of existing addresses. One such approach is Carrier-Grade NAT (CGN), where many end-users in a network share a single public IPv4 address. There is limited data about the prevalence of CGN, despite the implications on performance, security, and ultimately, the adoption of IPv6. In this work, we present passive measurement-based techniques for detecting CGN deployments across the entire Internet, without the requirement of access to machines behind a CGN. Specifically, we identify patterns in how client IP addresses are observed at M-Lab servers and at the UCSD network telescope to infer whether those clients are behind a CGN. We apply our methods on data collected from 2014 to 2016. We find that CGN deployment is increasing rapidly. Overall, we infer that 4.1K autonomous systems are deploying CGN, 6 times the number inferred by the most recent studies.

I. INTRODUCTION

IPv4 addresses are rapidly running out. In 2011 the Internet Assigned Numbers Authority assigned its last available IPv4 addresses to the Regional Internet Registries (RIRs); moreover, RIRs are currently allocating from their last /8 block [1]. This scarcity and the slow uptake of IPv6 [2] – the long term solution to IPv4 depletion – have prompted Internet Service Providers (ISPs) to search for alternatives to meet their address needs. One alternative is to deploy large scale network address translation (NAT) mechanism known as Carrier-Grade NAT (CGN), which allows ISPs to put many customers behind a single public IP address.

CGN is built on the IPv4 address sharing principle used by the traditional NAT mechanism, where a single IPv4 address is shared among a number of users (end devices) [3]. In NAT44 (traditional home NAT) – deployed in home networks and small businesses – the address translation typically occurs at the users’ Customer-Premise Equipment (CPE) device. In the case of CGN, the address translation typically occurs inside the service provider’s network.

Surprisingly, there is little information available about the prevalence of CGN. Knowing which ISPs implement CGN could assist in: developing IP reputation systems such as blacklists, analyzing web usage from logs, capacity planning by CDNs or content providers, troubleshooting degradations in application performance [4], [5], and understanding the implications on regulatory tasks like lawful traffic interception and anti-spoofing [6]. Most existing CGN inference approaches require custom measurements from hosts within the ISP,

which prevents them from scaling to a comprehensive Internet-wide study. An exception is Richter *et al.*’s recent work [7], which (in addition to measurements from hosts within the ISP) analyzed internal address space leakage to identify few hundred ASes deploying CGN. However, it is unlikely this leakage-based technique will extend to other datasets.

Our main contribution is two generalizable methods for inferring CGN deployment using passive measurements, that could be applied to many datasets, e.g., CDN logs or Web access logs. The first method infers /24 blocks used for CGN by analyzing traffic with client identifiers. We apply this method to BitTorrent packets collected at the UCSD network telescope [8]. The second method achieves the same goal by investigating how frequently the same IP address accesses a web service. We develop this method using tests against the Measurement Lab (M-Lab) infrastructure [9].

There are limitations of our methodology, which we attempt to overcome by filtering our final results. First, our methods identify cases of large-scale IP address sharing, which is an inherent property but not a concrete proof of CGN. Thus, we may detect middleware solutions other than CGN (e.g., proxies). Moreover, there are diverse and complicated relationships between service providers and their customers. Without pinpointing the location of the CPE device, there is ambiguity as to where address translation occurs. For example, a non-CGN AS in our validation set announces prefixes for universities; these universities may deploy CGN themselves.

Using data collected from July 2014 to September 2016, we infer CGN deployments in 4.1K of the 17.4K measured autonomous systems (ASes). We validate our methods against recent CGN detection methodologies [7], [10], and detect more than 85% of the networks they detect as deploying CGN. We find that the number of inferred CGN ASes is increasing over time. However, we find no evidence that CGN deployment negatively impacts IPv6 adoption: ASes with CGN deployments are 17% more likely to originate IPv6 prefixes than the average AS. Our findings suggest that a large fraction of inferred CGN networks are deploying CGNs as a solution to IP address shortage, and do not acquire IPv4 addresses from the IPv4 transfer markets.

II. DATASETS

A. Primary datasets

We infer CGN deployment with IBR and M-Lab data.

IBR data: The UCSD Network Telescope (UCSD-NT) [8] passively collects unsolicited traffic, called Internet Background Radiation (IBR), sent to an unused contiguous /8 address block. We extract BitTorrent KRPC packets from IBR. KRPC is the protocol implementing BitTorrent’s distributed hash table (DHT) – the mechanism through which BitTorrent nodes discover torrent locations [11]. UCSD-NT receives KRPC packets because some DHT nodes spread misinformation – either accidentally (e.g., bit flips, programmatic bugs) or intentionally (e.g., to inhibit torrent downloads) [12]. When the misinformation states that a UCSD-NT address can assist in a torrent download, genuine clients send it KRPC packets.

These packets provide a sample of BitTorrent users, but there are several challenges with this dataset. First, there is considerable fluctuation in the data: notably, starting with July 2015 the number of observed BitTorrent packets increased from about 100M to 10 billion per month [13]. We further refer to the months prior to July 2015 as *low-volume*, and other months as *high-volume*. Additionally, there may be some false information. According to the KRPC specification, every packet contains the node’s globally unique, randomly generated 160-bit ID. This should imply a one-to-one correspondence between machines and IDs. However, not all nodes follow the KRPC specification.

M-Lab data: The M-Lab project provides a set of measurement tools that users can run to test their network connection [9]. From the M-Lab dataset, we use server-side logs generated by the Network Diagnostic Tool (NDT) test to extract the public IPv4 address of the client running each test, then group these runs according to the origin Autonomous Systems (AS). We use BGP data from RouteViews [14] and RIPE NCC [15] to find the origin ASes of the IP addresses. We leverage the intuition that within a period of time an IP address is likely to appear more often if it is one of the public addresses of a CGN, than if it comes from a network that does not deploy CGN (see Sec. III-B).

B. Filtering the data

Our data is collected externally from the networks for which we infer CGN deployment. However, such networks may route prefixes for other entities (i.e., edge networks), that could deploy within their networks large-scale IP address sharing solutions. To avoid such cases, we filter the data to only include /24 blocks owned by access/transit networks. We further summarize the datasets used to filter our results, and describe the filtering process.

AS Classification dataset: Researchers from the Center for Applied Internet Data Analysis (CAIDA) developed a method that classify ASes based on their business type [16] into: “Enterprise” networks, “Content” and “Transit/Access” providers.

WHOIS data: RIRs maintain databases that contain information regarding the registered Internet resources within their own region. We leverage WHOIS data to identify routed IPv4 address blocks that are registered to “Transit/Access” ASes.

RIRs extended delegation files: The RIRs publish daily files that summarize the current allocation and assignment of their

Internet resources (i.e., IPv4 and IPv6 addresses, and AS numbers) [17]. For each resource, the RIR provides: type, country, date of the allocation/assignments of the resource. Recently, the RIRs started to provide for each resource also an organization identifier that corresponds to the resource holder [18]. This identifier is the same within a single file (i.e., for each day), but is not guaranteed to be consistent over time. Using the organization identifier, we map the IPv4 space allocated/assigned to “Transit/Access” ASes.

Filtering process: Using the AS classification dataset, we identify 5.1M /24s routed by 17.4k “Transit/Access” ASes. We then use RIRs extended delegation files to select IPv4 address blocks that share the same *organization identifier* as the AS that announces them. Next, we examine WHOIS data: we query Team Cymru’s service [19] to extract the *name* of each AS; we then compare this value to the *name* and *description* of the blocks advertised by the AS in the bulk WHOIS data. That is, for each AS we identify the IPv4 space that matches in either name or description. In total, 4.2M (81.1%) /24 address blocks matched the ASes in at least one of the two analyzed datasets (i.e., RIRs extended delegation file or WHOIS data).

C. Validating our inferences

To validate our CGN detection techniques presented in Sec. III, we construct lists of “Transit/Access” networks (ASes) that were deploying CGNs and that were not deploying CGN during July 2015. We leverage (i) information extracted from a survey carried by CAIDA [20], (ii) email confirmations and online resumes (LinkedIn profiles indicating experience with CGN at an Internet provider), (iii) reverse DNS names of IP addresses coming from a set of networks for which we manually collected the number of subscribers, and (iv) CGN networks detected with client-side measurements by Lutu *et al.* [21]. Our list of validation networks is comprised of 22 CGN ASes and 15 non-CGN ASes. Networks in both categories include mobile and fixed operators and are spread across four RIRs: 9 from ARIN, 8 from APNIC, 18 from RIPE and 2 from LACNIC.

III. INFERRING CGNS USING PASSIVE MEASUREMENTS

In this section, we describe two methods for inferring CGN. The first method leverages traffic containing unique client identifiers to identify networks where many clients share external IP addresses. We apply this method to BitTorrent IBR. The second method estimates how frequently public-facing CGN IP addresses appear in logs generated by users when accessing a web service. We apply this method to logs of M-Lab tests. Both methods aim to identify CGN networks with 100 or more users per external facing IP address, as this was reported as a popular configuration in the CAIDA survey [20].

A. Inferring CGNs using IBR data

Many types of Internet traffic contain unique *client identifiers* (IDs) – a packet field or combination of fields that uniquely identify a machine. Observing multiple IDs with the same external (public) IP address indicates that the address

is shared. Unfortunately, technologies besides CGN, namely home NAT and DHCP, also enable IP address sharing.

To differentiate CGN from home NAT, we leverage CGN implementation characteristics. We expect that more clients share an external IP address with CGN than home NAT. Moreover, we expect that external IP addresses that the ISPs configure for CGN are both numerous and contiguous. Specifically, we look for /24 blocks where many IP addresses in the block show evidence of address sharing. This requirement helps eliminate single IP addresses used as web proxies. We work at the /24 granularity as this is typically the smallest block announced in BGP and likely similarly configured.

To differentiate CGN from DHCP, we leverage temporal differences in traffic patterns. In CGN, many machines simultaneously use the same external IP address. In DHCP, a device uses an external IP address until the lease expires or it is relinquished. Intuitively, the observed sequence of IDs from a given IP address should be more mixed in CGN than in DHCP configurations. To capture the notion of mixed, we call a sequence of time-ordered packets from the same source IP address *interwoven* if there exists some ID whose associated packets do not all appear consecutively.

From a temporal perspective, we also look for persistence: in CGN configurations, clients collectively send traffic throughout the observation period, resulting in a persistent signal. Such a signal helps differentiate from home NAT and DHCP,¹ and guards against temporary misinformation.²

Finally, we evaluate additional traffic attributes to increase our confidence in our inferences (e.g., previous work fingerprints machines using TTL [23] or TCP options [24]). Attributes corroborate that diverse, heterogeneous clients generated the signal. In our analysis, we use the client version contained in BitTorrent packets (e.g., uTorrent).

Methodology: We infer that a network deploys CGN in a contiguous /24 block if the individual IP addresses send traffic that is persistent, from heterogeneous clients, and associated with many interwoven IDs. We score a contiguous /24 block, B , using the following metric:

$$S(/24) = \sum_{i \in /24} ids_i \times interwoven_i \times persistent_i \times diversity_i$$

where ids_i , $interwoven_i$, $persistent_i$, and $diversity_i$ are values assigned for an IP address, i , between 0 and 1 indicating non-CGN-like and CGN-like behavior respectively. Higher scores indicate higher confidence that a network deploys CGN.

In general, setting the values that compose our metric depends both on the specific type of ID used and on the dataset to which our method is applied (e.g., on the popularity of the application generating IDs and the fraction of the traffic

observed). We use a combination of domain knowledge (how popular is the program generating IDs?) and understanding of the dataset (how common are interwoven IDs?) to effectively set ids_i , $interwoven_i$, $persistent_i$, and $diversity_i$. For our IBR dataset, we calculate our metric, on a monthly basis, for a /24 block, as follows:

- ids_i : Since we are interested in the magnitude of the number of IDs, relative to CGN deployments with 100 devices per external IP address, we use a logarithmic scale to compute ids_i as $\min(1, \log(\text{num IDs})/\log(100))$.
- $interwoven_i$: We set this value to 1 if the packets are interwoven (as defined above). Otherwise, we penalize the IP address by assigning a value of 0.5.³
- $persistent_i$: We calculate the fraction of days in the month that we observe packets from the IP address.
- $diversity_i$: If we observe multiple client versions (e.g., uTorrent and Vuze), we set this value to 1. Otherwise, we assign a value of 0.5.³

While chosen somewhat arbitrarily, these values represent a reasonable starting point for datasets with unique client IDs. With additional ground truth we could optimize the constants and relative importance of the parameters. Moreover, we found that slight alterations to the definitions did not substantially alter the set of /24 blocks with high scores.

What score indicates a CGN deployment? Our threshold for inferring CGN depends on the underlying traffic. To show this dependence, we simulate various network configurations while varying p , the probability that a BitTorrent host sends traffic to the UCSD-NT on a given day. Specifically, we compare CGN with 10, 100, and 1000 devices per IP address to home NAT (with 5 devices [25]) and various DHCP lease times. With a 30% probability [26] we determine if a device runs BitTorrent; for each device running BitTorrent we determine which days (out of 30 days) it sends traffic to UCSD-NT by conducting a Bernoulli trial with probability p of success.

Figure 1 shows the result of 100 runs with 10 different values of p . The median score for a /24 block used in CGN increases with the probability of observing BitTorrent traffic. For most of the tested probability levels, we note a clear separation of scores between networks in CGN with 100 or more devices sharing an IP address and non-CGN deployments.

Although the results of a small experiment⁴ suggest there is a low probability of a BitTorrent host contacting to UCSD-NT, it is unclear how to determine the exact probability. Absent a comprehensive set of ground truth networks, we set the threshold empirically – each month – by finding outliers in the score distribution. We expect the scores from non-CGN deployments

¹For moderately popular protocols or websites, we do not expect devices behind home NAT/DHCP to collectively generate continual traffic as instead happens for devices behind CGN.

²E.g., we observed an apparent programmatic error in an implementation of the DHT security extension [22] that caused some clients to send a single burst of many properly-formed-according-to-the-security-extension IDs. Without evaluating temporal attributes we would have mistaken these IDs as generated by many hosts.

³In low-volume datasets scoring IP addresses as zero because they failed to meet a criterion (e.g., not interwoven) resulted near-zero scores for all /24 blocks, including networks known to deploy CGN. To overcome this limitation and conduct a broad (though potentially less accurate) study, we instead moderately penalize these IP addresses.

⁴We constantly ran two BitTorrent clients (but did not torrent any files) for two months; one client attempted to contact UCSD-NT in 4 days, the other attempted to contact UCSD-NT in 24 days.

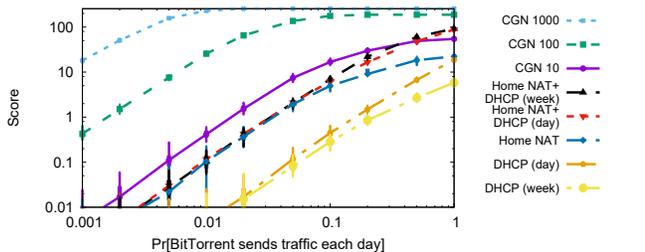


Fig. 1. Score for /24 address block used in different network configurations. CGN configurations with at least 100 devices per IP have a higher score than configurations involving home NAT and DHCP.

to follow the Poisson distribution in the Generalized Linear Model⁵. We use the scores for all /24 blocks with clients sending BitTorrent IBR (e.g., all 3.7M /24 blocks in July 2015) to estimate the parameters of the Poisson distribution. By including CGN networks in the parameter estimation, we conservatively set our threshold, but the result should not be egregiously high as we expect non-CGN networks to dominate our dataset (in terms of /24 blocks). Specifically, we set the CGN threshold to the 99.99th percentile, which is 0.95 and 5.33 for January and July respectively.

Do low scores indicate that CGN is not deployed? A low score does not imply that CGN is not deployed. We need to observe significant traffic to determine that an operator deploys CGN (e.g., in our data, /24 blocks with scores greater than 5 sent at least 700 packets). We are unable to detect CGN deployments if the operator prohibits or limits BitTorrent usage. Additionally, our method is more likely to identify CGN deployments when many hosts share an external address. This limitation prevents us from conclusively determining which networks deploy CGN, but does not hinder our primary objective of improving our understanding of CGN prevalence.

In our analysis, we are interested in how an operator configures their network. In these cases, we can argue that there is evidence that CGN is not deployed when a low score is accompanied by either significant traffic volume or other /24 blocks with high scores in the same AS. However, there may be exceptions to this argument (e.g., a single host could generate significant BitTorrent IBR traffic; file-sharing policies could differ throughout a network).

Sensitivity analysis: To understand the influence of each component of our metric, we generalize the formula for $S(/24)$ as follows:

$$S(/24) = \sum_{i \in /24} ids_i^a \times interwoven_i^b \times persistent_i^c \times diversity_i^d$$

The values a , b , c , and d are artificial weights used to manipulate the importance of our parameters. That is, by changing the values of a , b , c , d we can increase (e.g., $a > 1$) or decrease (e.g., $a < 1$) the importance of the respective component. Table I reports the percentage of ASes we inferred as deploying CGN in January and July 2015 (i.e.,

⁵The Poisson distribution counts the number of *events* (IP addresses in non-CGN configurations showing evidence of sharing) in a *fixed interval* (/24 block), and is appropriate when events occur *independently* (the behavior observed for hosts with non-CGN address does not depend on other hosts) and with a *known average rate* (our simulation shows the average score produced by an IP address is close to zero in all non-CGN configurations).

TABLE I
PERCENTAGE OF ASes THAT APPEAR TO DEPLOY CGN UNDER THE MODIFICATION OF OUR STANDARD INFERENCE ($a=b=c=d=1$). THE RELATIVELY HIGH PERCENTAGES SHOW THAT OUR METRIC IS ROBUST TO MINOR METRIC VARIATIONS.

		Jan. 2015 - x						
		0	0.125	0.25	0.5	2	4	8
ids:	$a=x, b=c=d=1$	77%	86%	87%	88%	88%	72%	54%
interwoven:	$b=x, a=c=d=1$	87%	88%	88%	90%	91%	88%	85%
persistent:	$c=x, a=b=d=1$	84%	84%	86%	86%	80%	65%	54%
diversity:	$d=x, a=b=c=1$	77%	86%	87%	88%	88%	72%	54%
		July 2015 - x						
		0	0.125	0.25	0.5	2	4	8
ids:	$a=x, b=c=d=1$	96%	98%	98%	99%	85%	72%	66%
interwoven:	$b=x, a=c=d=1$	97%	97%	98%	99%	95%	91%	90%
persistent:	$c=x, a=b=d=1$	89%	90%	92%	95%	87%	77%	69%
diversity:	$d=x, a=b=c=1$	97%	98%	98%	99%	98%	98%	98%

$a = b = c = d = 1$) that also meet the CGN thresholds under the various parameters. In general, changing one parameter does not exclude CGN ASes – especially in the high-volume dataset (July 2015). The attributes of *ids* and *persistent* have the largest effect on our results. We believe that this finding stems from our dataset’s sparsity – both in the number of hosts contributing traffic and how often they contribute.

B. Inferring CGNs using M-Lab data

M-Lab logs IP addresses of users that run performance tests against its infrastructure. We leverage these IP addresses and how frequently they appear in a given time window for detecting CGN deployment. The underlying assumption of our approach is that CGN public-facing IPs are likely to appear more frequently than IPs from non-CGN configurations.

Methodology: We divide our measurement period, T , into N windows of length t , which we call *detection windows*. Let p be the probability that a user runs an M-Lab test at least once in T , i.e., in at least one detection window. The probability p captures M-Lab popularity and user’s behavior, and can vary from one network to another. Our goal is to estimate the likelihood of observing an IP address in more than one detection window for different addressing configurations. We start by discussing the likelihood of observing multiple IP addresses from non-CGN configurations. Then we proceed to identifying false negatives caused by M-Lab’s popularity.

Likelihood of observing IPs from non-CGN configurations. As defined earlier, p is the probability that a user accesses M-Lab at least once in T . However, once the user accesses M-Lab, the probability that the same user accesses M-Lab again in T is no longer p . This is because the user now knows about M-Lab and thus her future access decisions do not depend solely on M-Lab’s popularity. They will depend on the reasons for accessing M-Lab the first time and the user’s satisfaction with her first experience. In the absence of information about these factors, we assume that the probability a user accesses M-Lab, given that she has accessed it before, is 0.5. We perform a sensitivity analysis (at the end of the section) to explore the impact of this parameter on our inference. For users with *static* IP addresses, this probability also represents the probability that an IP address appears in more than one detection window. Consequently, the probability of observing μ IP addresses coming from *static* configurations in multiple detection windows is $Q_{\mu}^s = (0.5)^{\mu}$.

In *DHCP* and *home NAT* configurations, an IP address is no longer assigned to one user. For these configurations, the likelihood of observing an IP multiple times in T depends also on the configuration parameters, i.e., DHCP time lease and the number of users behind the NAT. Recent work showed that a DHCP lease is typically a few days long [27], which is longer than our detection window t . Recall that the probability that a DHCP-assigned or NAT IP appears in our logs once is p . The probability this IP appears again in T is: p if the IP is reassigned/reused by another user or 0.5 if the same user repeats the test. Hence, the probability of observing μ IPs used in *DHCP* or *home NAT* configurations is between p^μ and $(0.5)^\mu$. Hence, *DHCP* and *home NAT* are not expected to increase the likelihood of observing several IPs on multiple detection windows beyond the static case.

Since Q_μ^s is the probability of observing μ IPs coming from a /24 block used in non-CGN configurations, its complement corresponds to the probability of μ IPs coming from a /24 block used in CGN configuration. That is, the complement of Q_μ^s corresponds to the level of confidence in the hypothesis that the observed μ is caused by a CGN configuration. Accordingly, we can set the desired confidence level and calculate the corresponding μ value. For e.g., if we observe 15 IPs or more from an address block in more than one detection window, we can be 99.99% confident that this block is not assigned statically, via DHCP, or behind a home NAT. Note that the approach above is independent of address block lengths.

Addressing false negatives. The task of distinguishing CGN from non-CGN configurations is prone to *false negatives*, i.e., CGN IPs that are classified as being in a non-CGN configuration. This happens when a public-facing CGN IP appears only in a single detection window. To control for false negatives, we derive an expression that estimates, for a public-facing CGN address block, the number of IPs that we expect to see in a single detection window in the measurement period T . Let x be the size of the address block and c be the CGN compression factor, i.e., the number of users that share a single public IP address. Further, let P' be the probability that an public-facing CGN IP address appears *once* in T , which follows a binomial distribution and is given by:

$$P' = \binom{N}{1} p_{cgn} (1 - p_{cgn})^{N-1} \quad (1)$$

where p_{cgn} is the probability that a public-facing CGN IP appears in at least one detection window, which is given by $p_{cgn} = 1 - (1 - p)^c$. Recall that p is the probability that a user accesses M-Lab and N is the number of detection windows in T . Hence, the expected number of false negatives, E , is:

$$E = (2^x - 2) \times P' \quad (2)$$

Given the block size and compression factor, we can use Eq. 2 to estimate the expected number of false negatives. We then reject (accept) the hypothesis that a non-statically configured block is used for CGN, if the observed false negatives is higher (lower) than this expected value.

Parameter estimation. The above methodology involves a number of key parameters. First, we need to pick an appro-

prate measurement period T , and detection window t . In the following, we set T to 90 days and t to one day. Setting T to 90 days allows us to account for the low popularity of the service. Also a one-day detection window makes our methodology less prone to false positives due to DHCP address reassignment. Essentially, t needs to be shorter than typical DHCP leases to reduce the number of times we see the same IP because of reassignments. Second, we need to determine the threshold for separating CGN and non-CGN configurations, i.e., the minimum number of IPs that we need to observe in more than one day. We set this threshold to 15 IPs which corresponds to a 99.99% confidence level. Third, we need to estimate the probability, p , that a user accesses M-Lab and use it to estimate the expected number of false negatives, E , per a given address block, i.e., using Eq. 1 and Eq. 2. For a network a , let c_a be the number of its customers (as estimated by APNIC Labs [28]), n_a^d be the number of IPs from this network seen on a day d . The probability p for this network can be estimated as $\hat{p}_a = \frac{1}{90} * \sum_{d=1}^{90} \frac{n_a^d}{c_a}$. To get an idea about the range of p , we estimate it for the *non-CGN* networks in the validation set. The estimated values vary from 0.0000039 to 0.00016, showing that M-lab has a very low popularity. Using the maximum estimated value of \hat{p}_a and assuming a compression factor of 100, we estimate that $E = 88$ for a /24. The choice of the compression factor is based on responses to CAIDA's IPv6 survey [20].

Sensitivity analysis: Using data from APNIC Labs, we are able to estimate the number of false negatives E for a /24 address blocks. However, we rely on different assumptions when choosing the values for the probability p that the same user accesses M-Lab multiple times in T and the minimum number of μ IP addresses observed from a /24 address block in multiple detection windows. Thus, we seek to determine the impact of p and μ on our inference of CGN configurations. We show in figure 2 the level of confidence that μ observed IPs come from /24 blocks behind CGNs, for different values of μ and p . We find that observing $\mu < 10$ IPs during our measurement period provides a high confidence level only for $p < 0.1$. For $0.1 < p < 0.9$, we need to observe at least 10 IPs in multiple detection windows in order to classify with a 90% confidence that an address block as used in CGN deployments. For $p > 0.9$, we obtain the same level of confidence only for $\mu > 44$ IPs. In the figure, the vertical solid line corresponds to our chosen threshold of $\mu = 15$ IPs, which provides a high confidence level for $p < 0.9$. We thus believe that observing 15 IPs from a /24 block in multiple detection windows is a strong indicator that the /24 is used in a CGN deployment.

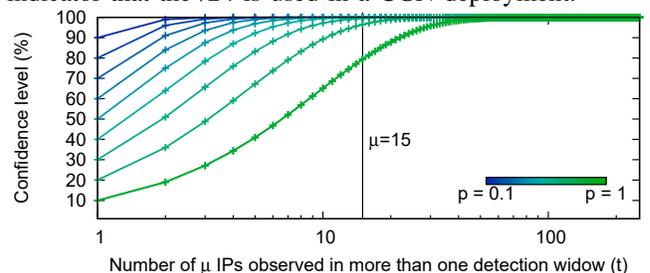


Fig. 2. Confidence level that μ IPs come from /24 blocks behind CGNs.

IV. CROSS-CHECKING WITH THE VALIDATION SET

We use the validation set of CGN and non-CGN networks to check our methodologies. For each AS from the validation set, we compare the IBR and M-Lab inferences at the /24 granularity during both the low-volume (January-March) and high-volume (July-September) periods of time in 2015; recall that BitTorrent IBR increased 100-fold between the two time periods [13]. The two datasets collectively (commonly) observe 111K (81K) and 144K (77K) /24s from 28 and 30 ASes during the former and latter periods, respectively. Figure 3 shows our inferences for the commonly observed /24 blocks. For each AS, we represent the percentage of /24 blocks inferred as used in CGN deployment by both methodologies (*Both*), by only one methodology (*Only M-Lab/IBR*), and the percentage of /24s that both IBR and M-Lab classify as not used in CGN deployments (*Neither*). The solid and hashed patterns distinguish between the two periods of time.

True Positives/False Negatives: A network operator may deploy CGN in only a portion of its network, keeping traditional deployments (home NAT, one IP address per machine, etc.) in the other parts (see Sec. V). Accordingly, we should infer CGN for at least one /24 block from the ASes known to deploy CGN in our validation set.

Our methods successfully inferred 1,233 CGN /24 blocks in 15 of the 19 ASes known to deploy CGN. There is a high degree of consensus between the two methods: in July-August 2015, 99% of /24 blocks inferred as CGN through the M-Lab method are also inferred as CGN by the IBR-based method. The M-Lab method is consistent across time periods, identifying 80 and 93 /24 blocks (11 and 14 ASes) in the two measurement periods. The IBR traffic increase resulted in an increase from 7 to 15 ASes that are true positives.

Four ASes known to deploy CGN are false negatives as they did not meet either IBR or M-Lab requirements. While our methods miss CGN deployments with few users per external IP address, these false negatives more likely reflect the unpopularity of BitTorrent and M-Lab. For example, two of the ASes, British Telecom (AS2856) and Sky (AS5607) are based in the UK, which is known to block popular torrent sites [29]. Our inability to infer CGN in this scenario is reasonable, and highlights the potential benefit using of multiple data sources. *True Negatives/False Positives:* Failing to meet the criteria set by our methods does not imply that CGN is not deployed. However, /24 blocks in ASes that use traditional deployments should not meet our CGN requirements.

In both time periods almost all of the 90K /24 blocks from traditional deployments fail to meet our CGN requirements. Moreover, we rarely find exceptions. Both methods identify /24 blocks in Orange (AS3215) used by mobile clients (according to WHOIS) as deploying CGN, which is consistent with a presentation by the company [30]. The remaining false positives, identified by the IBR-based method, come from Hungarian universities with larger student populations than allocated IP addresses. We believe this is a case where a large-scale IP address sharing is implemented by the university

and not their upstream provider, Hungronet (AS1955). Consequently, these false positives reflect the difficulties involved in curating a set of ground-truth ASes and not a misclassification of traditional deployments as CGN.

V. CGN DEPLOYMENT ANALYSIS

In this section, we analyze our inference by assessing the number of ASes and /24s used in CGN deployment.

Inferred CGN networks

We apply our methodologies to five three-month periods of IBR and M-Lab⁶ data collected from 2014 to 2016, and report our findings for “Transit/Access” ASes and the advertised /24s space owned by such networks.

Analysis of the inferred CGN networks: Figure 4 shows the CGN deployment in terms of percentage and number of inferred ASes and /24s. We split the results into three categories: /24s and ASes inferred by both methodologies (*Intersection(IBR,MLab)*) and by just one methodology (*Only(IBR/MLab)*). From the second half of 2015, we observe a significant increase both in the number of inferred ASes and /24 blocks. This finding is the result of the hundred-fold increase in BitTorrent IBR volume, which enabled the analysis of significantly more networks. Overall, we infer 4,191 ASes and 154,098 /24 blocks involved in a CGN deployment. These correspond to 23.9% and 3.64% of the measured ASes and /24 blocks. Three-quarters of the inferred ASes are detected only by the IBR-based method, whereas 22% are inferred by both methodologies. Networks comprised in the former category fail to meet the M-Lab method’s CGN requirements most likely due to the low popularity of the service. For these networks, we observe M-Lab tests coming on average from 4 IPs per /24 blocks, which is significantly lower than the minimum number of IPs we impose to observe multiple times from /24s used in CGN deployments. However, the Only(IBR) values are not consistent across the periods.

The overall number of inferred networks increased significantly over time – from 1.2K in 2014 to 3.4K in 2016. While our visibility improved due to an increase in BitTorrent volume in July 2015, we still see an increase from late 2015 to 2016. Digging deeper into the set of inferred CGN networks, we progressively infer 2,920 ASes; i.e., networks that start deploying CGN at some point during the study period and continue deploying it until the end. A large fraction of these ASes (2,177) are inferred starting with the July-September 2015. 634 ASes are inferred throughout study period (2014-2016). Our analysis suggests that CGN has been used to solve IPv4 address scarcity. This is consistent with the observation that the heavy-hitter addresses (e.g., those used by CGN) received an increased share of the bytes served by a CDN [31]. *Types of networks that deploy CGN:* Using WHOIS data we classify the inferred CGN ASes from the perspective of regional registry, and whether the network is a mobile operator or not. The majority of networks (66.52%) are from the RIPE

⁶We curate the M-Lab data by filtering 556 IPs responsible for a disproportionately large number of tests, which we call *heavy hitters*. We confirm that these IPs correspond to users that run periodic measurements against M-Lab.

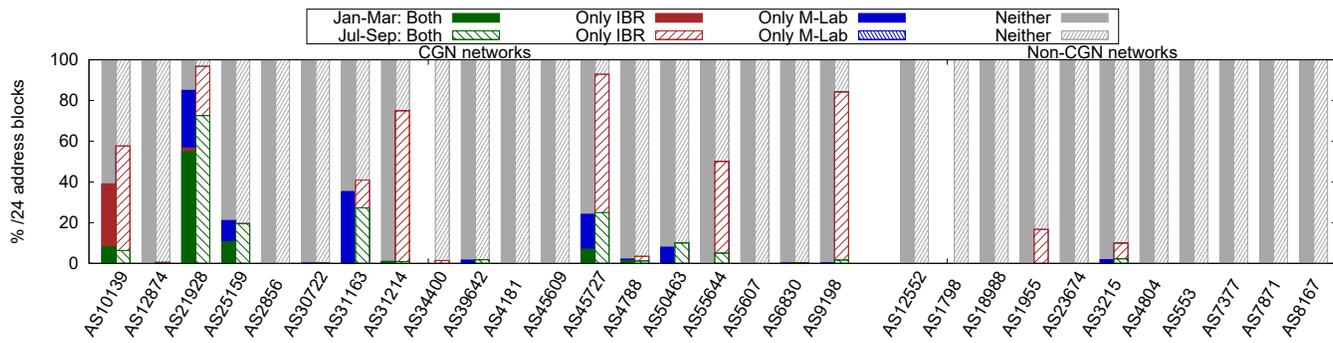


Fig. 3. Classification of the /24 blocks in non-CGN and CGN ASes from the validation set during January-March, and July-September 2015.

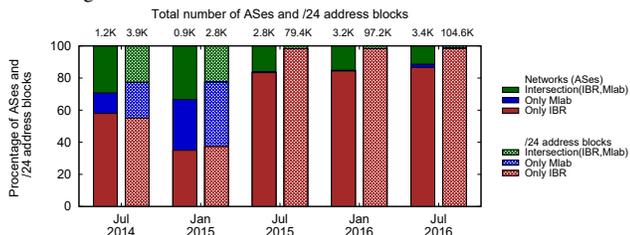


Fig. 4. Percentage of ASes and /24s inferred as deploying CGN by both methodologies (*Intersection(IBR,MLab)*) and just one methodology (*Only(IBR/MLab)*). Starting with July-September 2015, the IBR results account for more than 98% and 80% of the inferred /24 blocks and ASes, respectively. The remaining networks are distributed among all other RIRs; 13.88% from APNIC, 11.23% from LACNIC, 4.74% from ARIN and 3.60% from AFRINIC. To classify a network as a mobile operator, we retrieve the name and description of the inferred /24 blocks in the WHOIS data, and search for keywords (e.g., “mobile”, “mobility”) and names of well-known mobile operators [32]. 28.85% of the inferred networks appear to be mobile operators.

Comparison with active inference methodologies

We compare our inference to the results reported by two recent active measurement studies [7], [10]. NAT Revelio [21] is an active *inside-the-network* technique for detecting CGN at the IP address granularity. Using this technique, Mandalari *et al.* [10] detected CGN deployment in 52 /24 blocks. We inferred 47 (90%) of those 52 /24 blocks as part of a CGN configuration in at least one time period. The most comprehensive prior study on CGN prevalence [7] infers CGN deployment through Netalyzr [33] clients and internal address space leakage exposed during crawls of the BitTorrent DHT. In their study, Richter *et al.* collected data during the second part of 2015 to discover 421 ASes deploying CGN. However, the authors supplied us with an updated list that included 609 CGN-positive ASes. Of these ASes, 582 are classified as “Transit/Access” networks. We inferred 507 (87.11%) ASes of the 582 as deploying CGN; most of the networks that we miss are inferred by Netalyzr.

Overall, we discover more than 85% of the CGN networks inferred by previous work. The missed CGN deployments may be due to a small number of devices sharing an external IP address (e.g., operators preemptively deploying CGN), or a lack of visibility of our datasets into the networks with NAT Revelio or Netalyzr clients.

Using our approach we detect six times more ASes than the inferred 609 ASes by Richter *et al.* By applying our generic techniques to large datasets, we achieved our goal of

a better understanding about the prevalence of concurrent IP address sharing. A portion of this increase may be attributed to less stringent requirements (Richter *et al.* only analyze ASes where their crawl of the BitTorrent DHT revealed at least 200 peers with unique IP addresses), our longer observation period which captures churn in the addresses used for CGN deployment, and errors in our inferences (e.g., both active measurement techniques leverage the presence of internal IP addresses which is a better indicator of CGN – as opposed to other middleware technologies – than concurrent clients).

CGN deployment configuration

Mixed configurations: Our data suggests that most ASes deploying CGN, use non-CGN network configurations in other portions of their address space. For example, only mobile address ranges in Orange (AS3215) have high IBR scores, while all other blocks have low scores. We say these networks have a *mixed* configuration. Mixed configurations most likely result from ISPs incrementally converting their addresses to CGN to meet the increased demand for IPs.

If we infer that an AS deploys CGN using BitTorrent IBR, then the AS does not universally block BitTorrent, and a low score for other /24 blocks is evidence of non-CGN deployment. For the ASes meeting our IBR CGN requirements, we conclude that the AS has a mixed configuration if there is at least one /24 block sending BitTorrent IBR with a below average score (in all three months). Overall, three-quarters of ASes inferred to deploy CGN using BitTorrent IBR have mixed configurations.

Which IP addresses are used for CGN? We find that some networks deploy CGN in only a portion of a /24 block. In particular, T-mobile (AS21928), had 12 /24 blocks where the lower half of each block appeared to be part of a CGN configuration (IBR score for IP address greater than 0.5), while we did not receive traffic from the upper half. A T-mobile engineer confirmed that this pattern was consistent with their implementation [34]. Using July 2015 IBR data, we found 187 /24 blocks in 88 ASes where the CGN deployment appeared to be limited to a single /25, /26, /27 or /28.

Pooling types: A CGN device maps internal addresses to external addresses in either an arbitrary or paired configuration [35]. In the arbitrary case, an internal IP address may map to multiple external addresses at the same time. In the paired case, the same external address is used for all sessions associated with the internal address. For proper UDP functionality, RFC 4787 recommends paired address pooling.

Using our IBR dataset, we say that there is evidence of

arbitrary pooling for an IP address i if the upper bound on its usage is less than 5 minutes. That is, given a sequence of time ordered packets with the same ID, there exists a subsequence of three packets, p_A, p_i, p_B from IP addresses $A, i,$ and B respectively ($A \neq i$ and $B \neq i$), such that the difference in timestamps of p_B and p_A is less than 5 minutes. We include a /24 block in our final count if more than five IDs using IPs in the /24 block provide evidence of arbitrary pooling.

In the dataset from July 2015, we identify 14k /24 blocks in 73 ASes that likely use arbitrary pooling. These 14k blocks account for 42% of all CGN blocks in the July 2015 IBR dataset. This percentage is skewed due to the behavior of two Chinese ASes (AS4134 and AS4812) that collectively account for 94% of the /24 blocks inferred to use arbitrary pooling. Outside of these ASes, most networks appear to follow the recommendations of RFC 4787.

Are networks that deploy CGN acquiring address space?

To understand whether inferred CGN networks also acquire more address space to satisfy their addressing needs, we collect routing tables [14] in July 2014 and September 2016, and compute for each network the size of the advertised IPv4 space. We conduct this analysis for 3,906 of the inferred CGN ASes that route IPv4 blocks in both months. For 44.87% of these ASes we do not detect any change in their address space; these networks seem to satisfy their addressing needs only by deploying CGN. 41% of the ASes increased their advertised address space; 32.27% of the ASes that increased their IPv4 space advertised IP blocks from the last /8s allocated to the four RIRs that entered the exhaustion phase. A closer analysis shows that for 49% of the inferred CGN ASes the IP address space change is by at most 256 /24 address blocks (i.e., one /16 block). We hypothesize that most CGN networks rely on their CGN deployment to satisfy their address space needs.

Given that the RIRs are now rationing the allocation of IPv4 addresses, our next question is whether inferred CGN networks are resorting to other means such as the IPv4 address transfer market. We use the list of reported transferred prefixes [36], [37], [38], to extract the organizations involved in the transfer (using the process described in [39]), and match them against the inferred CGN networks. We find 208 of the overall 3,050⁷ inferred networks participated in the IPv4 transfer market. Thus, most of the inferred networks seem to satisfy their address space needs without going to the IPv4 transfer market.

IPv6 adoption

Finally, to understand whether the inferred CGN networks deploy CGN as a “stop-gap” measure during the transition to IPv6, we first examined whether they also originated IPv6 prefixes. As of July-September 2016, 39.48% of the inferred ASes deployed IPv6. For the same period, however, 22% ASes from the IPv4 AS graph had deployed IPv6, i.e., the inferred CGN networks were *more likely* to have deployed IPv6 than the average AS. Second, we consider our measurement period (July-September 2014 to July-September 2016) and analyze

⁷Our analysis does not comprise networks inferred in 2016 as the study period of the IPv4 transfer markets stops in September 2015.

whether the inferred dual-stacked CGN ASes start advertising IPv6 prefixes prior to the period when we first detect them as deploying CGN; 46.55% of the dual-stacked CGN ASes appear to deploy CGN prior to advertising IPv6 prefixes. Our results indicate that CGN deployment does not have a negative impact on IPv6 adoption. However, given the rate at which the number of CGN deployments is increasing, this analysis deserves future revisiting.

VI. RELATED WORK

Many existing NAT analysis techniques rely on active measurements [40], [41], [33], [42], [43], [44], including methods for inferring CGN [21], [45], [7]. Researchers can learn detailed information about NAT deployments by initiating measurements from custom tests running on clients behind the NAT (e.g., the implementation of NAT444 [21], the topology of cascaded NATs [45], distance from the NAT device, or mapping timeouts [7]). Coverage can improve by conducting active measurements from outside the network (e.g., Casado *et al.* induced any client communicating with their custom web server to run active network characterization tests [44]; Richter *et al.* crawled BitTorrent’s DHT to identify cases of internal address leakage [7]).

Alternatively, applying passive NAT detection algorithms to datasets containing Internet-wide traffic requires little overhead. Existing techniques fingerprint machines based on specific packet fields such as TTL [23], operating system-specific TCP options [24], ephemeral ports [46], or the HTTP user-agent string [47]. These algorithms infer NAT when multiple fingerprints are observed with the same publicly routed IP address. While these algorithms are applicable to any packet trace, we are unaware of any extensions to the CGN setting.

VII. CONCLUSIONS AND DISCUSSION

We have presented two methods for detecting CGN deployment using existing passive measurement datasets collected from *outside* the target network. We have developed and validated our using IBR and M-Lab data, and expect that our methodology generalize to similar passively collected datasets.

To the best of our knowledge, our methods are the first to detect CGN with passive measurements collected outside the network, which comes with three advantages. First, our techniques do not require clients to install and execute any tests, and can thus achieve a larger coverage in terms of measured networks. Second, our techniques are simple and general. We can apply our IBR-based method to any traffic with unique client identifiers, and our M-Lab-based method to any traffic where we can model the likelihood of seeing an IP address as a function of user’s behavior. Finally, our methodology can be applied to data collected in the past, to analyze the evolution of CGN deployment.

In total, we have inferred that 4.1K ASes and 154K /24s are deploying CGN for the period from July 2014 to September 2016. During this period, we find a significant increase in the number of CGN networks. Our analysis shows that CGN deployment does not negatively impact the IPv6 adoption. Half

of the inferred networks seem to deploy CGN to prolong the lifespan of their IPv4 address space. Moreover, the remaining networks also seem to rely on CGN as they appear to acquire a small number of /24s blocks from the RIRs. Given that IPv4 depletion is still ongoing, each of these findings will need to be reassessed. Also, we find evidence of arbitrary pooling usage for CGN deployments. Our findings highlight the challenging nature of inferring CGN.

Acknowledgments

We thank Philipp Richter and Andra Lutu for sharing their CGN inference results. This work was supported by National Science Foundation grants CNS-1730661, CNS-1528148 and CNS-1111449, and by the Norwegian Research council grants number 209954 (Resilient Networks 2) and 240850 (DOMINOS). This material is based on research sponsored by the Department of Homeland Security (DHS) Science and Technology Directorate, Homeland Security Advanced Research Projects Agency, Cyber Security Division (DHS S&T/HSARPA/CSD), BAA 11-01-RIKA and Air Force Research Laboratory, Information Directorate under agreement number FA8750-12-2-0326. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of Department of Homeland Security, Air Force Research Laboratory or the U.S. Government. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

- [1] NRO, "Internet Number Resource Report (September 2017)," <https://www.nro.net/statistics/>.
- [2] IPv6 Statistics, <https://www.google.com/intl/en/ipv6/statistics.html>.
- [3] "RFC 1918: Address Allocation for Private Internets," 1996, <http://tools.ietf.org/html/rfc1918>.
- [4] "RFC 6269: Issues with IP Address Sharing," 2011, <http://tools.ietf.org/html/rfc6269>.
- [5] "RFC 7021: Assessing the impact of Carrier-Grade NAT on Network Applications," 2013, <http://tools.ietf.org/html/rfc7021>.
- [6] Europol, "Closing the online crime attribution gap: European law enforcement tackles Carrier-Grade NAT (CGN)," <https://www.europol.europa.eu/newsroom/news/closing-online-crime-attribution-gap-european-law-enforcement-tackles-carrier-grade-nat-cgn>, 2017.
- [7] 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. of ACM IMC*, 2016.
- [8] CAIDA/UCSD, "The UCSD Network Telescope," http://www.caida.org/projects/network_telescope/. Dates used: 2014-2016.
- [9] "Measurement Lab Project," <http://www.measurementlab.net/>.
- [10] A. Mandalari, A. Lutu, A. Dhamdhare, M. Bagnulo, and k. claffy, "Tracking the Big NAT across Europe and the U.S." in *CAIDA Technical Report*, Apr 2017.
- [11] A. Loewenstern and A. Norberg, "DHT Protocol," Jan. 2008. [Online]. Available: http://www.bittorrent.org/beps/bep_0005.html
- [12] J. Liang, N. Naoumov, and K. W. Ross, "The Index Poisoning Attack in P2P File Sharing Systems," in *Proc. of INFOCOM*, 2006.
- [13] K. Benson, "Leveraging Internet Background Radiation for Opportunistic Network Analysis," Ph.D. dissertation, UC San Diego, 2016.
- [14] David Meyer, "Route Views Project," <http://www.routeviews.org/>.
- [15] RIPE, "Routing Information Service (RIS)," <http://www.ripe.net/ris/>.
- [16] CAIDA, "The CAIDA UCSD AS Classification Dataset," august 2015. <https://www.caida.org/data/as-classification/>.
- [17] RIPE NCC, "Allocation and assignment resource file," <ftp://ftp.ripe.net/pub/stats>.
- [18] RIPE, "RIR Statistics Exchange Format," <ftp://ftp.ripe.net/ripe/stats/RIR-Statistics-Exchange-Format.txt>.
- [19] T. Cymru, "IP TO ASN MAPPING: WHO IS," <http://www.team-cymru.org/IP-ASN-mapping.html#whois>.
- [20] I. Livadariu, A. Elmokashfi, A. Dhamdhare, and K. Claffy, "Analysis of IPv6 deployment survey responses," *CAIDA Technical Report*, 2013, http://www.caida.org/~amogh/papers/ipv6survey2012_analysis.pdf.
- [21] A. Lutu, M. Bagnulo, A. Dhamdhare, and K. Claffy, "NAT Revelio: Detecting NAT444 in the ISP," in *Proc. of PAM*, 2016.
- [22] A. Norberg, "BitTorrent DHT security extension." [Online]. Available: http://www.libtorrent.org/dht_sec.html
- [23] S. M. Bellovin, "A technique for counting natted hosts," in *Proc. of 2nd ACM SIGCOMM Workshop on Internet measurement*, 2002.
- [24] R. Beverly, "A Robust Classifier for Passive TCP/IP Fingerprinting," in *Proc. of PAM*, 2004.
- [25] S. Grover, M. S. Park, S. Sundaresan, S. Burnett, H. Kim, and N. Feamster, "Peeking Behind the NAT: An Empirical Study of Home Networks," in *Proc. of ACM IMC*, 2013.
- [26] "Illegal file sharing on the wane in Sweden," Jun 2014, <http://www.thelocal.se/20140603/illegal-file-sharing-on-the-wane-in-sweden>.
- [27] R. Padmanabhan, A. Dhamdhare, E. Aben, K. Claffy, and N. Spring, "Reasons Dynamic Addresses Change," in *Proc. of ACM IMC*, 2016.
- [28] APNIC, "Visible ASNs: Customer Populations (Est.)," <http://stats.labs.apnic.net/aspop>.
- [29] T. Newton, "ISP Traffic Management: BT vs Virgin vs Sky vs TalkTalk vs EE," 2015, https://recombu.com/digital/article/isp-traffic-management-bt-sky-virgin-media-ee-talktalk_M11045.html.
- [30] C. Jacquenet, "An Introduction to Orange IPv6 Strategy," Jun 2013, <https://colloque-ipv6.greyc.fr/Slides/J2-01.pdf>.
- [31] P. Richter, G. Smaragdakis, D. Plonka, and A. Berger, "Beyond Counting: New Perspectives on the Active IPv4 Address Space," in *Proc. of ACM IMC*, 2016.
- [32] "List of mobile network operators," https://en.wikipedia.org/wiki/List_of_mobile_network_operators.
- [33] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson, "Netalyzr: Illuminating the Edge Network," in *Proc. of ACM IMC*, 2010.
- [34] C. Byrne, personal communication, April 2015.
- [35] F. Audet and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP," RFC 4787 (Best Current Practice), Internet Engineering Task Force, Jan. 2007, updated by RFCs 6888, 7857. [Online]. Available: <http://www.ietf.org/rfc/rfc4787.txt>
- [36] ARIN, "Inter-RIR and specified transfers of Internet number resources," <https://www.arin.net/knowledge/statistics/transfers.html>.
- [37] APNIC, "IPv4 Transfers," <ftp.apnic.net/public/transfers/apnic/>.
- [38] RIPE, "IPv4 Transfer Statistics," <https://www.ripe.net/manage-ips-and-asns/resource-transfers-and-mergers-transfer-statistics>.
- [39] I. Livadariu, A. Elmokashfi, A. Dhamdhare, and K. C. Claffy, "A first look at IPv4 Transfer Markets," in *Proc. ACM CoNEXT*, Dec. 2013.
- [40] B. Ford, P. Sriuresh, and D. Kegel, "Peer-to-peer communication across network address translators," in *Proc. of USENIX*, 2005.
- [41] S. Guha and P. Francis, "Characterization and Measurement of TCP Traversal Through NATs and Firewalls," in *Proc. of ACM IMC*, 2005.
- [42] A. Knutsen, R. Frederick, J. Mahdavi, Q. Li, and W. Yeh, "TCP Option for Transparent Middlebox Discovery," Internet Draft: draft-knutsen-tcpm-middlebox-discovery-04.txt, May 2010.
- [43] J. Rosenberg, R. Mahy, P. Matthews, and D. Wing, "RFC 5389: Session Traversal Utilities for NAT (STUN)," <http://tools.ietf.org/search/rfc5389>.
- [44] M. Casado and M. J. Freedman, "Peering Through the Shroud: The Effect of Edge Opacity on IP-Based Client Identification," in *Proc. of NSDI*, 2007.
- [45] A. Miller, F. Wohlfart, and G. Carle, "Analysis and topology-based traversal of cascaded large scale NATs," in *Proc. of HotMiddlebox*, 2013.
- [46] G. J. Armitage, "Inferring the Extent of Network Address Port Translation at Public/Private Internet Boundaries," CAIA, Swinburne University of Technology, Tech. Rep., 2002.
- [47] G. Maier, F. Schneider, and A. Feldmann, "NAT usage in Residential Broadband Networks," in *Proc. of PAM*, 2011.