IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received 14 April 2023, accepted 22 May 2023, date of publication 31 May 2023, date of current version 20 June 2023.

Digital Object Identifier 10.1109/ACCESS.2023.3281333

RESEARCH ARTICLE

On Improving the Accuracy of Internet Infrastructure Mapping

PAUL MCCHERRY, VASILEIOS GIOTSAS, AND DAVID HUTCHISON
School of Computing and Communications, Lancaster University, LA1 4WA Lancaster, U.K.

Corresponding author: Paul McCherry (p.mccherry@lancaster.ac.uk)

This work was supported by the National Cyber Security Centre, U.K., under Grant 4211350/RFA 15844.

ABSTRACT This study identifies a method to create fine-grained multilayer maps of the Internet's structure, which are currently lacking. We begin with an investigation of current techniques for geolocating hosts using passive, active, and hybrid methods. This is followed by a survey of the fundamental problems that IP geolocation techniques must address. The survey points to the obvious difficulties in using Delay-Distance models and suggests that the use of Return Trip Times can lead to highly misleading results. We therefore develop a new procedure that combines state-of-the-art methods to avoid many of the fundamental problems in Internet topology mapping, whilst creating finer-grained internet maps than those currently available. This procedure is tested on the UK infrastructure by conducting a series of tests using distributed measurement points provided by the RIPE Atlas platform. Our results show that we can accurately geolocate routers between two endpoints to create a fine-grained map of the internet infrastructure involved in our measurements. Researchers have long recognized the scarcity of ground truth datasets where IP geolocation is a concern. As a byproduct of our new method reported in this paper, we create a validation dataset that maps hundreds of IP addresses to geo-coordinate landmarks or vantage points, which is highly desirable for IP geolocation research. Finally, we discuss some limitations of this method, and we summarise the next steps toward accurate and complete Internet infrastructure maps.

INDEX TERMS Co-location facility, IP geolocation techniques, internet mapping, interconnection services, internet topology,

I. INTRODUCTION

Knowledge of the geographical locations of the Internet infrastructure is a necessary requirement for cyber situational awareness and can allow us to understand and mitigate risks related to topological vulnerabilities and design more resilient networks and routing policies [1]. For instance, the ability to predict what would happen if a co-location facility or Internet eXchange (IX) fails, can inform better fallback policies and more efficient resource allocation. To develop appropriate prediction techniques, we must measure the relevant routing paths and infer the traversed interconnection points. Analysis of those paths can provide clues about connectivity changes that will prevent choke points, single points of failure, or serious performance degradation owing to the failure of a facility. Consequently, researchers and network engineers can design and evaluate new protocols and services or analyse the vulnerability of the network infrastructure.

The first step in developing improved infrastructure maps is to investigate and assess the fundamental limitations of state-of-the-art Internet cartography. The scarcity of valid ground truth data sources is a classic problem in IP Geolocation. Motamedi et al. [2] remarks on the notoriety of a lack of ground truth data sets. This scarcity of valid data sources means that researchers must rely on incomplete or coarse-grained abstractions of Internet topology. These abstractions miss many interconnection details and render them largely irrelevant to real-world Internet engineering problems [3], and many of the findings that rely on simplistic models are controversial or misleading, owing to the incompleteness and inaccuracies of the produced

The associate editor coordinating the review of this manuscript and approving it for publication was Byung-Seo Kim.

maps [4]. A recent report by the UK National Cyber Security Centre (NCSC) [5] highlighted the importance of the stability of IP-based networks and argued that the need for increased security of routing information that underpins the delivery of Internet services has increased dramatically. It is therefore paramount to develop appropriate methods and practices to make the Border Gateway Protocol (BGP) more secure, and thus maintain the integrity of the routing system which relies almost exclusively on BGP. The European Network and Information Security Agency (ENISA) concluded in a 2015 report [6] that the current lack of structural transparency is the biggest obstacle to tackling the inherent vulnerabilities and architectural shortcomings of the Internet routing system.

In this paper we research the current techniques in IP geolocation, and we build on that research to develop a new method in IP infrastructure mapping with the ultimate goal of creating fine-grained multilayer maps of the Internet infrastructure that are currently lacking. We exhibit the utility of our technique by mapping the UK Internet interconnection facilities after conducting a large-scale distributed measurement campaign using the RIPE Atlas platform [7].

The aims of this work are as follows:

- To extend DNS-based geolocation from city-level to facility-level and address shortcomings of the state-of-the-art with respect to their limited geographical coverage.
- To introduce a new technique to create constraints in DNS geohints inference. While past work has relied on RTT measurements, our work uses traceroute-derived constraints by combining IXP datasets with forward and reverse traceroute measurements to observe the bi-direction interfaces.
- To construct a dataset of facility-level landmarks that can be used in future research work to improve RTT-based geolocation.
- To illustrate the applicability of our work by geolocating a number of IPs at the level of colocation facilities, and then show that our method can create detailed maps of interconnection infrastructures at large metropolitan Internet hubs like London.
- To evaluate our inferences and estimate its success using a carefully curated dataset obtained by two of the largest London IXPs.

We review the related work in Section II, focusing on methods for geolocating hosts and describing the issues and challenges surrounding these methods. In Section III, we provide an overview of our innovative approach towards geolocating the Internet infrastructure and explain our contributions in terms of tools and methods. In Section IV, we describe this procedure in detail. In Section V, we provide an example of this procedure and the results that can be obtained along with the validation methods. Section VI provides details of the rules and methods used to automate the procedure and the results obtained. Section VII discusses these results while Section VIII discusses to what extent the study has met its aims. Section IX summarises what has been achieved and gives an indication of the next steps that are envisaged, including further research.

II. RELATED WORK

IP-based geolocation maps an IP address to the geographic location of a real-world internet-connected device. IP geolocation can attempt to map an IP address to different granularities, including latitude and longitude, interconnection facility, metropolitan area or country. IP geolocation methods can be broadly classified into three types: passive, active, and hybrid.

A. PASSIVE IP GEOLOCATION

Passive methods involve the collection and synthesis of geolocation information from databases and websites. For example, Domain Name Service (DNS) LOC records are DNS records proposed in 1996 in RFC1876 [8] which are designed to hold the geographical coordinates of the IP address host. However, they are rarely created by administrators [9].

Another source of passive geolocation data is the WHOIS protocol [10], which stores information on the owners of Internet resources, including IP addresses. Among this information is often the address of the organization or individual to which an IP address is assigned. WHOIS servers are operated by the five Regional Internet Registers (RIRs), which are also responsible for the allocation and registration of Internet resources. However, it is often left to network administrators to update the information which can become outdated without timely maintenance. In addition, WHOIS maps IP addresses to a registered administrative location, which may not reflect their actual location.

Geofeeds, another example of passive IP geolocation, are self-published IP geolocation data that provide geolocation coordinates and are described in the Internet Engineering Task Force (IETF) self-published IP geolocation data RFC8805 [11]. Finally, several commercial geolocation services that use proprietary methods provide location data to subscribers, such as Maxmind [12], IP2Location [13], and Neustar [14]. However, past research on the accuracy of those databases shows that especially for router and infrastructure IPs commercial databases can be highly inaccurate [15], [16], [17].

B. ACTIVE IP GEOLOCATION

Active IP geolocation relies on network-level latency measurements between a node with a well-known location (landmark) and the IP which we want to geolocate. Assuming that the Speed of Internet (SoI) is known, the latency can be then translated to distance from the landmark [18], [19], [20], [21], [22], [23], [24], [25], [26], [27]. While active geolocation tends to be more accurate than passive geolocation, it incurs a much higher measurement overhead and it is hard to scale to geolocation of millions of IPs. Additionally, the SoI is not fixed but depends on transmission medium and the network conditions. Geoping is

one of the earliest active geolocation techniques introduced in 2001 by Padmanabhan et al. [18]. Geoping measures the latency between multiple landmarks and creates a latency vector for each landmark. It then measures the latency from all landmarks to the target IP and geolocates it to the landmark with the most similar latency vector. In 2006, Gueye et al. [19] proposed 'Constraint-Based Geolocation' (CBG) as an enhancement of Geoping. CBG also employs measurements from multiple landmarks, but it combines the measured delays using multilateration which can geolocate IPs not only in the locations of the landmarks but also in the area between them. Zu et al. [20] propose a new algorithm based on router error training, which requires an exhaustive mapping of the Metropolitan Area Network (MAN) of the city where the target IP should be located in order to infer its street-address location. While this technique achieves high accuracy; it is limited to cities with a suitably large number of measurement vantage points.

C. HYBRID IP GEOLOCATION METHODS

Hybrid IP Geolocation techniques aim to combine passive and active geolocation to alleviate their individual limitations. To depart from oversimplified models, it has been argued that it is necessary to identify the geolocation of Internet infrastructure [5], which would provide a useful tool for detecting poor routing structures and understanding why damaging routing events occurred. A method was proposed in which a combination of data sources could be used, such as crowdsourcing, reverse DNS records, tagged naming schemes, Return Trip Time (RTT) delay-distance models, and Internet Exchange Points.

In "Topology Based Geolocation" (TBG), Katz-Basset et al. [26] argue that the directness of a network path from a landmark to a particular target cannot be predicted, and a single conversion factor for the entire network is not sufficient to capture the intricate details of the network topology and routing policy. This method also uses multilateration, as used in CBG, but issues traceroute measurements instead of pings to map the entire IP path between a landmark and the target IP. The intermediate IP hops are geolocated using location hints in their reverse DNS records, allowing more detailed knowledge of the network, and the traversed locations.

Spotter [27] uses a probabilistic approach to derive a generic model of the relationship between network delay and geographic distance instead of using either a pre-determined IoS value or separate calibration data for each landmark. This delay distance model was then used to geolocate an IP.

Wang et al. [28] refined the granularity of CBG to achieve street-level geolocation. To this end they mine web-based geolocation hints for locally hosted web servers to significantly expand the list of passive landmarks. They observed that they try to leverage the observation that "*many entities host their Web services locally*", but since then, the trend of cloud-hosted services and resource centralization certainly inhibit the applicability of their technique.

Octant [29] also considers the locations of intermediate routers as landmarks to geolocate the target. Furthermore, Octant considers both positive and negative constraints, which define where the node can and cannot be. Then it tries to geolocate the target IP as an error minimization constraint satisfaction problem. Even though Octant achieves better accuracy that CGB, the authors pointed out that extracting useful positive and negative information is a challenge. In contrast to CBG and TBG, the authors allowed for circuitous routes and employed intermediate routers as secondary landmarks to reduce the latency errors caused by this issue. Octant refers to a proprietary database of router DNS names for geographical locations, to use routers as secondary landmarks. Their conclusion was that, in many cases, the closer the landmark, the greater the accuracy, which is a common finding in all active geolocation methods.

Scheitle et al. developed a method called Hints-Based Geolocation (HLOC) [30], which extracts geo-hints from router DNS names, similar to Octant. It then validates these hints by selecting several RIPE Atlas probes based on the extracted geo-hints and measuring the RTT values between them and the domain. This solution compares a previously compiled database of router DNS names and codes with target DNS names. Interestingly, the authors investigated and proposed a latency delay of 9ms over a maximum distance of 900 km to accommodate the packet buffering, processing, and scheduling delays. If the total latency is considered low, the target geocoordinates are assumed to be those of the router, and the hint provided by the router's DNS name is verified.

RIPE IPMAP [7] is a multi-engine geolocation platform operated by RIPE NCC that uses active IP geolocation as well as passive methods to locate the geographical coordinates of the targets. One of the IPMAP geolocation engines uses a method called Single Radius, which first finds the AS that announces the prefix that contains the target IP, and then locates the RIPE Atlas probes that are close to the target IP. Pings are then sent from these probes, and any delays of over 10ms are discounted. The probe with the minimum latency to the target IP is then selected, and the distance is calculated using the the signal transmission speed through the optical fibre of 0.66c. All cities within this distance from the probe are then ranked by numerous factors, such as population density, and the highest ranked one is inferred as the location of the IP. A major problem is that RIPE probes are heavily biased toward Europe and North America and become quite sparse in Latin America, Asia, and Africa. This may indicate that other methods, such as the shortest ping or CBG, yield comparable results in these regions.

Livadariu et al. [31] identified that DNS names do not accurately map geolocations without improved lookup tables and proposed the use of Looking Glass servers[1] as additional landmarks. They also investigated the accuracy of RIPE

[1]Looking Glass servers provide non-privileged interfaces to BGP routers that allow execution of basic commands, such as the querying of a route, or the exestuation of traceroute and ping measurements.

IPMAP against various methods, such as WHOIS, DNS, geolocation databases, and HLOC. They find that various approaches can disagree even at the country level and raise the point that organisations may be unaware of the countries through which their traffic is routed. They also found that geolocation databases fail to accurately locate IPs that belong to international ASes on many occasions, and that commercial geolocation databases appear to use information from WHOIS, which can often be wrong, as their primary source of data.

Luckie et al. [32]. demonstrated significantly improved DNS to geolocation lookups by compiling an extensive list of regular expressions. Dan et al. [33] applied Machine Learning to the task of learning DNS names and their locations, showing that their work significantly outperformed previous academic baselines, and was complementary and competitive with commercial databases.

Dan et al. [34] proposed an IP geolocation technique that exploits the concept of IP interpolation, according to which if at least two IPs within a /24 prefix are in the same location, then all IPs in that prefix are also in that location. Additionally, they exploit the observation that there is strong correlation between delay differences along a traceroute path, physical distance.

Giotsas et al. [35] developed a method known as Constrained Facility Search (CFS), which combines data from various sources such as Internet Exchange websites, PeeringDB [36] and traceroute measurements to infer the connection facility of a specific IP address. Using this method, they were able to locate 71% of the router interfaces to a specific facility. Motamedi et al. [2] extended the geolocation of interconnection facilities to private and cloud interconnections using the Belief Propagation algorithm on a specially defined Markov Random Field graphical model.

D. FUNDAMENTAL PROBLEMS

Several past works determined that techniques that try to measure the Internet topology and geolocation IP addresses using traceroutes suffer several problems [1], [4], which we summarize in this section.

Layer 2 clouds are largely opaque to tools using traceroutes. Willinger and Roughan found that Internet connections that appear to have trivial or simple IP layer topologies can have complex layer-2 topologies. Technologies such as Software-Defined Networking (SDN) and Multi-Protocol Label Switching (MPLS) can further complicate this situation by creating logical layer-2 and layer-3 networks without physical devices. Measurements often see only one layer, creating misunderstandings regarding the true resilience of a network. Furthermore, traceroute-based measurements can return the RTT of a proxy server, which may be several miles away. In fact, Padmanabhan and Subramanian [18] observed that a significant fraction of a proxys clients were located several hundred to thousands of kilometers from the location of the proxies. Network Delay measurements are oblivious to this and would incorrectly return the location of the proxy server.

Traceroute RTT includes both application-layer and network-layer delays, and if a measurement device is overloaded or under-resourced then the RTT times may be inflated. This is a problem that RIPE Atlas probes may encounter (especially older versions) [37].

The RTT can also be inflated by circuitous routes, which happen when the network path between two endpoints does not follow the shortest geographical path. For instance, Figure 1 shows that an ICMP packet travels from Blackpool to Lancaster, through London and Manchester. Blackpool to Lancaster is approximately 40 km apart: however, this packet travels approximately 800 km one-way.

Generally, the RTT is divided by 2 to give a delay approximation on the one-way journey; in this case, RTT is 32ms, therefore, the one-way journey from Blackpool to Lancaster took approximately 16ms. The signal transmission speed through the optical fiber is estimated as.66 $\times$ speed of light (c) where the speed of light is approximately 3×10^8 m/s. Dividing the distance of 40 km by $0.66c$ gives the time that the packet should have taken on a direct one-way journey is 0.2 of a millisecond. However, we know that the packet travelled approximately 800 km on its one-way journey from Blackpool to Slough and returning north to Lancaster, which should have taken 4 milliseconds over this circuitous route. Therefore, the remaining 12 milliseconds should be allocated to packet scheduling, packet processing, interface delays, and other factors. Indeed, as pointed out by the authors of HLOC, they include a latency delay of 9ms to account for these issues.

Another complication of traceroute-based measurements discovered by Candela et al. [38] is the infrastructure diversity in different regions worldwide, leading to different delay coefficients. These delay coefficients are not only hard to estimate but also very dynamic, as the infrastructure and the related network phenomena can change very frequently.

E. IP GEOLOCATION SUMMARY

While active IP geolocation can provide real-time updates and requires no administrative upkeep, many active IP geolocation solutions employ active measurements, that is, traceroute, to discover network interfaces and topology. However, the traceroute tool was designed primarily for troubleshooting, and its use in network discovery was not what it was designed for. Therefore, the results cannot always be trusted.

Problems such as asymmetric forward and return paths, circuitous routes, different router configurations, route congestion, and faster-than-fiber technologies such as microwave links, SDN, ATM, and MPLS clouds can have varying effects on delay-based geolocation techniques. Passive methods suffer from out-of-date or completely incorrect information; therefore, it appears that both active and passive methods have their own strengths and weaknesses. A hybrid mix of

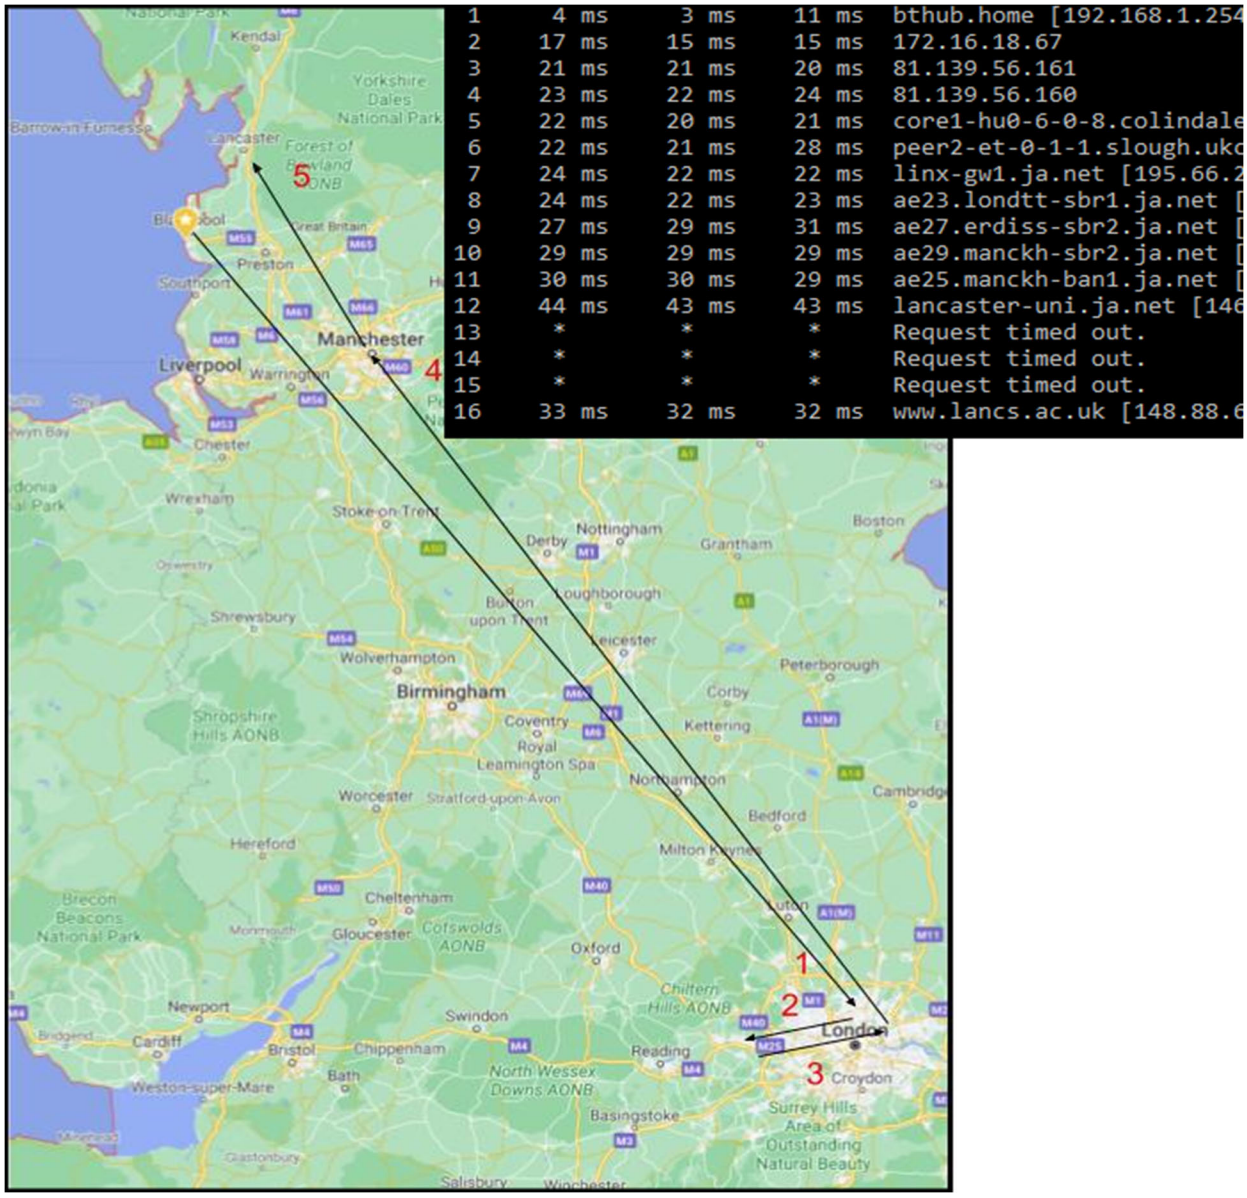

FIGURE 1. Example of traceroute measurement demonstrating circuitous route between two IP addresses.

active and passive techniques has the potential to alleviate those weaknessess and offer the most accurate IP geolocation solutions.

To depart from oversimplified models, it has been argued that it is necessary to identify the geolocation of Internet infrastructure [39], which would provide a useful tool for detecting poor routing structures and understanding why damaging routing events occurred. A method was proposed in which a combination of data sources could be used, such as crowdsourcing, reverse DNS records, tagged naming schemes, RTT delay-distance models, and Internet Exchange Points.

The key contribution of this paper is that it extends DNS-based geolocation from city-level to facility-level and addresses shortcomings of the state-of-the-art in respect to their limited geographical coverage. We showcase the applicability of our work by geolocating over a thousand IPs at the level of colocation facilities. While the dataset is small, to the best of our knowledge it is the first working prototype at this level of granularity and illustrates that our method can create detailed maps of interconnection infrastructures at large metropolitan Internet hubs like London.

III. OVERVIEW OF MAPPING METHOD

According to Motamedi et al. [1], Point of Presence (PoP) is the ideal resolution for geographical mapping of network infrastructure. Motamedi described PoP as a concentration of routers that belong to an AS; however, for the purpose of this method, PoP is simply a facility where a router and its interconnections are housed. This study proposes a

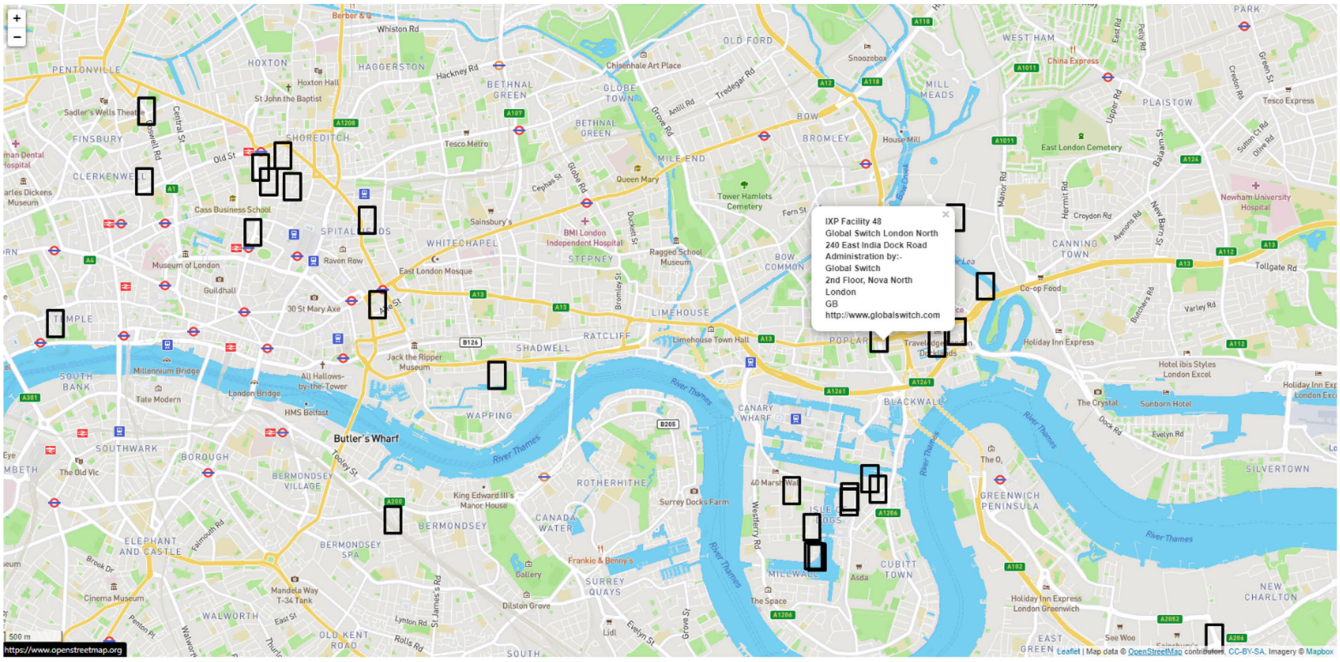

FIGURE 2. Geo-mapping interconnection facilities using PeeringDB and OpenStreetMap.

new method to map the Internet infrastructure at this PoP layer, starting with the discovery of the geolocation of each interconnection facility in the UK. PeeringDB is a freely available, user-maintained database of networks, facilities, and internet exchanges (IX) that provides comprehensive details on each facility's address, geolocation, and other useful information.

Step 1 of this process involves querying the PeeringDB REST API to obtain a list of UK-based facilities and their geographical coordinates. These facilities are then mapped onto OpenStreetMap (OSM), a collaborative project to create a freely editable geographic database of the world. PeeringDB identifies each facility through a unique identification number, which we use to reference facilities in this paper.

In **Step 2** of this process, we query again the PeeringDB API to find in which of these UK facilities Internet Exchanges have deployed their switching equipment to build an OSI layer 2 map of the IX network infrastructure. Additional information is downloaded from each Internet Exchange website, such as the connection speeds of the peering ports.

Step 3 involves the execution of traceroute measurements using the RIPE ATLAS platform, which has over 600 probes in the UK, allowing traceroutes to be created in both directions, to and from each probe, creating a mesh of over 350,000 measurements.

Step 4 maps each hop to a facility where possible using a combination of DNS lookups, Internet Exchange website information, and PeeringDB data. This information also creates a list of valuable Vantage Point (VP) information that will be useful for future research. To map these intermediate hops, a tool was created, which reads the data from the traceroute measurements created in step 3, and queries various sources, such as PeeringDB, DNS, and the Internet Exchange websites, to locate the position of the router where these hops are interfacing, considering the previously discovered facility and IXP information.

IV. THE METHOD IN ACTION

A. PREAMBLE

The first step in developing a technique for mapping Internet infrastructure involves mapping interconnection facilities to their geophysical coordinates. Internet exchange directories are publicly available in many locations, such as Packet Clearing House website (PCH) [40], IXPDB website [41], and the PeeringDB website. Among these directories, PeeringDB has the most comprehensive list [42]. A simple data extraction can be performed due to PeeringDB being a freely available database of networks that contains a well-updated list of IXs (Internet eXchange), facilities, and their geolocations, as well as a REST API. Peeringdb also facilitates the global interconnection of networks in Internet Exchanges, data centres, and other interconnection facilities. However, as Kloti et al. [42] points out PeeringDB is also incomplete, the data from some IX's are not included in the Peeringdb Database. This causes additional failures in the code to recognise the geographical location of those IP addresses registered to those IX's.

OpenStreetMap (OSM) [43] is an open-source project that creates a freely editable geographical database in which tags can be created to provide information about elements, as shown in Figure 2. A list of UK-based facilities was extracted from PeeringDB, along with the geographical coordinates of each facility. Where facility records have

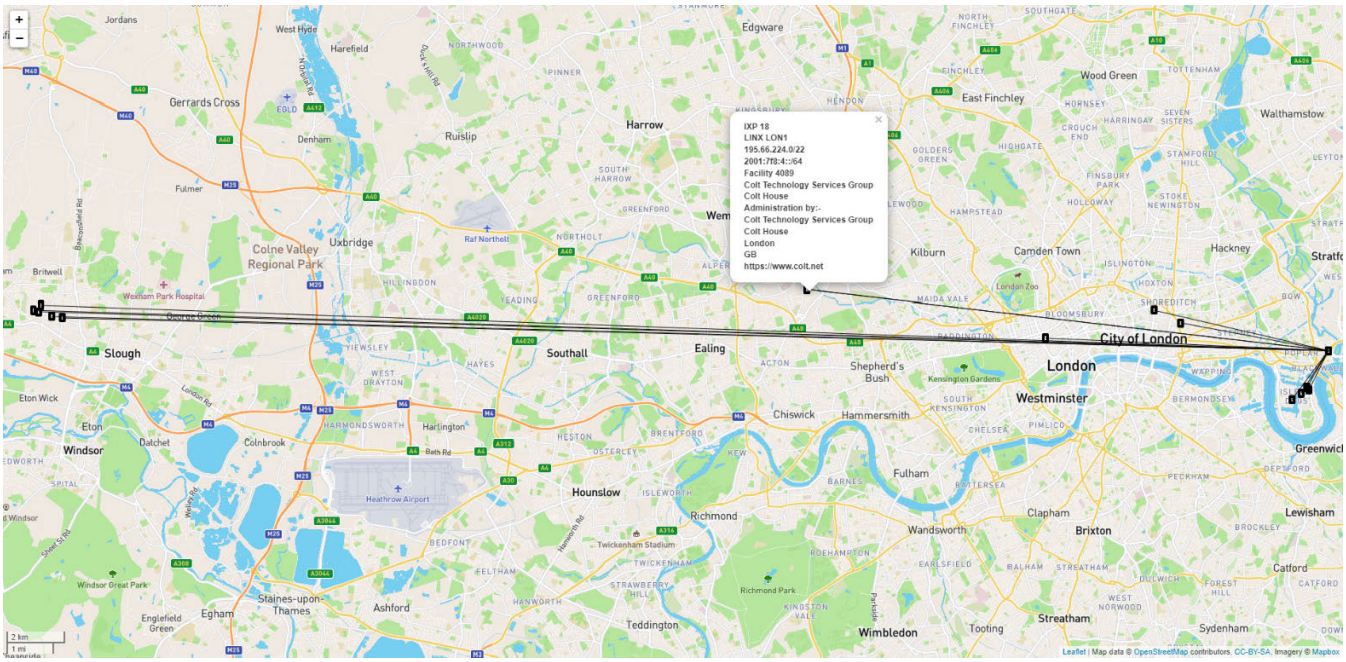

FIGURE 3. LINX LON1 IX showing its public peering facilities and its layer 2 network.

no geolocation information available, the address of the facility is input into Nominatim [44], which is a tool for searching OSM data by name and address (geocoding) and generating synthetic addresses of OSM points (reverse geocoding). There are occasions when addresses do not return any geocoding data; in this case, the address of the facility must be entered manually. Of the 235 UK facilities listed by PeeringDB (as of 16/2/2023), only sixteen facilities had to be manually geolocated. Figure 2 shows the number of facilities (black rectangles) geolocated in the London area using PeeringDB, Nominatim, and OSM.

Internet Exchanges (IX) consist of one or more switches to which participant Internet service providers, transit providers, and content delivery networks (CDN) exchange data. They are housed in co-location facilities, are generally located close to large populations, and are therefore essential to the internet network infrastructure.

IXs connect the facilities where they interconnect at layer 2; therefore, data entering an IX's network at one facility can traverse the IX network and exit at any other facility where the IX interconnects. There are two methods by which an organisation may wish to connect to an IX: direct or remote peering. Direct peering requires an organisation to have physical presence at a co-location facility where the IX also has presence, while remote peering allows an organisation to peer with the IX using one of the IX's partners over layer 2 MPLS clouds.

According to PeeringDB, there are currently 21 Internet exchanges in the UK, although two are listed with no connected networks. Packet Clearing House (PCH) lists 15 active IXs, whereas the IXPDB website lists nine IXs with connected networks. To map an Internet Exchange, PeeringDB is queried to discover the facilities at which each IX publicly interconnects, and these layer 2 networks are mapped to OpenStreetMap, as shown in Figure 3. Public information is not available to map the actual physical cables; thus, while point-to-point connections are depicted in this Figure, it is possible that connections forming a mesh network, where each point is connected to every other point, may be in use. However, the principle of the Layer 2 logical network remains the same.

Once the public peering points from PeeringDB have been mapped, we can refer to the IX's website to collect any additional IX public peering facilities that may have been missed by PeeringDB. For example, LINX London is one of the largest Internet Exchanges in the world with one of the highest numbers of participants. The information available on the LINX Internet Exchange website is comprehensive and includes the ASN, IP address, connection location, connection speed, relevant routers, ports, and port type. The location and IP information allow us to geolocate the interconnection with great accuracy, whereas the service speed allows us to understand the maximum bandwidth that a connection can use, perhaps allowing for future investigation of any cause of congestion. Additionally, many ports are marked with a port type of 'Connexions', which are LINX's reseller partners, and provide information on clients that connect using remote peering. According to the LINX website, there are eight UK facilities to which LINX LON1 interconnects, which PeeringDB has failed to list. These are connections to other LINX IXs such as LINX Manchester, LINX Wales, and LINX Scotland.

Probes and anchors on the RIPE Atlas platform were chosen to create measurements across the UK infrastructure

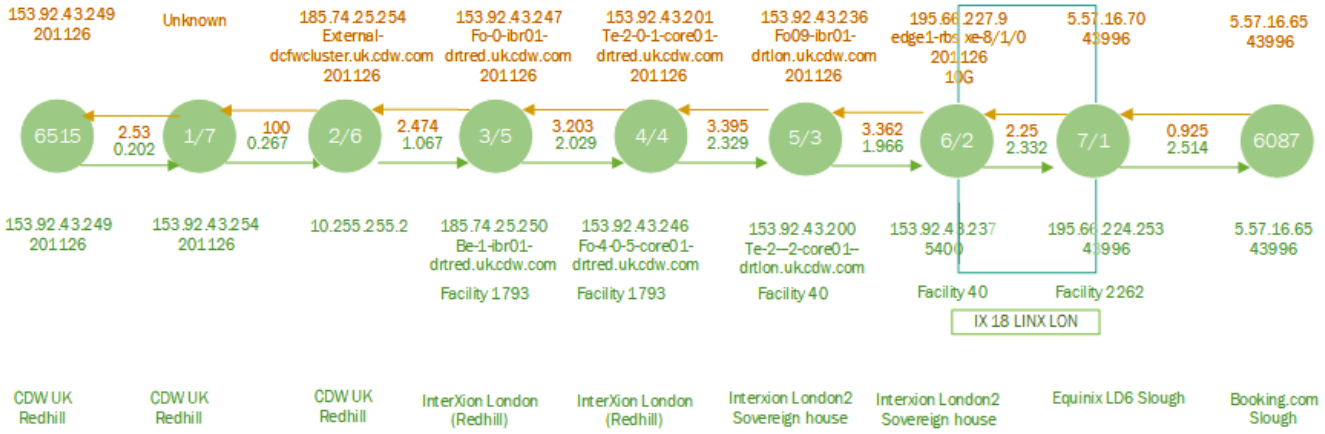

FIGURE 4. Forward and reverse traceroute measurements between two probes demonstrating facility and IX mapping.

to build a snapshot of the connections between UK facilities. The RIPE Atlas has over 600 active probes and anchors located throughout the UK, which can be used as a bootstrap to create a detailed infrastructure map. Traceroutes using CAIDA's ARK platform [45] and Looking Glass (LG) servers, where available, can also add details to the overall picture.

One problem with using a traceroute is that a packet may take any one of the possible routes where load balancing is involved. Paris Traceroute [46] avoids this problem by adapting the probe packet header fields in a manner that allows all probes toward a destination to follow the same path; per-flow load balancing is an option. The RIPE ATLAS platform uses the Paris Traceroute as a default. Although it cannot enumerate paths in all situations, it has been shown to perform considerably better than the classic traceroute.

B. THE METHOD IN ACTION

In putting the method into action, we carried out 1190 traceroutes in both the forward and reverse directions between 35 ATLAS Anchors. The following formula (1) was used to calculate the credit cost of a user-defined measurement:

$$\text{Traceroute credit cost} = 10 \times N((S/1500) + 1) \quad (1)$$

where:

- N = Number of packets per traceroute (default is 3)
- S = packet size (default is 40)
- S is effectively 1 for all intents and purposes, meaning each traceroute costs $(10 \times 3 \times 1)$ credits.

Thus, the total cost is $1190 \times 30 = 35700$ credits.

However, measurements do not necessarily have to be limited to the RIPE ATLAS platform but can be carried out whenever access to both ends of the traceroute is possible. In addition, most RIPE ATLAS measurements are publicly available for read access without requiring credits, and a search of two existing measurements, where the target and source become the source and target, can be used.

Where possible, each hop in a traceroute was mapped to a facility using a combination of DNS lookup, Internet Exchange website information, and PeeringDB information. The combined information creates new Vantage Points (VP) enroute to destinations; these will be invaluable in further research, especially when there is a dearth of ground-truth data, as has been recognised by many researchers.

V. EXAMPLE

An example of this logical network mapping is shown in Figure 4, where a traceroute is first carried out from probe 6515 toward probe 6087.

Each hop's IP address corresponds to an ingress port interface on a router located at a specific location. The hop timings for this measurement are listed in Table 1.

A DNS lookup for each IP address on the route can provide useful information regarding the location of the port/router combination. For example, in Figure 4, we find that the first hop has no DNS name assigned to it, and is most likely the probe owner (CDW UK) default gateway. The RTT value of 0.202 ms suggests that the router is co-located with the probe. The second hop uses a private 10.0.0.0 subnet range, which could be a Local Area Network (LAN), Virtual Private Network (VPN), or Multiprotocol Label Switching (MPLS) connection, and the RTT value suggests that this is also local.

At the third hop, we now have an IP address with a DNS name of 'be-1-ibr01-drtred.uk.cdw.com', which would indicate that it is at the Redhill Facility. Cross-checking with PeeringDB, we find that CDW does, in fact, peer at Facility 1793 Interxion Redhill, so we can be confident that because we know the geolocation of the facility, we now have a new IP address/geolocation combination and, therefore, a new Vantage Point (VP) for use in future IP geolocation work. The hop 4 DNS name shows that we are still in Redhill, but DNS suggests that we have now moved from an edge router to the core router, this IP address/geolocation combination is a new VP. The hop 5 DNS name shows that

TABLE 1. Traceroute measurement table.

Hop	RTT milliseconds	IP Address
1	0.202	153.92.43.254
2	0.267	10.255.255.2
3	1.067	185.74.25.250
4	2.029	153.92.43.246
5	2.329	153.92.43.200
6	1.966	153.92.43.237
7	2.332	195.66.224.253
8	2.514	5.57.16.65

TABLE 2. Reverse traceroute measurement table.

Hop	RTT milliseconds	IP Address
1	0.925	5.57.16.70
2	2.25	195.66.227.9
3	3.362	153.92.43.236
4	3.395	153.92.43.201
5	3.203	153.92.43.247
6	2.474	185.74.25.254
7	unknown	185.74.25.254
8	2.53	153.92.43.249

we are now at a core router in London, and PeeringDB states that the only facility where CDW interconnects in London is Facility 40 Interxion, thus creating another VP that will be verified on the return traceroute.

Hop 6 also displays a London DNS name that must still be in the same facility as hop 5, which provides another IP/geo combination or VP. The diagram shows the route through an Internet Exchange Point. The only indication is that the DNS name appears to show a possible LINX gateway interface; later results on the return traceroute will eventually prove this. Hop 7 has an IP address within the LINX LON1 IX's assigned prefix range of 195.66.224.0/22. We know that the traceroute is now exiting the Internet Exchange, and by cross-checking the LINX Internet Exchange website, we are given the facility name for this IP address, viz. Facility 2262 Equinix LD6 along with other secondary information such as port speed (10G), organization, Ipv6 information, and router/port name. Cross-referencing the facility name in PeeringDB provided vital geolocation data, and another VP was added to the ground-truth dataset. We can also surmise that the previous hop was an Internet Exchange entry point. Finally, the traceroute ended at the target address. The next step was to create a traceroute measurement in the reverse direction from probe 6087 to probe 6515, which is shown in green from right to left in Figure 4. The traceroute timings are listed in Table 2.

Hop 1 provides little information regarding its location, and at this stage, we can only assume that it is a gateway router. Hop 2's IP address is within the LINX LON1 prefix range, so we know that the packet is now exiting the Internet Exchange. Checking the LINX Internet Exchange website, we are given the facility for this IP address, which is facility 40 at Interxion London, along with the other secondary information mentioned earlier, such as the connection's 10Gb service speed. The DNS name closely resembled the DNS name from hop 6 on the forward leg. Therefore, it is safe to assume that they belong to the same router. In addition, we can surmise that hop 1 must have been the entry point for the Internet Exchange, which we already located at Facility 2262 Equinix LD6 in Slough. Therefore, two additional VPs (Vantage Points) can be added to the ground truth dataset. Hop 3 has a London DNS name, stating that it is a port on the core01 router, similar to hop 5 on the outward leg. Therefore, this must be performed at facility 40 in London, adding another VP to the ground truth dataset. Hop 4 has a Redhill DNS name similar to hop 4 on the outward leg. Therefore, we know that the packet has now traveled to the Redhill 1793 Interxion facility, adding another VP to the Table. Hop 5's DNS name shows that the packet is still in Redhill, but has now moved to an edge router, adding a further VP to our VP Table, as shown in Table 3.

Hop 6 has the DNS name of 'external-dcfw-cluster.uk.cdw.com', which does not provide any clues regarding its location. Hop 7's IP address in this direction is unknown; however, the forward traceroute shows that hop 2 ends at a private IP address of 10.255.255.2; therefore, the remote end of this connection must also be in this private subnet range. This coincides with the unknown IP address in the reverse traceroute at hop 7, and it is assumed that this interface does not reply to ICMP packets. Another verification of this assertion is to examine hop 3 on the forward traceroute with hop 6 on the reverse traceroute, both of which are in the 185.74.25.x subnet range. This indicates that we can be confident we are not dealing with asynchronous routes. Because this is the last hop, we can safely conclude that this is the initial gateway router that connects to the probe.

The results of this method allowed us to build a detailed picture of the infrastructure between these two probes by combining information from our three sources (DNS, PeeringDB, and LINX websites), as shown in Figure 5. This diagram shows a traceroute from RIPE Probe 6515 to RIPE Probe 6087, which first passes through three routers (blue circles) on its way to the Interxion and LINX LON1 interconnection facility at the Interxion Sovereign House in London.

Colour coding was used in Figure 5 only as a visual cue to denote the approximate speeds of transmission over these hops. The approximate speeds were calculated by dividing the distance between hops by the difference in time between the previous hop and this hop. It should be noted, however, that each router may prioritise the ICMP packets differently depending on their target, and timings

TABLE 3. Vantage points table.

IP Address	PeeringDB Facility ID	Long	Lat	DNS Name	Port Name	Speed
185.74.25.250	1793	51.2476	-0.1571	be-1-ibr01-drtred.uk.cdw.com.	none	Unknown
153.92.43.246	1793	51.2476	-0.1571	fo-4-0-5- core01-drtred.uk.cdw.com.	none	Unknown
153.92.43.200	40	51.4998	-0.0107	te-2-0-1- core01-drtlon.uk.cdw.com	none	Unknown
153.92.43.237	40	51.4998	-0.0107	fo-0-0-0-20- ibr01-drtlon.uk.cdw.com.	none	Unknown
195.66.224.253	2262	51.5243	-0.6380	None	edge5-eq4 xe-3/0/3	10G
195.66.227.9	40	51.4998	-0.0107	None	edge1-rbs xe-8/1/0	10G
153.92.43.236	40	51.4998	-0.0107	fo-4-0-5- core01-drtlon.uk.cdw.com.	none	Unknown
153.92.43.201	1793	51.2476	-0.1571	te-2-0-1- core01-drtred.uk.cdw.com.	none	Unknown
153.92.43.247	1793	51.2476	-0.1571	fo-0-0-0-20-ibr01-drtred.uk.cdw.com	none	Unknown
185.74.25.254	1793	51.2476	-0.1571	External-dcfwcluster.uk.cdw.co	none	Unknown

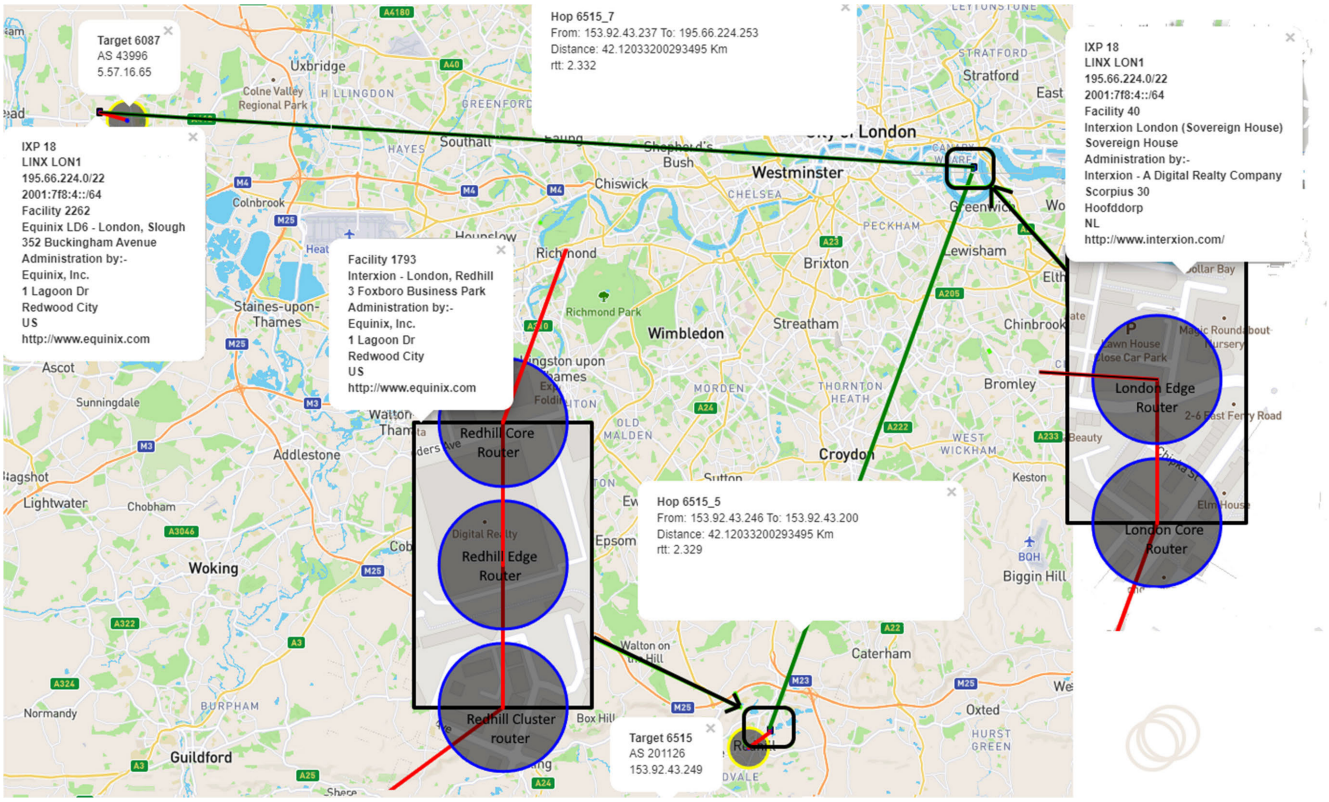

FIGURE 5. Example of UK infrastructure mapping incorporating LINX LON1 internet exchange point.

can also suffer from packet forwarding decisions, circuitous routes, different router configurations, and congestion, and the time taken does not always reflect distances. A different method can divide the overall RTT time of the hops by the distance from the source to the intermediate router. However, this also has its own problems. For example, delays owing

to administrative packet forwarding decisions, circuitous routes, different router configurations, and router congestion will multiply timing errors depending on the number of intermediate routers between the source and hop.

In Figure 5, the green lines indicate relatively fast connections. Red lines denote slow speeds, that is, less than 100 km/ms, and yellow lines denote medium speeds, that is, 100 km/ms to 200 km/ms, whereas green lines are used for anything greater than 200 km/ms. However, it should be emphasised that this is only a rough indication of transmission speeds, regardless of the method used. The ICMP packets then pass through two further routers before entering the Internet Exchange Layer 2 network on their way to Slough. The packets exit the LINX LON1 Internet Exchange at the Slough Equinix Facility and are routed to probe 6087.

The reverse measurement from RIPE Probe 6087 to RIPE Probe 6515, shown in Figure 5, follows the same path in this case, but it is highly likely that many measurements will follow alternative reverse paths. Indeed, both forward and reverse measurements may even take different paths at contrasting times of the day, depending on congestion, providing further details about the network infrastructure and additional Vantage Points. It would be useful to test this hypothesis in future studies. It is interesting to note that in this example and in this reverse direction, ICMP packets seem to be given low priority at hops 3,4,5 and their RTT timings reflect this. It seems unlikely that this is due to congestion or other problems, as the RTT timings at the rest of the interfaces are in line with outward-bound measurements.

Improved validation of the results of the IP address locations within the VP Table could involve contacting the various facilities or AS organisations to confirm the location; however, this would be a very slow method but would guarantee 100% confidence in the result. This method provided good validation in several cases. For example, at hop 3 in the results, the hop IP address has a DNS address, indicating that it is hosted in Redhill. PeeringDB confirms that probe owners (CDW UK) interconnect at the Redhill Interxion Facility and there are no other options. This would suggest that without contacting the facility or CDW UK, we can be confident that we have the correct location. Hop 2's location in the results is a little more obtuse; its DNS name, 'external-dcfw-cluster.uk.cdw.com', refers to an external cluster but does not provide a city name. However, it is connected to hop 1, which appears to be CDW's gateway router at 153.92.43.254 via a private network, perhaps a VPN, or a point-to-point connection. Hop 3's IP address is 185.74.25.250, whereas hop 2's egress IP address is 185.74.25.254, indicating that it is either on the same LAN or possibly on a point-to-point link. The RTT difference in timing between hop 2 (0.267ms) and hop 3 (1.067ms) indicates that it is likely that the router at hop 3 is not local to hop 2. An educated guess would be that this external cluster router is situated at either CDWs Redhill offices or at the Redhill facility, but full validation would have to be confirmed by contacting the facility or CDW UK. In the meantime, the two offices are only 1 mile apart, and either geolocation would provide a useful Vantage Point. So, from these various confidence levels we could add a confidence column to the VP Table as shown in Table 4 where a 1 is fully confident, a 2 is probable, a 3 is likely and a 4 is "best guess." A score of 1 indicated that validation was confirmed by the facility or company. A score of 2 indicates where the DNS name corresponds to a facility location and there are no other possible facilities. A score of 3 would indicate where various other factors such as LAN IP addresses link two hops, as between hops 2 and 3, or perhaps RTT times between the two hops make it impossible for the router to be geolocated elsewhere. A score of 4 was assigned if only minor evidence indicated its location. Further reinforcement of these IP geolocations could result from additional traceroute measurements from RIPE probes located within the AS that owns the hop's IP Address.

VI. AUTOMATING THE SOLUTION

In this section, we describe the automation of this method. Thirty-five probes were used, each targeting the other 34 probes, to create a mesh of 1190 traceroutes across London and South England. The results of each hop from a traceroute are first passed through a filter that tests the hop results against five assumptions (or rules).

A. RULE 1

If this is the first hop, it is assumed that the first router encountered will most likely be the gateway router for the source probe. This may or may not be in the same location as that of the probe.

To provide a sanity check, a test is conducted to discover whether the RTT to this router is less than 1ms; if less than 1ms it is assumed that the gateway router is in the same location as the source probe. This is the only use of the delay-distance model. However, this could be verified through further tests. We begin by testing to ensure that the hop's IP address is not the target, as it has been found that in some cases, the RTT responses are blocked, or the packets are discarded by some or all of the intermediate routers on the way to the target (occasionally, the first IP address encountered is the target address). Once a valid IP address has been determined, a reverse IP lookup is made. If the IP address returns a DNS name, then this is put through a series of search patterns to discover the likely town or city where the IP address is located. If a town or city name was discovered, we checked which facilities were in that town or city. If there are multiple facilities in the town, we compare the AS interconnections of each facility with those of the ASN of the previous hop, which will hopefully reduce the list to a single facility.

If multiple facilities or no facilities were returned, we attempted the reverse traceroute method. In this case, we created a traceroute probe from the target back to the

TABLE 4. Vantage points table with confidence column.

IP Address	Facility ID	Long	Lat	DNS Name	Port Name	Speed	Confidence
185.74.25.250	1793	51.2476	-0.1571	be-1-ibr01-drtred.uk.cdw.com	none	Unknown	2
153.92.43.246	1793	51.2476	-0.1571	fo-4-0-5-core01-drtred.uk.cdw.com	none	Unknown	2
153.92.43.200	40	51.4998	-0.0107	te-2-0-1-core01-drtlon.uk.cdw.com	none	Unknown	2
153.92.43.237	40	51.4998	-0.0107	fo-0-0-0-20-ibr01-drtlon.uk.cdw.com	none	Unknown	2
195.66.224.253	2262	51.5243	-0.6380	None	edge5-eq4xe-3/0/3	10G	2
195.66.227.9	40	51.4998	-0.0107	None	edge1-rbsxe-8/1/0	10G	2
153.92.43.236	40	51.4998	-0.0107	fo-4-0-5-core01-drtlon.uk.cdw.com	none	Unknown	2
153.92.43.201	1793	51.2476	-0.1571	te-2-0-1-core01-drtred.uk.cdw.com	none	Unknown	2
153.92.43.247	1793	51.2476	-0.1571	fo-0-0-0-20-ibr01-drtred.uk.cdw.com	none	Unknown	2
185.74.25.254		51.2395	-0.1525	External-dcfwcluster.uk.cdw.com	none	Unknown	3

source and compared the first hop of the forward tracker with the last hop of the reverse traceroute. If both IP addresses are in the same subnet prefix, we can assume that the penultimate interface on the reverse traceroute is an interface on the forward traceroute's first-hop router. A reverse DNS lookup of this interface's IP address is made, and any resulting DNS address is again put through a series of regular expression (REGEX) search patterns in an attempt to discover its location by comparing various parts of the DNS address with the UK (United Kingdom) town or city names where facilities are known to be located.

Figure 6 shows an example traceroute from RIPE Probe 6087 to Ripe Probe 6843, and the reverse path where the first router encountered at IP address 5.57.16.70 only has an RTT of.369ms, which is a good indication that this router is in the same location as the source probe. However, we can attempt to verify this by examining the other side of this router by carrying out a reverse traceroute from RIPE Probe 6843 back to probe 6087. In this case, we can examine the penultimate incoming interface and compare the IP address prefixes, where it is found that the incoming IP address 195.66.224.253 is in the same prefix range as the outgoing IP address of 195.66.224.108, indicating that we are dealing with the same first-hop router on both the outward and return journeys. Therefore, DNS clues to the location of the incoming interface also provide us with the location of the outgoing interface.

A lookup of the LINX IX membership database shows that IP address 195.66.224.253 belongs to booking.com and is located at the Equinix LD4 facility. A database lookup at PeeringDB provides the geo-coordinates of the Equinix LD4 facility, which shows that it is on a specific street in Slough. This return traceroute in Figure 6 has provided us with the location of the outgoing traceroute's first hop interface because the outgoing interface of this router with an IP address of 5.57.16.70 is an interface on the same router as that of the geolocated interface with an IP address of 195.66.224.253; both of these IP addresses along with the geo-coordinates of Equinix Facility LD4 can be used in our Vantage Points table.

We need to compare this traceroute with that shown in Figure 7, where the penultimate incoming hop has not returned an IP address because the packet is discarded or blocked, whereas the forward hop after the initial router has an IP address of 195.66.224.234. In this case, we cannot ascertain that this is the same router, and as can be seen in the figure, it would appear that the forward and reverse paths are asynchronous, using different Internet Exchanges on the forward (IX 18 LINX London) and reverse routes (IX 321 LINX London2). An additional point to make here is that the packet route seems to follow a somewhat circuitous route from the Slough Equinix LD6 facility to the London Telehouse West facility and then back to the Slough Equinix LD4 facility (Equinix advertises local cross-connects between LD6 and LD4). This example shows that these methods may offer Internet Exchanges with some opportunities to improve the network speed and reduce congestion.

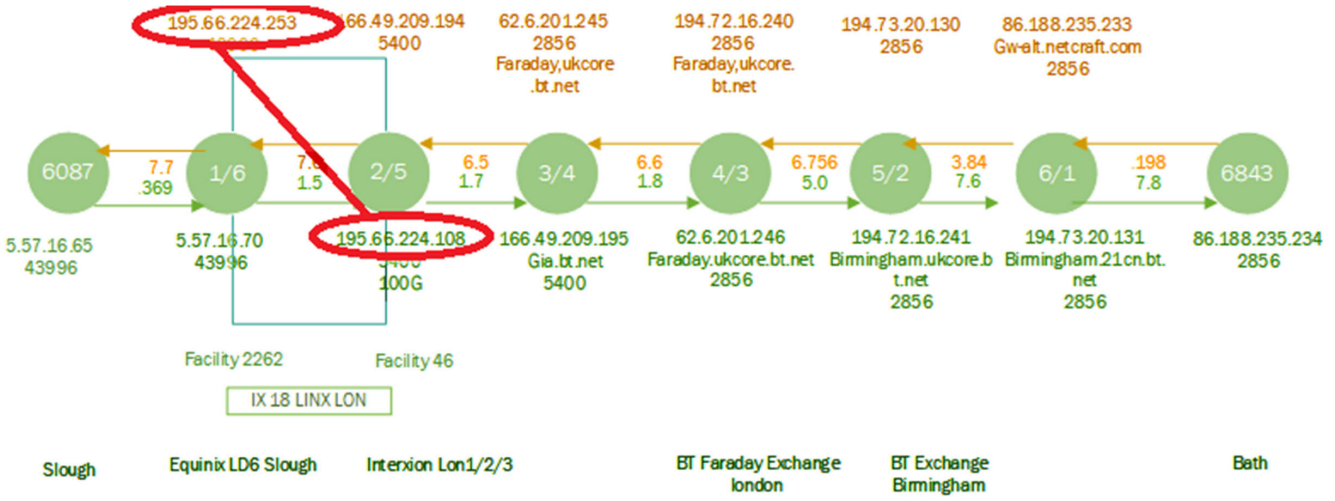

FIGURE 6. Forward and reverse traceroutes from IP address 5.57.16.65 to 86.188.235.234, showing outgoing and incoming IP addresses in same prefix range.

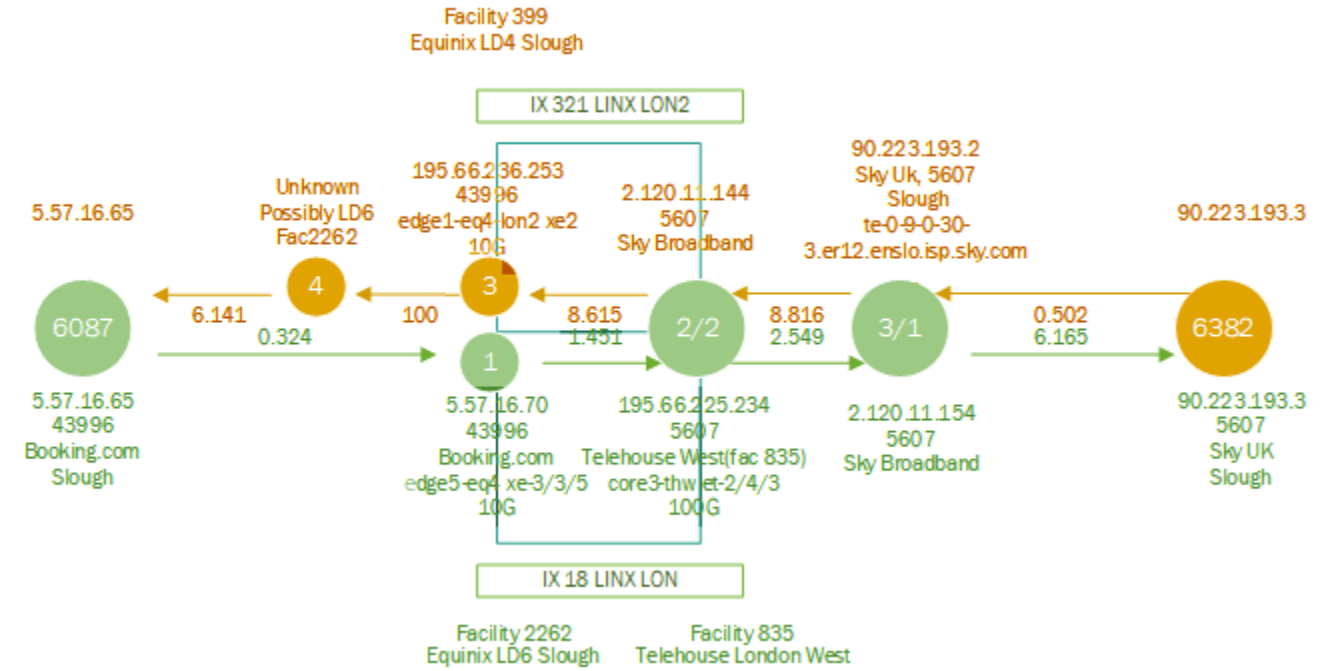

FIGURE 7. Forward and reverse traceroutes from IP address 5.57.16.65 to 90.223.193.3, showing incoming penultimate IP address blocked and outgoing IP addresses at 195.66.225.234. Also showing asynchronous forward and reverse routes using different Internet Exchanges on the forward and return legs.

In the first scenario, it was fortunate that the penultimate return hop was across an Internet Exchange, where a list of IP addresses and their facility locations was readily available. However, the penultimate return hop may be another connection, as shown in Figure 8. In this case, the reverse lookup is 'Birmingham.21cn.bt.net', which our REGEX search script would normally locate to Birmingham. However, this first-hop router cannot possibly be located in Birmingham because the initial probe is located in Bath; with a.198ms RTT to this router, the sanity check locates the router in Bath, which contradicts the reverse DNS lookup. In this case, the result of the RTT sanity check is prioritised over the results of the reverse traceroute method.

If the list cannot be reduced to a single facility, the central coordinates of the town or city are returned for use as a general location. A list of UK towns and cities and their central coordinates was downloaded from the Office for National Statistics (ONS), which provides free and unrestricted access to a definitive source of geographical products.

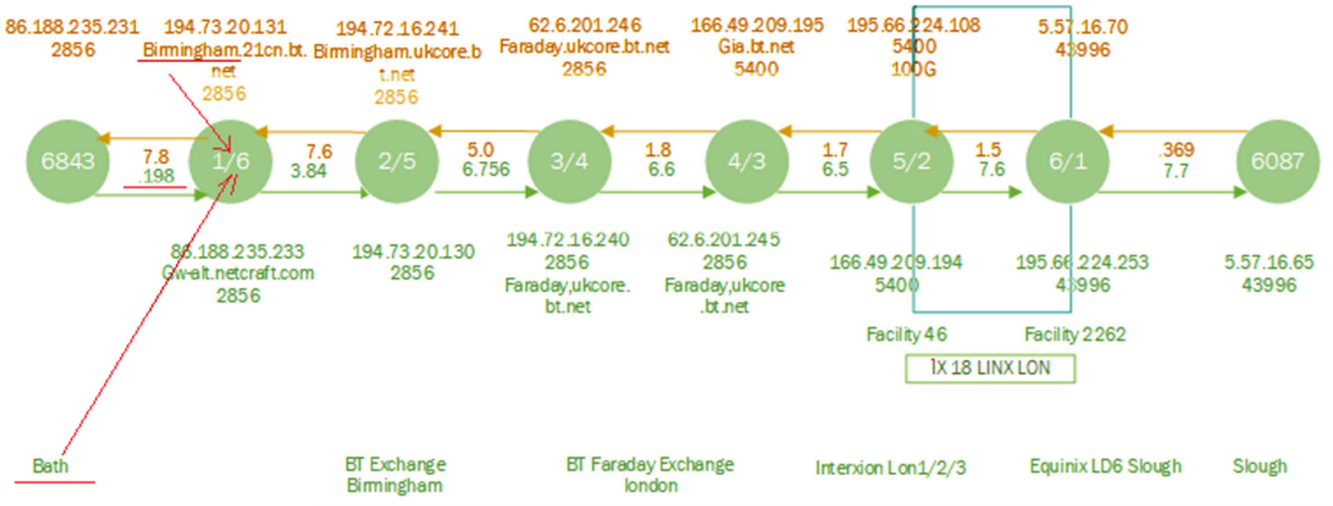

FIGURE 8. Forward and reverse traceroutes from IP address 86.188.235.231 to 5.57.16.65, showing false positive in finding IP address location.

B. RULE 2

All IPv4 addresses can be divided into two major groups: **global** (or public, external), which are those used on the Internet, and **private** (or local, internal), which are those used in local area networks (LANs) and may also be used in VPNs (Virtual Private Network) (Virtual Private Network), cross-connections, or site-to-site links.

If the IP address of this hop is within a private subnet range and the previous hops are not, then the packet has crossed a LAN, VPN, or site-to-site link. Therefore, it is possible that it is still in the same location as the previous hop or has traveled to a new location. However, because it is a private IP address, there is little benefit in locating its true coordinates, as private IP addresses cannot be added to a Vantage Point table because of their possible use in multiple locations. If the location of this router is considered important, a comparison of the difference between the RTT values of the previous and successive RTTs provide clues as to whether the location of this router is local or remote to the previous router. Some private IP addresses have been geolocated in this fashion.

C. RULE 3

If the hop under consideration corresponds to the target IP address, its coordinates are those of the target probe, which can be discovered within the RIPE ATLAS database.

D. RULE 4

This rule determines whether a hop's IP address is registered on an Internet Exchange. A test was conducted on each hop's IP address to determine whether it falls within any IX-registered prefix. The initial design would then look at the previous hop to discover the ASN entering this IX and compare it with a list of IXs and their ASN peers from PeeringDB in an attempt to discover the sole facility where the IX and previous ASN are peers. Therefore, it would provide an entry facility. The same method was then applied to the existing facility using the successive hop ASN.

However, finding the entry facility in this manner becomes unnecessary, as Rule 5 (which applies to all other hops that do not fall into any previous rule) attempts to discover the correct facility. Additionally, the IX entry facility is not as important as the exit facility, which can be readily used as an ideal Vantage Point to discover the possible location of a given IP address. Furthermore, in the initial design, the exit facility was found using a similar process, where it looked at the successive hop to discover the ASN exiting this IX and compared it with a list of IXs and their ASN peers in an attempt to discover the sole facility where the IX peers with the successive ASN. However, this becomes unnecessary after discovering that the IXPDB database (an authoritative, comprehensive, and public source of data related to IXPs) provides a list of IX-registered IP addresses along with the facility where they are located. The IXPDB database also integrates the data from third-party sources. The website provides a comprehensive and corroborated view of the global interconnection landscape. The combined data can be viewed, analysed, and exported through a web-based interface or API.

It is this IXPDB database, which is now initially interrogated in an attempt to discover whether a hop's IP address is registered at an Internet Exchange. It was discovered that the IXPDB database is not as comprehensive as advertised; some IXs mark their IP to facility information as private, so no information is uploaded to IXPDB. One of these IXs is Equinix, which plans to release this information in the near future. In cases where the IP address cannot be found in the IXPDB database, but we have discovered that the IP address of the hop falls within an IX-registered IP prefix range, we revert to the previously discussed method of comparing the successive ASN with each Internet Exchange's known peers.

E. RULE 5

If no other rule is applicable, then a DNS lookup is made on this hop's IP address, and any resulting DNS address is passed through a series of REGEX search strings to extract the town or city name from the reverse DNS address. Each part of a DNS address is compared with a list of known towns and cities in which facilities are located. If part of a DNS address matches the beginning of a town or city name, then the facilities for that town are extracted. If this list of facilities contains only one facility, the hop's IP address is located successfully. However, if no facilities are located, we can attempt a reverse traceroute to discover the outgoing interface of the router. This is similar to the procedure described in Rule 1, where both the outgoing and incoming IP addresses are in the same prefix range, and we can then make a safe assumption that the incoming and outgoing interfaces are on the same router. If this traceroute is successful in discovering the outgoing IP address, we can attempt to carry out a reverse DNS lookup on that IP address and pass any results through our REGEX search to discover the town and facilities, as described previously. If all of these methods fail, we must classify this as a failure to find the IP address; therefore, we are unable to add this IP address to our Vantage Point table.

Referring back to our previous example in Section V and Figure 4, at hop 2 on the forward route we have a private IP address that would normally be impossible to geolocate. However, it will complete the routing diagram for this particular traceroute if we know the exact geolocation of this particular hop. If we examine the reverse traceroute, we can see that the reverse hop provides us with a DNS address of 'External-dcfwcluster.uk.cdw.com'.

Combining this with our rules regarding sanity checks and incoming/outgoing IP addresses on the same subnet range, we can safely assume that the router with an incoming hop of 10.255.255.2 on the forward traceroute and an incoming hop of 185.74.25.254 on the reverse incoming hop is indeed the same router, and therefore we can geo-locate this 10.255.255.2 and the reverse 185.74.25.254 IP address to Redhill.

VII. RESULTS

To demonstrate our proposed method, 1190 traceroutes were created using the RIPE ATLAS platform with an API tool specifically written to create the necessary measurements on the platform.

Once the measurements are completed using RIPE ATLAS, the second API tool reads all the measurements from it and creates a local JSON file. The third tool reads the JSON file and discovers the likely geocoordinates of each hop, creating another JSON file and a Vantage Point table as the outputs. The fourth optional tool maps these results to an OpenStreetmap, as shown in Figure 5.

The results of applying the described rules and methods through an automated process are presented in Table 5. This table also shows the effect of IXPDB data on the final results. The first two columns show that the complex "Common_Fac" method was used to discover IP address locations; however, once the LINX and LONAP Internet exchange data were added from IXPDB, this method became almost redundant. It is expected that once all IX datasets become available, the Common_Fac method will not need to be used.

Although the rules have already been described above and the methods have been discussed at several points in this paper, we now provide a summary list for convenience.

Regex is a process in which a hop's DNS address is filtered through a series of regular expressions to find a town or city name. It was found that in-depth knowledge of the network region is required to provide the correct tests for the Regex method.

Reverse traceroute is the process of discovering the IP address of the outgoing interface to discover the location of a router via a second interface located on the same router, which may provide better clues regarding the router's location.

Reverse DNS is where an IP address is looked up in an attempt to discover its DNS address. This method is typically used in conjunction with the **Regex** method and can also be combined with the Reverse Traceroute method.

Common Facility is one of the earliest processes used in our work and was designed to geo-locate a router where a packet enters and exits an Internet Exchange. This was done by comparing the ASN entering a facility with the ASN peer at each Internet exchange. This has been largely superseded by the **Facility to IP Table** described below.

Facility to IP Table consists of an API lookup of the IXPDB website, which holds a comprehensive list of Internet exchange-registered IP addresses and their locations.

It should be noted that many of the IP addresses were tested multiple times using different traceroutes; hence, a much larger number of successes and failures occurred compared with the discrete number of IP addresses. The first two columns of Table 5 show the original success and failure rates for each rule and method used. The second two columns show the success and failure rates once the IXPDB data that apply to the LINX Internet Exchange are added. The software stops relying on the complex method of determining the facility by comparing the ingoing/outgoing ASNs with the IX facilities, as shown in red, and begins by using the IXPDB database, as shown in green. Finally, the third set of columns shows the results when the LONAP Internet Exchange data from the IXPDB database is added. The Reverse Traceroute method shows zero successes, but this is not a test in itself, it must be combined with reverse DNS in order to provide a result.

Of 1190 traceroutes employed in the test, 1047 individually discrete IP addresses were found, of which 372 were geolocated to a confidence level of 3 and above. However, three of those IP addresses are anycast, which means that they are shared by devices in multiple locations. This discounts them from being able to be geolocated to a single location. Therefore, we were able to geolocate 369 IP addresses.

Of these 369 successfully geolocated IP addresses, 102 were geolocated without a DNS lookup. When analyzing the individual contribution of each geolocation rule described in the previous section we find the following:

- 12 IPs were geolocated by Rule 1 using a sanity check on the RTT value. Due to these 12 IP addresses being at hop 1 in each traceroute we can confidently use the RTT value to predict the delay-distance values;
- 33 IPs were geolocated by Rule 3, which are the target IP addresses;
- 102 IPs were geolocated by Rule 4, using the facility's location;
- 211 of the successful IP addresses were located by Rule 5, where the geolocation is discovered using a combination of REGEX and previous rules.

Table 6 summarizes these statistics according to the geolocation rules.

If we rule out the 26 private IP addresses from our formulae due to private addresses only providing a location pertaining to that specific traceroute, we end up with a total of 343 vantage points out of a possible 1021 distinct IP addresses, which gives a 33.6% success rate. Luckie et al. downloaded 1.39 million IP addresses from the CAIDA ITDK 2020/2021 datasets. Of these he discovered 220,000 had geohints and from these 220,000 they geolocated 183,000 which works out at 7.1% of the original 1.39 million. Whilst our initial dataset only contains 1047 discrete IP addresses, the geolocation of 369 of them represents 33.6%, which is a significantly higher percentage.

VIII. DISCUSSION

The five rules presented in this solution have evolved over time, and as new processes have been discovered, some have become more relevant while others are less relevant. Some rules use a similar program flow, such as:

1. Attempt to discover the DNS address of this IP.
2. The DNS address is fed through a REGEX solution to discover a possible city or town name.
3. Find all the facilities in that city or town and narrow them down to one facility using sanitcy checks, ASN lookups, and RTT values.
4. Otherwise, perform a reverse traceroute, if available, and subject any reverse DNS to the same REGEX filter.
5. Failing all of this carries out a Common Facility comparison to determine the facilities at which the previous ASN and current ASN. interconnects.
6. If none of the methods proved successful, we failed to locate the IP address.

In Step 2, DNS parsing uses a regular expression script similar to that developed by Luckie et al. [32] and Dan et al. [33]. In many cases, generic regular expressions automate the discovery of a facility and its coordinates. However, it should be noted that the success of a regular expression script is highly dependent on the local infrastructure knowledge. The regular expression script employed was developed purely for the UK, where detailed information can also be hard coded. For example, British Telecom uses its own telephone exchanges as facilities, and these are not listed in PeeringDB. However, the locations of BT's DNS addresses are easily identified when the script is provided with the necessary expression. Some of the BT DNS names provide the telephone exchanges town such as 'acc1-te0-0-0-0. **kingston**.ukcore.bt.net' which is in Kingston-upon-Thames. Others are slightly more obtuse such as 'core2-hu0-7-0-3. **southbank**.ukcore.bt.net' which is situated at Columbo House, London. Others are listed after the name of the property such as 'core3-hu0-6-0-0. **faraday**.ukcore.bt.net' which corresponds to Faraday House in London. Therefore, many of these locations must be added to the list of regular expressions, and many other companies have equally obtuse DNS addresses. NTT, for example, appears to have misspelled London in all their DNS names, for example, 'ae-2.r21. **londen**12.uk.bb.gin.ntt.net'. Faelix has identified facilities with names, such as an IP address with a DNS name of 'eth5. **aebi.**m.faelix.net', which corresponds to PeeringDB's facility number 46, which is the Interxion facility in London.

The results of the use of traceroutes in measurements should be interpreted with caution. Although this method avoids many of the fundamental problems described in Section II-D, there are still limitations that need to be addressed. ICMP echo packets are often given second-class treatment by routers and target hosts. This means that ICMP echo requests and responses may be given a lower priority than traffic, which is considered more important. The end result indicates that the return trip time reflected by the traceroute can easily be different from that experienced by other higher-priority traffic types. In addition, because routers may consider ICMP traffic to have a small packet size, they can experience different routing paths compared to fully laden TCP or UDP packets. However, the goal of this method is to create maps of the Internet infrastructure and not to be overly concerned about packet timings. With the exception of using RTT values as a secondary check, RTT values are not a major part of this method.

It should also be noted that this method works well because of the abundance of RIPE probes located in the UK, and it is likely that the use of this method will not be as effective in regions where RIPE probes are sparse, such as Africa, Russia, or China. In these cases, other traceroute platforms, such as CAIDA's ARK platform, could be employed, where the IP address and geolocation are already known to be used as sources and targets for initiating traceroutes.

While building up this detailed visualisation of the UK Internet infrastructure, the method additionally creates a dataset of IP addresses to geolocations: the Vantage Points or VPs. With over 600 probes in the UK, it is theoretically possible to create a mesh of over 300,000 traceroutes, each discovering on average 1–10 IP address/geocoordinate combinations, providing a dataset of over one million VPs from which future research on IP geolocation can be based. In addition, it should be noted that the IXPDB dataset is seem-

TABLE 5. Rules and methods success /failure table.

Rules & Methods	Before adding LINX data		after adding LINX data		after adding LONAP data	
	failures	successes	failures	success	failures	success
regex	3244	2760	3244	2760	3218	2760
reverse_tr	3085	0	3065	0	3064	0
reverse_dns	35	17	35	17	10	0
common_fac	52	855	52	40	10	19
fac_to_ip_table	0	0	92	815	29	878
Rule1	634	2003	634	1986	634	1987
Rule2	126	236	126	236	126	236
Rule3	0	1243	0	1243	0	1243
Rule4	35	855	35	855	10	897
Rule5	1254	32	1254	33	1254	33

TABLE 6. Rules total successful geolocations.

RULE 1 – Hop 1 Sanity Checks	12
RULE 2 – Private IP addresses	26
RULE 3 – Target IPs	33
RULE 4 – Facility Location	90
RULE 5 – REGEX and other methods	211 (3 excluded due to being anycast)
TOTAL	372 (3 excluded due to being anycast)

ingly an untouched source of Vantage Points/Landmarks, which can be used in future research, additionally, the IP address-to-geolocation pairs that can be derived from this data are naturally located close to population centres.

IX. CONCLUSION

The aim of this study was to investigate the current methods used by Internet mapping techniques to determine the optimum method for developing finer-grained infrastructure maps. We built on these methods by developing tools and techniques that can help create more fine-grained infrastructure maps. The new method developed in this study uses four Python tools that gather all the UK's facility geolocations and map them onto OpenStreetMap. The tools also locate all the UK's IXs and map the network structure and interconnection facilities. The tools then create measurements from the selected RIPE ATLAS probes to create a UK infrastructure map by geolocating every hop within each traceroute, where possible. As a by-product of this method, a useful IP address to geolocation dataset was created.

OpenStreetMap is not capable of effectively displaying all fine-grained information regarding Internet infrastructure; therefore, research into improved methods for visualising this information would prove useful. For example, Virtual Reality may provide better methods for visualising interconnections and geographical data.

A REGEX filter is one of the main components of this solution. First, the DNS name of each IP address is discovered and then fed through REGEX in an attempt to discover the town or city where it is located. While several researchers have already worked on this issue, such as Luckie et al. [32] and Dan et al. [33], a new solution has been created and reported here, specifically designed for the UK Internet infrastructure. The limited number of UK towns (75) where 179 facilities are located makes the process of geolocating Internet infrastructure slightly easier, allowing for some amount of brute force techniques to be used, for example, searching for specific facility names. However, there is much room for improvement here; an investigation into the 3000+ failures of this REGEX technique would lead to more comprehensive results. The machine learning techniques from Dan et al. and learning geographic naming conventions from Luckie et al. could also significantly improve these results.

Future work could also involve creating the same traceroute measurements over extended time periods and using different routes, adding alternative hops as backup paths,

or finding completely new paths, allowing new infrastructure details to be realised and further vantage points to be created.

The various methods and rules presented in this study evolved over time. At the beginning of this research, discovering the entry and exit points of a packet crossing an Internet Exchange was a complicated process for finding the common peers of an IX against a preceding or succeeding ASN. However, the IXPDB website has made this task much easier by providing the facility name and cross-reference to the PeeringDB number for each IX-registered IP address. The IXPDB website is a previously untouched mine of useful Vantage Points that are close to population centres.

The solution uses a confidence level mechanism to provide some idea of the accuracy of the methods explained in this paper. Of 1190 traceroutes employed in the test, 1021 individually discrete IP addresses were found, of which 343 were geolocated using the procedure and methods described above. This gives a success rate of 33.6% in geolocating the 1021 IP addresses to a confidence level of 3 or above.

Relatively little research has been carried out on geolocating IP addresses to the facility level. Luckie et al [32] were able to geolocate 7.1% of IP addresses to a city level when comparing their dataset against the CAIDA ITDK [45] dataset. When attempting to compare our dataset against the CAIDA ITDK datasets it was only possible to compare city locations as the CAIDA dataset resolves to city-level resolution rather than at facility level. While the CAIDA dataset has 1.7 million UK nodes, our dataset has 492 overlapping IP addresses for comparison. Out of the 1047 distinct IPs in our dataset, 288 were geolocated to a city level by CAIDA and 369 to a facility level by our method; this includes 89 new IPs which were not in the CAIDA dataset. Given that these promising results are preliminary, we next need to demonstrate scaling the solution to millions of traceroutes and different regions.

In summary, the work reported here has successfully achieved its aim of discovering and developing a new method to create finer-grained maps of the UK's Internet infrastructure.

REFERENCES

[1] R. Motamedi, R. Rejaie, and W. Willinger, "A survey of techniques for internet topology discovery," *IEEE Commun. Surveys Tuts.*, vol. 17, no. 2, pp. 1044–1065, 2nd Quart., 2015.

[2] R. Motamedi, B. Yeganeh, B. Chandrasekaran, R. Rejaie, B. M. Maggs, and W. Willinger, "On mapping the interconnections in today's internet," *IEEE/ACM Trans. Netw.*, vol. 27, no. 5, pp. 2056–2070, Oct. 2019, doi: 10.1109/TNET.2019.2940369.

[3] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (In)Completeness of the observed internet AS-level structure," *IEEE/ACM Trans. Netw.*, vol. 18, no. 1, pp. 109–122, Feb. 2010.

[4] W. Willinger, M. Roughan, "Internet topology research redux," in *Recent Advances in Networking*, H. Haddadi and O. Bonaventure, Eds., 2013, pp. 1–59.

[5] NCSC. *Technical Report: Responsible Use of the Border Gateway Protocol (BGP) for ISP Internetworking*. Accessed: Jan. 10, 2020. [Online]. Available: https://www.ncsc.gov.uk/files/border-gateway-protocol-technical-paper.pdf

[6] (Jan. 15, 2015). *ENISA, Threat Landscape of Internet Infrastructure*. [Online]. Available: https:// www.enisa.europa.eu/publications/iitl

[7] R. N. Staff, "RIPE atlas: A global internet measurement network, RIPE," *Internet Protocol J.*, vol. 18, no. 3, pp. 2–26, 2015.

[8] C. Davis, P. Vixie, T. Goodwin, and I. Dickinson. (1996). *A Means for Expressing Location Information in the Domain Name System*. Request for Comments: 1876. [Online]. Available https://www.rfc-editor.org/rfc/rfc1876

[9] G. Cumming. (2014). *The Weird and Wonderful World of DNS LOC Records*. [Online]. Available: https://blog.cloudflare.com/the-weird-and-wonderful-world-of-dns-loc-records/

[10] L. Daigle. (2004). *WHOIS Protocol Specification. Request for Comments: 3912*. [Online]. Available https://www.rfc-editor.org/rfc/rfc3912

[11] E. Kline, K. Duleba, and Z. Szamonek. *Self-Published IP Geolocation Data*. Accessed: Jun. 5, 2020. [Online]. Available: https:// tools.ietf.org/html/draft-google-self-1281 published-geofeeds-02

[12] MaxMind, Inc. *Detect Online Fraud and Locate Online Visitors*. Accessed: Mar. 18, 2021. [Online]. Available: https://www.maxmind.com/en/home

[13] IP2Location.com. *Geolocate IP Address Location Using IP2Location*. Accessed: Sep. 9, 2021. [Online]. https: //www.ip2location.com/

[14] Neustar, Inc. *IP Intelligence*. Accessed: Sep. 9, 2021. [Online]. https://www.security.neustar/digitalperformance/ip-intelligence

[15] M. Gharaibeh, A. Shah, B. Huffaker, H. Zhang, R. Ensafi, and C. Papadopoulos, "A look at router geolocation in public and commercial databases," in *Proc. Internet Meas. Conf.*, Nov. 2017, pp. 463–469.

[16] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye, "IP geolocation databases: Unreliable?" *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 2, pp. 53–56, Apr. 2011.

[17] Y. Shavitt and N. Zilberman, "A geolocation databases study," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 10, pp. 2044–2056, Dec. 2011.

[18] V. N. Padmanabhan, L. Subramanian. *An Investigation of Geographic Mapping Techniques for Internet Hosts*. Accessed: Jul. 7, 2020. [Online] Available: https://homes.cs.washington.edu/ arvind/cs425/doc/geomap.pdf

[19] B. Gueye, A. Ziviani, M. Crovella, and S. Fdida, "Constraint based geolocation," *IEEE/ACM Trans. Netw.*, vol. 14, no. 6, p. 1219, Dec. 2006.

[20] S. Zu, X. Luo, and F. Zhang, "IP-geolocater: A more reliable IP geolocation algorithm based on router error training," *Frontiers Comput. Sci.*, vol. 16, no. 1, Feb. 2022, Art. no. 161504.

[21] G. Ciavarrini, M. S. Greco, and A. Vecchio, "Geolocation of internet hosts: Accuracy limits through Cramér–Rao lower bound," *Comput. Netw.*, vol. 135, pp. 70–80, Apr. 2018.

[22] B. Eriksson, P. Barford, J. Sommers, and R. Nowak, "A learning-based approach for IP geolocation," in *Proc. Int. Conf. Passive Act. Netw. Meas.* Cham, Switzerland: Springer, 2010, pp. 171–180.

[23] R. Jayant and E. Katz-Bassett, "Toward better geolocation: Improving internet distance estimates using route traces," Pennsylvania State Univ., Harrisburg, PA, USA, Tech. Rep., 2004.

[24] I. Youn, B. L. Mark, and D. Richards, "Statistical geolocation of internet hosts," in *Proc. 18th Int. Conf. Comput. Commun. Netw.*, Aug. 2009, pp. 1–6, doi: 10.1109/ICCCN.2009.5235373.

[25] D. Li, J. Chen, C. Guo, Y. Liu, J. Zhang, Z. Zhang, and Y. Zhang, "IP-geolocation mapping for moderately connected internet regions," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 2, pp. 381–391, Feb. 2013.

[26] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe, "Towards IP geolocation using delay and topology measurements," in *Proc. 6th ACM SIGCOMM Conf. Internet Meas.*, Oct. 2006, pp. 71–84.

[27] S. Laki, P. Mátray, P. Hága, T. Sebok, I. Csabai, and G. Vattay, "Spotter: A model based active geolocation service," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 3173–3181.

[28] Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang, "Towards street level client-independent ip geolocation," in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2011, pp. 365–379.

[29] B. Wong, I. Stoyanov, and E. G. Sirer, "Octant: A comprehensive framework for the geolocalization of internet hosts," in *Proc. NSDI 4th USENIX Symp. Netw. Syst. Design Implement.*, 2007, pp. 1–14.

[30] Q. Scheitle, O. Gasser, P. Sattler, and G. Carle, "HLOC: Hints-based geolocation leveraging multiple measurement frameworks," 2017, *arXiv:1706.09331*.

[31] I. Livadariu, T. Dreibholz, A.S. Al-Selwi, H. Bryhni, O. Lysne, S. Bjørnstad, and A. Elmokashf. *On the Accuracy of Country-Level IP Geolocation*. Accessed: Oct. 14, 2021. [Online] Available: https:// www.simulamet.no/sites/default/files/publications/files/anrw2020.pdf

[32] M. Luckie, B. Huffaker, A Marder, Z. Bischof, M. Fletcher, and K. Claffy, "Learning to extract geographic information from internet router hostnames," in *Proc. CoNEXT*. Berlin, Germany: Association for Computing Machinery, Dec. 2021, pp. 440–453, doi: 10.1145/3485983.3494869.
[33] O. Dan, V. Parikh, and B. D. Davison, "IP Geolocation through Reverse DNS," *ACM Trans. Internet Technol.*, vol. 22, no. 1, pp. 1–29, Oct. 2021, doi: 10.1145/3457611.
[34] O. Dan, V. Parikh, and B. D. Davison, "IP geolocation using traceroute location propagation and IP range location interpolation," in *Proc. Companion Proc. Web Conf.*, Apr. 2021, pp. 332–338, doi: 10.1145/3442442.3451888.
[35] V. Giotsas, G. Smaragdakis, B. Huffaker, M. Luckie, and K. Claffy, "Mapping peering interconnections to a facility," in *Proc. 11th ACM Conf. Emerg. Netw. Exp. Technol.*, Dec. 2015, pp. 1–13, doi: 10.1145/2716281.2836122.
[36] *PeeringDB*. Accessed: Apr. 8, 2021. [Online] Available:https://peeringdb.comxd
[37] T. Holterbach, C. Pelsser, R. Bush, and L. Vanbever, "Quantifying interference between measurements on the RIPE atlas platform," in *Proc. Internet Meas. Conf.*, Oct. 2015, pp. 437–443.
[38] M. Candela, E. Gregori, V. Luconi, and A. Vecchio, "Using RIPE atlas for geolocating IP infrastructure," *IEEE Access*, vol. 7, pp. 48816–48829, 2019.
[39] (2013). *Router Geolocation, E. Aben, RIPE67*. [Online]. Available: https://ripe67.ripe.net/presentations/341-2013-10-ripe67-router-geoloc-emile-aben.pdf
[40] *Packet Clearing House*. Accessed: Apr. 8, 2021. [Online] Available: https://www.pch.net/
[41] *The IXP Database*. Accessed: Mar. 25, 2022. [Online] Available: https://ixpdb.euro-ix.net/en/
[42] R. Klöti, B. Ager, V. Kotronis, G. Nomikos, and X. Dimitropoulos, "A comparative look into public IXP datasets," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 46, no. 1, pp. 21–29, Jan. 2016.
[43] OpenStreetMap Contributors. *OpenStreetMap*. Accessed: Jun. 7, 2020. [Online] Available: https://www.openstreetmap.org
[44] *Nominatim—OpenStreetMap Wiki*. Accessed: Jun. 7, 2020. [Online] Available: https://nominatim.org/
[45] (Jun. 22, 2022). *The CAIDA UCSD Macroscopic Internet Topology Data Kit (ITDK)*. Accessed: Apr. 15, 2023. [Online]. Available: https://www.caida.org/catalog/datasets/internet-topology-data-kit/
[46] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with Paris traceroute," in *Proc. 6th ACM SIGCOMM Conf. Internet Meas.*, Oct. 2006, pp. 153–158, doi: 10.1145/1177080.1177100.

PAUL MCCHERRY received the B.Sc. degree (Hons.) in network engineering and the M.Sc. degree in cyber security from the University of Lancaster, in 2017 and 2019, respectively. He is currently pursuing the Ph.D. degree in internet security. He joined the Royal Air Force, in 1979, to become an Electrical Technician Servicing a Lightning Fighter Aircraft, attaining the rank of a Junior Technician before leaving the RAF to build his career in computer and network engineering, in 1984. In 1994, he started Syscom Ltd., a local IT sales and support company subcontracting to organizations, such as Camelot, National Lottery, Blackpool Council, and many local Northwest England businesses. In 2019, he became a Research Assistant with Lancaster University. His research interests include IP geolocation and the security of the border gateway protocol (BGP).

VASILEIOS GIOTSAS received the Ph.D. degree from the University College London (UCL). He is currently a Lecturer with Lancaster University, where he leads the Networks Area Research of the Security Institute. His research interests include network measurements and the analysis of the internet routing systems.

DAVID HUTCHISON is currently an emeritus Professor in computing with Lancaster University, U.K., and the Founding Director of InfoLab21. His work is well known internationally for contributions in a range of areas, including quality of service, active and programmable networking, multimedia and content distribution networks, and testbed activities. His current research interests include the resilience of networked computer systems and protection of critical infrastructure and services.