



Revisiting the Privacy Implications of Two-Way Internet Latency Data

Brian Trammell^(✉)  and Mirja Kühlewind

Networked Systems Group, ETH Zurich, Zurich, Switzerland
trammell@tik.ee.ethz.ch

Abstract. The Internet measurement community is increasingly sensitive to the privacy implications of both active and passive measurement. Research into the drawbacks of network data anonymization has led the community to investigate data sharing techniques, as well as to focus on active measurements and active measurement datasets. A key metric in these datasets is round-trip-time (RTT) as measured e.g. by `ping` or `traceroute`. This paper examines the assumption that the analysis of Internet RTT data is safe for open research by posing the question: what potentially-private inferences can be made about a remote target given periodic latency measurements from known vantage points under one's control? We explore the risks to end-user privacy both through a review of diverse literature touching on the subject as well as on the analysis of RTT data from fixed and mobile Internet measurement infrastructure. While we find that the common assumption of safety generally holds, we explore caveats and give recommendations for mitigation in those cases where it may not.

1 Introduction

The Internet measurement research community has long been concerned with the privacy impact of its measurements on the end users of the Internet. The personally-identifiable nature of IP addresses, for example, as it is linkable to end-user activity, is well-understood, and even a subject of current regulation¹, with a body of literature on anonymization techniques [1] and the effectiveness thereof [2] to protect this information. However, other information that can be gleaned from passive observation of traffic at multiple layers [3] can be used to track end users as well. Encryption of application-layer payload does not necessarily provide protection from tracking [4]. As the inferences that we make as Internet measurement researchers are inextricably related to the inferences necessary to perform user tracking, ethical standards are necessary to minimize end-user harm [5].

There is, however, a common understanding that certain types of data are safer than others. Simple round-trip time or two-way delay information between

¹ e.g. the European General Data Protection Regulation (GDPR); see <http://www.eugdpr.org>.

two infrastructure addresses, for example as widely used in diagnostics and operations using `ping` or `traceroute`, and as publicly available at scale via active measurement platforms such as RIPE Atlas², is taken to be unthreatening. Even a latency time series gleaned from user traffic says more about the dynamics of the network paths that traffic took than anything about the user’s behavior.

In this paper, we examine that assumption by considering the components of end-to-end round trip time, defining possible threat models for RTT privacy and evaluating the utility of latency data for the defined attackers. There are two broad concerns here. First, since RTT is related to distance, RTT measurements from a set of distributed vantage points could be used to determine the location of an endpoint and its associated end-user. Second, since RTT has a component of far-end delay, RTT measurements over a period of time could be used to glean information about the relative level of activity on some remote endpoint. Depending on the resolution of this information, different inferences could be made. Even low-resolution information from a home network could be used to guess whether someone is at home during a given time period, for example. We examine both of these concerns in this work.

We conclude that RTT information is generally safe to use, but should be treated as sensitive in specific circumstances, and provide guidance to mitigate privacy risk when handling this data in Sect. 5.

This paper is, in part, an answer to a related question raised in the IETF QUIC working group. As QUIC’s transport layer headers are encrypted, passive RTT measurement as available with TCP [6] is not available in QUIC. A proposal to add explicit RTT measurement to QUIC’s wire image in the spirit of IPIM [7] was met with concern that passive RTT measurement might pose a privacy risk. Though this paper is more concerned with active RTT measurement, its insights are applicable to passive measurement as well, with the caveat that an entity in position to perform passive traffic measurement is in a position to gain more information about a given target than a random endpoint in the Internet armed only with `ping`.

This paper is designed to be easily reproducible: the Jupyter notebooks used in the analyses in this paper are available online³, including code and/or instructions for the retrieval of the source data we used.

2 Components of End-to-End Latency

We begin with an examination of the components of end-to-end latency as can be observed at either endpoint of a transport-layer connection, the sender of an ICMP Echo Request, or the observer of a TCP flow with full information about sequence and acknowledgment numbers and timestamps in both directions of a flow. This observable RTT RTT_{obs} is given by Eq. 1, for f hops in one direction and r hops in the opposite direction, where D_{prop} is propagation delay on a link, D_{queue} is queueing delay at a forwarding node, D_{proc} is processing delay at a

² <https://atlas.ripe.net>.

³ <https://github.com/mami-project/rtt-privacy-paper>.

forwarding node, D_{stack} is stack delay at the remote endpoint (the time it takes for a packet to make it from the network interface to the application and back, including acknowledgment delay [8] when traffic is unidirectional), and D_{app} is application delay at that endpoint.

$$RTT_{obs} = \sum_{n=0}^f (D_{prop_{n \rightarrow n+1}} + D_{queue_n} + D_{proc_n}) + \sum_{m=0}^r (D_{prop_{m \rightarrow m+1}} + D_{queue_m} + D_{proc_m}) + D_{stack} + D_{app} \quad (1)$$

This equation illustrates the confounding effect of end-to-end RTT measurement, which we will explore in more detail later. Each potential threat to privacy uses only one component of delay measured in the observable RTT, but all components are mixed together in a given RTT sample. The challenge in exploiting this information is then to reduce the irrelevant components to a known constant. For example, in the geolocation case, the desired RTT would be (a) perfectly symmetric and (b) made up of only propagation delay (c) in a straight line between endpoints, which would allow a distance measurement as in Eq. 2, where $c_{internet}$ is the speed of light in the Internet, assuming a known and constant factor for refraction in optical fiber and/or propagation in other physical media. $dist$ is an inequality because even in an ideal case (c) does not hold: the light path following the great circle between two points and the light path actually followed by physical Internet infrastructure differ.

$$dist < \frac{\sum_{n=0}^f D_{prop_{n \rightarrow n+1}} + \sum_{m=0}^r D_{prop_{m \rightarrow m+1}}}{2} \times c_{internet} \quad (2)$$

On the flip side, if light distance could be known, and processing and queuing delay were zero, these terms could be subtracted out from yielding only stack and application delay, turning RTT observations into “load” observations as in Eq. 3.

$$load \propto D_{stack} + D_{app} \quad (3)$$

The utility of RTT measurements to various geolocation and activity fingerprinting tasks, then, is directly related to the separability of these terms. This is the question we address in the rest of this work.

3 Latency and Geoprivacy

We first examine the geoprivacy question. The threat model here is one of an attacker armed with RTT measurements between a target with unknown location and distributed vantage points with known location, who wants to know the location of the target with arbitrary accuracy.

There is a wide array of recent literature related to this subject. Much of this focuses on “exclusion” based approaches, which uses assumptions about $c_{internet}$ to successively determine where a node or endpoint with unknown location *cannot* be. For example, Cicalese et al. [9] used active RTT measurement to discover and heuristically geolocate anycast infrastructure in the Internet: IP addresses whose RTT-derived circles of exclusion from known vantage points do not overlap must be anycasted.

In any case, much of the literature focuses explicitly on improving the accuracy of latency-based geolocation techniques; i.e., on the attacker’s side of the question we pose. Indeed, this underpins the provision of location-based services. Latency has been used to improve IP geolocation accuracy [10–12], and uncover potential fraud [13].

A common theme here is that the vantage points must be optimally selected, since unwanted error terms in RTT_{obs} increase with distance, so accuracy depends highly on the distance of the vantage points to the targets. In the case of passive or opportunistic RTT measurement, one must instead be lucky to be able to observe low-latency paths to the target. Katz-Bassett et al. [14] showed that accuracy under 100 km was possible by augmenting delay measurements with known topology information, and Gueye et al. [15] extended this approach with the use of multilateration, bringing the median error distance down to 25 km for the region of Western Europe.

These proposals are based on previously proposed techniques: IDMaps [16], a multicast-based service for geolocating, and IP2Geo [17] that is based in the first step on the location information of the closest DNS resolver. Model-based approaches for predicting the distance between measured network nodes [18] further refine this, and have achieved a median errors on the order of 30 km [19, 20].

More recently a method to utilize crowd-sourcing has been proposed to use smart-phones as landmarks and leveraging their GPS and WiFi-based location information [21]. Their measurement shows an median error of several hundred kilometers, reflecting both the use of a mobile dataset (where D_{proc} is generally higher) and the variability of real-world data, as compared to the testbed-based measurements of earlier works. Unsurprisingly, this work also confirms that the accuracy highly depends of the distance of the selected landmark to the measurement target.

Other work in location-based services [22–25] focuses on locating nodes in a virtual coordinate system (VCS), as opposed to physical space, following the argument of Ratnasamy et al. [26] that such high accuracy location in physical space is unnecessary for common location-aware services such as content server selection. These approaches are ideal for providing selection of distributed services without necessarily enabling geolocation of endpoints.

Our work is also related to location attacks against low-latency anonymity networks such as TOR. Ries et al. [27] investigate how virtual network coordinate systems can be utilized for timing attacks and exploitation of timing information and also conclude that small changes in latency can have a high influence on the accuracy. The availability of large numbers of latency samples has been

shown to increase the success of such attacks as well [28,29], and these become even more powerful if timestamp information is available and samples can be correlated based on clock skew [30]. However, given their architectural peculiarity, observations about these networks do not translate well to the impact on privacy in the non-anonymized Internet. Therefore, we do not address the case of anonymity networks in this work.

3.1 Measurements with Atlas and MONROE

We revisit the question by examining latency measurements taken with the RIPE Atlas and MONROE [31] measurement platforms. Both provide us with latency measurements between vantage points with known location toward targets with known location: Atlas through its anchoring measurements, and MONROE via periodic pings toward the MONROE collection infrastructure.

Exclusion. We start by using Atlas anchoring measurements to attempt geolocation by exclusion, exploiting the inequality for $dist$, the observation that the RTT between two endpoints cannot be less than the speed of light in the medium of the Internet multiplied by twice the distance between those endpoints. We looked at all anchoring ICMP traceroute measurements on Monday 2 October 2017 to a set of 39 anchors, and filter out any measurements from probe-anchor pairs with reported locations less than 500 m from each other⁴. We assume that each of our RIPE Atlas Anchor targets is unicast. This yields a total of 9.61 million individual measurements over 22,072 probe-anchor pairs. We then took the minimum end-to-end RTT measurement for each probe-anchor pair, taking this to be the best measurement for exclusion purposes.

We then draw exclusion circles corresponding to each probe’s minimum RTT at that probe’s location, and examine the intersection of these circles. We note



(a) Glattdbrugg, Switzerland

(b) Nairobi, Kenya

(c) Brno, Czechia

Fig. 1. Exclusion circles around selected anchors (red dot) and associated MaxMind Geolite City geolocation result (green cross) (Color figure online)

⁴ Here, the reasoning is that such pairs are either colocated in the same rack, or possible connected to the same local- or metropolitan-area network, and as such do not accurately reflect Internet RTT measurement.

that for 35 of the 39 anchors, intersection gives no additional information; i.e. the closest probe’s exclusion circle is completely covered by that of the next closest probe. RTT location via exclusion is therefore largely a matter of luck of the location of the known vantage point. An illustration of this most common case is shown in Fig. 1(a). In the other cases, either the refinement to the exclusion area is insignificant, or the location estimate covers a large region with or without intersection. Figure 1(b) shows an example of this case; note here that the both the location estimate and IP geolocation yield national-scale results.

Though sometimes comparatively remote probes can refine each others’ exclusion circles, in no case did we find such a refinement resulting in a reasonably accurate location estimate: the uncertainty in RTT simply grows too quickly with distance. Figure 1(c) illustrates this. Here, estimates from Prague and Vienna yield an area roughly the size of Czechia, but do exclude Prague.

When the IP address of the target is known, IP address geolocation can also be used to estimate its location. We therefore attempt to geolocate each anchor based on its IPv4 address in the freely-available MaxMind GeoLite City database⁵, which we take as a worst-case IP geolocation result, noting that better IP geolocation databases will yield better results [32]. We compare the error between the geolocation result and the anchor’s declared location with the uncertainty circle for each probe. Here we find that only 140 of 22079 of our anchor-probe pairs have less uncertainty than IP geolocation error, and only 14 of the 39 anchors, generally in areas with a very high probe density, have a measurement from at least one such probe. This further underscores the role of luck in vantage point selection.

We take the RIPE Atlas anchoring measurement dataset to be representative of Internet RTT measurements. Given that opportunities for location area reduction by intersection are not significant in this dataset, we now make a simplifying assumption that the best estimate for the location of an anchor is the center of the uncertainty circle of the probe with the lowest minimum RTT, and therefore that the error in the best estimate is simply the distance from that probe to the anchor. Median error in the Atlas dataset is 39 km while the median IP geolocation error is 16 km. We note that even though our methodology is far simpler than those described in the literature, it achieves comparable accuracy, underscoring the finding that skill (or luck) in vantage point placement is the dominant factor in accuracy in geolocation by exclusion using RTT measurements.

Atlas measurements are largely from residential or infrastructure networks toward infrastructure networks. Recent work by Bajpai et al. [33] shows our findings also to be applicable to the location of residential subscribers. This analysis of Atlas and SamKnows measurements of last-mile latency finds latency to depend on provider, technology, and point of presence, with median (two-way) latencies per provider between 5 and 20ms. Last-mile latency is therefore responsible, on its own, for an exclusion radius between about 500 km and 2000 km.

⁵ As retrieved from <https://stat.ripe.net> on 10 October 2017.

We therefore take the location of residential endpoints to be more challenging than location by exclusion of Atlas anchors.

Note that while landmark selection is a challenge for active measurement, when RTT information is observed passively, e.g. during a transport or application-layer handshake, or using passive TCP measurement [6], the increased flatness of the Internet topology [34] implies that there is a decent chance to observe active communications between a client and a nearby content server.

Linear Distance Modeling. We also attempted trilateration through the creation of a linear model relating RTT to distance; i.e. $dist_{est} = f(RTT_{obs})$, based both on Atlas and MONROE measurements. The linear models we derived from our measurements (Atlas: $RTT = 0.0190 \times dist + 22.317$ with $r = 0.86$; MONROE⁶: $RTT = 0.0154 \times dist + 37.0735$ with $r = 0.78$, for RTT in milliseconds and distance in kilometers) are too imprecise to use as a basis for trilateration, with variance and last-mile latency making distance estimation even less feasible on mobile networks.

However, in examining the absolute (Fig. 2(a)) and relative (Fig. 2(b)) error in these models, a guideline for using RTT measurement for location estimation emerges. Restricting RTT data in ways that are possible using only simple inference or measurement can lead to better models with less error. Figure 2 also shows error results for models based on subsets of the Atlas RTT data, considering only pairs with a minimum RTT less than 50 ms, or considering only short paths (with less than 6 hops).

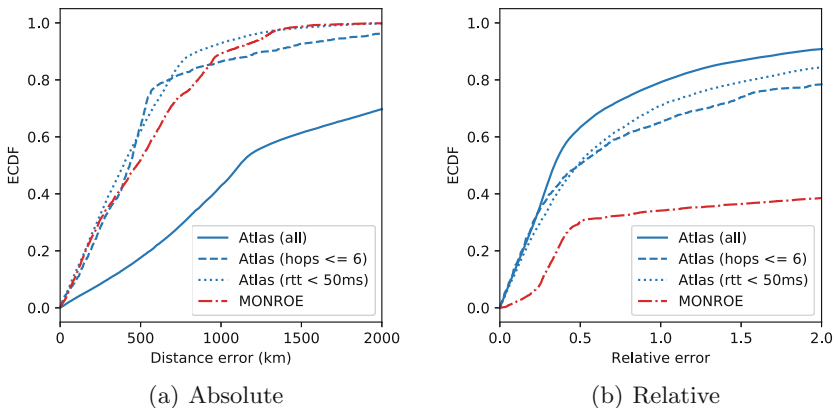


Fig. 2. Distance error for linear models

⁶ MONROE nodes provide GPS metadata for mobile nodes for location ground truth. We split MONROE data from 1 September 2017 into 5 min bins (300 pings) and associated the geographic average GPS location with the minimum RTT in each bin to yield 3,863 samples from 45 nodes.

4 Load Telemetry

RTT measurement is of interest to in-network operations precisely because it can be used to gain insight into the functioning and malfunctioning of network devices. An unusual D_{queue_n} or D_{proc_n} is often indicative of a fault to be corrected. This is the basic insight behind the measurement of in-network buffering performed by Netalyzr [35] and other measurement platforms, and has been used more recently in the inference of congestion at network interconnects [36]. Load on the endpoint can also be visible in RTT measurements, as shown by Holterbach et al. [37], who showed in a study of the load dependent accuracy of the Atlas platform that several milliseconds of RTT error could be induced and measured by varying the load on RIPE Atlas probes.

This utility, however, has a flipside, as it necessarily exposes information about D_{queue_n} or D_{proc_n} to any device on path which can use active or passive measurement of RTT. More precisely, we now consider a threat model where the attacker knows an IP address associated with a given target, and wants to estimate activity on that target’s network. Here, the attacker can leverage four assumptions about common characteristics of residential access networks to successfully determine activity on a residential customer’s network:

- The access link is usually the bottleneck link for residential access, so latency variation is due to D_{queue} on the two directions of this link.
- Residential access links are frequently “bufferbloomed” [35, 38]; i.e., modems have overdimensioned buffers that lead to high D_{queue} under load.
- In many markets, a single customer has a single public IP address at a single point in time, so an ICMP ping to a given address will traverse the access link. Note that this assumption does not hold when carrier-grade NAT is used to conserve IP addresses [39].
- ICMP packets, if not blocked, will generally share a queue with other packets, and can therefore be used to measure D_{queue} induced by other traffic.

In other words, remote buffer measurement is possible, and can be used to infer activity on residential networks. We ran a simple proof-of-concept activity inference attack against one of the authors’ home networks to illustrate this point.

We pinged an author’s public IP address in Zurich ten times a second from vantage points in Amsterdam and Singapore, while subjecting the inbound network link to varying load through TCP downloads using `curl` at various rate limits. The results are shown in Fig. 3. Idle and active periods are clearly visible due to RTT change from baseline from both the near and far vantage point, even down to a 300 kB/s downstream rate limit, one tenth of the capacity of the link. This author’s home access network is also home to a RIPE Atlas probe. Indeed, examining the 5-minute RTT series from second-hop latency measurements from this probe on a different day show clear diurnal peaks in maximum RTT measured during morning and evening weekday network activity, indicating that load telemetry is even possible with very limited RTT data.




Fig. 3. RTT time series from Amsterdam and Singapore toward a typical residential cable access network in Zürich with a 30 Mbit download link, during downloads from that network with various rate limits.

The success of load telemetry via RTT data is dependent on each of these assumptions holding. Indeed, in a similar test against a virtualized endpoint in a datacenter, where the bufferbloat assumption does not hold, we saw no significant difference in RTT regardless of generated load.

We have made an automated version of this load telemetry measurement available as an online tool⁷. Initial analysis of data collected during beta testing of this tool from 26 access networks shows that remote load telemetry is possible on a minority of examined networks: 13 (50%) on which ICMP ping is always blocked, 9 (35%) which are pingable but on which there is no apparent correlation between RTT and load, and 4 (15%) on which RTT is correlated with load.

The ability to perform load telemetry of a remote network using RTT data illustrates the well-known utility of RTT data to queue delay measurements. However, variance in queueing delay is often indicative of other kinds of activity, indicating that RTT data may be privacy-sensitive in limited circumstances.

5 Conclusions and Recommendations

We have confirmed the network operations rule of thumb that 1ms of RTT is 100 km of distance, and the findings of previous studies that RTT measurement can provide location accuracy on the order of 30 km to 100 km. That a very basic exclusion-based methodology can perform as well as techniques with higher complexity using publicly available datasets shows that luck in vantage point selection is the dominant factor in accuracy. The sensitivity of RTT measurements for geoprivacy is therefore related to the minimum RTT represented by those measurements. However, RTT measurement is less accurate than IP geolocation using even the most basic, free databases. We therefore recommend care in dissemination of RTT measurement datasets in those cases where the datasets themselves are dominated by samples on the order of less than 10ms, and/or where one (but not both) IP addresses are anonymized.

As for load telemetry, the ability for RTT measurements to provide insight into network queueing delay can be used to infer human activity on networks

⁷ <https://pingme.pto.mami-project.eu>.

where certain assumptions hold. While the ongoing reduction of bufferbloat in residential access networks will mitigate the utility of RTT measurements for the inference of residential activity long equipment replacement cycles in these networks mean that bufferbloat will be with us for some time to come. In any case, high-resolution, long-duration RTT datasets collected from networks where bufferbloat is likely should be treated with care.

Acknowledgments. Many thanks to RIPE for making Atlas available to the research community, and to the MONROE project for access to the mobile dataset used in this work. Thanks to the anonymous reviewers and our shepherd, Ramakrishna Padmanabhan, for comments improving the organization and focus of this paper. Thanks also to the members of the IETF QUIC Working Group RTT Design Team for the discussions leading to this paper. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 688421, and was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0268. The opinions expressed and arguments employed reflect only the authors’ views, and not those of the European Commission or the Swiss Government.

References

1. Moore, D., Claffy, K.C: Summary of anonymization best practice techniques, December 2008. <https://www.caida.org/projects/predict/anonymization/>
2. Burkhart, M., Schatzmann, D., Trammell, B., Boschi, E., Plattner, B.: The role of network trace anonymization under attack. *SIGCOMM Comput. Commun. Rev.* **40**(1), 5–11 (2010)
3. Coull, S., Wright, C., Monroe, F., Collins, M., Reiter, M.: Playing devil’s advocate: inferring sensitive information from anonymized network traces. In: Proceedings of the 14th Annual Network and Distributed Systems Security Symposium, San Diego, CA, USA (2007)
4. Barnes, R., Schneier, B., Jennings, C., Hardie, T., Trammell, B., Huitema, C., Borkmann, D.: Confidentiality in the face of pervasive surveillance: a threat model and problem statement. RFC 7624, RFC Editor, August 2015
5. Partridge, C., Allman, M.: Ethical considerations in network measurement papers. *Commun. ACM* **59**(10), 58–64 (2016)
6. Strowes, S.D.: Passively measuring TCP round-trip times. *Commun. ACM* **56**(10), 57–64 (2013)
7. Allman, M., Beverly, R., Trammell, B.: Principles for measurability in protocol design. *SIGCOMM Comput. Commun. Rev.* **47**(2), 2–12 (2017)
8. Ding, H., Rabinovich, M.: TCP stretch acknowledgements and timestamps: findings and implications for passive RTT measurement. *SIGCOMM Comput. Commun. Rev.* **45**(3), 20–27 (2015)
9. Cicalese, D., Joumblatt, D.Z., Rossi, D., Buob, M.O., Augé, J., Friedman, T.: Latency-based anycast geolocation: algorithms, software, and data sets. *IEEE J. Sel. Areas Commun.* **34**(6), 1889–1903 (2016)
10. Grey, M., Schatz, D., Rossberg, M., Schaefer, G.: Towards distributed geolocation by employing a delay-based optimization scheme. In: 2014 IEEE Symposium on Computers and Communications (ISCC), pp. 1–7, June 2014

11. Hillmann, P., Stiemert, L., Rodosek, G.D., Rose, O.: Dragoon: advanced modelling of IP geolocation by use of latency measurements. In: 2015 10th International Conference for Internet Technology and Secured Transactions (ICITST), pp. 438–445, December 2015
12. Wang, Z., Mark, B.L.: Robust statistical geolocation of Internet hosts. In: 2015 IEEE Globecom Workshops (GC Wkshps), pp. 1–6, December 2015
13. Abdou, A., Matrawy, A., van Oorschot, P.C.: CPV: delay-based location verification for the internet. *IEEE Trans. Dependable Secure Comput.* **14**(2), 130–144 (2017)
14. Katz-Bassett, E., John, J.P., Krishnamurthy, A., Wetherall, D., Anderson, T., Chawathe, Y.: Towards IP geolocation using delay and topology measurements. In: Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement, IMC 2006, pp. 71–84. ACM, New York (2006)
15. Gueye, B., Ziviani, A., Crovella, M., Fdida, S.: Constraint-based geolocation of internet hosts. *IEEE/ACM Trans. Networking* **14**(6), 1219–1232 (2006)
16. Francis, P., Jamin, S., Jin, C., Jin, Y., Paxson, V., Raz, D., Shavitt, Y., Zhang, L.: IDMaps: a global Internet host distance estimation service. In: Proceedings of IEEE INFOCOM, pp. 210–217 (2000)
17. Padmanabhan, V.N., Subramanian, L.: An investigation of geographic mapping techniques for internet hosts. *SIGCOMM Comput. Commun. Rev.* **31**(4), 173–185 (2001)
18. Laki, S., Mátray, P., Hága, P., Csabai, I., Vattay, G.: A model based approach for improving router geolocation. *Comput. Netw.* **54**(9), 1490–1501 (2010)
19. Wong, B., Stoyanov, I., Sirer, E.G.: Geolocalization on the internet through constraint satisfaction. In: Proceedings of the 3rd Conference on USENIX Workshop on Real, Large Distributed Systems, WORLDS 2006, vol. 3, p. 1. USENIX Association, Berkeley, (2006)
20. Dong, Z., Perera, R.D., Chandramouli, R., Subbalakshmi, K.: Network measurement based modeling and optimization for IP geolocation. *Comput. Netw.* **56**(1), 85–98 (2012)
21. Ciavarrini, G., Luconi, V., Vecchio, A.: Smartphone-based geolocation of internet hosts. *Comput. Netw.* **116**(Supplement C), 22–32 (2017)
22. Ng, T.S.E., Zhang, H.: Global network positioning: a new approach to network distance prediction. *SIGCOMM Comput. Commun. Rev.* **32**(1), 73–73 (2002)
23. Dabek, F., Cox, R., Kaashoek, F., Morris, R.: Vivaldi: a decentralized network coordinate system. *SIGCOMM Comput. Commun. Rev.* **34**(4), 15–26 (2004)
24. Chen, Y., Xiong, Y., Shi, X., Deng, B., Li, X.: Pharos: a decentralized and hierarchical network coordinate system for Internet distance prediction. In: IEEE GLOBECOM 2007 - IEEE Global Telecommunications Conference, pp. 421–426, November 2007
25. Lim, H., Hou, J.C., Choi, C.H.: Constructing internet coordinate system based on delay measurement. *IEEE/ACM Trans. Networking* **13**(3), 513–525 (2005)
26. Ratnasamy, S., Handley, M., Karp, R., Shenker, S.: Topologically-aware overlay construction and server selection. In: Proceedings, Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies, vol. 3, pp. 1190–1199 (2002)
27. Ries, T., State, R., Engel, T.: Measuring anonymity using network coordinate systems. In: 2011 11th International Symposium on Communications Information Technologies (ISCIT), pp. 366–371, October 2011
28. Hopper, N., Vasserman, E.Y., Chan-Tin, E.: How much anonymity does network latency leak? *ACM Trans. Inf. Syst. Secur.* **13**(2), 13:1–13:28 (2010)

29. Serjantov, A., Sewell, P.: Passive attack analysis for connection-based anonymity systems. In: Sneekenes, E., Gollmann, D. (eds.) ESORICS 2003. LNCS, vol. 2808, pp. 116–131. Springer, Heidelberg (2003). https://doi.org/10.1007/978-3-540-39650-5_7
30. Murdoch, S.J.: Hot or not: revealing hidden services by their clock skew. In: Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, pp. 27–36. ACM, New York (2006)
31. Alay, O., Lutu, A., Garcia, R., Peon-Quiros, M., Mancuso, V., Hirsch, T., Dely, T., Werme, J., Evensen, K., Hansen, A., Alfredsson, S., Karlsson, J., Brunstrom, A., Khatouni, A.S., Mellia, M., Marsan, M.A., Monno, R., Lonsethagen, H.: Measuring and assessing mobile broadband networks with MONROE. In: 2016 IEEE 17th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM), pp. 1–3, June 2016
32. Gharaibeh, M., Shah, A., Huffaker, B., Zhang, H., Ensafi, R., Papadopoulos, C.: A look at router geolocation in public and commercial databases. In: Internet Measurement Conference (IMC), November 2017
33. Bajpai, V., Eravuchira, S.J., Schönwälder, J.: Dissecting last-mile latency characteristics. SIGCOMM Comput. Commun. Rev. **47**(5), 25–34 (2017)
34. Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., Willinger, W.: Anatomy of a large European IXP. In: Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication. SIGCOMM 2012, pp. 163–174. ACM, Helsinki (2012)
35. Kreibich, C., Weaver, N., Nechaev, B., Paxson, V.: Netalyzer: illuminating the edge network. In: Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, IMC 2010, Melbourne, Australia, pp. 246–259 (2010)
36. Luckie, M., Dhamdhere, A., Clark, D., Huffaker, B., claffy, k.: Challenges in inferring Internet interdomain congestion. In: Proceedings of the 2014 Conference on Internet Measurement Conference, IMC 2014, Vancouver, BC, Canada, pp. 15–22 (2014)
37. Holterbach, T., Pelsser, C., Bush, R., Vanbever, L.: Quantifying interference between measurements on the RIPE Atlas platform. In: Proceedings of the 2015 Internet Measurement Conference, IMC 2015, Tokyo, Japan, pp. 437–443. ACM (2015)
38. Gettys, J., Nichols, K.: Bufferbloat: dark buffers in the internet. Queue **9**(11), 40:40–40:54 (2011)
39. Lutu, A., Bagnulo, M., Dhamdhere, A., Claffy, K.C.: NAT revelio: detecting NAT444 in the ISP. In: Karagiannis, T., Dimitropoulos, X. (eds.) PAM 2016. LNCS, vol. 9631, pp. 149–161. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-30505-9_12