



HAL
open science

GeoGiant: Vers une géolocalisation d'adresses IP à l'échelle grâce aux géants d'Internet

Hugo Rimlinger, Kevin Vermeulen, Timur Friedman, Olivier Fourmaux

► To cite this version:

Hugo Rimlinger, Kevin Vermeulen, Timur Friedman, Olivier Fourmaux. GeoGiant: Vers une géolocalisation d'adresses IP à l'échelle grâce aux géants d'Internet. CoRes 2024: 9èmes Rencontres Franco-phones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication, May 2024, Saint-Briac-sur-Mer, France. à paraître. hal-04566957

HAL Id: hal-04566957

<https://hal.science/hal-04566957v1>

Submitted on 2 May 2024

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.

GeoGiant: Vers une géolocalisation d'adresses IP à l'échelle grâce aux géants d'Internet

Hugo Rimlinger^{1,2}, Kevin Vermeulen³, Timur Friedman^{1,2} et Olivier Fourmaux¹

¹ Sorbonne Université, CNRS, LIP6, 4 place Jussieu, 75005 Paris

² LINCS, 19 place Marguerite Perey, 91120 Palaiseau

³ LAAS-CNRS, 7 avenue du Colonel Roche, 31031 Toulouse

Malgré l'importance de la géolocalisation des adresses IP, les chercheurs sont souvent limités par l'utilisation de bases de données privées, opaques et peu fiables. RIPE Atlas offre une solution potentielle en fournissant des milliers de points de mesure dans le monde, permettant la création d'une base de données géolocalisée à l'échelle d'Internet. Malgré d'importants efforts pour y parvenir, nous ne disposons pas pour le moment de méthodes qui nous permettent de réaliser cet objectif. En nous appuyant sur l'hypothèse selon laquelle deux adresses IP géographiquement proches seront redirigées vers un même point de présence d'un CDN, nous sommes en mesure de sélectionner un nombre limité de points de mesure, suffisants pour assurer une précision similaire à celle obtenue en utilisant l'ensemble des points disponibles. Cette approche, basée sur l'option *EDNS Client Subnet* (ECS), permet de passer à l'échelle la géolocalisation par mesures actives, indépendamment de la taille de la plateforme de mesure considérée.

Mots-clefs : IP addresses geolocation, Internet, Network measurements

1 Introduction and related work

Determining the location of an IP address is vital for a variety of sectors. It helps the content distribution industry to enhance user experience through efficient delivery and to enforce region-specific intellectual property rights. It helps network operators in their efforts to trace cybersecurity attacks. And it helps scientists to gain a deeper understanding of the structure of the Internet. Despite years of effort, achieving geolocation on an Internet scale, which is to say being able to geolocate all active IP addresses, continues to be a significant challenge. As a result, researchers often rely on geolocation databases provided by private companies such as MaxMind, which are generated using opaque techniques and which can produce incorrect results [GSH⁺17]. The research community now stands on the brink of a significant shift towards independence, thanks to the growth of the RIPE Atlas infrastructure. This system now boasts around 10,000 probing agents, enabling the launch of active measurements like ping and traceroute from vantage points around the world. Given that the accuracy of ping-based geolocation hinges on the proximity of a vantage point to the IP address being located [DRD⁺23], the size and coverage of this infrastructure paves the way for free and open geolocation.

Besides commercial solutions, latency based solutions (CBG, shortest ping [GZCF04]) are usually privileged by researchers because they provide a way to constrain the geolocation of an IP address. Previous studies were able to scale those techniques to the entire IPv4 address space by leveraging the fact that only a subset of closest vantage point to the source is sufficient to maximize geolocation precision [HHP12]. Previous efforts relied on probing from all vantage points available to find the best set of vantage point to geolocate any IP address [HHP12]. Although this approach was possible on smaller platforms, it fails to scale to the immensity of RIPE Atlas [DRD⁺23]. Other techniques like RIPE IMap, determine the set of closest vantage point based on network information [DCH⁺] [RIP]. Nevertheless, this technique is constrained to IXPs and lacks scalability for Internet-wide geolocation.

In this paper, we take a step towards Internet-scale geolocation by framing it as a search problem. The primary question we address is : given the large universe of RIPE Atlas vantage points, how can we identify

a small subset from which to conduct measurements towards a specific IP address? We can significantly reduce the load on the RIPE Atlas infrastructure if each of the billions of IP addresses is geolocated from its own small subset. We use ECS, an optional extension to DNS, for this search.

2 Efficient vantage point selection with ECS

The EDNS Client Subnet option (ECS) was designed to enable a CDN operator to consider the client’s network when choosing which server IP address should resolve a particular hostname for that client. Whereas a standard DNS query can be paraphrased as : “What is the IP address for hostname X?”, with ECS it becomes : “What is the IP address for hostname X, when asked on behalf of a client in network Y with prefix Z?” Here, Z is 24 or less for IPv4 and 56 or less for IPv6. ECS queries are not required to originate from the client’s network, so we can issue queries for all active /24 and /56 networks from a single machine. This work only focuses on IPv4.

Our intuition is that a CDN will typically return a server IP address geographically close to the client network, so as to reduce client latency. Consequently, two client networks that are directed to the same server address are likely to be geographically close. Major CDNs such as Google, Meta, and Amazon use ECS. For a given target IP address, we can formulate a ‘fingerprint’ based on several hostnames from these CDNs : this is the vector of server IP addresses to which the target’s /24 is directed for each of those hostnames. We can also generate such fingerprints for the IP addresses of the RIPE Atlas vantage points. Based on the similarity between target and vantage point fingerprints, we can construct a small subset of vantage points, presumably close to the target, from which to issue geolocation pings.

We identify two challenges : identifying CDN hostnames for which ECS is used to reduce latency, and formulating a metric for assessing fingerprint similarity.

Hostname selection : To determine whether a hostname supports ECS, we can look at whether the *source-scope* field of the ECS option in the DNS response is non-zero. To decide if ECS is used to minimize latency, we use a simple heuristic : if two client networks in different regions are mapped to the same server, it is likely that ECS is not used to reduce latency, so we do not consider the hostname.

Fingerprint similarity : We defined the similarity $S_{t,v,h}$ of two ECS fingerprints for a given hostname as the size of their respective fingerprint intersection, divided by the size of the smallest fingerprint. The overall similarity $S_{t,v}$ between two ECS fingerprint is defined as the average similarity over all hostnames :

$$S_{t,v} = \text{avg}_{h \in H} \left(\frac{\text{len}(F_{t,h} \cap F_{v,h})}{\min(\text{len}(F_{t,h}), \text{len}(F_{v,h}))} \right) \quad (1)$$

where t is the target’s subnet, v is a vantage point’s subnet and $F_{t,h}$ and $F_{v,h}$ denote their respective ECS fingerprint for the hostname h (H being the entire set of hostnames). We then rank the vantage points based on their similarity, the higher the better.

There are some refinements that are needed to build the fingerprint. For instance, a client can be served by multiple front-end servers that are located at a single CDN point of presence, meaning that a difference in IP addresses does not necessarily indicate a difference in location.

3 Preliminary results

We evaluate our methodology on a validation dataset of 772 IP addresses that we have chosen to represent target ground truth, as we consider them to be reliably geolocated. These are the RIPE Atlas ‘anchors’, which are larger servers to which the infrastructure’s administrators pay particular attention. Despite the dataset’s limited size, it allows us to test our assumptions across diverse locations, spanning 93 countries and 6 continents. It is, however, biased towards western Europe and North America. We employ the lighter-weight RIPE Atlas ‘probes’, of which there are approximately 10k, as vantage points.

We select 60 hostnames hosted by CDNs listed in the Google CrUX top domain dataset [CrU]. After our filtering step we obtain a set of 15 hostnames, owned by five large CDNs (Google, Meta, Amazon, Apple, and CDNetworks).




FIGURE 1 – CDFs of the proportion of target function of : Figure 1a the distance error and Figure 1b the measured latency

For each target, we perform the ECS measurements, compute the fingerprint, and issue pings from the top N vantage points with the highest similarity. We then extract the vantage point with the minimum RTT, and assign its geolocation to the target. To compute a reference of what should be achievable with ping measurements, we also issue pings from the 1,000 geographically closest vantage points to the target and assign the geolocation of the one with the minimum RTT to the target as the reference geolocation.

Figure 1a shows the cumulative distribution of the distance between the assigned geolocation and the actual geolocation of a target, for different numbers of vantage points (VPs) selected to issue the pings, from 1 to 500 VPs, and the reference. The median error varies from 6 km to 119 km for 500 VPs and 1 VP, while the reference is also at 6 km, showing that the optimal VP is within the 500 first VPs according to our ranking. Taking the 40 km error threshold defining an accurate city level geolocation, we can geolocate between 37% and 78% of targets at city level.

In absence of ground truth for the geolocation of the RIPE Atlas ‘probes’ which are providing the VPs, one must rely on the RTT obtained from the ping measurements to bound the distance error. Theoretically, a 1 ms RTT can correspond to up to 100 km between the VP and the target, taking $\frac{2}{3}c$ as the speed of packets over the Internet where c is the speed of light, but prior work used less conservative thresholds to say that two IP addresses were located in the same city, such as 2 ms [GKF⁺20], to take path inflation into account. Figure 1b shows the cumulative distribution of the minimum RTT obtained from the ping measurements, for different numbers of selected VPs. Using only the VP with the most similar fingerprint (orange curve), 23% of the IP addresses have a minimum RTT of less than 2 ms, while it is 76% when using 500 VPs.

4 Challenges

Although the validation dataset that we used to evaluate our methodology is larger than those found in previous studies [HHP12], it is important to acknowledge its limitation in validating our methodology at Internet scale. This is because the distribution of RIPE Atlas anchors tends to be biased towards western Europe and North America, two regions with highly efficient network infrastructure. We might well encounter worse results were we to consider more targets elsewhere in the world. Additionally, the validation dataset does not adequately represents the broad diversity of IP addresses. Trying to geolocate the IP address of a router, a server, or an end-user host presents different issues. For instance, the last-mile delay effect on end hosts can inflate latency of 10 ms, making latency constraint approaches impractical. The challenge to expanding our validation dataset is to find targets outside of RIPE Atlas that are reliably geolocated.

Furthermore, we rely on the smallest granularity permitted by ECS to determine the geographical proximity of two IP addresses. Consequently, we assume a proximity of all IP addresses within the same /24 subnet. While this assumption has been employed in previous studies [HHP12], we acknowledge its limitation. Notably, it does not hold true for certain cases, such as tier-1 networks or trans-Atlantic links with

continent-spanning /24 subnets.

Finally, while our methodology is effective in identifying a restricted set of valuable vantage points for geolocating an IP address, it falls short of determining the singular closest one. Consequently, the vantage point with the highest ECS similarity is often not the closest in geographical proximity. Our ultimate objective is to pinpoint this closest vantage point, rendering our methodology independent of pings and suitable even for IP addresses that do not respond to ping requests.

5 Conclusion

In this paper, we introduced a novel approach for choosing a small set of vantage points to geolocate any IP address, leveraging the EDNS Client Subnet option (ECS). Unlike previous techniques, our methodology exclusively depends on publicly accessible data, and it is both lightweight and scalable. This scalability lends itself to Internet-scale measurements from the RIPE Atlas infrastructure. Our results demonstrate that the hypothesis of considering two IP addresses with similar ECS fingerprints as geographically close is valid and useful for limiting the number of pings needed to geolocate an IP address.

Références

- [BGP] Bgptools. <https://github.com/bgptools/anycast-prefixes/>.
- [CrU] Crux-top-lists. <https://github.com/zakird/crux-top-lists>.
- [DCH⁺] Ben Du, Massimo Candela, Bradley Huffaker, Alex C Snoeren, and KC Claffy. RIPE IP-Map active geolocation : Mechanism and performance evaluation. *ACM SIGCOMM Computer Communication Review*.
- [DRD⁺23] Omar Darwich, Hugo Rimlinger, Milo Dreyfus, Matthieu Gouel, and Kevin Vermeulen. Replication : Towards a publicly available internet scale ip geolocation dataset. In *Proc. IMC*, 2023.
- [GKF⁺20] 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. IMC*, 2020.
- [GSH⁺17] 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. IMC*, 2017.
- [GZCF04] Bamba Gueye, Artur Ziviani, Mark Crovella, and Serge Fdida. Constraint-based geolocation of internet hosts. In *Proc. SIGCOMM*, 2004.
- [HHP12] Zi Hu, John Heidemann, and Yuri Pradkin. Towards geolocation of millions of ip addresses. In *Proc. IMC*, 2012.
- [RIP] Ripe ipmap. <https://ipmap.ripe.net/>.
- [rou] routeviews. <https://www.routeviews.org/routeviews/>.
- [SBA⁺20] Raffaele Sommese, Leandro Bertholdo, Gautam Akiwate, Mattijs Jonker, Roland van Rijswijk-Deij, Alberto Dainotti, KC Claffy, and Anna Sperotto. Manycast2 : Using anycast to measure anycast. In *Proc. IMC*, 2020.
- [SBC⁺13] Florian Streibelt, Jan Böttger, Nikolaos Chatzis, Georgios Smaragdakis, and Anja Feldmann. Exploring EDNS-client-subnet adopters in your free time. In *Proc. IMC*, 2013.