



HAL
open science

Replication: Towards a Publicly Available Internet scale IP Geolocation Dataset

Omar Darwich, Hugo Rimlinger, Milo Dreyfus, Matthieu Gouel, Kevin
Vermeulen

► To cite this version:

Omar Darwich, Hugo Rimlinger, Milo Dreyfus, Matthieu Gouel, Kevin Vermeulen. Replication: Towards a Publicly Available Internet scale IP Geolocation Dataset. ACM Internet Measurement Conference (IMC 2023), ACM, Oct 2023, Montréal, Canada. 10.1145/3618257.3624801 . hal-04215113v2

HAL Id: hal-04215113

<https://hal.science/hal-04215113v2>

Submitted on 25 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Replication: Towards a Publicly Available Internet Scale IP Geolocation Dataset

Omar Darwich
LAAS-CNRS, Université de Toulouse,
CNRS, Toulouse, France

Hugo Rimlinger
Sorbonne Université

Milo Dreyfus
Sorbonne Université

Matthieu Gouel
Sorbonne Université

Kevin Vermeulen
LAAS-CNRS, Université de Toulouse,
CNRS, Toulouse, France

ABSTRACT

IP geolocation is one of the most widely used forms of metadata for IP addresses, and despite almost twenty years of effort from the research community, the reality is that there is no accurate, complete, up-to-date, and explainable publicly available dataset for IP geolocation. We argue that a central reason for this state of affairs is the impressive results from prior publications, both in terms of accuracy and coverage: up to street level accuracy and locating millions of IP addresses with a few hundred vantage points in months. We believe the community would substantially benefit from a public baseline dataset and code. To encourage future research in IP geolocation, we replicate two geolocation techniques and evaluate their accuracy and coverage. We show that we can neither use the first technique to obtain the previously claimed street level accuracy, nor the second to geolocate millions of IP addresses on today's Internet and with publicly available measurement infrastructure. In addition to this reappraisal, we re-evaluate the fundamental insights that led to these prior results, as well as provide new insights and recommendations to help the design of future geolocation techniques. All of our code and data are publicly available to support reproducibility.

CCS CONCEPTS

• **Networks** → **Network measurement**;

KEYWORDS

Active Internet Measurements; IP geolocation; Replicability

ACM Reference Format:

Omar Darwich, Hugo Rimlinger, Milo Dreyfus, Matthieu Gouel, and Kevin Vermeulen. 2023. Replication: Towards a Publicly Available Internet Scale IP Geolocation Dataset. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC '23)*, October 24–26, 2023, Montreal, QC, Canada. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3618257.3624801>

1 INTRODUCTION

IP geolocation is one of the most widely used forms of metadata for IP addresses to help Internet measurements [29], and it supports

many network research applications, such as path troubleshooting [20, 22, 43], path prediction [18], or cybersecurity [19], but also commercial applications, such as geowalls, location based advertising, or fraud detection. But despite almost twenty years of effort from the research community, we have to face the reality: there is no accurate, complete, and explainable publicly available dataset for IP geolocation, and we are still far from the objective. The only large scale publicly available datasets are ITDK [16], RIPE IPMap [40], which are research efforts to provide the geolocation of router IP addresses of the Internet topology, and commercial databases. If they all provide a city level accuracy, which should satisfy the need of most applications [17, 18, 28, 47], none of these datasets fill the three criteria. On one hand, ITDK and RIPE IPMap focus on router IP addresses, so they are not complete. On the other hand, the commercial databases are not explainable, and prior work showed that the provided city level accuracy was overclaimed [26, 39].

We argue that one of the central reasons why we are far from our objective is that a significant body of prior work has been published in IP geolocation, and that this is a disincentive for producing new work, especially when, among this prior work, there are papers that appear to have solved two major problems: accuracy, up to street level [46], and coverage, with nearly geolocating the full range of IP addresses in a few months [32]. Wang *et al.* [46] uses a combination of latency measurements and locally hosted websites that serve as landmarks to geolocate IP addresses at street level, while Hu *et al.* [32] provide a methodology to scale classic latency-based geolocation techniques [31], by showing that a small subset of ten close vantage points to a target can obtain almost the same performance as hundreds.

A new technique should therefore compare its performance to those techniques in order to prove that it provides a significant improvement. As these techniques came with no publicly available code, we cannot easily evaluate or compare them. Further, as ten year old techniques, the results may not hold in the current (or future) Internet, motivating our imperative to produce code that provides up-to-date baselines to re-evaluate the techniques. To encourage future research on IP geolocation, we propose in this paper to fill this gap and replicate the two papers cited above [32, 46]. Our work makes several contributions:

- A new baseline of performance for the two replicated techniques for future geolocation techniques to compare against.
- The reappraisal of the fundamental insights that led to prior results.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

IMC '23, October 24–26, 2023, Montreal, QC, Canada

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0382-9/23/10...\$15.00

<https://doi.org/10.1145/3618257.3624801>

- A set of recommendations for future research on IP geolocation based on both old and new insights found in our study.
- All of our code and data that is publicly available to facilitate reproducibility at the following url: <https://github.com/dioptra-io/geoloc-imc-2023>

Our main results are: For the paper of Hu *et al.* [32], we are unable to replicate the geolocation of millions of IP addresses in a few months: To obtain this result, Hu *et al.* designed a vantage point selection algorithm to reduce the number of vantage points necessary to geolocate a target: To find which vantage points to use for a target, the algorithm first probes representative IP addresses that are likely to be close to the target (*i.e.*, in the target's /24 prefix) from all the vantage points, and then use the subset of vantage points with low latency to the representatives to probe the target. We find that although the algorithm still works in theory, probing all the /24 from all the vantage points produces too much measurement overhead and hinders the possibility of deploying it on the current biggest publicly available measurement platform, RIPE Atlas. We propose an extension of the algorithm to reduce its overhead and find that our extension uses as little as 13.2% of the measurements used by the original algorithm while obtaining the same accuracy. Nonetheless, even if we are not able to replicate the main result of the paper, we find that the main hypothesis that lead to this result still holds: A few geographically close vantage points to the target and even a single one, can obtain the same accuracy as the full set of 10K RIPE Atlas vantage points, the biggest publicly available measurement platform in terms of geographic coverage and number of vantage points.

For the paper from Wang *et al.* [46], we find that the claimed street level geolocation technique performs only as well as the classic constrained based geolocation (CBG) technique [31] on our dataset, with a median error of 28 km, versus 29 km for CBG - far from the 690 m of the original paper. We show that we could not obtain the same results because the two insights of the original paper do not hold. First, whereas the authors found that there exist street level landmarks for most targets, we find that only 28% of our targets have a landmark located at less than 1 km away. This number provides an upper bound on the number of targets that we can geolocate at street level, as the final step of the geolocation is to map a target to a landmark. Second, they found that the relative order of the latency-based measured distances between the landmarks and the target was preserved in the physical distance. In other words, if a landmark has a lower latency-based distance to the target than another landmark, this landmark is actually physically closer to the target than the other landmark. If this idea seems sound, we could not replicate these results as we found that the measurements to compute the latency-based distances are often unusable.

Finally, in addition to the replication, we had the opportunity to share and discuss our results with a geolocation database company that outperformed CBG and the street level technique on our dataset. We demystify partly how this geolocation base works by sharing the main results of our discussion, explaining that it mainly uses publicly available data accessible to researchers, which is a step towards more explainability of these commercial databases.

Gathering all these results, we provide a new baseline that on our dataset of 723 RIPE Atlas anchors as targets, 73% of them could be

geolocated at city level and 11% of them at less than 1 km away from their real geolocation. We also provide a set of recommendations for future work in IP geolocation, including the challenges to use and improve latency-based and landmark-based techniques and insights to develop new techniques.

2 MOTIVATION

Our goal is to reinvigorate the community to pursue efforts toward an accurate, complete, and publicly available IP geolocation dataset and open methodology for the entire routable space. A reasonable granularity that would satisfy the community is a city level accuracy, as it is enough for most applications, such as knowing if the geolocation of a proxy is accurate [47], predicting the geographic path of a measurement [18], or geolocating an outage [20, 22], so our paper is written with this granularity in mind. Without this dataset, one needs to perform their own measurements, or rely on free or paid geolocation databases that are inscrutable black boxes: 58 papers published in network conferences used the MaxMind [4] geolocation database during the 2016-2020 period [29]. By replicating two ten year old techniques that previously achieved high performance, we hope to establish a proper baseline for future work in the area.

2.1 Which papers do we replicate?

We replicate the two papers, “Towards Geolocation of Millions of IP addresses”, published at IMC in 2012 [32], and “Towards Street-Level Client-Independent IP geolocation”, published at NSDI in 2011 [46]. They are called million scale and street level papers from now. In the million scale paper, the authors were able to geolocate 35% of the IPv4 space in a few months, while in the street level paper, the median error of the technique was 690 meters. With these results, it is hard to propose a major contribution in the domain, and one can be discouraged to invest efforts in the field, as publishing in a major venue is hard considering these prior works.

Our work focuses on IPv4: For the million scale paper, the vantage point selection algorithm relies on finding representative IP addresses in every /24 prefix to use fewer vantage points to geolocate an IP address (§3.1). Due to the sparsity of the IPv6 space, finding representatives of a target is not straightforward. For the street level paper, there is no fundamental reason why the technique cannot be adapted. We leave the replication of these papers in IPv6 as future work.

2.2 What does the community stand to learn from the replication?

2.2.1 A new baseline for comparing geolocation techniques. We want to provide a new baseline of the results obtained with the two techniques in the current Internet with publicly available measurement platforms, such as RIPE Atlas. RIPE Atlas [41] is the largest geographically and topologically distributed public measurement platform providing more than 10K vantage points in 172 different countries, 5 continents, and 3,494 ASes. People can run measurements including pings and traceroutes on this platform with a system of credits. Ten years ago, the evaluation was made on small datasets of tenths of targets or proprietary datasets of unknown size, using a few hundred vantage points. Replicating the results

with larger datasets of targets and vantage points enabled by the RIPE Atlas platform will not only give a new baseline, but also give more robust and more representative results. We want this baseline and our results to be easily extended and updated, so we make our code publicly available and only use publicly available datasets.

2.2.2 Revisiting the fundamental insights that led to earlier results. Beyond the new baseline, we want to re-evaluate whether the insights that led to the results of the techniques are still valid because they could serve for future geolocation techniques: In the million scale paper, the main insight was that a few vantage points close to a target can perform as well as many vantage points. In the street level paper, there were two main insights: (1) There exists locally hosted websites close to a target that one can use as landmarks; (2) The relative order of the geographic distances between these landmarks and a target is preserved by the measured distance, so mapping the target to the landmark with the lowest RTT to it gives the best results.

2.2.3 Giving recommendations for IP geolocation usage and research. Finally, we want this replication study to serve for both future usage of IP geolocation and research in IP geolocation. Re-evaluating the insights and maybe finding new ones will help us to give useful recommendations about which ideas are more likely to work in the future.

3 METHODOLOGY

Given the goals described in Section 2, our methodology to replicate the papers should: (1) Be faithful to the original, to be able to provide a sound baseline of comparison with the two techniques (2) Provide updated results on the end to end results of the techniques, but also on the insights that led to these results (3) Provide, when possible, additional results to help future research using IP geolocation.

To follow these guidelines, for each paper, we give a background of the technique, the insights that were used and found, and the details of the replication, *i.e.*, whether there is any change in the methodology, the results that we replicate, and some additional results that could help future work.

To be clear, we replicate the papers, and do not reproduce them. We do not use the same set of targets, vantage points, or the same mapping service for the street level paper. These differences are explained in their own section (§4).

Before delving into the details of how we replicate the two papers, we give some background on Shortest ping and Constrained Based Geolocation (CBG) [31] which are two classic latency-based IP geolocation techniques on which the two replicated papers rely.

Latency based geolocation: To geolocate a target IP address, latency-based geolocation consists in issuing measurements from distributed vantage points with known locations to the target in order to obtain RTTs between them and the target. Then, to transform these measurements into a geolocation: (1) Shortest Ping maps the target IP address to the geolocation of the vantage point with the lowest RTT. (2) CBG transforms the RTTs into a maximum distance between the vantage points and the target by converting time to distance via speed of light, forming several circles centered at the vantage points, providing a set of constraints where the IP address can possibly be located. CBG estimates the target’s location as the

centroid of the intersection of the different constrained regions. Figure 1a shows an example.

3.1 Million scale paper

3.1.1 Background. This paper describes how to scale Shortest ping and CBG, by showing that a set of well chosen vantage points (VPs) close to the target can be as accurate as the set of full VPs. To show this result, they start by showing three hypotheses (Section 4 of the million scale paper):

- A few VPs can be as accurate as many VPs.
- Certain small subsets have good accuracy.
- The closest VPs generally maximize accuracy.

The intuition behind these hypotheses is that only the VPs with small RTTs to the target contribute to the constraints. Typically, a VP with an RTT of 100ms to the target results in a constrained region with a radius of 10,000 km, and almost never serves to accurately geolocate a target, as some smaller circles are included within this one. If these hypotheses are valid, then one could only use a few VPs per target, drastically reducing the active probing overhead from Shortest Ping and CBG, enabling faster geolocation.

Once the authors have shown that their hypotheses are reasonable, the challenge is then to know which VPs to select per IP address. Their idea is to probe representatives for each target: the representatives of a target are three responsive (if they exist) IP addresses in its /24. The intuition is that IP addresses in the same /24 should be close, so VPs close to the representatives should be also close to the target. They probe these representatives from all the VPs to obtain RTTs between the VPs and then select the ten VPs with the lowest RTTs to the representatives to probe the target.

With this technique, they achieve a comparable performance to CBG with all vantage points, showing that their VP selection algorithm works, with a median error decreasing from 231 km to 208 km, while only using 2% of the VPs. These savings allow the authors to geolocate 35% of the IPv4 routable space (in 2012) in a few months.

3.1.2 Replication.

Methodology and changes: We replicate the methodology without any change.

Insights and results that we re-evaluate: First, we evaluate whether the insight that a few close VPs can perform as well as many VPs is still valid, by looking at the three hypotheses (Sec. 4.1, Figure 2 and Figure 3 of the million scale paper) and the VP selection algorithm (Sec. 4.3, 4.4, and Figure 5 of the million scale paper). Then, as the main goal of the paper is to scale geolocation techniques, we provide an evaluation about the applicability of the VP selection algorithm on RIPE Atlas.

Results that we do not replicate: We do not provide the hilbert curve that shows the visualization of the geolocated IPv4 address space, we could not geolocate millions of IP addresses with the technique on RIPE Atlas (§5.1).

3.1.3 New insights. We extend the results from ten years ago, showing that a single well chosen VP per target performs as well as many VPs. We also extend the original VP selection algorithm, which is

not deployable on RIPE Atlas because of its measurement overhead, to make it more scalable.

3.2 Street level paper

3.2.1 Background. This paper presents a methodology to accurately geolocate a target IP address, up to street level, by leveraging two insights: (1) There exists some locally hosted websites landmarks close to the target that can serve as third-party VPs although they are not controlled by the authors, and (2) The relative order of geographical distances between the landmarks and the target is preserved by the measured distances with RTTs, *i.e.*, a landmark geographically closer to the target should have a smaller RTT between it and the target than a geographically further landmark. With these two insights, to obtain a very accurate geolocation, one has to find the landmark with the lowest RTT to the target and map the target to this landmark’s coordinates.

The authors present a three-tier system to geolocate a target IP address using these insights, shown in Figure 1.

Tier 1, Section 2.1 of the street level paper: The first tier (Figure 1a) goal is to obtain a first approximation of the geolocation of the target, by running ping measurements from the VPs to the target and perform CBG. The centroid of the CBG region is extracted and given as input to the tier 2.

Tier 2, Section 2.2 of the street level paper: The second tier’s (Figures 1b and 1c) goal is to obtain landmarks from locally hosted websites. The idea is that websites owned by entities such as university, the government, or a company, might be located at the postal address given by a mapping service (*e.g.*, OpenStreetMap or Google Maps). For instance, when typing our laboratory on a mapping service, one finds its postal address, and its website `laboratory.com`¹. Our website is locally hosted in the laboratory, so it can serve as a landmark. If we can obtain multiple websites in the CBG region and compute delays between them and the target, we will likely reduce the region where the target must be.

To obtain those websites, the technique uses a mapping service to map some points in the CBG region to their zip code and extracts the websites of the points of presence that are close to these zip codes. To sample the points in the CBG region, one draws concentric circles centered at the centroid of the region obtained with CBG, increasing the radius of the circle by step of R , here $R = 5$ km. The process stops when no points of a circle are within the CBG region. For each circle, one extracts points from it by rotating the point at 0 degree from an angle of α degree, here $\alpha = 36$, so 10 points are extracted from each circle.

A website is not necessarily locally hosted (*e.g.*, it can be hosted by a CDN), so the technique performs multiple checks to assess whether the website is locally hosted or not, consisting in (1) comparing the zip code of the postal address of the entity given by the mapping service with zip code of the coordinates of the point of the circle (2) checking that the content is not served by a CDN (3) checking whether a website appears in multiple zipcodes. If one of the three tests fails, the website is not used as a landmark. Section 3.2 of the street level paper provides more details on these tests.

Then, to obtain the delay between the landmarks and a target, one runs traceroutes from each VP to each landmark and to the target (Figure 1c). For each pair of traceroutes from a VP to a landmark and the target, one extracts the delay between the last common hop in the traceroutes, here $R1$ for the traceroutes from $V1$ and $R2$ for the traceroutes from $V2$. $D1$, the delay between $R1$ and the landmark, and $D2$, the delay between $R2$ and the target, are then obtained from the traceroutes from $V1$, and similarly $D3$ and $D4$ are obtained from the traceroutes from $V2$. Then, the minimum of $D1 + D2$ and $D3 + D4$, here $D1 + D2$, is selected to be an upper bound of the delay between the landmark and the target. We describe in Section 3.2.2 how we compute the D_i , as the street level paper lacks explanations. With these delays between the landmarks and the target, one converts these delays into distances to obtain a new CBG region formed by the intersection of the circles centered at landmarks.

Tier 3, Section 2.3 of the street level paper: The last tier consists in repeating the tier 2 technique but with the new region obtained from the tier 2, and with a finer granularity (Figure 1d). Now the concentric circles are drawn by increasing the radius by steps of $R=1$ km and using a rotation angle α is 10 degrees, obtaining 36 points per circle. After repeating the steps of finding locally hosted websites and of performing the traceroutes, the geolocation of the target is set to the geolocation of the landmark with the smallest delay Figure 1e. The key here is that the authors use the insight that the order of the delays is preserved in the order of the distances, so selecting the landmark with the smallest delay should result in the smallest error distance.

3.2.2 Replication.

Methodology and changes: In the tier 1, the only thing that differs with the million scale paper is that the speed of Internet used to convert RTT into distance is $\frac{2}{3}c$, whereas it is $\frac{4}{9}c$ in the street level paper, as they say that $\frac{2}{3}c$ is too conservative. We are faithful to their methodology and use $\frac{4}{9}c$ for the replication of the street level paper.

We replicate the tier 2 and the tier 3, with one modification to reduce the measurement overhead: for each landmark of a target, instead of running traceroutes from all the VPs to the landmark, we only use the 10 closest VPs to the target to run traceroutes to the landmark, as our results show that adding more VPs does not bring useful information (§5.1).

We also add a statement about how the delay between a landmarks and a target is computed. In Figure 1c, $D1$ and $D2$ represent the delays between $R1$, the last common router between the two traceroutes, and the landmark and the target. In the paper, it is simply stated: “Next, we calculate the latency between the common router and the landmark [...] and the latency between the common router and the target”. However, obtaining these values is not possible without information about the reverse paths (appendix B), and the values of $D1$ and $D2$ are very noisy (§5.2). As the authors did not mention the usage of a technique to get information about the reverse paths [35], we assume that the authors just subtracted the RTTs between ($V1$, landmark) and ($V1$, $R1$) to obtain $D1$ and subtracted the RTTs between ($V1$, target) and the landmark ($V1$, $R1$) to obtain $D2$, so we do the same.

¹anonymized for submission

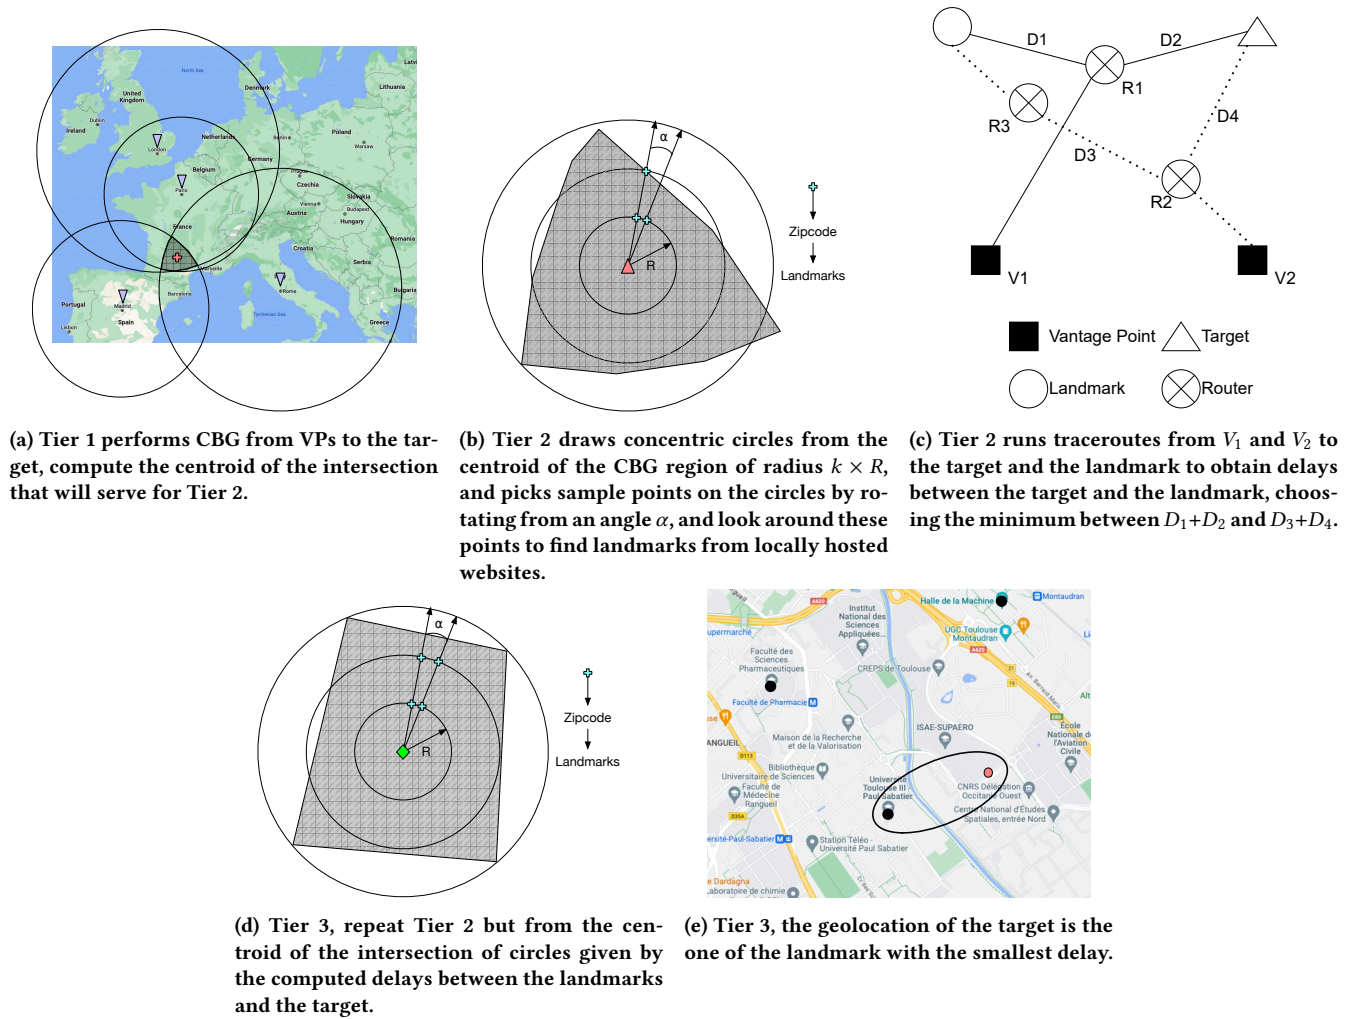


Figure 1: Three tier technique to geolocate an IP address in the street level paper [46]

Insight and results that we re-evaluate: We first replicate the results on the overall accuracy of the technique (Figure 8 of the street level paper), to show whether the technique can still geolocate IP addresses with a street level precision.

Then, we re-evaluate the insights that led to these results, *i.e.*, whether there exists some locally hosted landmarks near the targets (Figure 9 of the street level paper), and whether the relative order of the geographical distances between the landmarks and the target is preserved in the measured distances (Figure 5 of the street level paper). We also replicate the results on the relationship between accuracy and population density (Figure 11 of the street level paper).

Finally, we also re-evaluate the scalability of the technique (Section 5 of the street level paper).

Results that we will not replicate: We do not replicate the results on relationship between the accuracy and the different ISPs.

4 DATASETS

For each replicated paper, we describe the datasets that were used, their limitations, and present our choices for the replication. In particular, we describe our choice of targets and vantage points, and the measurements run between them. Other datasets include the mapping service used to perform reverse geocoding and to obtain websites that serve as landmarks in the street level paper. We then describe how we sanitize the geolocation of our vantage points and targets, and finally discuss the bias of our datasets. Table 1 gives a recap of the datasets and APIs used in the two replicated papers and in our replication. All the datasets were collected during the March-April-May 2023 period.

4.1 Million scale paper

4.1.1 Vantage points.

What has been done: The authors used 400 PlanetLab nodes.

	Million scale paper [32]	Street level paper [46]
Original targets	PlanetLab nodes (25)	PlanetLab nodes (88), Residential dataset (72) Driving service dataset (?)
Replication targets	RIPE Atlas anchors (723)	
Original vantage points	PlanetLab nodes (400)	Ping servers (163) Traceroute servers (136)
Replication vantage points	RIPE Atlas probes (10k)	RIPE Atlas anchors (723)
Original Other datasets	None	Geonames [2]
Replication Other datasets	None	Nominatim [5] OpenStreetMap [6] Overpass API [7]

Table 1: Targets, vantage points, and other datasets used in the replicated papers and in our replication. Numbers in parenthesis for the targets and vantage points rows correspond to the number of elements. All datasets and APIs used in our replication are public.

Limitations: PlanetLab nodes were mainly located in universities in research and education networks, which are not representative of the global Internet connectivity. Nowadays, universities are considered to be located at the edge of the Internet and not well connected, in comparison to vantage points in IXPs, for instance. Finally, PlanetLab is no longer an option as it was decommissioned several years ago.

Our choices: With more than 10K vantage points in 172 countries and 3,494 ASes, RIPE Atlas is the publicly available platform with the best coverage to run geolocation measurements. We choose to use the 10K vantage points to obtain an upper bound of the accuracy that one can obtain with the technique with the entire coverage of the RIPE Atlas platform. To be clear, this work used a lot of RIPE Atlas credits (hundreds of millions) and has only been made possible because RIPE Atlas generously gave us an upgraded RIPE Atlas account with increased rate limits to perform our measurements in a reasonable time.

4.1.2 Targets.

What has been done: The authors used 25 servers in universities around the world.

Limitations: This dataset is small, and lacks of diversity. As stated earlier, universities are in networks that are not representative of the entire Internet, so, as stated by the authors, the accuracy might be different for these targets than for other targets in other networks.

Our choices: We use the set of the 723 RIPE Atlas anchors because their geolocation is more accurate than RIPE Atlas probes (but we do perform some sanity checks though, see Section 4.3). This dataset is best effort: although the number can appear as being small, there exists no publicly available ground truth dataset for IP geolocation.

These 723 targets are located in 441 cities, 96 countries, and 561 ASes, and non uniformly spread over the continents: we have 133 targets in Asia, 16 in Africa, 18 in Oceania, 125 in North America, 399 in Europe and 27 in South America.

4.1.3 Measurements. We run pings from all the RIPE Atlas probes to each target. For the VP selection algorithm, we use the ISI hitlist [25] to select the three representatives of the /24 prefix of each target with the highest score of responsiveness. For 8 targets there were fewer than three responsive representatives, so we select random IP addresses in their /24 prefix to till the missing representatives. We then run pings from all the RIPE Atlas probes to these representatives.

4.2 Street level paper

4.2.1 Vantage points.

What has been done: The authors used 163 publicly available ping and 136 traceroute servers geographically distributed in the US. The ping servers are used for the tier 1, while the traceroute servers are used for the tiers 2 and 3.

Limitations: The set of vantage points is restricted to the US. In addition, from a more practical point of view, even if prior work tried to automate measurements from ping and traceroute servers [27], these servers are not made to perform automated measurements, and often explicitly mention in their policy that automating measurements is a violation.

Our choices: We use the RIPE Atlas platform, but this time we restrict ourselves to the set of anchors to run the ping and traceroute measurements. We cannot use the full set of probes for this paper, because in addition to the ping measurements from the tier 1, we have to issue one traceroute from each vantage point to each landmark: the median number of landmarks per target is 111, so it is both unrealistic in terms of credits needed and in terms of measurement overhead put on the RIPE Atlas platform.

4.2.2 Targets.

What has been done: The authors used three datasets: 88 PlanetLab nodes, 72 IP addresses from a residential dataset obtained with crowd-sourcing, and another dataset containing three months of user’s search logs for driving directions from a “popular online maps service”, where each record of the dataset includes the user access IP address and the driving sequence represented by two pairs of latitude and longitude points. They extract IP addresses which appear in multiple driving sequence as source or destination and set the geolocation of these IP addresses to the source or the destination, depending on a set of heuristics (Section 4.1.3 of the street level paper). This last dataset has an unknown number of IP addresses (the lack of details is justified by the authors by a non disclosure agreement set with the company). All the IP addresses of the three datasets are located in the US.

Limitations: The first limitation is one of the ones described for the first million scale paper: the dataset the target IP addresses is rather small, even if this time, the three different dataset bring diversity in the type of networks. Another limitation is that the geolocation of the targets in the second dataset is based on crowd

sourced measurements and can be subject to errors. Finally, the last dataset is opaque and its validation relies on handcrafted heuristics.

Our choices: For the same reasons as given for the first million scale paper, we choose to use the RIPE Atlas anchors as targets. As in the Figure 7 of the street level paper, our dataset covers both rural and urban areas (appendix C).

4.2.3 Measurements. We run pings from all the RIPE Atlas anchors to the targets to compute the tier 1. For the tiers 2 and 3, for each landmark associated with a target, we select the ten VPs with the lowest RTTs to the target using the measurements from the tier 1 and run traceroutes to the landmark and to the target.

4.2.4 Services to obtain the websites and landmarks. The tier 2, which consists in obtaining landmarks from locally hosted websites found in the region determined by the tier 1 CBG, is divided into several steps: how to perform reverse geocoding (map a point with latitude and longitude coordinates into a postal code), how to extract points of interests in each postal code, and how to obtain the websites from the points of interest, if they exist.

What has been done: The authors used the Geonames [2] publicly available service to perform the reverse geocoding queries and obtain the websites that can serve as landmarks if they pass the tests of being locally hosted. They used the websites of entities qualified as “business”, “university” and “government office” by Geonames.

Limitations: The Geonames API allows a limited number of calls, 1000 calls per hour, and 20,000 calls per day, unless one subscribes for a premium tier. The median number of reverse geocoding queries that we had to do per target is 878, so the Geonames API was not an option to replicate the technique in a reasonable amount of time.

Our choices: To perform the reverse geocoding queries, we use a local instance of Nominatim [5], an engine built on top of OpenStreetMap [6], which allows us not to be restricted in terms of the number of queries that we perform on the server. To find the places of interest in a postal code, we use the overpass API [7] and query for all the amenities with a website, so it is not restricted to keywords used in the street level paper. We use a public instance allowing an unlimited amount of queries, although we experienced some rate limiting when trying more than 8 simultaneous requests [7]. All these datasets and services are publicly available.

4.3 Sanitizing RIPE Atlas geolocation

As we are using RIPE Atlas probes and anchors as vantage points and anchors as targets, it is important to have high confidence in their geolocation. If the common thought is that the geolocation of the anchors reported by RIPE Atlas can be trusted [10], it is less certain for RIPE Atlas probes, and it cannot hurt to perform additional checks.

Our sanitizing process starts by verifying anchors, and consists in counting the number of speed of Internet violations per anchor. We extract RIPE Atlas meshed measurements between anchors and count, for each anchor, how many of its RTTs from and to other anchors violates the speed of Internet (SOI) constraint. SOI is set to $\frac{2}{3}c$ where c is the speed of light [31]. We iteratively remove the anchor with the highest number of SOI violations, update the

number of SOI violations, until there is no anchor with any violation. With this process, we remove 9 anchors from both datasets of vantage points and targets. We do the same for probes: we perform ping measurements from all the probes to each anchor correctly geolocated, and remove 96 probes from our dataset of vantage points.

4.4 Potential bias of the RIPE atlas platform

Before describing the limitations of our dataset of targets and vantage points, we mention that these datasets are best efforts given the publicly available ground truth that we have. Moreover, our methodology is fully replicable and our code publicly available, so someone has a better set of targets and vantage points, it can re-run our methodology with this set and obtain a new baseline.

4.4.1 Geographic and topological bias. Our targets and vantage points are limited by the coverage of the RIPE Atlas platform, which is non uniformly geographically spread, with a denser presence in Europe. In addition to the geographic bias, we also look at a potential AS type bias. PlanetLab were mainly located in R&E networks and were not representative of the different types of network composing the Internet. Table 2 shows the distribution of the AS category of our targets according to the CAIDA AS classification dataset [1]. We see that the targets are located in different types of networks, with most of them being in content providers, access networks, and transit/access networks. In addition, we look at the AS type of our targets according to the ASDB dataset [50], finding that they fall into 16 categories, with 72% falling into the “Computer and Information Technology” category. The second most represented category being “R&E” with 5%. All of the other categories are below 5%. The classification according to these two different datasets show that even if our targets are not representative of the whole Internet, they contain more network diversity than the datasets used in the replicated papers.

We also show in Table 2 the distribution of the RIPE Atlas probes + anchors that serve as our vantage points for the million scale paper. We see that if the dataset is dominated by the access network category, there are nonetheless 25% of the probes in other categories, showing the topological diversity of the RIPE Atlas platform.

4.4.2 Last mile delay. Our targets, the RIPE Atlas anchors, are well-connected servers that typically not suffer from the last mile delay [37]. However, targets in access networks might suffer from it, adding precious milliseconds to latency measurements and make latency-based geolocation techniques less precise. As a result, geolocating the targets in access networks impacted by the last mile delay is even more challenging.

5 EVALUATION

For each replicated paper (§5.1 and §5.2), we first give the main takeaways in a leading paragraph, and then support each of the takeaways in a dedicated paragraph with the analysis supporting the takeaway.

5.1 Million scale paper

Takeaways: The number of VPs, the parameter used to evaluate the accuracy by the first two hypotheses of the million scale paper,

Dataset	Content	Access	Transit/Access	Enterprise	Tier-1	Unknown
Anchors	229 (31.7%)	211 (29.2%)	197 (27.2%)	55 (7.6%)	6 (0.8%)	25 (3.5%)
Probes	1112 (9.2%)	9124 (75.2%)	1005 (8.3%)	410 (3.4%)	166 (1.4%)	312 (2.6%)
Probes + Anchors	1361 (10.5%)	9347 (72.4%)	1221 (9.5%)	472 (3.7%)	174 (1.3%)	339 (2.6%)

Table 2: AS type of the RIPE Atlas probes, anchors, and probes + anchors according to the CAIDA AS classification dataset [1].

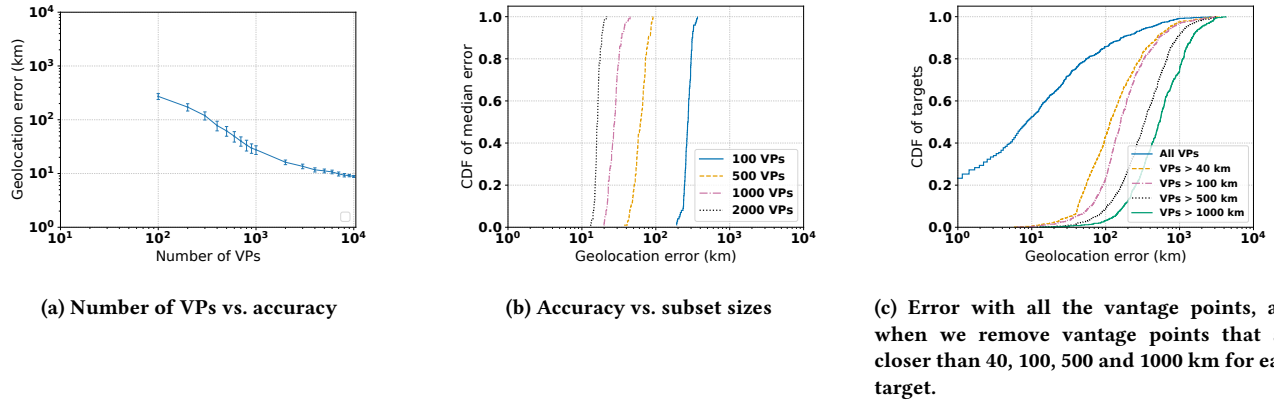


Figure 2: Replication of the three hypotheses on how using subsets of VPs affect accuracy.

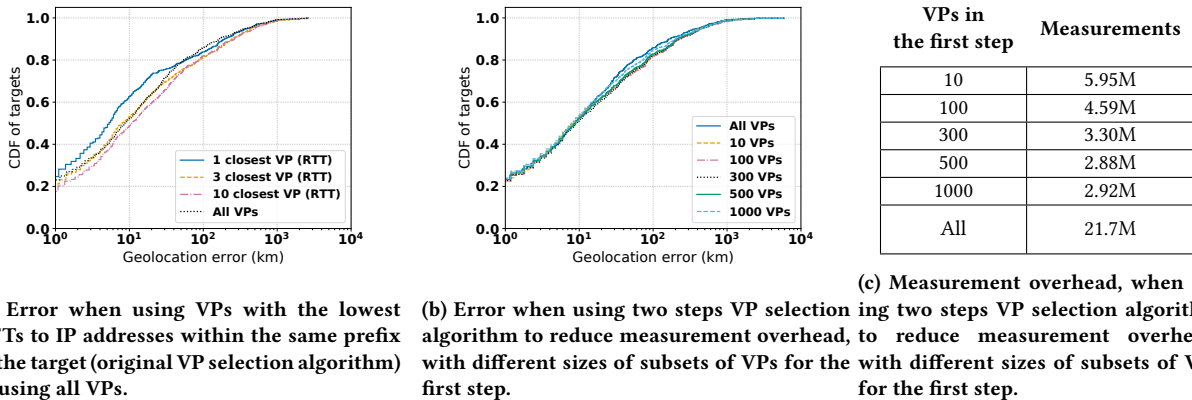


Figure 3: CBG performance with original VP selection algorithm and new VP selection algorithm.

is not the right metric to measure accuracy. Instead, what matters is the access to VPs close to the target (§5.1.1). In particular, we extend the prior result that a few close VPs were enough to geolocate a target: a single well chosen VP is enough (§5.1.2). Beyond the replication, we show that VP selection algorithm to find the closest VPs to a target cannot be deployed on RIPE Atlas because of its measurement overhead (§5.1.3), but that it is possible to reduce it to 13.2% of the original measurements, while achieving a similar performance, by decoupling the VP selection algorithm in two steps (§5.1.4). Finally, we have a look at whether the accuracy depends on the geographical coverage, finding that the better coverage of the RIPE Atlas platform does not necessarily translate into a better accuracy (§5.1.5). All the results are given for CBG, but results with shortest ping are similar.

5.1.1 *The number of VPs is not the right metric to evaluate accuracy.* The hypotheses that we need to re-evaluate are that: (1) “a few VPs can be accurate”; (2) “certain small subsets have good accuracy”; (3) “The closest VPs generally maximize accuracy.”

A few VPs can be accurate: We replicate the experiment that looks at the median geolocation error for subsets of 10 to 10K VPs. For each subset size, we run 100 trials with a random subset of VPs and we compute the median error for each trial. Figure 2a shows the error bars of the distribution of the median error (Figure 3a of the million scale paper). First, our results are qualitatively different from the million scale paper, where the best median of the median error was a few hundred kilometers. Here, the median of the median errors is as low as 8 km with 10K VPs. Then, whereas in the replicated work, beyond 60 VPs, the median error was stable,

the median error still decreases when we add VPs, even beyond thousands, although the gains are smaller.

Certain small subsets have good accuracy: We replicate the experiment that focus on specific sizes of subsets. Figure 2b shows the CDF of the median error for sizes of subsets of 100, 500, 1000, and 2000 VPs (Figure 3b of the million scale paper). Our results are also different: in the million scale paper, the authors found that some subsets were performing poorly, and some were doing well. Our distributions vary much less than theirs: for instance, for 100 VPs, where the median error is supposed to vary the most, the median error varies from 191 to 366 km, whereas in the million scale paper, for 50 VPs, where the median error varies the least, the median error varies from a few hundred km to almost a thousand km.

The closest VPs generally maximize accuracy: Figure 3c of the million scale paper showed on one example target that the geolocation error was dependent on the distance to the closest VPs. We generalize this figure: Figure 2c of our paper shows the geolocation error when, for each target, we remove VPs that are closer than a certain distance. We see that when we remove VPs that are closer than 40 km (basically, in the same city [26]), the median error increases from 8 km to 120 km, and only 6% of the targets have an error of 40 km or less when removing the VPs in the same city, versus 73% with all the VPs.

Do the hypotheses still hold? If the goal is to provide a city level accuracy for each target, one should at least have one VP in each city where there is a target. So, if the third hypothesis stating that the closest VPs maximize accuracy is still definitely true, the two first are not really relevant, as small subsets of VPs can only be accurate for geolocating targets close to each of them.

5.1.2 A well chosen single VP per target works well. As the key to be accurate is to have access to VPs geographically close to the target, we re-evaluate whether the VP selection algorithm (§3.1) can find these VPs, if they exist. Recall that this algorithm selects the VPs with the lowest RTTs to three representatives per prefix so that it can use this subset of VPs to geolocate a target in the prefix. As low RTTs imply that the VPs are geographically close to the target, we expect the VP selection algorithm to still work well.

Figure 3a shows the replication of the Figure 5 of the million scale paper, where we compute the CDF of the geolocation error if we take 1, 3, and 10, and all VPs with the lowest RTTs to the target (the figure in the million paper has only the line for 10 VPs). The results are a little bit surprising: for errors lower than 40 km, the single closest VP outperforms the other alternatives, with, for instance, 62% of the IP addresses located at 10 km or less, against 52% for all VPs. This shows that it is hard to obtain a better accuracy than city level (40 km) for CBG.

5.1.3 The VP selection algorithm cannot be deployed on RIPE Atlas because of its measurement overhead. Even if the VP selection algorithm provides a way to reduce the number of VPs needed to geolocate an IP address, each VP still needs to probe three representatives per prefix. To be able to geolocate 35% of the unicast allocated IPv4 address space in a few months, the authors used VPs which had a probing rate of 500 packets per second (pps). This is not feasible with RIPE Atlas, as the 10k VPs cannot have a probing

rate of 500 pps only for geolocation. Indeed, on average, an anchor has a probing rate between 200 and 400 pps, whereas a probe has a probing rate between 4 and 12 pps [11].

5.1.4 Towards scaling the VP selection algorithm. As we cannot use all the VPs to measure the three representatives of an IP address, we need to reduce the number of VPs without losing too much accuracy.

To reduce the number of VPs, we modify the VP selection algorithm to proceed in two steps: instead of using all the VPs to probe the representatives, we use a subset of VPs as a first step, compute CBG, and select one VP per AS/city in the CBG region. We then run pings to the representatives from this subset of VPs, and select the VP with the lowest median RTT to the representatives to geolocate the target. The subset of VPs for the first step is greedily selected to cover a maximum of the earth with a minimum number of VPs: at each iteration, we select the VP which maximizes sum of the logarithmic distances to the other VPs. This technique is similar to what has been done in prior work [15].

Figure 3b shows the error with this new VP selection algorithm using different sizes of subsets for the first step. We see that using this technique does not degrade the performance, even when using 10 VPs. Now that we know that this technique does not degrade performance, we have a look at by how much it reduces the probing overhead: Figure 3c shows the number of measurements needed depending on the number of VPs used for the first step. We see that the better tradeoff is realized with 500 VPs, which brings a total of 2.88M ping measurements, which represents 13.2% of the 21.7M needed by the original algorithm.

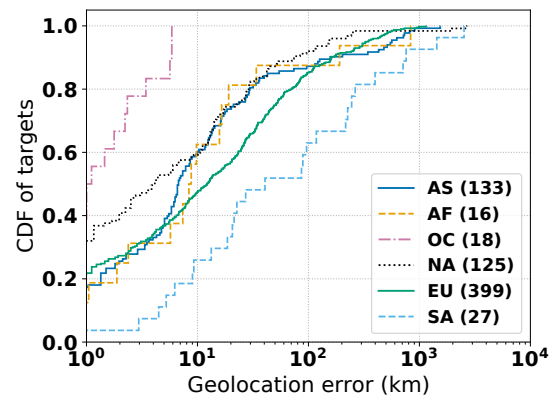


Figure 4: Error per continent

5.1.5 . Figure 4 shows the geolocation error when we split the targets by continent. The numbers in parenthesis indicate the number of targets in the continent. We observe that the accuracy does not necessarily match the coverage of the RIPE Atlas platform: Africa has an overall better performance than Europe, whereas there are way fewer VPs in Africa. Digging into these results, we find that the access to a close vantage point is not the reason why the results for Africa are better: 94% of the targets do have a VP at less than 40km, and this number grows to 99% for Europe. Inspecting the

26 targets in Europe with a high error above 300km, we find that the probes close to the targets did not give a small RTT, with a median over the targets of 7.96 ms, suggesting that the probes close to the targets could suffer from the last mile delay, or that their geolocation is not accurate.

5.2 Street level paper

Takeways: The street level technique has a similar accuracy to CBG on our dataset, with a median error of 28 km vs 29 km for CBG, far from the 690 m of the street level paper (§5.2.1). We explain this low performance by re-evaluating the insights of the street level paper, finding that they do not hold: (1) At most, 28% of targets have a landmark within 1km radius around them (§5.2.2) (2) The relative order of the geographical distances between the landmarks and a target is not preserved in the measured distances: the median correlation between the measured distance and geographical distance across the targets is 0.08 (§5.2.3). We also find that the accuracy of the technique does not depend on the population density as was previously found (§5.2.4). Finally, with our best effort setup, the median time to geolocate an IP address is 1,238 seconds (20 minutes), far from the few seconds that the authors of the street level paper gave as a theoretical time to geolocate an IP address (§5.2.5).

5.2.1 The technique has a similar performance to CBG. Figure 5a shows the CDF of the error per target for three techniques: street level, CBG, and closest landmark. The closest landmark technique, as its name indicates, selects the closest landmark to each target. This technique simulates an oracle where we know the geolocation of the target and is a lower bound of the error that the street level can obtain, for two reasons: it assumes that all the websites passing the tests of Section 3.2 are actually locally hosted and can be used as landmarks, and that the technique always select the closest landmark to the target. For 46 targets, we are not able to find a landmark for the target, so we assume that the street level technique and the closest landmark technique return the same geolocation as CBG. Finally, there are also 5 targets where the value of $\frac{4}{9}c$ for the speed of the Internet did not bring any intersection, so we used a value of $\frac{2}{3}c$ for them.

There are two results from the graph. First, the street level and CBG technique are close, with a median error of 28 km for the street level technique vs 29 km for CBG. In the street level paper, the authors found a median error of 690 m on the PlanetLab dataset, so two orders of magnitude of difference. Second, the closest landmark technique shows that at most, if we consider than 1 km is a street level precision, only 33% of the targets could eventually be geolocated at street level. We detail these two results in the next sections.

5.2.2 Most targets do not have a street level landmark. Figure 5b shows two things: (1) the second column shows the number targets with a locally hosted landmark at less than a certain distance. It is an upper bound as it corresponds to a scenario where the landmarks that passed the tests to be locally hosted are actually locally hosted. (2) The third column shows that same number but with additional latency checks on the landmarks. These checks consist in running pings from each target, which is a RIPE Atlas anchor, to all the

landmarks that passed the tests to be locally hosted at less than 40 km from this target. We only keep the landmarks with a RTT of less than 1 ms. These checks do not prove the exact geolocation of the landmark, but we have higher confidence that these landmarks are actually locally hosted.

We see that 28% of the targets have a landmark at less than 1 km away, and 76% of the targets have a landmark at less than 40 km away for the optimistic view, and these numbers decrease to 19% and 72% when we perform additional latency checks on the landmarks. We conclude that most of our targets cannot be geolocated at street level, but most of them could be at least geolocated at city level if a technique could select the right landmark.

In addition to these results, we mention that a lot of websites did not pass the tests to be classified as locally hosted: only 65,325 landmarks out of 2,584,527 passed the tests (2.5%).

5.2.3 Relative order of the geographic distances is not preserved by the measured distances. Figure 5c replicates the Figure 5 of the street level paper. It shows the scatter plot of measured vs geographic distances for targets that were geolocated with an error of 1 km, 5 km, 10 km and 40 km. For each target, the landmark selected corresponds to the point closest to the X axis. We can see that our results are noisier than the one found in the street level paper: the street level technique selects the closest landmark only for the target with an error within 1 km. More generally, we compute the Pearson correlation coefficient between the measured distance and the geographical distances for all targets and find a median value of 0.08, so almost no correlation.

The authors of replicated work said that the relative distance preserved the geographical distance because “the network paths that are used to estimate the distance between landmarks and the target share vastly common links”. We verify this property by looking at whether we can find many landmarks that are in the same network as a target, and whether the relative distance is preserved if we only consider landmarks in the same network as the target: 140 targets (19%) have at least one landmark within the same AS as the target, 59 targets (8%) have at least one landmark in the same BGP prefix retrieved from RouteViews, and only 3 targets have at least one landmark within the same /24. Even for the 57 and 24 targets with at least two landmarks within the same AS or the same BGP prefix, the Pearson correlation coefficient between the measured distance and the geographical distance is 0.39 and 0.19, showing that the insight does not really hold even with this property.

Finally, we show that the $D1 + D2$ value used to compute the measured distance between a landmark and a target is noisy. Figure 6a shows the CDF of the percentage of landmark for which the $D1+D2$ value is negative and therefore unusable to compute a distance. For 50% of the targets, at least 28% of the landmarks are unusable, so a significant number, questioning the computation of the $D1 + D2$ value. More details about this computation are given in Appendix B.

5.2.4 Geolocation accuracy does not depend on population density. Figure 6b replicates the results of the Figure 11 of the street level paper and shows a scatter plot of the error distance vs the population density, with a best fit linear regression. We obtain the population density data via the “Gridded Population of the World 2020” dataset, which gives the density of the population with a one




Figure 5: Performance of the street level technique (left) and re-evaluation of the two insights: is there a close landmark for most targets (middle)? Is the relative order of the geographical distance preserved by the measured distance (right)?

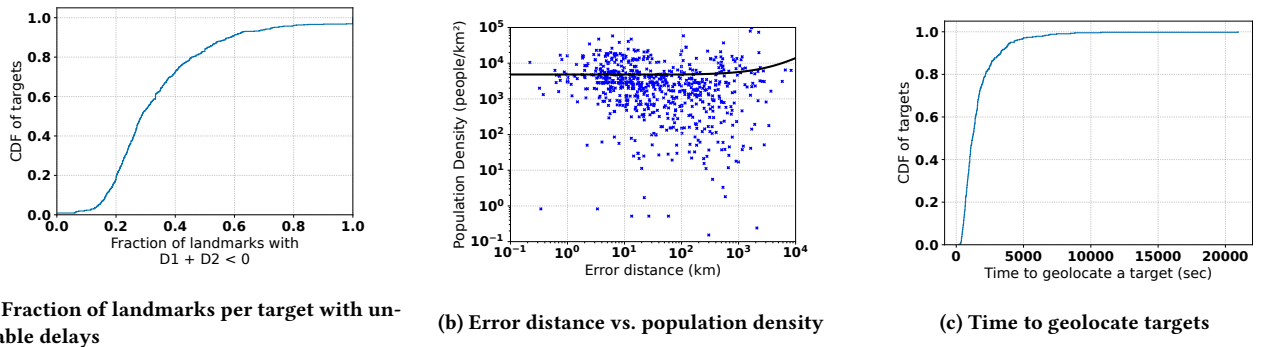


Figure 6: Delays between landmark and targets are noisy (left), accuracy does not depend on the population density (middle), the technique hardly scales with our setup (right).

kilometer square precision [8]. We see that the error is not really better for denser areas, as they found in the street level paper.

5.2.5 We could not scale the technique. In the Section 5 of the street level paper, the authors mention two things to show the scalability of the technique: (1) They can cache the results of the query to the mapping service to obtain the landmarks and the results of the test checking whether a landmark is locally hosted or not (2) They only need 8 RTTs to geolocate an IP address, where an RTT either corresponds to a master node asking a vantage point to run a measurement (a ping or a traceroute) or the measurement itself, resulting in a 1-2 second to run the measurements per target.

Although the authors are right for their first point on caching, this first step still represents a significant overhead. First, for 723 targets in 441 cities, we had to run 753,428 queries to our mapping service, where we observed a rate limiting of our queries to approximately 8 per seconds. Second, we ran 2,755,315 tests on the websites involving one DNS query and two wget per test.

On the second point, our RTTs are way bigger for us as our master node correspond to calling the RIPE Atlas API, and it generally takes a few minutes to get the results of a measurement. But more importantly, in terms of measurement overhead put in the network, the original technique ran a traceroute per landmark and per target

from all the vantage points, so for the same reasons as given in Section 5.1.3, the technique cannot be deployed on RIPE Atlas.

Finally, Figure 6c shows the time needed to geolocate a target using the street level technique. The median time to geolocate an IP is 1,238 seconds (20 minutes) using a multiprocessed code on a 32-core processor with 62 GB of RAM, so far from the theoretical 1-2 seconds that the authors of the street level paper found, so at least we can say that with our setup, we cannot use the technique at scale.

6 GEOLOCATION DATABASES

Commercial geolocation databases are often used by the community [29] as we lack a publicly available IP geolocation dataset from the community. We compare two geolocation databases, MaxMind (the free version) [4] and IPinfo (the free API) [3] to CBG with all the RIPE atlas probes (§5.1). These databases provide a mapping between IP prefixes (up to /32 for IPv4) and their geolocation. We query the data for both MaxMind and IPinfo in May 2023. Figure 7 shows that IPinfo is better than the two other datasets: 89% of the targets have an error of less than 40 km (correct at city level) for IP info, against 73% for all the RIPE Atlas VPs and 55% for MaxMind. We sent our results to IPinfo to better understand why they were

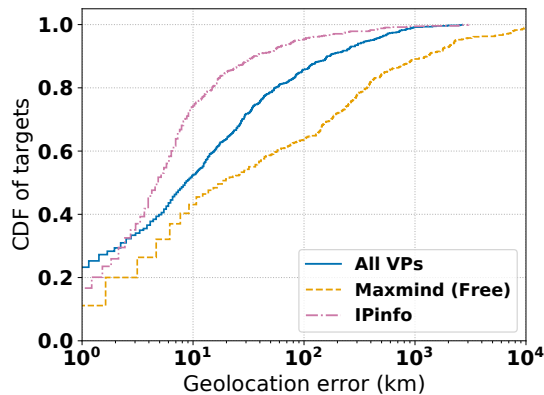


Figure 7: Geolocation of CBG with all the RIPE Atlas VPs versus geolocation databases.

doing better than the other techniques. This exchange² demystified a bit the geolocation databases usually considered as black boxes: They told us for 20% of the targets, their latency measurements gave an error of 42 km or less, and 70% of the targets with an error of 137 km, whereas it is 74% and 87% for these errors for CBG with all RIPE Atlas probes. Then, to further refine the geolocation, they told us that they were using hints extracted from DNS, WHOIS, geofeeds [9], or other proprietary data. To be clear, these results are specific to our dataset of targets, and IPinfo does not claim to be able to geolocate 70% of the Internet targets with an error of 137 km with latency measurements.

These explanations suggest that our goal to obtain an Internet scale IP geolocation is realistic if one combines the latency measurements from RIPE Atlas with hints from these different sources. Some work has been made in this direction using DNS names and latency measurements but was limited to routers [21, 36, 42] so future work could be to extend these techniques to the entire Internet. One of the big challenges is to manage to geolocate the IP addresses that do not have DNS names or are unresponsive.

7 LESSONS LEARNED

7.1 A new baseline

Our evaluation has set a new baseline for geolocation techniques: (1) On accuracy, 73% of the IP addresses could be geolocated at city level with both the street level and CBG techniques, and 11% of them could be geolocated at less than 1 km away, which we can consider as being street level. (2) On coverage, no technique is able to geolocate millions of IP addresses in a few months with RIPE Atlas (§5.1). In the future, if one is able to significantly beat the results of our replication on this dataset or another one by re-running our methodology, it should be considered as a substantial contribution.

²we publish these numbers with IPinfo approval

7.2 Insights for future work

7.2.1 We need inferential techniques when we have no close vantage points to use latency based techniques. We learned two lessons about latency based techniques (§§ 5.1 and 5.2): (1) For classic latency based techniques such as CBG, the accuracy only depends on the access to geographically close vantage points to the target. (2) Latency techniques based on RTT difference in traceroutes to obtain the latency between the landmarks and a target are noisy and are not trustworthy without information on the reverse path.

These two results suggest that we need to develop inferential techniques to geolocate IP addresses in cities where we have no vantage points.

7.2.2 There exists some landmarks at city level. We retain from Section 5.2 two lessons about the landmark-based techniques. The first one is that, surprisingly, in an era dominated by the cloud providers where one might think that it would be hard to find locally hosted websites, we found at least one locally hosted website in the same city as 72% of our targets, so that one can still use these locally hosted websites as landmarks for a city level accuracy, as we saw that street level geolocation is too optimistic.

The second lesson is that although these landmarks exist, choosing the closest one to a target is challenging, as one measuring the latency between the landmarks and a target is hard (§5.2.3). For future usage, one should develop new techniques to determine how to map a landmark to a target.

7.2.3 Round based geolocation is one key to scale. In Section 5.1, we have seen that the proposed VP selection algorithm was not deployable because of its measurement overhead. However, to reduce the overhead, we have proposed to decouple the VP selection in two steps, showing that we can use only 13.2% of the measurements needed by the original VP selection algorithm while obtaining the same accuracy. This principle could be easily extended to multiple rounds instead of two, and attain a number of rounds for which the measurement overhead is minimum. The tradeoff is that multiple rounds take more time than a single round, as each round needs the results of the previous round as input. As a result, multiple rounds need multiple calls to the RIPE Atlas API instead of one, so it would take a few minutes more in practice, but this is not really an issue as we do not expect the geolocation of IP addresses to quickly change over time.

8 RELATED WORK

Replicability of IP geolocation techniques: IP geolocation has been a longstanding goal of the community: Multiple types of techniques have been developed, such as latency based techniques [31, 32, 34, 38, 46, 48], DNS based combined with latency based measurements [21, 33, 36, 42], statistical techniques [49]; or machine learning techniques [24]. It is not easy to compare with all of these techniques, as only a few of them come with available code [21, 36]. And even for those, it is not easy to replicate them, as they both use proprietary data to train their model. Our paper is a step towards providing a new baseline to compare geolocation techniques, as it presents a replication of two techniques that achieved high performance in terms of coverage and accuracy [32, 46] using only publicly available data.

Towards a publicly available dataset of the routable space:

The authors of the million scale paper [32] are the last ones to have published a dataset of all the routable IP addresses [12], and this dataset is a single snapshot dated of 2012. In parallel, some efforts have been made on the router IP addresses: (1) CAIDA ITDK [16] provides a dataset of geolocated router IP addresses collected with the Ark platform. The geolocation is made using a combination of sources, including DNS based techniques [36], peering information [13, 14], and MaxMind [4]. (2) RIPE IPMap is the service providing a daily dataset of the router IP addresses that the single radius technique [23] could geolocate at city level. A snapshot of March 2023 contained 419K IP addresses, which is significant, but still far from the 3.3M router IP addresses seen by state-of-the-art topology mapping systems [30, 44].

Prior work looked at the performance of the geolocation databases [26, 39], and how dynamic they are [29], finding some differences between the databases and stating that these databases should not be trusted. Our work shows a more nuanced story: At least one geolocation database, IPinfo, outperformed CBG with all the RIPE Atlas vantage points on our dataset, and we know now that IPinfo mostly uses standard geolocation techniques. In the future, we could consider using geolocation databases for research if they provide more explainability, such as which geolocation technique, and eventually which dataset were used per IP address.

None of these public datasets are accurate, complete, and explainable, as ITDK and RIPE IPMap focus on router IP addresses, while geolocation databases are not explainable. In addition, our paper shows that no current techniques are capable of providing a dataset meeting the three criteria, setting up a new baseline for future contributions.

9 CONCLUSION

In this paper we presented a replication of two ten year old high performance techniques for IP geolocation, showing that no technique can either geolocate millions of IP addresses in a few months with the available measurement platforms, nor geolocate IP addresses at street level precision in the current Internet with the most powerful publicly available active measurement infrastructure (RIPE Atlas). In addition to providing a new baseline for future work, we revisited the fundamental insights of these prior techniques and found new insights in our study to inform considerations for the design of future IP geolocation techniques. This paper thus serves as a step towards the community’s goal of obtaining a publicly available Internet scale dataset and code for IP geolocation. All our code is available: <https://github.com/dioptra-io/geoloc-imc-2023>.

ACKNOWLEDGEMENTS

A research grant from the French Ministry for the Armed Forces has made this work possible. We thank: Robert Beverly and Timur Friedman, for their careful reading and feedback on earlier versions of this paper; the RIPE Atlas team which provided us with an upgraded RIPE Atlas account allowing us to run many measurements in parallel to realize this study in reasonable time. The anonymous reviewers from the IMC TPC, and our shepherd, Philipp Richter, for their careful reading of this paper and suggestions for its improvement. Omar Darwich and Kevin Vermeulen are associated with the LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France. Hugo Rimlinger, Milo Dreyfus, and Matthieu Gouel are associated with Sorbonne University, CNRS, Laboratoire d'informatique de Paris 6, LIP6, F-75005 Paris, France. Hugo Rimlinger and Matthieu Gouel have partially carried out the work presented in this paper at LINCS (www.lincs.fr).

A ETHICS

This work does not raise any ethical issues.

B LATENCY BETWEEN THE LANDMARKS AND THE TARGETS

We show that the step of computing the latency between the landmarks and the targets in the tier 2 of the street level paper (Figure 1c) is not straightforward without information on the reverse path.

Given the two traceroutes from the VP to the landmark L and the target T and the common router R_1 , what we can obtain is $RTT(VP, R_1)$, $RTT(VP, L)$, and $RTT(VP, T)$.

Let us decompose these RTT into one way delays, written $D(X, Y)$ for the one way delay from X to Y , to make the terms $D1$ and $D2$ appear. We have:

$$RTT(VP, L) = D(VP, R_1) + D1 + D(L, VP)$$

$$RTT(VP, T) = D(VP, R_1) + D2 + D(T, VP)$$

With no assumption on the reverse paths from L to VP and T to VP , we cannot isolate the $D1$ and $D2$ terms. If we do the assumption that the last link (R_1, L) and (R_1, T) is symmetric and then the reverse paths from R_1 are the same (which should be the case because of destination based routing [45]), we have now:

$$RTT(VP, L) = D(VP, R_1) + D1 + D(L, R_1) + D(R_1, VP)$$

$$RTT(VP, T) = D(VP, R_1) + D2 + D(T, R_1) + D(R_1, VP)$$

which simplifies into:

$$RTT(VP, L) = RTT(VP, R_1) + 2D1$$

$$RTT(VP, T) = RTT(VP, R_1) + 2D2$$

and so we can compute $D1$ and $D2$. The assumption of symmetry between two hops has been shown to be valid in 90% of the cases when the two hops belong to the same AS [45], so we see that further analysis needs to be done when one wants to compute $D1$ and $D2$. But in the street level paper, the authors do not give any explanation about this question.

C TARGETS DATASET POPULATION DENSITY

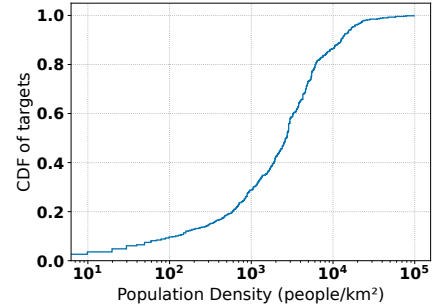


Figure 8: The distribution of the population density of our targets dataset

Figure 8 shows the CDF of population density per target. Like in the replicated work, we have both targets in rural and urban areas.

REFERENCES

- [1] CAIDA AS classification dataset. URL <https://www.caida.org/catalog/datasets/as-classification/>.
- [2] Geonames API. <https://www.geonames.org/export/ws-overview.html>.
- [3] IP info website. <https://ipinfo.io>.
- [4] MaxMind website. <https://www.maxmind.com/>.
- [5] Nominatim website. <https://nominatim.org>.
- [6] OpenStreetMap website. <https://www.openstreetmap.org/>.
- [7] Overpass API. https://wiki.openstreetmap.org/wiki/Overpass_API.
- [8] Dataset Center for International Earth Science Information Network - CIESIN - Columbia University, 2018 Gridded Population of the World, Version 4 (GPWv4): Population Density, Revision 11 Palisades, New York NASA Socioeconomic Data and Applications Center (SEDAC). <https://sedac.ciesin.columbia.edu/data/set/gpw-v4-population-density-rev11/data-download>.
- [9] RFC 9092 on geofeeds. <https://datatracker.ietf.org/doc/rfc9092/>.
- [10] RIPE Atlas anchors information. <https://atlas.ripe.net/about/anchors/>.
- [11] RIPE probe traffic API. https://stat.ripe.net/data/atlas-probe-traffic/data.json?probe_id=?,
- [12] USC/ISI ANT datasets. URL <https://ant.isi.edu/datasets/all.html>.
- [13] PCH, 2023. URL <https://www.pch.net/ixp/data>.
- [14] Peering DB, 2023. URL <https://www.peeringdb.com/>.
- [15] Malte Appel, Emile Aben, and Romain Fontugne. Metis: Better Atlas vantage point selection for everyone. In *Proc. IFIP/IEEE TMA*, 2022.
- [16] CAIDA. The CAIDA ark IPv4 Internet topology data kits dataset. <https://www.caida.org/catalog/datasets/internet-topology-data-kit/>.
- [17] Patricia Callejo, Marco Gramaglia, Ruben Cuevas, and Angel Cuevas. A deep dive into the accuracy of IP geolocation databases and its impact on online advertising. *IEEE Transactions on Mobile Computing*, 2022.
- [18] Ítalo Cunha, Pietro Marchetta, Matt Calder, Yi-Ching Chiu, Bruno VA Machado, Antonio Pescapè, Vasileios Giotsas, Harsha V Madhyastha, and Ethan Katz-Bassett. Sibyl: a practical internet route oracle. In *Proc. USENIX NSDI*, 2016.
- [19] David Dagon, Guofei Gu, Christopher P Lee, and Wenke Lee. A taxonomy of botnet structures. In *Proc. ACSAC. IEEE*, 2007.
- [20] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapè. Analysis of country-wide internet outages caused by censorship. In *Proc. ACM IMC*, 2011.
- [21] Ovidiu Dan, Vaibhav Parikh, and Brian D Davison. IP geolocation through reverse DNS. *ACM Transactions on Internet Technology*, 22(1):1–29, 2021.
- [22] Amogh Dhamdhere, David D Clark, Alexander Gamero-Garrido, Matthew Luckie, Ricky KP Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex C Snoeren, and Kc Claffy. Inferring persistent interdomain congestion. In *Proc. ACM SIGCOMM*, 2018.
- [23] Ben Du, Massimo Candela, Bradley Huffaker, Alex C Snoeren, and Kc Claffy. RIPE IPmap active geolocation: Mechanism and performance evaluation. *ACM SIGCOMM Computer Communication Review*, 50(2):3–10, 2020.
- [24] Brian Eriksson, Paul Barford, Joel Sommers, and Robert Nowak. A learning-based approach for IP geolocation. In *Proc. PAM*, 2010.

- [25] Xun Fan and John Heidemann. Selecting representative IP addresses for internet topology studies. In *Proc. ACM IMC*, 2010.
- [26] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Han Zhang, Roya Ensafi, and Christos Papadopoulos. A look at router geolocation in public and commercial databases. In *Proc. ACM IMC*, 2017.
- [27] Vasileios Giotsas, Amogh Dhamdhere, and Kimberly C Claffy. Periscope: Unifying looking glass querying. In *Proc. PAM*, 2016.
- [28] Vasileios Giotsas, Thomas Koch, Elverton Fazzion, Ítalo Cunha, Matt Calder, Harsha V Madhyastha, and Ethan Katz-Bassett. Reduce, Reuse, Recycle: Repurposing existing measurements to identify stale traceroutes. In *Proc. ACM IMC*, 2020.
- [29] Matthieu Gouel, Kevin Vermeulen, Olivier Fourmaux, Timur Friedman, and Robert Beverly. IP geolocation database stability and implications for network research. In *Proc. IFIP/IEEE TMA*, 2021.
- [30] Matthieu Gouel, Kevin Vermeulen, Maxime Mouchet, Justin P Rohrer, Olivier Fourmaux, and Timur Friedman. Zeph & Iris map the internet: A resilient reinforcement learning approach to distributed IP route tracing. *ACM SIGCOMM Computer Communication Review*, 52(1):2–9, 2022.
- [31] Bamba Gueye, Artur Ziviani, Mark Crovella, and Serge Fdida. Constraint-based geolocation of internet hosts. *IEEE/ACM Transactions On Networking*, 14(6):1219–1232, 2006.
- [32] Zi Hu, John Heidemann, and Yuri Pradkin. Towards geolocation of millions of IP addresses. In *Proc. ACM IMC*, 2012.
- [33] Bradley Huffaker, Marina Fomenkov, and kc claffy. DRoP: DNS-based router positioning. *ACM SIGCOMM Computer Communication Review*, 44(3):5–13, July 2014.
- [34] Ethan Katz-Bassett, John P. John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. Towards IP geolocation using delay and topology measurements. In *Proc. ACM IMC*, 2006.
- [35] Ethan Katz-Bassett, Harsha V Madhyastha, Vijay Kumar Adhikari, Colin Scott, Justine Sherry, Peter Van Wesep, Thomas E Anderson, and Arvind Krishnamurthy. Reverse traceroute. In *Proc. USENIX NSDI*, 2010.
- [36] Matthew Luckie, Bradley Huffaker, Alexander Marder, Zachary Bischof, Marianne Fletcher, and K Claffy. Learning to extract geographic information from internet router hostnames. In *Proc. ACM CoNEXT*, 2021.
- [37] Gregor Maier, Anja Feldmann, Vern Paxson, and Mark Allman. On dominant characteristics of residential broadband internet traffic. In *Proc. ACM IMC*, 2009.
- [38] Venkata N Padmanabhan and Lakshminarayanan Subramanian. An investigation of geographic mapping techniques for internet hosts. In *Proc. ACM SIGCOMM*, 2001.
- [39] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. IP geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2):53–56, 2011.
- [40] RIPE IMap ftp. RIPE Atlas daily dumps. <https://ftp.ripe.net/ripe/ipmap/>.
- [41] RIPE NCC. RIPE Atlas, 2019. <https://atlas.ripe.net/>.
- [42] Quirin Scheitle, Oliver Gasser, Patrick Sattler, and Georg Carle. Hloc: Hints-based geolocation leveraging multiple measurement frameworks. In *Proc. IFIP/IEEE TMA*, 2017.
- [43] Kevin Vermeulen, Stephen D Strowes, Olivier Fourmaux, and Timur Friedman. Multilevel MDA-Lite Paris traceroute. In *Proc. ACM IMC*. ACM, 2018.
- [44] Kevin Vermeulen, Justin P Rohrer, Robert Beverly, Olivier Fourmaux, and Timur Friedman. Diamond-miner: Comprehensive discovery of the internet’s topology diamonds. In *Proc. USENIX NSDI*, 2020.
- [45] Kevin Vermeulen, Ege Gurmericiler, Ítalo Cunha, Dave Choffnes, and Ethan Katz-Bassett. Internet scale reverse traceroute. In *Proc. ACM IMC*, 2022.
- [46] Yong Wang, Daniel Burgener, Marcel Flores, Aleksandar Kuzmanovic, and Cheng Huang. Towards street-level client-independent IP geolocation. In *Proc. USENIX NSDI*, 2011.
- [47] Zachary Weinberg, Shinyoung Cho, Nicolas Christin, Vyas Sekar, and Phillipa Gill. How to catch when proxies lie: Verifying the physical locations of network proxies with active geolocation. In *Proc. ACM IMC*, 2018.
- [48] Bernard Wong, Ivan Stoyanov, and Emin Gün Sirer. Octant: A comprehensive framework for the geolocalization of internet hosts. In *Proc. USENIX NSDI*, 2007.
- [49] Inja Youn, Brian L Mark, and Dana Richards. Statistical geolocation of internet hosts. In *Proc. IEEE ICC*, 2009.
- [50] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. ASdb: a system for classifying owners of autonomous systems. In *Proc. ACM IMC*, 2021.