

# Statistical Network Monitoring: Methodology and Application to Carrier-Grade NAT

Enrico Bocchi<sup>1,2\*</sup>, Ali Safari Khatouni<sup>1</sup>, Stefano Traverso<sup>1</sup>, Alessandro Finamore<sup>3</sup>,  
Maurizio Munafò<sup>1</sup>, Marco Mellia<sup>1</sup>, Dario Rossi<sup>2</sup>

1 - Politecnico di Torino, Italy

2 - Telecom ParisTech, France

3 - Telefonica Research, Spain

---

## Abstract

When considering to passively collect and then process network traffic traces, the need to analyze raw data at several Gbps and to extract higher level indexes from the stream of packets poses typical big data-like challenges. In this paper, we provide a methodology to extract, collect and process passive traffic traces. In particular, we design and implement analytics that, based on a filtering process and on the building of empirical distributions, enable the comparison between two generic collections, e.g., data gathered from two different vantage points, from different populations, or at different times. The ultimate goal is to highlight statistically significant differences.

After introducing the methodology, we apply it to assess the impact of Carrier-Grade NAT (CGN), a technology that Internet Service Providers (ISPs) deploy to limit the usage of expensive public IP addresses. Since CGN may introduce connectivity issues and performance degradation, we process a large dataset of passive measurements collected from an ISP using CGN for part of its customers. We first extract detailed per-flow information by processing packets from live links. Then, we derive higher level statistics that are significant for the end-users, e.g., TCP connection setup time, HTTP response time, or BitTorrent average download throughput. At last, we contrast figures of customers being offered public or private addresses, and look for statistically significant differences. Results show that CGN does not impair quality of service in the analyzed ISP deployment. In addition, we use the collected data to derive useful figures for the proper dimensioning of the CGN and the configuration of its parameters in order to avoid impairments on end-users' experience.

*Keywords:* IP networks; Computer Network Management; Network Address Translation; Big Data; Network Measurements; Performance

---

## 1. Introduction and Motivation

Traffic measurements have always played a central role to guide traffic management, to improve network and application design, and, in general, to understand the Internet. As result, several tools are available for both active and passive measurements. The former let the network administrator run on-demand specific tests at the expense of an increased network load. The latter permits a continuous monitoring by simply observing traffic, a challenging task given the several Gbps currently carried by backbone links. Collected measurements can be gathered to form a big data-like repository, and later leveraged to extract further knowledge, e.g., to contrast performance before and

after an upgrade or to monitor performance of applications being accessed from different parts of the network. Often, the network administrator needs “to compare” measurements collected at different times and places, or, generally speaking, different population subsets (e.g., fiber vs copper cable, fixed vs mobile, etc). In case of significant and unexpected differences in measurement results, countermeasures can be taken to identify, and possibly fix, the root cause of these differences.

In this paper, we present a methodology that accomplishes the above process, and apply it to a specific use-case, namely the deployment of Carrier-Grade NAT (CGN). Abstracting from the use-case for the time being, we leverage the passive monitoring technologies recently developed by the mPlane project, which offers a scalable architecture to deploy, collect and analyze Internet measurements. Referring to Fig. 1, we form a *measurement* layer by instrumenting several Points of Presence (PoPs) of an ISP with Tstat [1], a high-performing passive probe. By observing packets exchanged by end-users, each probe builds detailed logs for TCP and UDP flows in real time. Logs are then moved to a central *repository*, where several

---

\*Corresponding author

*Email addresses:* enrico.bocchi@polito.it (Enrico Bocchi<sup>1,2</sup>), ali.safari@polito.it (Ali Safari Khatouni<sup>1</sup>), stefano.traverso@polito.it (Stefano Traverso<sup>1</sup>), alessandro.finamore@telefonica.com (Alessandro Finamore<sup>3</sup>), maurizio.munafò@polito.it (Maurizio Munafò<sup>1</sup>), mellia@tcl.polito.it (Marco Mellia<sup>1</sup>), dario.rossi@enst.fr (Dario Rossi<sup>2</sup>)

gigabytes of raw data are collected every hour. To obtain valuable information from the logs, we design *analytics* to extract the subsets of data of interest and compute high level performance indexes. In particular, we focus on defining a methodology that allows us to detect whether statistically significant differences are present in measurements comparing different user populations or periods of time. While several metrics [2] allow one to compare two empirical distributions, ingenuity is needed to engineer a robust system capable of conveying simple yet telling differences in a compact way.

While the methodology is generic and would allow the comparison of generic populations (e.g., IPv4 vs IPv6, HTTP/1.1 vs HTTP/2, Android vs iPhone, etc.), in this work we apply it to quantify the impact of CGN the ISP has deployed. Network Address Translation (NAT) techniques have become a viable cheap solution to alleviate public IPv4 exhaustion. In a nutshell, a router implementing NAT functionality remaps the IP address space of a private network into one (or more) public IP address(es). CGN technologies extend this concept by masking a whole ISP network using NAT [3]. In this scenario, customers' home routers are assigned private IP addresses. When communicating with hosts in the public Internet, the CGN router temporarily maps the private, edge-facing IP address of the customer to one available public, Internet-facing IP address. This approach enables the ISP to mask part of its network as a large private network, significantly reducing the total amount of public IP addresses to use. Indeed, ISPs are more and more looking into these solutions as the price of a public IP address has now reached 10\$/year per IP.<sup>1</sup>

However, NAT and CGN break the end-to-end paradigm of the Internet communication model. On the one hand, NAT-ed hosts cannot be directly addressed from the Internet, which is unsuitable for applications that require reachability from the public Internet. On the other hand, the NAT mapping operations may add delay to packets or cause loss. Despite a large body of work focusing on NAT technologies and NAT traversal techniques (which we overview in Sec. 7), little effort has been devoted to assess CGN impact on actual user experience.

We study both aspects in this work, whose main contributions can be summarized as follows:

- We define analytics to assess statistical differences between measurement aggregates, which allow one to either highlight spatial (e.g., different populations) or temporal (e.g., same population at different times) discrepancies (Sec. 2).
- We particularize these analytics to the case of a real CGN deployment for populations of users with either private or public IP addresses (Sec. 3). For the purpose, we consider key performance indicators that are

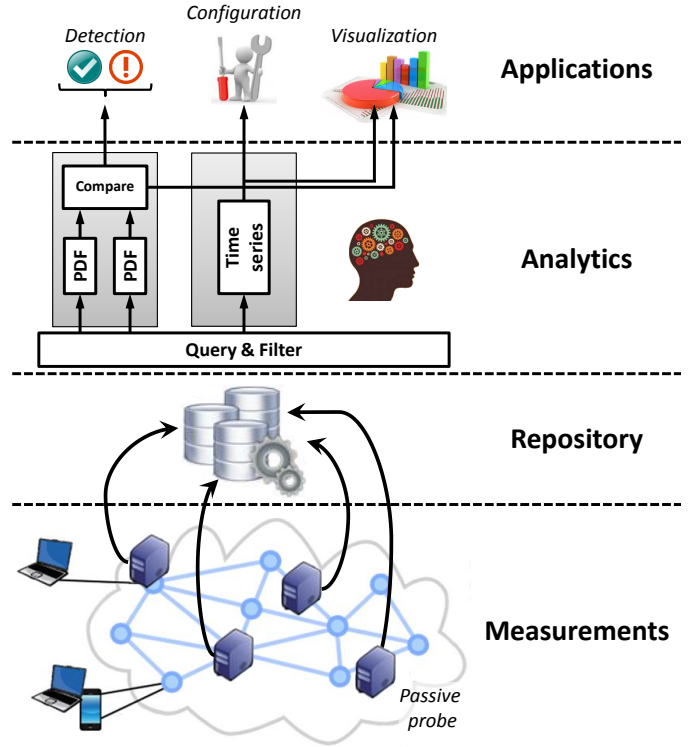


Figure 1: Illustration of the measurement framework.

relevant for user quality of experience, such as connection setup time for web traffic, average transfer rate for BitTorrent traffic, etc. (Sec. 4).

- We apply these analytics to assess the impact of CGN from an actual ISP deployment, processing the data collected by monitoring about 17,000 residential customers for one month. Results show that no statistically significant difference can be observed between the two populations for the considered performance indicators (Sec. 5).
- We leverage actual usage patterns to provide statistical figures that allow the network operator to properly design and dimension the CGN deployment while avoiding impairments to the end-users (Sec. 6).

## 2. Analytics to highlight and quantify statistical differences

In this work, we focus on analytics to compare measurements referring to different datasets, i.e., that help us to quickly pinpoint eventual performance differences between different populations of users. In this section, we provide an overall view and the necessary elements to understand the potential applications of the framework, deferring detailed statistical considerations and sensitivity analysis to Appendix A.

<sup>1</sup><http://www.ipaddressnews.com/2014/04/07/343>

### 2.1. Empirical distribution estimation

Fig. 1 illustrates our workflow. Several layers are visible. From the bottom, the *Measurements* layer consists of passive Tstat probes installed in an operational network. Probes are responsible for extracting traffic summaries based on the continuous observation of packets generated by end-users: at flow completion, Tstat logs more than 100 metrics whose schema is strictly defined. As such, the output is a tabular database, where each row represents a flow and each column is a specific key performance indicator (KPI). More details about KPI definitions are given in Sec. 4. Data is temporarily stored at the probe premises, and then moved asynchronously to the central *Repository* based on Hadoop and Hive, which is located in the Big Data Laboratory of Politecnico di Torino.

Once measurements are stored in the repository, a “Query & Filter” engine allows us to easily extract the measurement samples of interest (e.g., select Round Trip Time measurements for TCP connections where application layer protocol is HTTP, server name matches *\*.google.com*, client IP address is private, and date is any day of October 2014). Thanks to the SQL-like interface offered by Hive, the “Query & Filter” module allows us to easily access a large dataset in a simple, intuitive and scalable way, and quickly output sets of samples.

We next estimate the empirical Probability Density Function (PDF) and Cumulative Distribution Function (CDF) using a simple module that, given the size of bins and support range, computes the frequency of samples falling in each bin, i.e., the probability  $p_i$  that the sample takes values in the  $i$ -th bin. Given the amount of data to process is typically limited (few millions of samples), and the lack of iterative processing, we opted to implement this module using Python.

### 2.2. Comparison and quantization functions

While extracting data involves scalability issues successfully solved by the Hadoop approach, the comparison of PDFs no longer needs big-data processing, but poses significant challenges the analyst has to cope with. For instance, one should ensure to have an adequate population of samples, carefully choose the binning size, consider the nature of the performance indicators on which PDFs are built, etc. We defer the formal discussion of such aspects to the Appendix A. By now, we limit the discussion to the description of a method to compare PDFs.

To this goal, we define a statistical method to compactly quantify the difference between two distributions. The output can be “soft”, i.e., a real value in a continuous range, or “hard”, i.e., a categorical output from a (small) set of possible values. For instance, in the first case, comparison tells the analyst how different the statistics are, while in the second case it just tells if they differ or not.

In formal terms, the comparison function has the form  $F(p, q) : (\mathbb{R}^2, \mathbb{R}^2) \rightarrow \mathbb{R}$ , while the quantization function can be defined as  $Q(F(p, q)) : \mathbb{R} \rightarrow \mathbb{N}$ , where  $p = p(x)$  and  $q = q(x)$  are two empirical distributions under analysis.

Without loss of generality, in this paper we define a simple quantization function that considers three possible levels, correlating with a no difference state (0), a definitively different state (2), and a possibly different state (1) requiring further, though likely not urgent, investigation. Such quantization function can be written as:

$$Q(F(p, q)) = \begin{cases} 0 & \text{if } F(p, q) < Q^- \\ 1 & \text{if } Q^- \leq F(p, q) < Q^+ \\ 2 & \text{if } F(p, q) \geq Q^+ \end{cases}$$

with states discriminated by the lower  $Q^-$  and upper  $Q^+$  thresholds. Selecting these thresholds requires careful attention since (i) there is a dependency between the thresholds  $Q^-, Q^+$  values and the comparison function  $F(p, q)$ ; (ii) the value of  $F(p, q)$  can be noisy when distributions  $p$  and  $q$  are computed over small population samples; (iii) the value of  $F(p, q)$  can be affected by class imbalance when population samples of distributions  $p$  and  $q$  are of different orders of magnitude; and (iv) the value of  $F(p, q)$  can be affected by the measurement process (e.g., binning strategy, number of bins, etc.).

### 2.3. Jensen-Shannon divergence

We make illustrative examples of one statistical distance measure (SDM) on controlled distributions  $p, q$ . As representative SDM in this class, we take the Jensen-Shannon divergence ( $JS_{div}$ ), which is defined as:

$$JS_{div} = \sum_i \left\{ \frac{1}{2} p_i \ln \left( \frac{p_i}{\frac{1}{2} p_i + \frac{1}{2} q_i} \right) + \frac{1}{2} q_i \ln \left( \frac{q_i}{\frac{1}{2} q_i + \frac{1}{2} p_i} \right) \right\}$$

where  $p_i$  and  $q_i$  are the empirical probabilities of samples taking values in the  $i$ -th bin.  $JS_{div}$  is a popular statistical measure based on the Kullback-Leibler divergence, over which it brings some notable improvements, adding symmetry, i.e.,  $JS_{div}(p, q) = JS_{div}(q, p)$ , and bounded support, i.e.,  $JS_{div} \in [0, \ln(2)]$ .  $JS_{div}$  is equal to 0 if  $p = q$ , while it saturates to  $\ln(2)$  for two completely disjoint distributions.

We focus on understanding how the  $JS_{div}$  varies when comparing two synthetic PDFs, with the aim at defining  $Q^-, Q^+$  thresholds to separate the areas into three states in a generic case. To this extent, we consider (i) negative exponential distributions with different mean; and (ii) Gaussian distributions with different mean and/or different standard deviation. These PDFs are representative of diverse properties that may appear in network data; e.g., packet inter-arrival times and packet size in VoIP calls can be approximated with Gaussian distributions [4]; requests generated by user-activities are well approximated with Poisson processes and have as such negative exponential inter-arrival times [5, 6]. For the sake of the example, we now quantify differences between controlled PDFs that are representative of fictional, yet plausible, processes.

Such analysis is useful both to visually tie the  $JS_{div}$  behavior to some well-known distributions, and to identify

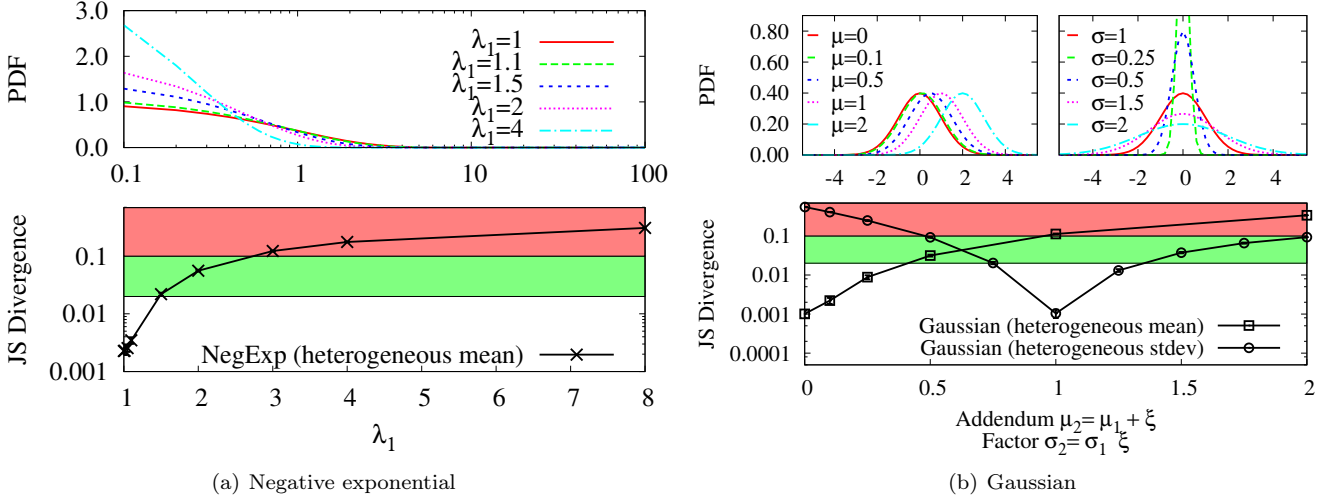


Figure 2: Illustrative examples of Jensen-Shannon divergence computed on: (a) negative exponential distributions with heterogeneous mean rates  $\lambda_1$  versus reference mean  $\lambda = 1$ ; (b) Gaussian distributions with heterogeneous mean  $N(\mu, 1)$  or standard deviation  $N(0, \sigma)$  versus the reference distribution  $N(0, 1)$ .

quantization thresholds that discriminate among significant (2), noticeable (1) and negligible (0) differences between PDFs. The purpose is to illustrate the methodology we followed in setting the quantization function  $Q(f(p, q))$ . A more in-depth analysis is given in Appendix A.

We consider a dataset where samples are extracted from synthetic  $p, q$ . We take care of avoiding any bias in the  $JS_{div}$  scores that can be tied to population size, imbalance or binning strategy (see Appendix A). Let consider first the comparison of  $p, q$  which are both negative exponential distributions:

$$NegExp(x, \lambda) = \lambda e^{-\lambda x} \text{ for } x \geq 0; 0 \text{ otherwise}$$

We consider  $p = NegExp(x, \lambda_0)$  as a reference, and choose  $\lambda_0=1$ , while  $q = NegExp(x, \lambda_1)$  is instead shaped according to a distribution of parameter  $\lambda_1$ , with  $\lambda_1 \in [1, 8]$  in our experiments. From both distributions, we extract  $10^6$  samples, obtain the empirical PDFs using 1000 bins in a  $[0, 100)$  support. This leads to bins of size  $\Delta b = 100/1000 = 0.1$ . For each bin  $i$ , we estimate  $p_i$  and  $q_i$  as the ratio between the number of samples falling in the  $i$ -th bin, i.e.,  $[i\Delta b, (i+1)\Delta b)$ , and the total number of samples.

Negative exponential PDFs  $p$  and  $q$  are depicted in the top portion of Fig. 2(a), whereas the bottom plot reports the  $JS_{div}$  versus  $\lambda_1$ . Without loss of generality, we select thresholds  $Q^- = 2/100$  and  $Q^+ = 1/10$ , so that a clearly visible changes in the distribution space (top) are visible in the  $JS_{div}$  space (bottom) as well. Intuitively, when  $JS_{div} \in [Q^+, \ln(2)]$ , the difference between the two PDFs is significant (red area). When  $JS_{div} \in [Q^-, Q^+)$  the difference is noticeable (green area), and negligible if  $JS_{div} \in [0, Q^-)$  (white area).

We repeat the experiment this time considering Gaus-

sian distributions, i.e.,

$$N(x, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

As before, we generate a reference sample  $p$  corresponding to  $(\mu, \sigma) = (0, 1)$ , and samples  $q$  with different  $(\mu, \sigma)$  parameters. The upper-left plot of Fig. 2(b) shows PDFs of  $q$  with parameters  $\mu \in \{0, 0.1, 0.5, 1, 2\}$  and  $\sigma=1$ , while the upper-right plot shows PDFs when  $\mu=0$  and  $\sigma \in \{0.25, 0.5, 1, 1.5, 2\}$ . Fig. 2(b) lower plot reports the  $JS_{div}$  values when comparing the above-mentioned distributions against  $p$ .

The previous threshold selection proves to be effective also in the case of Gaussian distributions: visible differences in the upper plots of Fig. 2(b) appear to be separated by the  $Q^- = 2/100$  and  $Q^+ = 1/10$  thresholds.

In real cases, the sensitivity of a domain expert can be used to set thresholds. In general, any threshold choice results arbitrary, which applies to any SDM of choice (and possibly being even more complicated for those SDMs with infinite support). We point out that the framework we propose is not limited to the use of  $JS_{div}$  measure. Rather, in Appendix A we consider a large set of SDMs and identify a set of those that are equivalent in this respect, as they share a number of desirable properties (e.g., symmetry and boundness). While we discuss these issues further in Appendix A, the information provided in this section allows us to understand the application of the general framework to the CGN use-case we focus on in the reminder of this paper.

### 3. Monitoring Scenario and Dataset

To characterize the implications of CGN, we rely on passive measurements obtained by instrumenting a non-

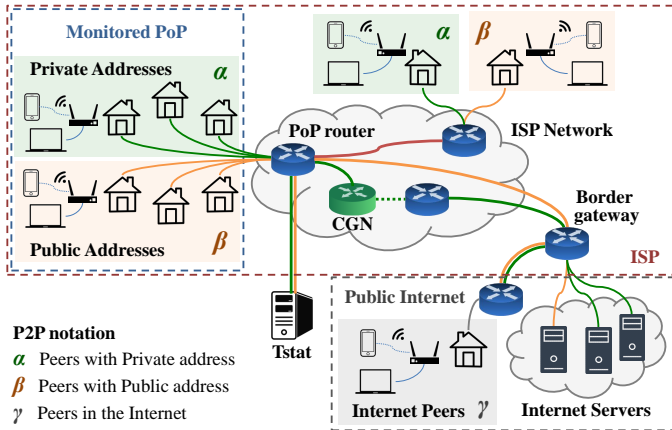


Figure 3: The monitoring scenario we consider in this study.

itoring probe in the operational network of an European country-wide ISP. Each customer device accesses the Internet via an ADSL home router. The ISP assigns either a public or private IP address to each home router according to the customer’s subscription type. Traffic directed to the Internet and coming from home routers with a public IP address (public home routers) is routed directly to the final destination, while traffic from home routers with a private IP address (private home routers) has to cross the CGN device first.

The CGN used by the monitored ISP is based on the NAT444 standard [7], which relies on *sessions* to translate the private, edge-facing IP address of a home router into a public, Internet-facing one. When the CGN receives the first packet from a private home router, it starts a new session, temporarily mapping the private address to the first available public address in a pool. It then converts the address of all subsequent packets according to the same mapping.<sup>2</sup> After a given inactivity time during which no packets are observed, the session expires and the public address is put back in the pool of available addresses.

### 3.1. Monitoring Setup

Fig. 3 depicts the monitoring scenario in which we operate. Three regions are identified: (i) a monitored PoP; (ii) the ISP network; and (iii) the public Internet. We deploy a passive probe inside the PoP and we instrument it to process the packets flowing through the PoP router. This router forwards the traffic coming both from private and public home routers, thus we are in the condition to monitor the traffic produced by the two populations of users.

In the case of web traffic, private and public home routers have to reach servers located in the public Internet. Therefore, the traffic they produce has to cross the

<sup>2</sup>The amount of public addresses available at the NAT is smaller than the number of customers provided with a private IP address. Consequently, the pool size of public addresses must be carefully set to minimize allocation costs, while guaranteeing satisfactory connectivity. See Sec. 6 for a thorough discussion.

	Private	Public
# of TCP flows	990M	767M
# of UDP flows <sup>4</sup>	2,676M	1,941M
# of failed-TCP flows	301M	347M
Traffic Volume	168TB	105TB

Table 1: Statistics for home routers with private and public IP addresses.

PoP router, the CGN if the traffic is generated by private home routers, and the ISP border gateway. In the case of Peer to Peer (P2P) traffic, peers can be located both inside the ISP network ( $\alpha, \beta$  in Fig. 3) and in the public Internet ( $\gamma$ ). The ISP assures end-to-end connectivity among customers within its network, independently on the type of IP address assigned to each home router.

### 3.2. Dataset description

We leverage a dataset collected during the month of October 2014. It consists of TCP, UDP and failed-TCP<sup>3</sup> logs carrying 1,757M, 4,617M and 648M records respectively, for a total of more than 273TB of network traffic. We split each of our logs in two subsets according to the IP address type of the customer’s home router.

Tab. 1 provides statistics about the traces we consider, separately for private and public home routers. In total, we find more than 17,000 active home routers. Out of these 60% (40%) are assigned a private (public) IP address.<sup>5</sup>

Unless stated otherwise, in the remainder of the paper we present the results obtained by focusing on October 2014. Analysis conducted on other periods show very similar results. Overall, the dataset at our availability is large enough to avoid biases due to population size or imbalance, and performance indicators (see Sec. 4) are not affected by quantization artifacts.

### 3.3. P2P description

For this work we consider only BitTorrent traffic as it is the most used P2P application in our dataset. To isolate peers that actively use the BitTorrent, we consider only those showing at least one flow carrying 1 MB in our dataset, and we call them *active*.

According to the type of IP address at the home router and to the location of the peer in the network ( $\alpha, \beta, \gamma$  in Fig. 3), peers can experience different reachability conditions. To assess this, we check if peers are able to receive incoming connections from their counterparts. We define a peer as *reachable* if its home router is properly configured

<sup>3</sup>Tstat labels as failed TCP connections for which the Three-Way Handshake is not completed (e.g., when the sole SYN message is observed).

<sup>4</sup>In Tstat, a UDP flow starts at the first seen packet exchanged between two endpoints and ends 65 seconds after the last seen packet.

<sup>5</sup>The home router IP address can be considered as an identifier of the household. It may hide several devices connected to the Internet.

		Private	Public
TCP	Reachable	631 (35%)	496 (50%)
	Unreachable	1188 (65%)	499 (50%)
	Total	1819	995
UDP	Reachable	891 (77%)	591 (95%)
	Unreachable	262 (23%)	33 (5%)
	Total	1153	624

Table 2: Number of active peers classified according to their reachability condition.

and ports are forwarded to the P2P application. In case the home router is misconfigured and the P2P application is non reachable, we define a peer as *unreachable*.

Four classes of peers emerge:

- *Private-Unreachable*: any peer behind private home routers that does not receive incoming connections;
- *Private-Reachable*: any peer behind private home routers that receives incoming connection from other peers in the ISP network ( $\alpha$  and  $\beta$  in Fig. 3). Reachability from peers in the Internet is not guaranteed as the CGN limits incoming connections;
- *Public-Unreachable*: any peer behind public home routers that does not receive incoming connections;
- *Public-Reachable*: any peer behind public home routers that receives incoming connections from both the ISP and the Internet. This is the only class of peers that is reachable by everyone.

Tab. 2 characterizes the number of active peers over TCP and UDP according to their reachability condition. Notice that only the 35% (50%) of peers with a private (public) IP have their home router properly configured and are thus reachable over TCP. This potentially owes to the scarce success of NAT traversal techniques for TCP flows.

In the case of UDP, instead, the ratio of reachable peers is higher both for private (77%) and public (95%) home routers. This is due to the fact that NAT traversal techniques like STUN [8, 9] are more effective over UDP, and enable peers to receive incoming connections.

#### 4. Key performance indicators

In this section we define the performance indicators we are interested in when considering the CGN impact. Among the many measurements provided by Tstat, we consider for each traffic flow: (i) The TCP Round-Trip-Time (RTT) between the probe and server; (ii) the Time-To-Live (TTL) seen at the probe of packets sent by the server; (iii) the total per-flow amount of bytes sent and received by the client; (iv) the application layer protocol (e.g., HTTP, HTTPS, BitTorrent); and (v) the timestamps of packets that are instrumental to obtain further

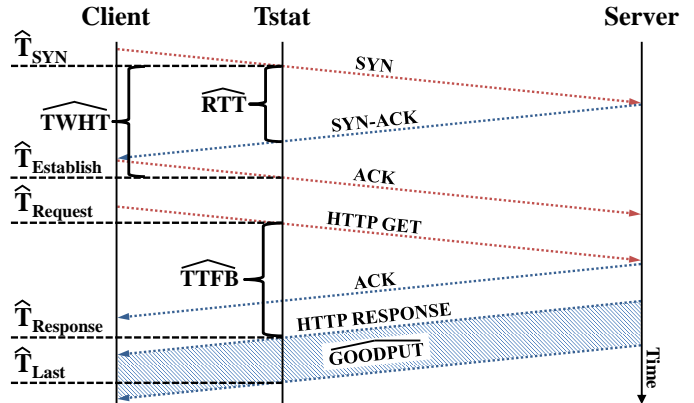


Figure 4: An example of HTTP transaction with metrics considered for our analysis.

indices.<sup>6</sup> Finally, we use the Fully Qualified Domain Name (FQDN) [10] of the server to split traffic according to the service generating it.

Fig. 4 shows Tstat observing a HTTP transaction. In a nutshell, Tstat correlates TCP data segments and acknowledgments, and records timestamps of significant packets. For instance, by correlating times of a data segment with the corresponding acknowledgment, it computes a sample of the RTT. Average and standard deviation of RTT is then obtained by considering all samples in the TCP flow. We defer the interested reader to [1, 11] to obtain the detailed description of how performance indicators are extracted from packet traces.

Since CGN may impact both network, transport and application layer performance, we detail how we combine basic metrics provided by Tstat to build higher level measurements that we use to contrast the impact of CGN at different layers.

##### 4.1. Network layer metrics

###### Hop count from Server to PoP (#Hops).

The minimum number of hops being traversed by packets transmitted from the server to the client. The operating system of the server sets the initial value of the TTL, with power of 2 values being the typical choice.<sup>7</sup> Each router along the path then decreases the TTL. The values observed at the probe is thus an indication of the number of hops on the path from the server to the probe located in the PoP. In more details, given a flow, we take the maximum server-to-client TTL observed by Tstat. We then choose  $x$  as the exponent minimizing  $\#Hops = 2^x - TTL$ ,  $\#Hops > 0$ . The resulting  $\#Hops$  is the minimum number of hops

<sup>6</sup>Notice that the probe measures the timestamps at a vantage point close to the customers. Therefore, for some metric  $X$  we can only gauge its estimated measure  $\hat{X}$ .

<sup>7</sup><http://subinsb.com/default-device-ttl-values>

that packets in the considered flow have traversed before reaching the probe. In our scenario we expect packets received by private home routers to traverse a possibly larger number of hops due to the presence of the CGN (one or more hops).

#### PoP to Server Round Trip Time ( $RTT$ ).

The average RTT Tstat measures in a flow ( $\widehat{RTT}$ ) on packets transmitted from the client to the server. Referring to Fig. 4, we consider only the RTT from the probe in the PoP to the server and backward, thus including only the backbone part of the path and ignoring the access portion. The RTT computation is complicated by the presence of packet retransmission, delayed acknowledgments, etc. We refer the reader to [11] for details. We expect packets transmitted by private home routers to experience a higher latency because of the CGN packet processing.

#### 4.2. Transport layer metrics

##### TCP Three-Way Handshake Time ( $TWHT$ ).

The amount of time measured by Tstat ( $\widehat{TWHT}$ ) required to successfully establish a TCP connection using the Three-Way Handshake (TWH). Referring to the upper part of Fig. 4, let  $\hat{T}_{SYN}$  be the timestamp of the SYN packet sent by the client to start the connection establishment procedure, and let  $\hat{T}_{Establish}$  be the timestamp of the packet carrying the ACK message ending the TWH. We define the  $\widehat{TWHT}$  as

$$\widehat{TWHT} = \hat{T}_{Establish} - \hat{T}_{SYN}$$

In our scenario we expect the  $\widehat{TWHT}$  to be higher for private home routers due to the time needed by the CGN to allocate the resources for the new communication session.

For the sake of completeness, we also consider some advanced specific TCP metrics that are directly computed by Tstat [11]: (i) The number of SYN messages observed during connection setup,  $SYN$ ; (ii) the number of out of sequence segments,  $OoS$ ; (iii) the number of duplicated segments  $Dup$ . These are measurements that we expect to be altered in case of connectivity issues introduced by the CGN. A large value of  $SYN$ , for instance, indicates that the client experienced difficulties in establishing the connection due to, e.g., exhaustion of NAT resources.

#### 4.3. Application layer metrics

##### Time to first byte ( $TTFB$ ).

Referring to Fig. 4, the amount of time that elapses between the first segment containing the HTTP request sent by the client ( $\hat{T}_{Request}$ ) to the first segment with payload sent by the server ( $\hat{T}_{Response}$ ). We define the  $\widehat{TTFB}$  as

$$\widehat{TTFB} = \hat{T}_{Response} - \hat{T}_{Request}$$

In HTTP flows, it represents a measure of the time span between the application request issued by the client and the consequent response by the server. Also in this case, we expect the CGN to eventually delay the response time due to NAT operations.

##### Per-connection Goodput ( $G$ ).

The average rate at which the server delivers information to the client. This is the paramount performance index for download services. Let  $\hat{T}_{Response}$  and  $\hat{T}_{Last}$  (see Fig. 4) be the timestamps of the first and the last data packet sent by the server, and let  $D_{down}$  be the size of the application payload sent by the server. We define the download goodput as

$$\hat{G}_{down} = \frac{D_{down}}{\hat{T}_{Last} - \hat{T}_{Response}}$$

It is similarly possible to evaluate the goodput in the upload direction by considering the amount of bytes sent by the client to the server ( $D_{up}$ ) and referring to the timestamps relative to the client traffic. To avoid the bias of short-lived flows, we evaluate the download goodput only on flows for which  $D_{down} \geq 1$  MB, and the upload goodput for flows where  $D_{up} \geq 500$  kB.

##### Average Throughput ( $Thru$ ).

$\hat{G}_{down}$  is a representative measure of performance when the download of a content is done using a single flow, e.g., when downloading some software from the web. For P2P applications however, the speed at which a peer downloads a content is more complicated to compute since multiple parallel connections are used by the application. For instance, BitTorrent typically downloads content from 5 to 10 peers at the same time, using both TCP or UDP at the transport layer. To measure the overall performance of a peer, we compute the average download (upload) throughput  $Thru$  considering all data received (sent) by a client in a time interval of duration  $\Delta T = 10$  mins. Only flows classified as BitTorrent are considered. Formally, given time interval  $i$ , we consider all TCP and UDP flows that Tstat classifies as BitTorrent, and terminated in the time interval,  $F(i) = \{f | \hat{T}_{Last}(f) \in i\Delta T\}$ . Let  $D_{tot}(i) = \sum_{k \in F(i)} D(k)$  the total amount of data those flows carried. Then

$$\widehat{Thru}(i) = \frac{D_{tot}(i)}{\Delta T}$$

## 5. Impact of CGN on users' traffic

The goal of this section is to check whether one of the two classes of customers experience worse performance than the other due to the type of IP address they have at their home router. To do so, we split flows into two subsets, based on if they are coming from private or public home routers. For each subset, we then compute the

Metric	Web Traffic				P2P	
	All Flows	www.google.com	TOP-50 Google	phobos.apple.com	Reachable	Unreachable
$\widehat{\#Hops}$	<b>0.223</b>	<b>0.666</b>	<b>0.682</b>	<b>0.689</b>	<b>0.184</b>	<b>0.162</b>
$\widehat{RTT}$	0.001	0.006	0.007	0.007	0.055	0.002
$\widehat{TWHT}$	0.002	0.010	0.011	0.016	0.029	0.008
$\#SYN$	<0.001	<0.001	<0.001	<0.001	<0.001	<0.001
$OoS$	<0.001	–	–	–	–	<0.001
$Dup$	0.001	0.001	0.001	<0.001	<0.001	<0.001
$\widehat{TTFB}$	0.002	0.006	0.008	0.006	0.031	0.005

Table 3: Jensen-Shannon divergence for considered metrics and different Internet services.

empirical PDF for each metric, and we finally evaluate the  $JS_{div}$  among the two PDFs. We remind the two thresholds identified in Sec. 2.3 for the  $JS_{div}$ :  $Q^- = 0.02$  and  $Q^+ = 0.1$ .

### 5.1. Impact on network and TCP layer metrics

We start our analysis by gauging the impact of CGN on network- and transport- layer metrics described in Sec. 4.1 and Sec. 4.2, respectively. We report the collected results in Tab. 3. We focus on the Web traffic first, as reported on the left-hand side of the table. We show the result of experiments considering flows directed to (i) any remote server (“all flows”); (ii) “www.google.com” servers (i.e., *Google Search*); (iii) TOP-50 most used IP addresses of Google servers (“TOP-50 Google”); and (iv) “phobos.apple.com” servers providing *iTunes Store* contents.<sup>8</sup>

As shown, the only metric that consistently overcomes the alarm threshold  $Q^+$  for both web traffic and P2P is the number of hops,  $\widehat{\#Hops}$ , which is highlighted in bold in the table. To validate the above finding, we directly compare the distributions of  $\widehat{\#Hops}$  in Fig. 5. For the ease of visualization, we report the CDF some services, as results are similar for any service. A clear offset between the  $\widehat{\#Hops}$  of private and public home routers appears, showing that private ones have to traverse more hops to reach the Internet. Such offset is present for all services. We verified this outcome with the ISP network administrators, who confirmed that the difference is due to some extra routers that packets sent/received by private home routers have to go through to reach the CGN. However, such routers are well dimensioned and not congested, with little to no implication on the performance, as testified by other metrics in Tab. 3.

In summary, the  $JS_{div}$  values for web traffic are below  $Q^-$ , meaning that the CGN configuration of our scenario does not induce any significant bias.

Let us focus on the time needed to establish a TCP connection  $\widehat{TWHT}$ . This is a typical metric one would expect to be affected by additional delay introduced by

<sup>8</sup>We focus on this selection of services as they appear to be popular on the monitored network, and the amount of TCP flows for each of them satisfies the requirements for a proper use of the  $JS_{div}$ .




Figure 5: CDFs of the hop count ( $\widehat{\#Hops}$ ) from the server to the client for private and public home routers against different web services. Clear differences are visible.




Figure 6: CDFs of time needed to complete the Three-Way Handshake ( $\widehat{TWHT}$ ) for private and public home routers against different web services. No significant differences are visible.

the CGN when private home routers try to establish new connections. Indeed, the CGN may require some time to initiate the session and translate addresses. Also in this case  $JS_{div}$  is very small for Web traffic. Fig. 6 shows details distributions for private and public home routers with respect to the same Internet services. Differences are practically negligible.


### 5.2. Impact on application layer metrics

We complement the above findings by applying the  $JS_{div}$  on the indices presented in Sec. 4.3. The last row of Tab. 3 shows the  $JS_{div}$  of the Time to First Byte,  $\widehat{TTFB}$ . Results for web traffic indicate that this metric is again not affected by the presence of the CGN, and that users




	Service	FQDN	$JS_{div}$
Download	All	*	0.001
	Facebook Video	fbcdn-video-*.akamaihd.net	0.004
	Tumblr	media.tumblr.com	0.021
	Phobos	phobos.apple.com	0.022
Upload	All	*	0.004
	Amazon S3	eu-irl-*.s3.amazonaws.com	0.007
	Whatsapp	mm*.whatsapp.net	0.033
	Dropbox	dl-*.dropbox.com	0.046

Table 4: Jensen-Shannon divergence for goodput distributions in download and upload directions.



(a) Download CDFs for flows carrying  $\geq 1$  MB.



(b) Upload CDFs for flows carrying  $\geq 500$  kB.

Figure 7: Normalized goodput CDFs for flows carrying Web traffic.

accessing the Internet from private or public home routers face similar delays.

Next, we perform the same analysis for the web traffic goodput  $\hat{G}$ . We consider several popular services that exchange a large amount of data, and for which  $\hat{G}_{down}$  is thus relevant, i.e., Facebook Video, Tumblr and Phobos. For  $\hat{G}_{up}$  we selected Amazon S3, Whatsapp and Dropbox. We report the results in Tab. 5.2, and draw the CDFs in Fig. 7.

Observe that the  $JS_{div}$  never overcomes the  $Q^+$  threshold, meaning that the CGN does not significantly harm the download/upload speed of private home routers. However, the  $JS_{div}$  values for Whatsapp and Dropbox in the upload direction, and for Tumblr and Phobos in the download direction, are higher than the  $Q^-$  threshold. Fig. 7(a) details the distribution of  $\hat{G}_{down}$  (we omit Facebook Video to ease the visualization).<sup>9</sup> The curves referring to private

<sup>9</sup>We normalize the measured throughput to not show the actual bandwidth provided by the monitored ISP.

	Reachable	Unreachable
Download	0.005	0.004
Upload	0.004	0.003

Table 5: Jensen-Shannon divergence for average throughput distributions in download and upload directions.

and public home routers show indeed very similar trends, justifying small  $JS_{div}$  values, as confirmed by Tab. 5.2. Fig. 7(b) reports results for  $\hat{G}_{up}$ . Also in this case the curves show very similar CDFs with the only exception of Whatsapp. In this latter case, the difference between the two distributions is confirmed by the  $JS_{div} = 0.033$ .

Interestingly, a relatively large amount of flows (13.98%) in Fig. 7(b) show almost zero throughput. By double-checking, we realize that those are long-lived flows with a duration higher than 10 min, and showing a number of uploaded bytes that slightly exceeds the 500 kB threshold. For some services, indeed, clients establish a single TCP connection with the remote server and keep sending tiny portions of data intermittently, de facto zeroing the upload throughput.

At last, we focus on the  $JS_{div}$  for BitTorrent traffic, distinguishing between (i) reachable peers (i.e., those who have port forwarding properly configured at their home router); and (ii) unreachable peers, as defined in Sec. 3.

We consider the average download throughput. Several works in the literature (see Sec. 7 for a detailed discussion) show how multiple factors impact P2P performance: content popularity, content availability, type of peers (e.g., seeders, leechers) involved in the transfer, peers cooperation techniques (e.g., tit-for-tat), etc. Let us assume population of peers in our dataset is not biased by any of these factors. To verify such assumption, we perform 10-fold validation: we randomly split the P2P dataset into 10 sub-groups, and we compare the  $\widehat{Thru}$  PDF of each group against all other groups. In none of these experiments we detected statistically significant differences, i.e.,  $JS_{div} > Q^-$ , validating that our population of peers shows homogeneous performance.

This allows us to safely contrast the  $\widehat{Thru}$  for *reachable* and *unreachable* peers. Tab. 5 shows the  $JS_{div}$  values. For both classes of peers, the computed values are one-order of magnitude below the  $Q^-$  threshold, proving that the CGN does not affect throughput for BitTorrent. In the nutshell, peers obtain the same performance, no matter if they have private or public addresses.

### 5.3. Discussion on P2P implications

Let us get back to the right-hand side of Tab. 3 which shows the  $JS_{div}$  values for BitTorrent traffic.  $JS_{div}$  values are below  $Q^-$  for *unreachable* peers, but they fall in the noticeable range ( $[Q^-, Q^+]$ ) for *reachable* peers. Let us consider first the results for  $\widehat{RTT}$ . Fig. 8 shows the




Figure 8: Round Trip Time for P2P traffic according to peers' reachability conditions.

$\widehat{RTT}$  of BitTorrent connections.<sup>10</sup> It is evident that the RTT for reachable peers with a private IP address appears to be lower than the RTT measured for all the other peers. To better understand this aspect, we characterize the reachability condition of peers inside the ISP, and of the counterparts with which they establish a connection.

To exemplify reachability conditions, refer to Fig. 3. It depicts peers inside the ISP ( $\alpha, \beta$ ) and peers outside ( $\gamma$ ). Consider a *reachable* peer with private IP ( $\alpha$ ): it can receive incoming connections from all the peers inside the ISP network ( $\beta$ ), but not from peers in the Internet ( $\gamma$ ). On the other hand, a *reachable* peer with public IP ( $\beta$ ) can receive incoming connections from both the peers inside the ISP ( $\alpha$ ) and from peers in the Internet ( $\gamma$ ).

Such reachability conditions have implications also on the distribution of contacted peers, as shown by the Table 6 which reports the percentage of contacted peers by *reachable* and *unreachable* peers. Reachable peers with a private IP establish more connections with other peers inside the ISP (52%) than in the Internet (42%). All other classes of peers are more prone to connect to peers in the Internet (>90%). This is due to peers in the ISP that contacted private but reachable peers, like  $\alpha$ .

As a consequence, private reachable peers experience a lower RTT since, they contact peers inside the ISP network, which are closer in space and exhibit a lower RTT. This also reflects in the other metrics,  $\widehat{TWHT}$  and  $\widehat{TTFB}$ , apparently showing noticeable differences. Being those metrics strictly related to the RTT (cfr. Fig. 4). This behavior is expected and is a direct consequence of lower RTT experienced by private reachable peers. Despite this bias no evident impact observed in download throughput, cfr. Table. 5.

## 6. Resource saving for different NAT policies

In this section we aim at providing some practical guidelines for the configuration of CGNs. In particular, we analyze different NATing policies and their saving in terms

<sup>10</sup>The measured  $\widehat{RTT}$  is inflated by the queuing delay of packets stacked in the upload queue of home routers. All measurements are equally biased by this phenomenon and it does not harm the reliability of the metric.

	Reachable	Unreachable
<b>Private</b>	58% ISP – 42% Internet	7% ISP – 93% Internet
<b>Public</b>	10% ISP – 90% Internet	6% ISP – 94% Internet

Table 6: Distribution of contacted peers




Figure 9: Fraction of active customers in different days.

of public IP addresses to be used to offer connectivity to the ISP customers. We consider two different cases: i) a simple NAT policy according to which a customer is given a public IP address for the period of time she is active. ii) a NAT and Port Address Translation (PAT) policy for which a customer is given a block of ports on given public IP address, for the time she is active. A *period of activity* of customer starts when the first packet is observed and ends when no packets is been observed for a period of time  $T_{out}$ , after which the resource is returned to the pool of available IP addresses of the CGN.

To conduct our analysis, we first must determine the number of active customers, and observe how their activity varies over the day. The analysis is conducted on the monitored PoP, but can be easily extended to the entire ISP customer population. We focus on the traffic directed to destinations outside the ISP network, ignoring the traffic internal to the ISP, which is not subject to the CGN. We consider TCP and UDP traffic. In particular for TCP, we take into account both successfully completed and failed connections, as in both cases the CGN has to allocate a public IP address and/or a block of port. We suppose that any TCP and UDP connection requires a dedicated (IP address, port) pair on the NAT, and this association must be maintained for the whole connection lifetime. The (IP address, port) pair will be released, freeing the resource, only later  $T_{out}$  minutes have passed.

For the experiments in this section we pick a workday (Wednesday) and an off day (Sunday), so to consider different activity patterns.

### 6.1. NAT based on simple address mapping

We first emulate the resource usage in the simple NAT scenario. We expect that in the worst case, i.e., when all customers are active, the ISP would need as many public IP addresses as the overall customer population inside the network.




Figure 10: Maximum fraction of active customers observed in different days vs. NAT's  $T_{out}$ .

From the NAT perspective, varying  $T_{out}$  influences the number of active customers in the network. We suppose  $T_{out}$  is 5 min. Fig. 9 shows the evolution, over 1-day, of the active customers, for both the weekday and the off day. The result shows that a simple NAT policy would turn into considering active approximately 40% of customers that are active at the same time, and hence the ISP would save roughly up to 60% of the public IP addresses. Finally, observe that the user activity is rather regular and similar for different days of the week. We conduct the same experiment on different days, and we observe very similar results (omitted for brevity).

For completeness, we check the impact of  $T_{out}$  on the estimation of active customers. In fact, the larger  $T_{out}$ , the longer the customer appears as active to the NAT, and the longer the time the NAT has to wait before redeem the public IP address. To this end, we show in Fig. 10 how the maximum fraction of active customers measured in the day (typically reached at the evening) changes when the  $T_{out}$  varies between 1 and 10 min.<sup>11</sup> As shown, the fraction of customers which have to be considered as active increases by a 10% only when increasing  $T_{out}$  up to 10 min. Observe also that there is no substantial difference between different week days.

## 6.2. CGN based on PAT policy

While above NATing technique might reduce the pool of public IP addresses to use, the actual savings are still limited, as the number of concurrent active customers is considerably large. Therefore, we investigate the resource requirements when NAT and PAT policy is in place, i.e., each active customer is given a block of ports on an public IP address.

For this policy, it is crucial to dimension the size of the block of ports the CGN shall allocate per customer. Hence, we have to count the per-customer number of concurrent connections. We expect the CGN to assign continuous

<sup>11</sup>Notice that RFCs suggest to set  $T_{out}$  to 2 min for UDP [12] and 2 h for TCP [13]. However, the suggested thresholds have been shown to be too long, and they lead suboptimal retention policies [14]. For this reason, we explore a threshold space closer to the order of tens of minutes.




Figure 11: Maximum, 99.9- and 99.9-percentile of the number of per-customer concurrent active connections (computed as the maximum between the numbers of UDP and TCP connections).

bulks of ports to each customer. The sizing of the block of ports should be based on the transport protocol, i.e., TCP or UDP, employing the largest number of ports. For instance, let  $p_{TCP}$  and  $p_{UDP}$  be the number of concurrent active TCP and UDP connections, respectively. The block size must then be larger than  $\max(p_{TCP}, p_{UDP})$ . We proceed as follows. We choose  $T_{out} = 5$  min, and we consider as concurrent the connections observed in 1 min long time bin. For each customer, we count the numbers of concurrent TCP and UDP connections. We then pick the maximum between the two, and use the result to build a per-minute distributions. In Fig. 11 we report the maximum, the 99.9- and the 99-percentiles obtained from the per-minute distributions and their evolution over time. As shown, the number of per-customer parallel connections rarely overcomes 20,000. In fact, we observe that customers employing a so wide number of ports are mostly users running P2P applications which open many parallel UDP flows. We see that 99% of customers never use more than 2,000 concurrent connections. Allocating a bulk of 2,000 ports for each customer would allow the ISP to use one public IP address for 32 costumers. Considering a more conservative approach, i.e., adopting the 99.9-percentile as a reference, we observe roughly 6,000 concurrent connections, leading in this case to allocate about 6 customers per public IP address. For the sake of completeness, we conduct the same experiment picking different days. The results are consistent and lead to very similar conclusions.

## 7. Related work

In this section we position our paper with respect to related studies. As the a thorough comparison of available statistical distance metrics is available in Appendix A, this section mostly discusses how our study complements the body of work present in the literature about NAT technologies.

In the last years, ISPs have deployed CGN in their networks to limit the utilization of public IPv4 addresses and postpone their final exhaustion day [15, 16]. Given their

strategic importance, CGNs have been matter of investigation in a large body of studies conducted by both standardization authorities and academia.

The IETF RFCs [12, 13, 17] standardize the requirements, implementations and behaviors for CGNs. A significant effort has been spent in standardizing mechanisms for NAT traversal, hole punching [18] and Interactive Connectivity Establishment (ICE) [19].

A remarkable amount of work has been dedicated to the task of identifying NAT deployment in residential networks [20, 21, 22, 23]. Similarly, but in a mobile scenario, [24] presents the results of an active measurement campaign to detect the presence of NAT middleboxes deployed in cellular networks.

Another branch of studies has focused on understanding the impact of CGNs on users' QoS and application-level experience [25, 26, 27, 28]. This paper falls in this category. Authors of [25] report a comprehensive classification of NATed scenarios and speculate about which impairments each of them could introduce. Authors of [26] presents a set of results obtained in a testbed. Specifically, they analyze the impact of network delays on the TCP connection establishment with and without CGN. In [27] authors evaluate the hole punching technique for NAT traversal, and how this impacts on the communication establishment in P2P applications. More similar in spirit to our work, [28] describes a case study conducted in controlled testbed where multiple CGN configurations are tested to evaluate their impact on several network applications and services. These include web, video streaming, P2P and gaming. These experiments are based on single sessions and do not consider actual performance testing. The results presented in [28] show that the presence of CGN has no substantial impact on users' browsing, thus confirming our observations. Differently, P2P applications like BitTorrent might be severely impaired. Our results, obtained in a real scenario and from a passive measurement perspective complement these observations. In fact, even if we can not discriminate leeches from seeders as authors of [28] do, we show that NAT444 has a deep impact on the peer selection.

A last family of work focuses on CGN dimensioning aspects, such as port allocation and retention, which we assess in Sec. 6. In [14] authors collect aggregate traffic traces from a real ISP network to investigate ports allocation and retention strategies in CGNs. The analysis shows that recommended timeout values in [13, 12] might be too long, resulting in suboptimal retention policies. In this paper we revisit the results presented in [14] by using more recent traces (2014 vs. 2009), collected from an actual residential ISP vs. campus network. Our results are different, given the evolution of Internet applications which tend to generate more concurrent connections.

Finally, to the best of our knowledge, this is the first work that specifically targets the problem of quantifying the impact of CGN on end-user experience from a passive perspective. We extend the results previously presented

in [29], by designing a more systematic methodology, and applying it to extend the study to P2P applications. Moreover, this extended version provides a thorough discussion (Sec. 6) about the possible resource saving which different CG-NAT policies would guarantee.

## 8. Conclusions

Network administrators lack effective tools to quickly pinpoint differences among several datasets obtained from traffic summaries. In this work, we aimed at filling this gap, and defined a methodology that builds on the statistical distance measures (such as Jensen-Shannon) divergence to assess statistical discrepancies between empirical distributions obtained from different populations or at different times. As an application example, we employed these analytics to study a large-scale CGN deployment, whereby ISP customers are split in two different populations, i.e., users assigned private IP addresses vs those assigned public IP addresses. For this particular scenario, we delineate several key performance indicators, relevant for user s' quality of experience. In particular, we gauge the impact of CGN deployment on the web browsing experience and on BitTorrent traffic.

Our results show that the CGN technology is stable and mature. As with any study based on passive measurement, results in the following are specific to the deployment that is under observation. Conditioned to measurement in our dataset, results suggest that if properly engineered and configured, CGN does not harm users' web browsing experience. Albeit the presence of the CGN has an evident impact on the neighborhood construction of BitTorrent, it does not affect the average transfer rate of peers. We conclude that the ISP we consider in our study may have no actual need to provide users with public IP addresses, when not specifically required. Finally, we analyzed our network traces to quantify the actual saving CGN policies could guarantee. In case of simple NAT policy, the considered ISP could save about 50% of public IP addresses. When NAT and PAT are combined, the saving can exceed one order of magnitude.

Generalizing such results is however an entirely different matter: in the case of CGN, admittedly there exist many different configurations and deployments, and it would be dangerous to project the lessons learned on our dataset across heterogeneous deployments. Of course, gains in terms of IP addresses savings are estimated based on the usage pattern of the ISP under study and cannot be generalized. For what instead concerns the performance implication, we point out that our study is limited to Web and BitTorrent traffic, but does not consider VoIP or Gaming traffic that have more stringent delay and jitter requirements: shall this traffic be especially important in another CGN deployment, the impact of CGN may be different in that case. Second, notice that our performance metrics are computed for successfully opened connections: as such,

we are not measuring if the deployment of CGN is changing the failure rate of connection attempts, which could possibly harm the performance at a session level. Third, changes in user/applications patterns can also heavily affect the results: e.g., currently browsers open many connections in parallel, which is unfavorable for CGN scenario while HTTP2 opens one per domain, so that these findings may need to be reassessed over long time periods with longitudinal studies.

In spite of these limitations, which are not peculiar to this work but that naturally arise whenever a specific dataset is considered, we believe that the methodology outlined in this work will survive the lessons learned by its application to the particular question about CGN we address here – and rather possibly enable such longitudinal investigations.

## Acknowledgements


This work has been funded by the *mPlane* project (grant agreement no. 318627) in the 7th European Framework Programme. We thank Valeria Di Gennaro for her initial help on this work, that she carried out in the context of her M.Sc. stage at LINC3 <http://www.lincs.fr>.

## References

- [1] A. Finamore, M. Mellia, M. Meo, M. Munafò, and D. Rossi, “Experiences of Internet Traffic Monitoring with Tstat,” *Network, IEEE*, vol. 25, pp. 8–14, 2011.
- [2] A. L. Gibbs and F. E. Su, “On Choosing and Bounding Probability Metrics,” *International statistical review*, vol. 70, no. 3, pp. 419–435, 2002.
- [3] S. Jiang, D. Guo, and B. Carpenter, “RFC 6264 - An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition,” Internet Engineering Task Force (IETF), Tech. Rep., 2011. [Online]. Available: {<https://tools.ietf.org/html/rfc6264>}
- [4] D. Bonfiglio, M. Mellia, M. Meo, D. Rossi, and P. Tofanelli, “Revealing Skype Traffic: When Randomness Plays with You,” *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 37–48, Aug. 2007.
- [5] R. Birke, M. Mellia, M. Petracca, and D. Rossi, “Understanding VoIP from backbone measurements,” in *INFOCOM 2007. 26th IEEE International Conference on Computer Communications*. IEEE, 2007, pp. 2027–2035.
- [6] D. Bonfiglio, M. Mellia, M. Meo, and D. Rossi, “Detailed Analysis of Skype Traffic,” *Multimedia, IEEE Transactions on*, vol. 11, no. 1, pp. 117–127, Jan 2009.
- [7] C. Donley, L. Howard, V. Kuarsingh, A. Chandrasekaran, and V. Ganti, “Assessing the Impact of NAT444 on Network Applications,” Internet Engineering Task Force (IETF), Tech. Rep., 2011. [Online]. Available: {<https://tools.ietf.org/html/draft-donley-nat444-impacts-01>}
- [8] Z. Hu, “NAT Traversal Techniques and Peer-to-Peer Applications,” in *Peer-to-peer technologies, networks and systems Seminar on Internetworking*, 2005.
- [9] J. Rosenberg, R. Mahy, P. Matthews and D. Wing, “RFC 5389 - Session Traversal Utilities for NAT (STUN),” Internet Engineering Task Force (IETF), Tech. Rep., 2008. [Online]. Available: {<https://tools.ietf.org/html/rfc5389s>}
- [10] I. N. Bermudez, M. Mellia, M. M. Munafò, R. Keralapura, and A. Nucci, “DNS to the Rescue: Discerning Content and Services in a Tangled Web,” in *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, ser. IMC ’12. New York, NY, USA: ACM, 2012, pp. 413–426.
- [11] M. Mellia, M. Meo, L. Muscariello, and D. Rossi, “Passive Analysis of TCP Anomalies,” *Computer Networks*, vol. 52, no. 14, pp. 2663–2676, 2008.
- [12] F. Audet, and C. Jennings, “RFC 4787 - Network Address Translation (NAT) Behavioral Requirements for Unicast UDP,” Internet Engineering Task Force (IETF), Tech. Rep., 2007. [Online]. Available: {<https://tools.ietf.org/html/rfc4787>}
- [13] S. Guha, K. Biswas, B. Ford, S. Sivakumar, and P. Srisuresh, “RFC 5382 - NAT Behavioral Requirements for TCP,” Internet Engineering Task Force (IETF), Tech. Rep., 2008. [Online]. Available: {<https://tools.ietf.org/html/rfc5382>}
- [14] S. Alcock, R. Nelson, and M. David, “Investigating the Impact of Service Provider NAT on Residential Broadband Users,” University of Waikato, Tech. Rep., 2010. [Online]. Available: {[http://www.wand.net.nz/~salcock/spnat/tech\\_report.pdf](http://www.wand.net.nz/~salcock/spnat/tech_report.pdf)}
- [15] A. Muller, F. Wohlfart, and G. Carle, “Analysis and Topology-based Traversal of Cascaded Large Scale NATs,” in *Proceedings of the 2013 Workshop on Hot Topics in Middleboxes and Network Function Virtualization*, ser. HotMiddlebox ’13. New York, NY, USA: ACM, 2013, pp. 43–48.
- [16] P. Richter, M. Allman, R. Bush, and V. Paxson, “A Primer on IPv4 Scarcity,” *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 2, pp. 21–31, Apr. 2015.
- [17] S. Perreault, I. Yamagata, S. Miyakawa, A. Nakagawa, and H. Ashida, “RFC 6888 - Common Requirements for Carrier-Grade NATs (CGNs),” Internet Engineering Task Force (IETF), Tech. Rep., 2013. [Online]. Available: {<https://tools.ietf.org/html/rfc6888>}
- [18] P. Srisuresh, B. Ford, and D. Kegel, “RFC 5128 - State of Peer-to-Peer (P2P) Communication across Network Address Translators (NATs),” Internet Engineering Task Force (IETF), Tech. Rep., 2008. [Online]. Available: {<https://tools.ietf.org/html/rfc5128>}
- [19] J. Rosenberg, A. Keranen, B. B. Lowekamp, and A. B. Roach, “RFC 6544 - TCP Candidates with Interactive Connectivity Establishment (ICE),” Internet Engineering Task Force (IETF), Tech. Rep., 2012. [Online]. Available: {<https://tools.ietf.org/html/rfc6544>}
- [20] L. DiCioccio, R. Teixeira, M. May, and C. Kreibich, “Probe and Pray: Using UPnP for Home Network Measurements,” in *Passive and Active Measurement*, ser. Lecture Notes in Computer Science, N. Taft and F. Ricciato, Eds. Springer Berlin Heidelberg, 2012, pp. 96–105.
- [21] S. M. Bellovin, “A Technique for Counting Natted Hosts,” in *Proceedings of the 2Nd ACM SIGCOMM Workshop on Internet Measurement*, ser. IMW ’02. New York, NY, USA: ACM, 2002, pp. 267–272.
- [22] G. Maier, F. Schneider, and A. Feldmann, “NAT Usage in Residential Broadband Networks,” in *Passive and Active Measurement*, ser. Lecture Notes in Computer Science, N. Spring and G. Riley, Eds. Springer Berlin Heidelberg, 2011, vol. 6579, pp. 32–41.
- [23] V. Krmicek, J. Vykopal, and R. Krejci, “Netflow Based System for NAT Detection,” in *Proceedings of the 5th International Student Workshop on Emerging Networking Experiments and Technologies*, ser. Co-Next Student Workshop ’09. New York, NY, USA: ACM, 2009, pp. 23–24.
- [24] Z. Wang, Z. Qian, Q. Xu, Z. Mao, and M. Zhang, “An Untold Story of Middleboxes in Cellular Networks,” in *Proceedings of the ACM SIGCOMM 2011 Conference*, ser. SIGCOMM ’11. New York, NY, USA: ACM, 2011, pp. 374–385.
- [25] 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.*, vol. 22, no. 2, pp. 391–404, Apr. 2014.
- [26] Y. Ohara, K. Nishizuka, K. Chinen, K. Akashi, M. Kohrin, E. Muramoto, and S. Miyakawa, “On the Impact of Mobile Network Delays on Connection Establishment Performance of a Carrier Grade NAT Device,” in *Proceedings of the AINTEC 2014 on Asian Internet Engineering Conference*, ser. AINTEC ’14. New York, NY, USA: ACM, 2014, pp. 1:1–1:8.

- [27] B. Ford, P. Srisuresh, and D. Kegel, "Peer-to-Peer Communication Across Network Address Translators," in *USENIX Annual Technical Conference*, 2005.
- [28] C. Donley, L. Howard, V. Kuarsingh, J. Berg, and J. Doshi, "RFC 7021 - Assessing the Impact of Carrier-Grade NAT on Network Applications," Internet Engineering Task Force (IETF), Tech. Rep., 2013. [Online]. Available: {<https://tools.ietf.org/html/rfc7021>}
- [29] E. Bocchi, A. Khatouni, S. Traverso, A. Finamore, V. Di Genaro, M. Mellia, M. Munafò, and D. Rossi, "Impact of Carrier-Grade NAT on Web Browsing," in *Wireless Communications and Mobile Computing Conference (IWCMC), 2015 International*, Aug 2015, pp. 532–537.

## Authors



**Enrico Bocchi** received his M.Sc. Degree in Telecommunications Engineering from Politecnico di Torino, Italy, in 2013. In 2014, he joined the Telecommunication Networks Group of Politecnico di Torino as Ph.D. Candidate. His research interests cover multiple aspects of Internet traffic monitoring, including cloud storage services measurements and benchmarking, network traffic characterization and classification, and applications to security. He is currently with the Laboratory of Information, Networking and Communication Sciences (LINCS), Paris, France, in the frame of a Joint-PhD degree between Politecnico di Torino and Telecom ParisTech, working on performance assessment of novel Internet protocols.

His research interests include Internet traffic monitoring, characterization and analysis.



**Ali Safari Khatouni** received his M.Sc. Degree in Computer and Communication Networks Engineering from Politecnico di Torino, Italy, in 2014. He joined the Telecommunication Networks Group of Politecnico di Torino as Ph.D. Candidate in 2015 and he is currently a participant of the mPlane Integrated Project.

His research interests include Internet traffic monitoring, characterization and analysis.



**Stefano Traverso** Ph.D. His research interests include privacy-preserving systems, network measurements and content delivery networks. During his Ph.D. and Post-doc he has been visiting Telefonica I+D research center (Barcelona, Spain), NEC Laboratories (Heidelberg, Germany) and Alcatel-lucent Bell Labs (Paris, France).

He is currently a Post-doc Fellow of the Telecommunication Networks Group of Politecnico di Torino.



**Marco Mellia** graduated from the Politecnico di Torino with Ph.D. in Electronic and Telecommunication Engineering in 2001, where he holds a position as Associate Professor. He has co-authored over 250 papers published in international journals and presented in leading conferences. He participated to the program committees of several conferences including ACM SIGCOMM, ACM CoNEXT, ACM IMC, IEEE Infocom, IEEE Globecom and IEEE ICC. He is Area Editor of ACM CCR, and ACM/IEEE Transactions on Networking. He is the coordinator of the mPlane Integrated Project, which focuses on building an Intelligent Measurement Plane for Future Network and Application Management.



**Maurizio Munafò** is Assistant Professor at the Department of Electronics and Telecommunications of Politecnico di Torino. He holds a Dr.Ing. degree in Electronic Engineering since 1991 and a Ph.D. in Telecommunications Engineering since 1994, both from Politecnico di Torino. He has co-authored about 70 journal and conference papers in the area of communication networks and systems. His current research interests are in simulation and performance analysis of communication systems and traffic modeling, measurement, and classification.


His current research interests are in simulation and performance analysis of communication systems and traffic modeling, measurement, and classification.



**Dario Rossi** is a Professor at Telecom ParisTech and Ecole Polytechnique. He received his M.Sc. and Ph.D. Degrees from Politecnico di Torino in 2001 and 2005 respectively, and his HDR Degree from UPMC in 2010. During 2003-2004, he held a visiting researcher position in the Computer Science division at University of California, Berkeley. He has coauthored over 9

patents and over 150 papers in leading conferences and journals, that received 3 best paper awards, a Google Faculty Research Award (2015), and an IETF Applied Network Research Prize (2016). He participated in the program committees of over 50 conferences including ACM ICN, ACM CoNEXT, IEEE Infocom of which he was also Distinguished Member (2015, 2016). He is a Senior Member of IEEE and ACM.

**Alessandro Finamore** received his Ph.D. in Electronics and Communication Engineering (2012), and M.Sc. (2008) from Politecnico di Torino. He has been an intern at University of Purdue, Lafayette, IL-USA in 2010, and Narus Inc., Sunnyvale, CA-USA in 2014. He coauthored over 30 publications, and participated in the



TPC of venues such as Infocom, PAM and TMA. His research interests are in the area of Internet traffic analysis, mobile systems, user quality of experience and mobility, CDNs services, and BigData frameworks. He is currently associate researcher at Telefonica Research in Barcelona.

Name	Abbrv	Formula	Support	Properties		
				Metric	Bounded	Symmetric
Jensen-Shannon	JS	$JS_{div}(p, q) = \sum_i \left\{ \frac{1}{2} p_i \ln \left( \frac{p_i}{\frac{1}{2} p_i + \frac{1}{2} q_i} \right) + \frac{1}{2} q_i \ln \left( \frac{q_i}{\frac{1}{2} q_i + \frac{1}{2} p_i} \right) \right\}$	$[0, \ln(2)]$		✓	✓
Kullback-Leibler	I	$KL_{div}(p, q) = \sum_i p_i \log \left( \frac{p_i}{q_i} \right)$	$[0, \infty)$			
Chi Square	$\chi^2$	$\chi_{dis}^2(p, q) = \sum_i \frac{(p_i - q_i)^2}{q_i}$	$[0, \infty)$			
Separation	S	$S_{dis}(p, q) = \max_i \left( 1 - \frac{p_i}{q_i} \right)$	$[0, 1]$		✓	
Total variation	TV	$TV_{dis}(p, q) = \frac{1}{2} \sum_i  p_i - q_i $	$[0, 1]$		✓	✓
Hellinger	H	$H_{dis}(p, q) = \left[ \sum_i (\sqrt{p_i} - \sqrt{q_i})^2 \right]^{\frac{1}{2}}$	$[0, \sqrt{2}]$		✓	✓
Kolmogorov	K	$K_{met}(P, Q) = \sup_x  P(x) - Q(x) $	$[0, 1]$	✓	✓	
Wasserstein	W	$W_{met}(P, Q) = \int_{-\infty}^{\infty}  P(x) - Q(x)  dx$	$[0, 1]$	✓	✓	
Discrepancy	D	$D_{met}(P, Q) = \sup_{\text{all closed balls } B}  p(B) - q(B) $	$[0, \text{diam}\Omega]$	✓	✓	

Table .7: Statistical Distance Measures. In the above formulas,  $p$  and  $q$  denote two empirical distributions on the measurable space  $\Omega$ , with  $p_i$  and  $q_i$  being their samples, and  $P$  and  $Q$  their cumulative distribution functions.

## Appendix A. Statistical Distance Measures

In this work we selected a specific Statistical Distance Measure (SDM) that we used as  $F(p, q)$ , namely the Jensen-Shannon divergence ( $JS_{div}$ ). The purpose of this section is thus to (i) contrast the broad set of SDMs from both a theoretic viewpoint, as well as making punctual examples to narrow down SDMs selection; (ii) show that, due to functional relationships between SDMs, it is possible to express the same methodology with multiple equivalent metrics, among which the Jensen-Shannon divergence ( $JS_{div}$ ); and (iii) assess robustness of the  $F(p, q)$  estimation as function of the  $p, q$  population size and binning strategy employed.

### Appendix A.1. SDM Comparison

The authors of [2] compare a variety of SDMs, shedding light on their properties and on the relationships among them. Without aiming at completeness, we report in Tab. 8 a list of 9 representative SDMs. Specifically, for each SDM the table reports its name, abbreviated notation, definition, co-domain and three relevant properties: (i) *Metric*, the SDM is a function defining a metric distance between each pair of elements in a set; (ii) *Bounded*, the SDM co-domain is finite; and (iii) *Symmetric*, the SDM is invariant to which of the two distributions is considered the reference, i.e.,  $F(p, q) = F(q, p)$ .

From Tab. 8 it is easy to see a rather heterogeneous picture. Most SDMs are divergence measures, with the exclusion of Kolmogorov (K), Wasserstein (W) and Discrepancy (D), which are metrics. With the exception of Kullback-Leibler (I) and Chi-Square ( $\chi^2$ ), all other SDMs have a bounded co-domain. Finally, only Jensen-Shannon (JS), Total Variation (TV) and Hellinger (H) are symmetric. None of the SDMs exhibits all three properties. As

we shall see later, these properties play an important role in the SDM selection.

In terms of provenance and use, JS and I are information theoretic measures. Loosely speaking, I expresses the amount of information that is required to encode  $q$  knowing  $p$ , while JS expresses the average amount of information carried by  $q$  which is not in  $p$ .  $\chi^2$ , H and K are often used for statistical tests. Metrics coming from both fields are apt in our framework.

We finally broaden the investigation by considering the 9 SDMs early introduced in Tab. 8, with the aim of both highlighting the relationships among them, as well as illustrating the behavior of each SDM in simple scenarios early considered for the  $JS_{div}$ . In principle, any of the SDMs in Tab. 8 can fit the purpose of our framework, so we illustrate here some relevant criteria to narrow down the SDM selection to a small set of equivalent functions.

To compactly represent these metrics, we arrange SDMs as a matrix of 9 blocks in Fig. 12(a). To show the intricate relationship among them, Fig. 12(a) represents dependency as arrows between block pairs, yielding an overall network of dependencies. It can be seen that, since each considered SDM is directly related to at least another one, their dependency graph consists of a single connected component. It also follows that it would be possible to adapt quantization functions  $Q(F(p, q))$  to the SDM in use, obtaining thus equivalent framework for different  $F(p, q)$ .

Fig. 12(b) and Fig. 12(c) show illustrative examples of these SDMs in action, by comparing as before synthetic probability distributions. The aim is here to visualize how the early illustrated SDM properties decline from practical viewpoint, and how these properties can be leveraged to narrow down SDM selection. For the sake of space, we focus here on the case of negative exponential distributions. As in Sec 2, we generate a reference sample with parameter

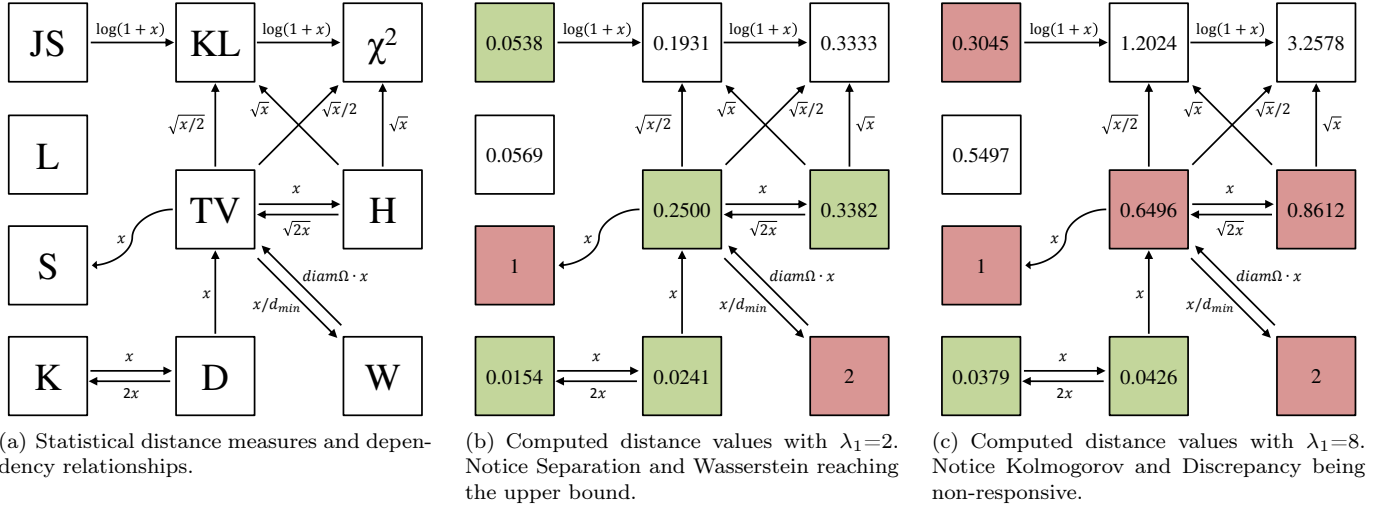


Figure A.12: Distance measures overview and computed values for negative exponential distributions.

$\lambda_0=1$ , and two other sample sets with parameters  $\lambda_1=2$ , and  $\lambda_1=8$ . These choices produce  $JS_{div}$  values falling in the green (noticeable difference) and red (significant difference) intervals, respectively, according to our  $Q^-$  and  $Q^+$  settings (cfr. Fig. 2(a)). We then compute the values of all considered SDMs comparing the two distributions of parameter  $\lambda_0$  and  $\lambda_1$ . Results are reported in Fig. 12(b) for  $\lambda_1=2$ , and in Fig. 12(c) for  $\lambda_1=8$ .

Several considerations hold contrasting Fig. 12(b) and Fig. 12(c). First, Separation (S) and Wasserstein (W) saturate to the upper bound already with  $\lambda_1=2$ , so that their fast varying dynamic is not able to express the greater difference among  $\lambda_0=1$  and  $\lambda_1=8$ .

Second, Kolmogorov (K) and Discrepancy (D) report a very low value both in case of  $\lambda_1=2$  and  $\lambda_1=8$ , with a slow dynamic that is opposite to the previous case. For all S, W, K and D, it is clear that the impact of  $Q^-$  and  $Q^+$  threshold selection in  $Q(F(p, q))$  becomes of paramount importance; additionally, no  $Q^-$  and  $Q^+$  selection would allow to express differences for SDMs such as S and W.

Third, I and  $\chi^2$  show a good sensitivity to the changes of the  $\lambda_1$  parameter. However, they not only are unbounded measures (which makes them not practical) but are not symmetric (which rules them out from our framework). Asymmetric metrics can be used when one wants to test a (suspect) population against a reference (well behaving), however the lack of symmetry makes it harder to use I and  $\chi^2$  in the general case where one has no a priori expectation about a population.

Fourth, we observe that the remaining measures, namely Jensen-Shannon (JS), Total Variation (TV) and Hellinger (H) are all good candidates. Not incidentally, JS, TV and H are the sole SDMs in Fig. 12(a) that are symmetric and bounded. In addition, they show to be consistently sensitive to different values of the  $\lambda_1$  parameter, as better detailed in Fig. A.13. Any of JS, TV and H are equivalent to our purpose. For practical purposes, it is pointless




Figure A.13: Trend of the three suitable distance measures.

to perform exhaustive analytics with each of these three metrics, which we show to be equivalent for our purpose. To avoid bringing redundant information and cluttering pictures and tables, we restrict our attention to  $JS_{div}$  as reference  $F(p, q)$  measure.

#### Appendix A.2. $JS_{div}$ Sensitivity analysis

We now assess the SDM robustness to factors that may affect the PDF estimation, as these may induce artificial bias leading to wrong conclusions. Indeed SDMs can be affected by various sources of bias, for instance, tied to (i) the binning strategy used to compute the samples of  $p$  and  $q$  distributions; (ii) the imbalance in the population size of  $p$  and  $q$ ; and (iii) the finitude of  $p$  and  $q$  populations.

Let us first start from the impact of the binning strategy. Taking  $JS_{div}$  as an example, we assess the operating conditions of the framework that ensure proper evaluation of the PDFs. We expect the binning adopted in estimating the PDF to play a role for continuous metrics with support in  $\mathbb{R}$ : intuitively, large bins smooth down differences ( $JS_{div}$  decreases, approaching 0 in the limit case where all samples fall in the same single bin). Small bins, in contrary, exacerbate differences ( $JS_{div}$  increases and approaches  $\ln(2)$  for rational bins of vanishing size, each of



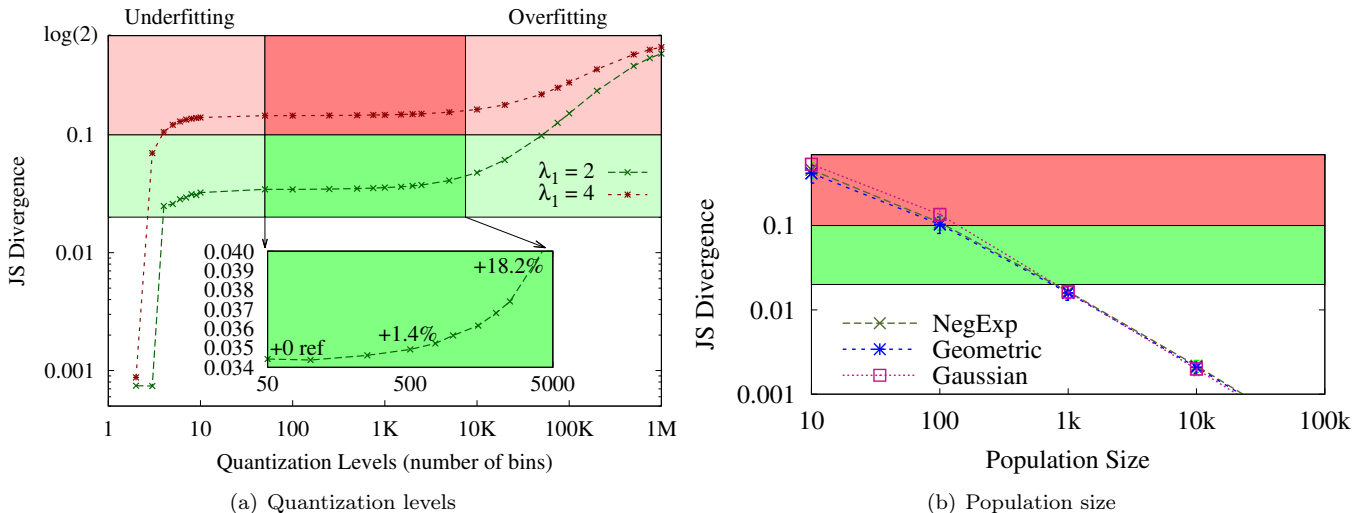


Figure A.14: Sensitivity analysis of Jensen-Shannon divergence for: (a) varying number of bins, (b) varying population size for two finite realization of the same process.

which contains a single or few samples).

We consider  $p$  and  $q$  as negative exponential distributions, with  $\lambda_0 = 1$ ,  $\lambda_1 \in \{2, 4\}$ . Given the previous  $Q^-, Q^+$  thresholds, we expect  $q = \text{NegExp}(x, 2)$  to be in the intermediate state, while  $q = \text{NegExp}(x, 4)$  to be significantly different from  $p$ . To avoid small population noise, we use finite sequences of  $10^6$  samples for each distributions. We then extract the empirical distributions from the two dataset by considering a number of bins which varies from 2 to  $10^6$ . Support is  $[0, 100)$ , thus  $\Delta b \in [0.001, 50]$ . We then compute the  $JS_{div}$  to compare  $p$  and  $q$ . For each value of the bin, we repeat 100 runs.

Fig. 14(a) show results, where the x-axis reports the number of bins (i.e., quantization levels) used, and the y-axis the corresponding  $JS_{div}$  value. Observe how the  $JS_{div}$  is quite stable for a number of bins in the 50-5,000 range. Conversely, when the number of bins is smaller than 50, a underfitting phenomenon shows up, so that the  $JS_{div}$  artificially drops to smaller values. Similarly, when the number of bins grows larger than 5,000, an overfitting phenomenon is visible, so that the  $JS_{div}$  artificially increases. Note the logarithmic scales. The inset details the relative error that occurs to  $JS_{div}$  with respect to the value obtained when using 50 bins, i.e., the reference. The relative error is below 19%. It follows that quantization oddities are easily controllable.


Notice that, for practical reasons, it is good practice to select a binning strategy that is tied to the physics of the metric: for example, use an unitary bin size for measurements that takes integer values (e.g., the Number of Hops), or relate the bin size to the unit of scale of interest (e.g., a ms accuracy for RTT and time-related metrics, or consider bins of 10 kbps when dealing with throughput). We thus expect binning policies to induce negligible distortion levels if properly selected.

Clearly, a specular question is in place: for a given bin size choice, which is the impact of the number of samples on the estimation of the PDF? Intuitively, while any finite sequence deviates from quantiles of the theoretic distribution, small population samples tend to exhibit larger deviations.

Taking two finite realizations of the same process, we estimate the empirical PDFs  $p$  and  $q$  and compute the  $JS_{div}$ . To avoid binning bias, we consider real-valued distributions (i.e., Gaussian, negative exponential) and an integer-valued distribution (Geometric). We then estimate the two (nominally identical) PDFs using a number of samples that varies from 10 samples only, to  $10^5$  samples. We compute the  $JS_{div}$  (nominally 0) considering quantization with 1,000 bins.

Fig. 14(b) show results. Irrespectively of the distribution,  $JS_{div}$  is strongly affected by the population size (linear slope in log-log plot). As expected, an excessively small population inflates the  $JS_{div}$  value. Specifically, having less than 1,000 (100) samples in the population causes the  $JS_{div}$  to exceed the warning threshold for noticeable (significant) differences for all the distributions considered. It is thus recommended to employ the  $JS_{div}$  on population larger than 1,000 samples, assumption verified in our dataset.

However, it is important to mention that artifacts caused by a limited population size may have an impact in case the methodology is used in real-time (e.g., on short time window) scenarios, or to compare the same population over different temporal samples. This possibly mandates a minimum duration of the observation period, especially in off-peak times, so to reach a minimum level of observation samples. We stress the importance of an adequate population size by reporting the example in Fig. A.15. We compute the per-hour  $JS_{div}$  on upload and



(a) Number of samples for public (filled) and private (dashed) upload and download goodput. household. (b) Computed  $JS_{div}$  values on upload and download goodput.

Figure A.15:  $JS_{div}$  per-hour on upload and download goodput badly biased by limited population size.

download goodput ( $\hat{G}$ ) as defined in Sec. 4.3 for one day. Fig. 15(a) shows the number of samples composing the population of customers with public or private IP address over the 24 hours considered. Notice the drop during the night time, and the consistently lower amount of samples for the upload direction. Now consider the corresponding  $JS_{div}$  values reported in Fig. 15(b), computed every hour. Notice how the  $JS_{div}$  is consistently higher for the upload goodput, and the peaks in both upload and download directions between midnight and 6am. This two phenomena are not driven by an effective difference in performance between the two populations, but owe to the finitude of available samples.