Public Review for
Dissecting Last-mile Latency Characteristics

Vaibhav Bajpai, Steffie Jacob Eravuchira, Jrgen Schnwlder

Measuring, quantifying and understanding access network performance has always been a challenge. In this work, authors exploit active measurements to investigate patterns and trends in last mile latency, digging into a longitudinal dataset from UK and US. Results present an in depth characterization of access delay, where some unexpected findings are observed.
Reviewers found the paper to provide a nice snapshot of the current state and diversity of access technologies, and appreciated the effort authors put in identifying interesting insights. As such, the paper is an interesting piece of work that will benefit the community at large.
The paper is fully reproducible. Authors make dataset and scripts accessible via a single github repository, that also includes installation instructions and user instructions. The dataset is SQLite, the scripts are plain Python 3 and Jupyter notebook Python 3 scripts. Data allows you to reproduce each result presented in the paper, from the fetching of RIPE Atlas data, lookup network types via peeringdb API, lookup AS information via RIPE stat API, to the processing and figures creation.

Public review written by
Marco Mellia
Politecnico di Torino

Artifacts Review for

Dissecting Last-mile Latency Characteristics

Vaibhav Bajpai, Steffie Jacob Eravuchira, Jürgen Schönwälder

In the paper *Dissecting Last-mile Latency Characteristics*, the authors investigate latency characteristics based on datasets obtained from two global Internet measurement platforms (RIPE Atlas, SamKnows). They provide insights on DSL networks including the usage of interleaving and its dynamic adaptation, latency, temporal, geographical, and ISP or broadband speed-related characteristics of access networks.

I would like to praise the fact that the authors have chosen to share the detail of their enriching methodology to the research community, providing an opportunity to increase the benefits of this article, or to further pursue their work to the willing researcher. Moreover, this article and the associated material can thoroughly serve as an example to young researchers wishing to pay special attention to the reproducibility of their own work.

Indeed, the datasets and scripts used to obtain the results are made publicly available. This allows to reproduce each result presented in the paper from data fetching (RIPE Atlas), using third-party public APIs (PeeringDB, RIPEstat), data processing and plotting.

I attribute to this article the following badges:

- **Artifacts Evaluated - Reusable**: All the results presented in this paper have been successfully replicated during the review process. The authors made a noticable effort in easing the task of re-running their calculations by splitting and documenting them. Moreover, the tools chosen by the authors are particularly suited for code exploration and reproducibility (jupyter, Python, pandas, public APIs).
- **Artifacts Available**: All required material is made publicly available.

Artifacts review written by
Korian Edeline
University of Liège

Dissecting Last-mile Latency Characteristics

Vaibhav Bajpai
TU Munich

Steffie Jacob Eravuchira
SamKnows Limited

Jürgen Schönwälder
Jacobs University Bremen

ABSTRACT

Recent research has shown that last-mile latency is a key network performance indicator that contributes heavily to DNS lookup and page load times. Using a month-long dataset collected from 696 residential RIPE Atlas probes and 1245 SamKnows probes, we measure last-mile latencies from 19 ISPs (RIPE Atlas) in the US and the EU, and 9 ISPs (SamKnows) in the UK. We show that DSL deployments not only tend to enable interleaving on the last-mile, but also employ multiple depth levels that change over time. We also witness that last-mile latency is considerably stable over time and not affected by diurnal load patterns. Unlike observations from prior studies, we show that cable providers in the US do not generally exhibit lower last-mile latencies when compared to that of DSL. We instead identify that last-mile latencies vary by subscriber location and show that last-mile latencies of cable providers in the US are considerably different across the US east and west coast. We further show that last-mile latencies vary depending on the access technology used by the DSL modem wherein VDSL deployments show last-mile latencies lower than ADSL1/ADSL2+ broadband speeds. The entire dataset and software used in this study is made available [2] to the measurement community.

CCS CONCEPTS

• **Networks** → Network monitoring;

KEYWORDS

SamKnows, RIPE Atlas, Last-mile Latency, Home Networks

1 INTRODUCTION

Srikanth Sundaresan *et al.* in [33] (2013), using the BISmark [4, 31] platform, have shown that latency becomes a critical factor impacting quality of experience in networks where downstream throughput exceeds 16 Mb/s. The effects of this observation are visible today with continuous efforts that attempt to move popular content as close [17] to the edge as possible. Yi-Ching Chiu *et al.* in [11] (2015) recently showed that popular paths to CDNs serving high volume client networks tend to be shorter than paths to other networks. This is taken even further by some large content providers that deploy content caches [8, 10] directly in service provider networks. Furthermore, new standards such as HTTP/2 [5] (2015) have been defined with a goal to improve webpage load times. Ongoing efforts such as QUIC [21] (2017) and TLS 1.3 [29] (2017) take this further to target operation on a much reduced latency (known as 0-RTT mode) overhead. In efforts to highlight confounding factors responsible for degraded webpage performance, Srikanth Sundaresan *et al.* in [33] (2013) recently showed that last-mile latency is a major contributor to end-to-end latency and it contributes heavily to DNS lookup and page load times. Last-mile latency is becoming a key broadband network performance indicator and factors affecting

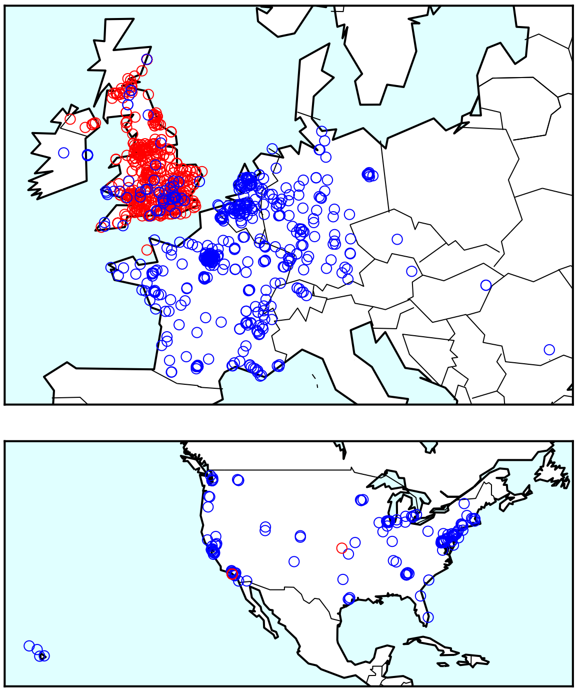

Figure 1: Distribution of 696 RIPE Atlas v3 (blue) and 1245 SamKnows (red) residential probes. RIPE Atlas probes span the EU (521) and the US (161), while SamKnows probes span the UK (1233) and the US (11).

last-mile latency need further investigation. We measure last-mile latencies using two datasets as shown in Fig. 1. The datasets have been obtained using 696 residential RIPE Atlas [4] probes deployed in 19 ISPs in the US and the EU and 1245 residential SamKnows [4] probes deployed in 9 ISPs in the UK.

Our **findings** −*a*) DSL service providers not only enable interleaving, but also dynamically adapt the depth levels (see § 4.1) with time, *b*) Last-mile latency is considerably stable over time (see § 4.2) and not affected by diurnal load patterns. Last-mile latencies for DSL center (see § 4.3) at ~16 ms, with cable at ~8 ms, and fibre deployments at ~4 ms, *c*) Subscribers of some US cable providers experience considerably different (see § 4.4) last-mile latencies across the US east (centered at ~24 ms) and west coast (centered at ~8 ms) and *d*) Last-mile latencies decrease with increase (see § 4.5) in broadband speeds. VDSL shows last-mile latencies lower than ADSL1/ADSL2+ deployments.

To the best of our knowledge, this is the first study that measures last-mile latency characteristics on multiple perspectives covering several service providers in the US and the EU. This is the first study to show interleaving depth levels, last-mile latency behaviour by time of day, last-mile latency by subscriber location and last-mile latency based on the access technology used by the DSL modem. To help with reproducibility [3], the entire dataset and software (see § 5 for details) used in this study is made available to the community.

2 RELATED WORK

Marcel Dischinger *et al.* in [15] (2007) inject packet trains and use responses received from home gateways to infer broadband link characteristics. They show that last-mile latencies are mostly affected by large modem queues and are higher for DSL when compared to cable networks. Srikanth Sundaresan *et al.* in [32] (2011) use the SamKnows platform to show that DSL networks enable interleaving on the last-mile which increases last-mile latencies for DSL users. Igor Canadi *et al.* in [9] (2012) show that end-to-end latencies to `speedtest.net` servers experienced by DSL users are higher in US markets. Our study extends this state of the art to show that DSL deployments not only enable interleaving, but also implement multiple interleaving depth levels (see § 4.1) and vary them over time. Furthermore, our study shows that last-mile latencies for cable users are *not* generally lower than that of DSL. We instead identify that last-mile latencies vary by subscriber location. We show that last-mile latencies of some cable providers in the US are considerably different (see § 4.4) across the US east and west coast. Subscribers of cable providers around the US east coast experience last-mile latencies similar to that of DSL. Prior research has measured last-mile latency as the latency to the first public IP hop and consequently included latencies within the home network. Our study shows that latencies within the home network have an impact and must not be included when measuring last-mile links.

Aaron Schulman *et al.* in [30] (2011) use PlanetLab [12] vantage points to send ICMP echo request packets to broadband hosts. They describe how physical factors (snow, wind, rain) affect the reliability of last-mile links. Zachary S. Bischof *el al.* in [6] (2012) run `traceroute` measurements from within a BitTorrent plugin to measure the effect of last-mile latencies on web performance. They show that while increasing bandwidth provides a trend of diminishing returns, high last-mile latency dramatically increases page rendering times. Daniel Genin *et al.* in [18] (2013) measure effects of congestion on access networks. They show that DSL links are mostly congested on the last-mile, while cable links usually experience congestion beyond the last-mile and show higher variability of such congestion events. Srikanth Sundaresan *et al.* in [33] (2013) show that last-mile latency is a bottleneck in high-throughput networks. They propose methods to perform DNS prefetching and TCP connection caching on the residential gateway to mitigate last-mile latency bottlenecks.

3 METHODOLOGY

RIPE Atlas has deployed ~22K [4] (with ~9.8K connected) and SamKnows has deployed ~250K [4] (with ~100K connected) dedicated hardware probes all around the globe as of June 2017. We begin by describing how we identified residential probes and provisioned month-long `traceroute` measurements from both platforms.

Table 1: Distribution of SamKnows (above) and RIPE Atlas (below) probes by service providers.

	ISP	ASN	TYPE	CC	##
1	BT	2856	DSL	UK	314
2	PLUSNET	6871	DSL	UK	271
3	VIRGINMEDIA	5089	CABLE	UK	201
4	OPALTELECOM	13285	DSL	UK	132
5	ORANGE	12576	DSL	UK	82
6	TISCALI	9105	DSL	UK	73
7	BSKYB	5607	DSL	UK	36
8	ZEN	13037	DSL	UK	35
9	TALKTALK	43234	DSL	UK	34
1	FREE	12322	DSL	FR	137
2	COMCAST	7922	CABLE	US	122
3	DTAG	3320	DSL	DE	61
4	ORANGE	3215	DSL	FR	60
5	TELENET	6848	CABLE	BE	30
6	XS4ALL	3265	DSL	NL	30
7	OVH	35540	DSL	FR	29
8	LDCOMNET	15557	DSL	FR	29
9	BELGACOM	5432	DSL	BE	25
10	UUNET	701	FIBRE	US	23
11	BT	2856	DSL	UK	23
12	LGI	6830	CABLE	EU	23
13	VIRGINMEDIA	5089	CABLE	UK	20
14	ZIGGO	9143	CABLE	NL	19
15	TWC	20001	CABLE	US	16
16	VIEWQWEST	18106	FIBRE	SG	14
17	TELEFONICA-DE	6805	DSL	DE	13
18	ZEN	13037	DSL	UK	12
19	VODAFONE	3209	DSL	DE	10

In the first step we clustered probes by their AS. For RIPE Atlas, we used the ASN revealed by the RIPE probe API [25], while for SamKnows, we used the RIPE data API [26] to map the public IP to the origin AS of the probe. We then ranked ASes by the number of probes and searched the literature for techniques that can classify ASes by network type. While Xenofontas Dimitropoulos *et al.* in [14] (2006) provide an approach to classify ASes using machine learning techniques, the dataset is outdated. In contrast, PeeringDB [23] which is a database holding peering information of participating networks serves as a living, viable alternative today. Therefore, we used PeeringDB to map ASes to their network type information. This mapping allowed us to select for ASes that belong to ISP networks. In the next step, we provisioned one-off `traceroute` measurements to identify residential probes. We define *residential probes* as probes that are directly wired to the home gateway. In order to achieve this, we searched for probes whose *hop*1 was in a private IPv4 address space [28], but their *hop*2 was in a public IPv4 address space. Going forward, we use the term probes to refer to residential probes.

We further classified probes by access technology based on service offers made on the website of the ISP. We also searched literature for techniques to validate the classification since neither dataset has ground-truth on access type used by the home gateway. For instance, UPnP discovery messages can be used to reveal access

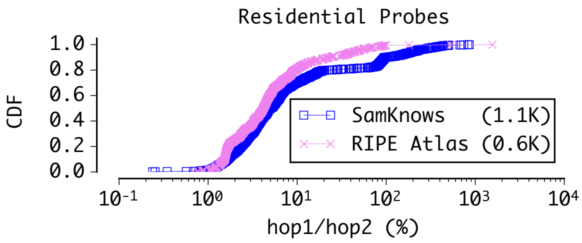

Figure 2: CDF of *hop*1 latency to that *hop*2. ~19% of RIPE Atlas and SamKnows probes show *hop*1 contributing to > 10% (but less than 100%) of *hop*2 latency. Home network latency should not be accounted when measuring last-mile latency.

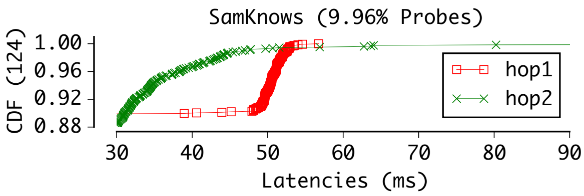

Figure 3: 9.95% of SamKnows probes show *hop*1 contributing to more than 100% of *hop*2 latency. The *hop*1 latency is stable at around 50 ms. We suspect, these probes are connected to home routers that prefer to rate limit ICMP responses to TTL expiry.

technology used on the WAN interface of a home gateway. However, since RIPE Atlas currently does not support a measurement that can perform UPnP queries and since this technique has been proven to be unreliable [13] (2012), we instead rely on reverse DNS entries derived from the public IP endpoints to validate the access type classifications with less than 1% mismatch error. We next provisioned month-long `traceroute` measurements towards RIPE Atlas anchors and SamKnows Measurement Lab [16] servers. Measurements were performed every 4 hours using `evtraceroute` busybox applet on RIPE Atlas and using `mtr` on SamKnows.

In this process, we discovered [1] (2015) that older versions of RIPE Atlas probes (~43.1% of all probes) experience load issues due to their hardware limitations. Recently, it has been further confirmed [20] (2015) that these delays are more pronounced in situations where older version of probes are loaded with concurrent measurements. We therefore base our measurements on the most recent hardware version (v3) only. The `traceroute` measurements were conducted every 4 hours over 35 days in (July-August) 2014. Note, since probes in either platform cannot associate to a wireless access point, these measurements do not get skewed by presence of wireless links in the home network.

This dataset consists of 135K last-mile latency data points captured from 696 residential v3 RIPE Atlas probes and 440K last-mile latency data points captured from 1245 residential SamKnows probes. Fig. 1 shows the geographical distribution of these probes. RIPE Atlas probes cover 19 different ISPs in the the EU (521 probes) and the US (161 probes), while SamKnows probes cover 9 ISPs in

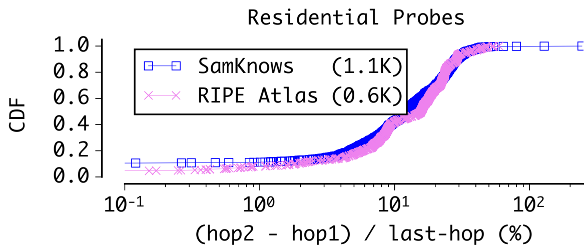

Figure 4: CDF of last-mile latency to end-to-end latency. More than half of RIPE Atlas and SamKnows probes experience last-mile latency contributing to > 10% of end-to-end latency.

the UK (1233 probes) and the US (11 probes). Table 1 further shows the number of probes broken down by ISP. We only consider ISPs in the analysis that have at least 10 probes.

4 DATA ANALYSIS

We begin by investigating the latency contributed by the home network (*hop*1) to that of the first hop in the service provider's network (*hop*2). A major portion of RIPE Atlas (~92%) and SamKnows (~80%) probes show expected *hop*1 latencies of less than 1.5 ms, while a discernible number of probes show more than expected *hop*1 latency. For instance, Fig. 2 shows the relative contribution of *hop*1 latency to that of *hop*2. We witness that 9.95% of SamKnows probes (and 0.4% of RIPE Atlas probes) show *hop*1 contributing to more than 100% of *hop*2. 69% of these probes are connected to PLUSNET home routers. Fig. 3 shows that *hop*1 latencies for these probes appear to be ~50 ms. We suspect that these probes are behind home routers that prefer to rate limit ICMP responses to TTL expiry and therefore have higher `traceroute` response times. We do not consider these probes as part of our dataset. Pruning these probes out, we witness that ~19% of both SamKnows and RIPE Atlas probes show *hop*1 latency contributing to 10% or more (but less than 100%) of *hop*2 latency. As such, latencies within the home network can have a discernible impact and must not be included when measuring last-mile latency. In order to circumvent effects of latencies induced within a home network, we calculate *last-mile latency* as the difference between the *hop*2 and *hop*1 latency. Last-mile latencies described beyond this point reflect this definition. Fig. 4 shows the contribution of last-mile to end-to-end latency. It can be seen that more than half of RIPE Atlas and SamKnows probes experience last-mile latencies that contribute to more than 10% of end-to-end latency with ~80% probes experiencing last-mile latency contributing to more than 5% of end-to-end latency. Thererfore, last-mile latency is a key broadband network performance indicator today and factors affecting last-mile latency need further investigation.

4.1 Interleaving depths in DSL networks

It is suspected that DSL networks enable interleaving on the last-mile to trade latency with lower packet loss rates [7]. An interleaving channel intersperses the payload between DSL frames to provide Impulse Noise Protection (INP) on the last-mile. This is usually

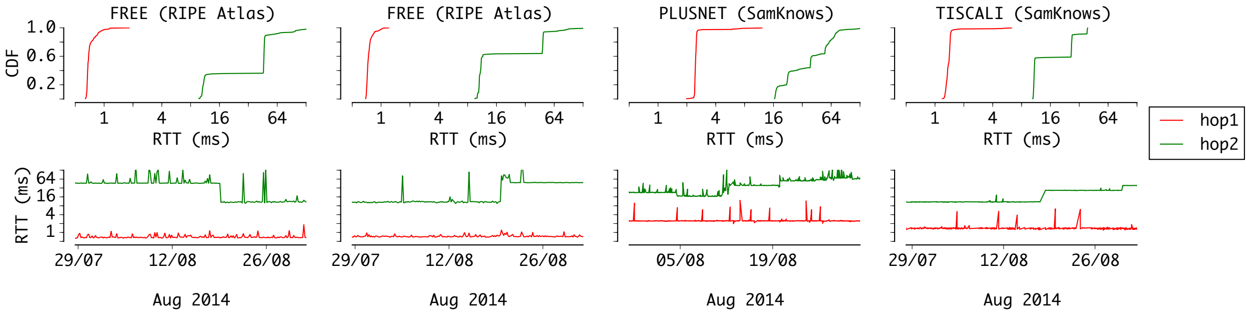

Figure 5: CDF (top) of *hop*1 and *hop*2 latencies from four probes connected to DSL networks. A step-wise change in *hop*2 latency exhibits an interleaving depth level change which matches with the timeseries (bottom).

implemented along with the Reed Solomon (RS) Forward Error Correction (FEC) technique to make the channel more resilient to packet loss. The number of RS codewords accumulated before transmitting the frame determines the depth of the interleaving channel. DSL deployments employ the Dynamic Line Management (DLM) technique to remotely monitor line characteristics such as the amount of packet loss encountered on the last-mile. They use this information to dynamically adapt interleaving depth levels. An increase in depth level increases latency. An increase in latency can directly impact applications leveraging congestion aware transport protocols such as TCP. An interleaving depth level of 1 is known as *fastpath* which is more suitable for real-time communication applications but only appropriate for links with low error rates. DSL operators tend to support both fastpath and higher depths, although not all operators allow fastpath on the last-mile. It is also unlikely that a deployment will only support fastpath.

In our pursuit to identify interleaving depths, we investigated latencies observed by both SamKnows and RIPE Atlas probes connected to DSL networks. A change in the interleaving depth level changes the *hop*2 latency by ~5 ms [7]. A step-wise transition on the CDF derived from *hop*2 latencies indicates a switch between such depth levels. Fig. 5 shows example probes that witnessed depth-level changes. These probes portray *hop*2 latencies distributed as step-wise functions. It can be seen that multiple depth level transitions occurred over the span of a month. The corresponding timeseries tends to match with the depth changes showing that DSL networks tend to vary interleaving depths over time. SamKnows probes perform measurements only in the absence of cross-traffic, as a result the second-hop transitions cannot be attributed to bufferbloat [19] on the home gateway. Each data point in the timeseries is an average of three queries, as a result, some spikes are also visible. We further refer the reader to [1] where we discuss the effects of averaging latencies over a single hop. In order to automate the discovery of probes experiencing such a behavior, we extracted relative maximas from the Kernel Density Estimation (KDE) derived from *hop*2 latencies witnessed by each probe. We used a sample threshold on the frequency of occurrence for each local maxima to ensure *hop*2 latencies remained stable for an extended period. We tagged probes with a depth-level transition in situations where the local maximas were at least 5 ms apart from each other. Fig. 6 shows the distribution

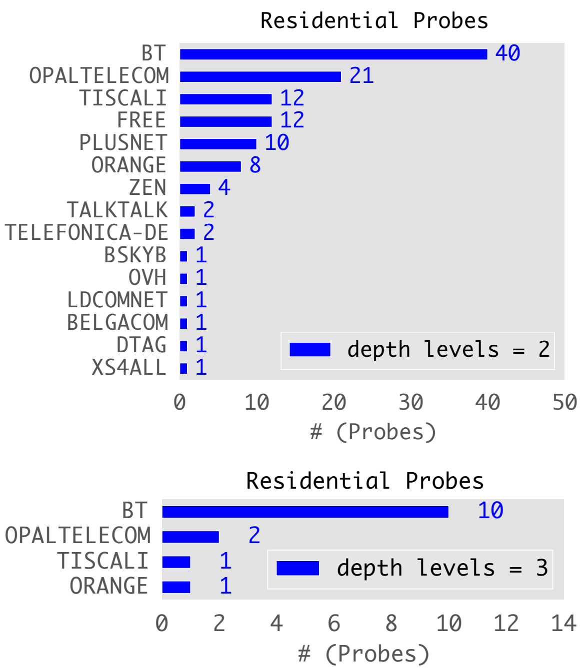

Figure 6: DSL probes that witnessed 2-levels (above) and 3-levels (below) of interleaving depth changes.

of both SamKnows and RIPE Atlas probes that experienced 2−levels and 3−levels of interleaving depth level changes. The observations were validated with one (BT) service provider. This analysis extends our understanding that DSL deployments *not* only enable interleaving, but also implement multiple interleaving depth levels and vary them over time.

4.2 Last-mile latencies by time of day

We investigated the distribution of last-mile latencies over 24 hour cycles for DSL, cable and fibre deployments. Fig. 7 shows boxplots of last-mile latencies observed over each hour in UTC. Note, our

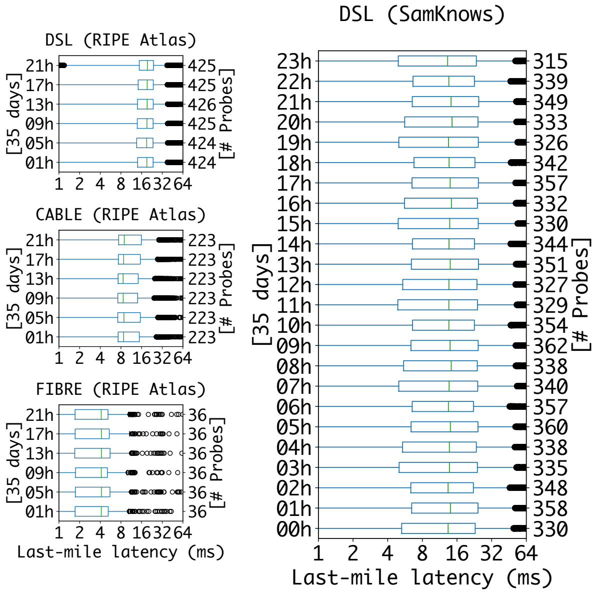

Figure 7: Last-mile latencies by time (UTC) of day. Last-mile latencies remain considerably stable by time of day.

measurements were taken every 4 hours over a 35 days period. The measurements were also provisioned using UTC. Since SamKnows tends to distribute probes within the frequency interval, measurements were spread over each hour of the day. RIPE Atlas only recently (since Nov 2015) introduced this feature [22] of controlling the spread. Given our dataset spans Aug 2014, RIPE Atlas measurements strictly occur on the 4 hour boundary. Since, SamKnows (unlike RIPE Atlas) probes do not perform measurements in presence of cross-traffic [1], the number of SamKnows probes running measurements change every hour unlike that of RIPE Atlas where all probes participate in the measurement. It can be seen that the last-mile latency is stable over time and is not affected by diurnal load patterns. Note that our measurement method has been designed to eliminate queuing delays such as delays caused by home gateways with bloated buffers [19] in front of an overloaded access line. As a such, this observation is in line with expectation. A DSL line is not shared with other customers (except indirectly via crosstalk impacting signal quality) and hence load should not affect DSL line behaviour in significant ways. For cable access networks, the situation is slightly different but it seems that deployments have enough capacity to sustain load such that the time-slotted approach makes them behave in a reasonably robust way.

4.3 Last-mile latencies by service provider

Fig. 7 shows that last-mile latencies for DSL center at ~16 ms, with cable at ~8 ms, and fibre deployments at ~4 ms. We further break down the last-mile latencies by service provider networks as shown in Fig. 8. We witness that last-mile latencies exhibited by DSL providers in the EU are higher (due to interleaving) when compared to cable providers with fibre deployments exhibiting relatively

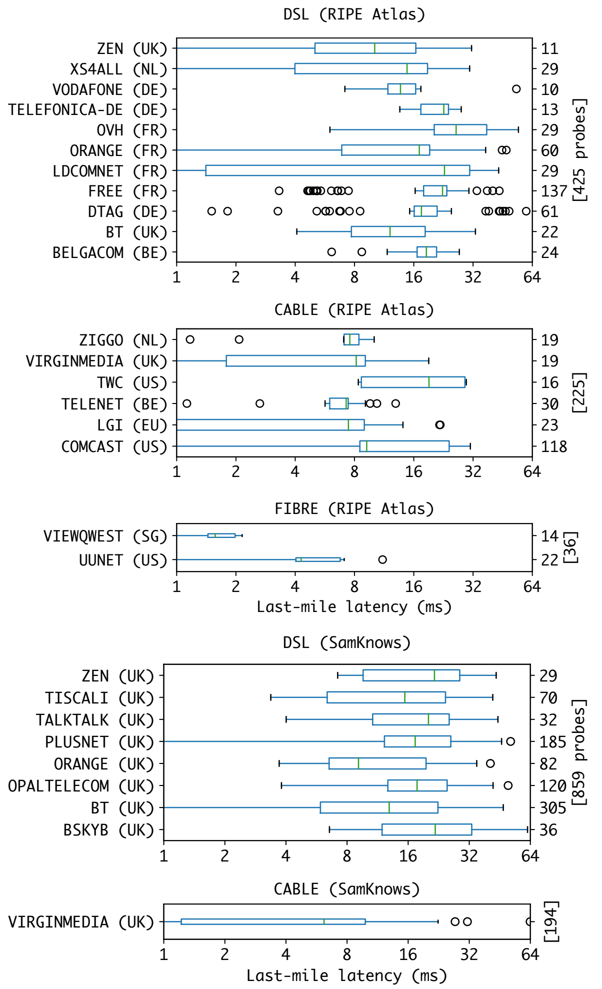

Figure 8: Last-mile latencies for DSL, Cable and Fibre ISPs. Last-mile latencies are ordered as DSL > Cable > Fibre for multiple ISPs in the EU.

lower last-mile latencies than that of DSL and cable deployments. The distribution shows higher variation in DSL networks due to multiple levels of interleaving depths enabled depending on the line characteristics and geographical location of the subscriber. The last-mile latencies for cable providers in the US (Comcast and TWC for instance) also appear to exhibit a multi-modal distribution. One of the clusters exhibit last-mile latencies similar to EU cable providers (centered at ~8 ms), while the other cluster exhibits last-mile latencies similar to EU DSL providers (centered at ~24 ms) which is discussed in the next section. Further, note that the sample of probes after splitting observations by ISP also goes down significantly. As such, it becomes difficult to reasonably discuss latency distributions for every ISP in greater detail. We captured some of these lessons in [1], where we show that the AS-based distribution of RIPE Atlas probes is heavily-tailed. As such, studies that require higher coverage of network origins tend to benefit more from RIPE Atlas than those that require high probe density within each network.

4.4 Last-mile latencies by subscriber location

We further investigated last-mile latencies by clustering probes of a service provider by their subscriber location. Given the RIPE Atlas dataset consists of probes located in both EU and US regions, the probes are located in different timezones. We use timezones since they provide a good granular separation by location (since countries are too coarse grained, while cities are too fine grained for the number of probes within each service provider). Fig. 9 shows the distribution of last-mile latencies grouped by location for selected ISPs where we have a higher sample (more than 100) of probes. This separation reveals the reason for the multimodal distribution (see Fig. 8) of last-mile latencies exhibited by ISPs. Fig. 10 shows that Comcast with last-mile latencies centered at ~8 ms are exhibited by probes in the LA region, while last-mile latencies centered at ~24 ms are exhibited by probes in the NYC region. Similar results are observed for TWC and LGI-UPC service providers. As such, unlike prior observations [15, 32], this analysis reveals that *not all* cable deployments show last-mile latencies lower than DSL. We instead identify that last-mile latencies vary by subscriber location. Last-mile latencies of cable providers within the EU are generally lower than that of DSL, but last-mile latencies of some cable providers in the US are considerably different across the US east and west coast. Subscribers of those cable providers around the US east coast experience last-mile latencies similar to that of DSL. However, the causes of this observed effect remain unclear. Further analysis is limited by the capabilities of the collected dataset and requires collaboration with service providers.

4.5 Last-mile latencies by broadband speeds

DSL technology has also evolved over the years. For instance, ADSL2 provides multichannel transmission capability that allows different latency characteristics to be applied to each channel over the last-mile. ADSL2+ uses higher frequencies to double bandwidth capacities. We further investigated the characteristics of last-mile latency based on the access technology used by the DSL modem. Fig. 11 shows last-mile latencies observed by DSL SamKnows probes separated by 4 broadband speed tiers. It can be seen that the last-mile latencies observed by probes behind ADSL1 and ADSL2+ speeds are similar. Although a cluster of probes behind ADSL2+ lines also center at ~8 ms and show last-mile latencies lower than ADSL1. On the other hand last-mile latencies for VDSL speeds tend to show considerably lower last-mile latencies when compared to ADSL1 and ADSL2+ speeds. Fig. 12 shows a decrease of last-mile latency with increase in broadband speeds. In DSL deployments, higher bandwidth capacities are made possible by using higher range frequencies on the physical link. These frequencies tend to dissipate over shorter distances. Therefore, ADSL2+ and VDSL deployments tend be closer to the traffic aggregation points. Although, a reduction in copper length does not have significant effects on last-mile latency. Furthermore, with an increase in line speeds, ADSL2+ and VDSL deployments allow frames to be transmitted faster. Higher transmission rates help reduce interleaving delays, which can significantly

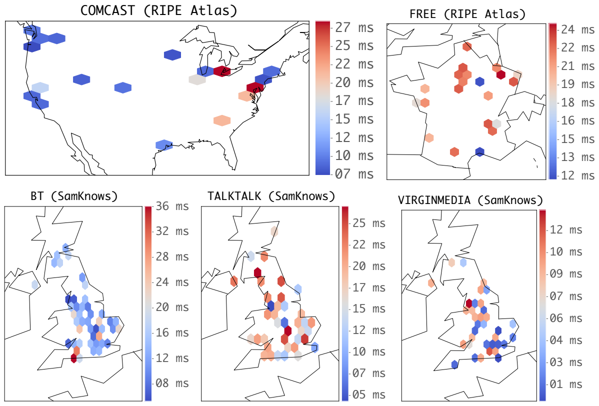

Figure 9: Last-mile latencies separated by location. Users witness different last-mile latencies depending on their location of subscription.

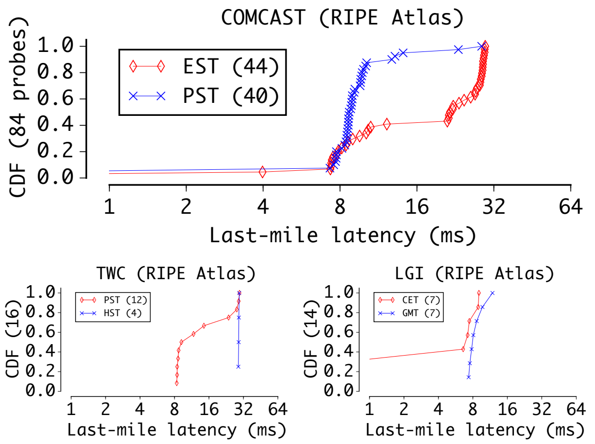

Figure 10: Last-mile latencies separated by timezone. Comcast and TWC users experience considerably different last-mile latencies across the US east and west coast.

reduce latencies experienced on the last-mile. This analysis reveals that last-mile latency is *not* the same for all subscribers of a DSL ISP, but it differs by access technology used by the DSL modem.

5 CONCLUSION

We leveraged the RIPE Atlas and SamKnows platform to measure last-mile latency. This is the first study that has measured last-mile latencies on such a scale from within multiple service providers networks in the US and the EU. We showed that DSL service providers not only enable interleaving, but some providers dynamically adapt interleaving depth levels. We witnessed that last-mile latency is considerably stable over time and not affected by diurnal load patterns. Last-mile latencies for DSL center at ~16 ms, with cable at ~8 ms,

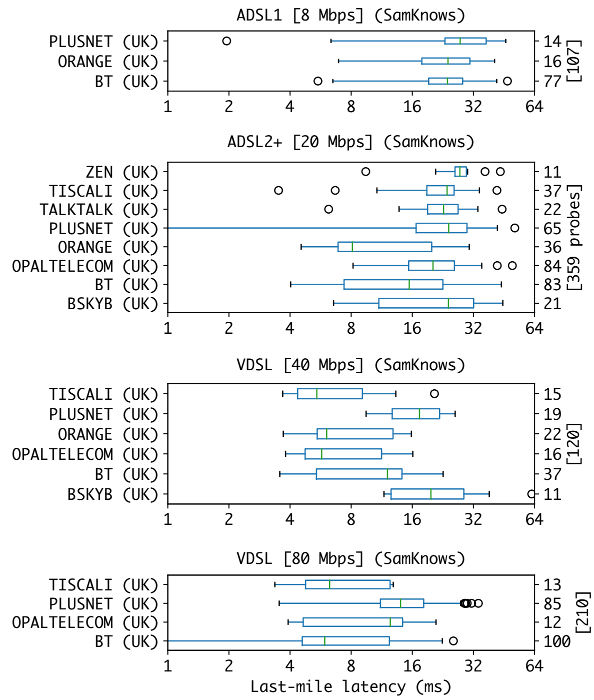

Figure 11: Box plots of last-mile latencies for DSL ISPs separated by broadband speeds. The last-mile latencies for ADSL/ADSL2+ > VDSL deployments.

and fibre deployments at ~4 ms. This observation will allow simulation studies to appropriately model DSL, cable and fibre links. We showed that last-mile latencies of a service provider can depend on the geographic location of a subscriber. We observed significant last-mile latency differences for US cable service providers across the east (centered at ~24 ms) and west (centered at ~8 ms) coast. Last-mile latencies of DSL deployments vary with broadband speeds. Last-mile latencies for VDSL are lower compared to that of ADSL1 and ADSL2+ broadband speeds.

This study extends our understanding of last-mile latency witnessed by home users. CDN providers that attempt to optimise content delivery towards the edge of the network will benefit from the identified characteristics of the last-mile. This work will also benefit ISPs since it promotes the possibility of caching popular content near to the home routers to further eliminate the bottlenecks induced by last-mile latency. This work serves as possible input for ongoing standardization efforts [21, 29] within the IETF that attempt to target operations in low latency modes. The methodology applied in this study is generally useful for broadband measurement studies [24, 27] using SamKnows and RIPE Atlas.

Reproducibility Considerations

The RIPE Atlas and SamKnows datasets are stored in a SQLite database (alongwith the SQL schemas) and released [2]. The software used in this study is also released [2]. This includes Jupyter

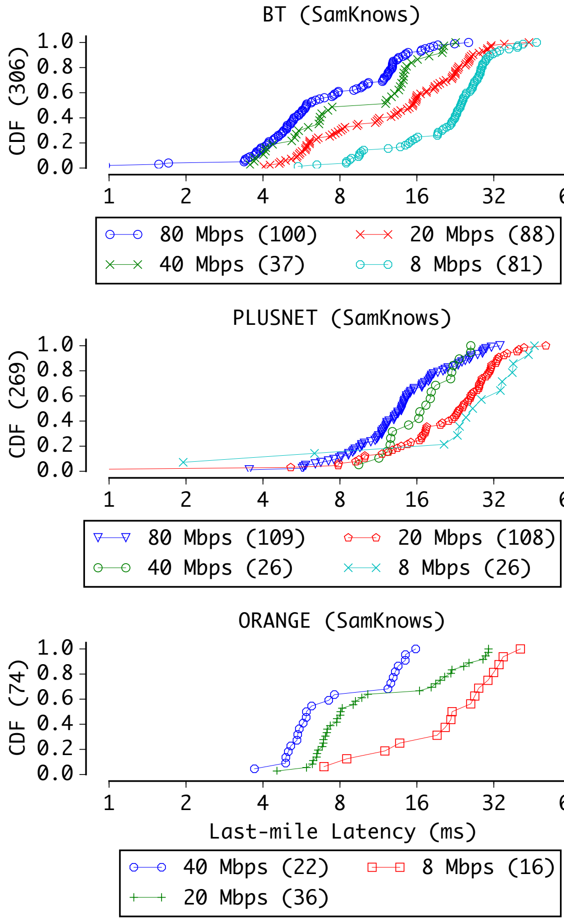

Figure 12: Last-mile latency decreases with increase in broadband speeds. Higher ADSL2+ and VDSL transmission rates help reduce last-mile latencies.

notebooks to provision measurements on RIPE Atlas, fetch measurements results, augment them using third-party datasets (such as PeeringDB [23] and RIPE stat APIs [26]) and compile datasets together into a SQLite database. The software used in the analysis to generate plots is also included. Guidance on how to repeat and reproduce these results is provided and reproducers are encouraged to contact the authors for further questions.

Acknowledgements

We like to thank Sam Crawford (SamKnows) for providing us `traceroute` datasets collected within UK access networks and giving us permission to release an anonymised version of the trace. We also thank Philip Eardley (BT) and Trevor Burbridge (BT) for reviewing our analysis. This work received funding from the European UnionâĂŹs Horizon 2020 research and innovation programme 2014-2018 under grant agreement No. 644866, Scalable and Secure Infrastructures for Cloud Operations (SSICLOPS). This work was also partially supported by the European CommunityâĂŹs Seventh Framework Programme in (FP7/2007-2013) Grant No. 317647 (Leone) and by Flamingo, a Network of Excellence project (ICT-318488) supported by the European Commission under its Seventh Framework Programme.

REFERENCES

[1] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. 2015. Lessons Learned From Using the RIPE Atlas Platform for Measurement Research. *Computer Communication Review* 45, 3 (2015), 35–42. https://doi.org/10.1145/2805789.2805796

[2] Vaibhav Bajpai, Steffie Jacob Eravuchira, and Jürgen Schönwälder. 2017. Dissecting Last-mile Latency Characteristics - Software and Dataset. (2017). Retrieved June 21, 2017 from https://github.com/vbajpai/2017-ccr-lastmile-analysis

[3] Vaibhav Bajpai, Mirja Kühlewind, Jörg Ott, Jürgen Schönwälder, Anna Sperotto, and Brian Trammell. 2017 (to appear). Challenges with Reproducibility. In *ACM SIGCOMM Reproducibility Workshop*.

[4] Vaibhav Bajpai and Jürgen Schönwälder. 2015. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *IEEE Communications Surveys and Tutorials* 17, 3 (2015), 1313–1341. https://doi.org/10.1109/COMST.2015.2418435

[5] M. Belshe, R. Peon, and M. Thomson. 2015. Hypertext Transfer Protocol Version 2 (HTTP/2). RFC 7540 (Proposed Standard). (May 2015). https://tools.ietf.org/html/rfc7540

[6] Zachary S. Bischof, John S. Otto, and Fabián E. Bustamante. 2012. Up, down and around the stack: ISP characterization from network intensive applications. *Computer Communication Review* 42, 4 (2012), 515–520. https://doi.org/10.1145/2377677.2377778

[7] Bob Briscoe, Anna Brunstrom, Andreas Petlund, David A. Hayes, David Ros, Ing-Jyh Tsang, Stein Gjessing, Gorry Fairhurst, Carsten Griwodz, and Michael Welzl. 2016. Reducing Internet Latency: A Survey of Techniques and Their Merits. *IEEE Communications Surveys and Tutorials* 18, 3 (2016), 2149–2196. https://doi.org/10.1109/COMST.2014.2375213

[8] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John S. Heidemann, and Ramesh Govindan. 2013. Mapping the Expansion of Google's Serving Infrastructure. In *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23-25, 2013*. 313–326. https://doi.org/10.1145/2504730.2504754

[9] Igor Canadi, Paul Barford, and Joel Sommers. 2012. Revisiting Broadband Performance. In *Proceedings of the 12th ACM SIGCOMM Internet Measurement Conference, IMC '12, Boston, MA, USA, November 14-16, 2012*. 273–286. https://doi.org/10.1145/2398776.2398805

[10] Fangfei Chen, Ramesh K. Sitaraman, and Marcelo Torres. 2015. End-User Mapping: Next Generation Request Routing for Content Delivery. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication, SIGCOMM 2015, London, United Kingdom, August 17-21, 2015*. 167–181. https://doi.org/10.1145/2785956.2787500

[11] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. 2015. Are We One Hop Away from a Better Internet?. In *Proceedings of the 2015 ACM Internet Measurement Conference, IMC 2015, Tokyo, Japan, October 28-30, 2015*. 523–529. https://doi.org/10.1145/2815675.2815719

[12] Brent N. Chun, David E. Culler, Timothy Roscoe, Andy C. Bavier, Larry L. Peterson, Mike Wawrzoniak, and Mic Bowman. 2003. PlanetLab: An Overlay Testbed for Broad-Coverage Services. *Computer Communication Review* 33, 3 (2003), 3–12. https://doi.org/10.1145/956993.956995

[13] Lucas DiCioccio, Renata Teixeira, Martin May, and Christian Kreibich. 2012. Probe and Pray: Using UPnP for Home Network Measurements. In *Passive and Active Measurement - 13th International Conference, PAM 2012, Vienna, Austria, March 12-14th, 2012. Proceedings*. 96–105. https://doi.org/10.1007/978-3-642-28537-0_10

[14] X. Dimitropoulos, D. Krioukov, G. Riley, and k. claffy. 2006. Revealing the Autonomous System Taxonomy: The Machine Learning Approach. In *Passive and Active Measurement Conference, PAM 2006*. http://arxiv.org/abs/cs/0604015

[15] Marcel Dischinger, Andreas Haeberlen, P. Krishna Gummadi, and Stefan Saroiu. 2007. Characterizing Residential Broadband Networks. In *Proceedings of the 7th ACM SIGCOMM Internet Measurement Conference, IMC 2007, San Diego, California, USA, October 24-26, 2007*. 43–56. https://doi.org/10.1145/1298306.1298313

[16] Constantine Dovrolis, P. Krishna Gummadi, Aleksandar Kuzmanovic, and Sascha D. Meinrath. 2010. Measurement Lab: Overview and an Invitation to the Research Community. *Computer Communication Review* 40, 3, 53–56. https://doi.org/10.1145/1823844.1823853

[17] Xun Fan, Ethan Katz-Bassett, and John S. Heidemann. 2015. Assessing Affinity Between Users and CDN Sites. In *Traffic Monitoring and Analysis - 7th International Workshop, TMA 2015, Barcelona, Spain, April 21-24, 2015. Proceedings*. 95–110. https://doi.org/10.1007/978-3-319-17172-2_7

[18] Daniel Genin and Jolene Splett. 2013. Where in the Internet is congestion? *CoRR* abs/1307.3696 (2013). http://arxiv.org/abs/1307.3696

[19] Jim Gettys and Kathleen M. Nichols. 2012. Bufferbloat: Dark Buffers in the Internet. *Communications of the ACM* 55, 1 (2012), 57–65. https://doi.org/10.1145/2063176.2063196

[20] Thomas Holterbach, Cristel Pelsser, Randy Bush, and Laurent Vanbever. 2015. Quantifying Interference between Measurements on the RIPE Atlas Platform. In *Proceedings of the 2015 ACM Internet Measurement Conference, IMC 2015, Tokyo, Japan, 2015*. 437–443. https://doi.org/10.1145/2815675.2815710

[21] Janardhan Iyengar and Martin Thomson. 2017. *QUIC: A UDP-Based Multiplexed and Secure Transport*. Internet-Draft. Internet Engineering Task Force. http://tools.ietf.org/html/draft-ietf-quic-transport-04 Work in Progress.

[22] Robert Kisteleki. 2015. RIPE Atlas Update: HTTP Measurements, CLI Tools, DomainMON. (2015). Retrieved June 04, 2017 from https://labs.ripe.net/Members/kistel/ripe-atlas-update-http-measurements-cli-tools-domainmon-and-more

[23] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, and kc claffy. 2014. Using PeeringDB to Understand the Peering Ecosystem. *Computer Communication Review* 44, 2 (2014), 20–27. https://doi.org/10.1145/2602204.2602208

[24] Andra Lutu, Marcelo Bagnulo, Amogh Dhamdhere, and Kimberly C. Claffy. 2016. NAT Revelio: Detecting NAT444 in the ISP. In *Passive and Active Measurement, PAM 2016, Heraklion, Greece, March 31 - April 1, 2016. Proceedings*. 149–161. https://doi.org/10.1007/978-3-319-30505-9_12

[25] RIPE NCC. 2017. RIPE Atlas Probe API: v2. (2017). Retrieved June 04, 2017 from https://atlas.ripe.net/api/v2/probes

[26] RIPE NCC. 2017. RIPE Stat API. (2017). Retrieved June 04, 2017 from https://stat.ripe.net

[27] Cristel Pelsser, Luca Cittadini, Stefano Vissicchio, and Randy Bush. 2013. From Paris to Tokyo: On the Suitability of ping to Measure Latency. In *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23-25, 2013*. 427–432. https://doi.org/10.1145/2504730.2504765

[28] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear. 1996. Address Allocation for Private Internets. RFC 1918. (Feb. 1996). https://tools.ietf.org/html/rfc1918

[29] Eric Rescorla. 2017. *The Transport Layer Security (TLS) Protocol Version 1.3*. Internet-Draft. IETF. http://tools.ietf.org/html/draft-ietf-tls-tls13-20

[30] Aaron Schulman and Neil Spring. 2011. Pingin' in the rain. In *Proceedings of the 11th ACM SIGCOMM Internet Measurement Conference, IMC '11, Berlin, Germany, 2011*. 19–28. https://doi.org/10.1145/2068816.2068819

[31] Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter de Donato. 2014. BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks. In *2014 USENIX Annual Technical Conference, USENIX ATC '14, Philadelphia, PA, USA, June 19-20, 2014*. 383–394. https://www.usenix.org/conference/atc14/technical-sessions/presentation/sundaresan

[32] Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescapè. 2011. Broadband Internet Performance: A View From the Gateway. In *Proceedings of the ACM SIGCOMM 2011 Conference, Toronto, ON, Canada, 2011*. https://doi.org/10.1145/2018436.2018452

[33] Srikanth Sundaresan, Nick Feamster, Renata Teixeira, and Nazanin Magharei. 2013. Measuring and Mitigating Web Performance Bottlenecks in Broadband Access Networks. In *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23-25, 2013*. 213–226. https://doi.org/10.1145/2504730.2504741