# Advanced Geolocation Techniques and Geopolitical Integration for a Resilient Internet Infrastructure

**Paul McCherry, BSc (Hons), MSc**

School of Computing and Communications

Lancaster University

A thesis submitted for the degree of

*Doctor of Philosophy*

June, 2025

# Abstract

**Governments and institutions are alarmed by the number of recent incidents that have compromised the confidentiality, availability, and integrity of critical infrastructure and services, and exposed the fragility of the Internet architecture. BGP offers limited performance and security mechanisms to protect the integrity of exchanged routing information and to provide authentication and authorisation of the advertised IP address space. Instead, each AS operator implicitly trusts that the routing information exchanged through BGP is accurate. As a result, the Internet backbone is potentially exposed. To better inform BGP administrators when choosing their routing paths, this thesis seeks to improve and advance current geolocation techniques, integrating geopolitical considerations into IP routing and introducing new IPv4 and IPv6 tools. By examining three distinct but interrelated aspects - improving current IP geolocation methods - enabling data routing for end users and network administrators - introducing a new IPv6 method of IP geolocation - this research aims to contribute to a more secure, efficient, and geographically aware Internet infrastructure. The thesis begins 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. The thesis builds on this current work by introducing new procedures and methodologies to create fine-grained multilayer maps of the structure of the Internet. Next, the thesis explores the additional benefits that IPv6 can bring to IP geolocation. IPv6 introduces a significant evolution in the area of Internet Protocols which resolves many of the issues with the limitations of IPv4 and provides**

**an improved framework for the future of the Internet. The concept of extension headers is a feature that enhances the IPv6 protocol's flexibility and functionality, and it is key among these advancements. The thesis conceptualises the design of a new IPv6 extension header, which aims to incorporate a geopolitical dimension into each data packet, optionally allowing network paths to be dynamically adjusted based on country codes of transit networks. The thesis builds on this tool by developing a new IPv6 tool to map network infrastructure, aiming to surpass current methodologies in accuracy, comprehensiveness, and utility. The tool provides a more precise and comprehensive mapping of the network's topology, including geolocation data and peer connections of network nodes. The thesis discusses how we can build on these foundational tools by combining them to produce new fault-finding techniques and a robust network analysis methodology. These methods and tools will benefit BGP administrators by informing them of better routing decisions, helping to avoid possible single points of failure, and enhancing overall network resilience. Finally, we discuss some limitations of the proposed approach and summarise some next steps needed towards accurate and complete Internet infrastructure maps.**

# Acknowledgements

This thesis marks the culmination of eleven years of dedication, research, and personal growth, a journey that would not have been possible without the steadfast support of those around me.

First and foremost, I would like to express my heartfelt gratitude to my wife, Christina, and to my family and friends, whose unwavering support, patience, and encouragement sustained me through every challenge. Their belief in me has been a constant source of strength.

I am deeply thankful to my supervisors for their guidance, insight, and encouragement throughout this research. In particular, I am profoundly grateful to Professor David Hutchison, whose mentorship has been instrumental from the outset. His thoughtful feedback, tireless support, and belief in the value of this work have shaped both this thesis and my academic journey. Without him, this project would not have reached completion.

I would also like to thank the National Cyber Security Centre (NCSC) Ed for funding this research and supporting its objectives. Their backing made this work possible. I am equally grateful to Lancaster University for providing an outstanding academic environment, comprehensive resources, and a community that fostered my growth as a researcher.

To my colleagues, collaborators, and all those who contributed in ways both large and small—whether through discussion, technical assistance, or moral support, thank you. Your contributions are sincerely appreciated and have left a lasting impression on this work.

# Declaration

I declare that the work presented in this thesis is, to the best of my knowledge and belief, original and my own work. The material has not been submitted, either in whole or in part, for a degree at this or any other university. This thesis does not exceed the maximum permitted word length of 80,000 words, including appendices and footnotes, but excluding the bibliography. A rough estimate of the word count is: 49340

Paul McCherry

The publication shown below has been created directly from the thesis, from which large portions of this published work is used:

P. McCherry, V. Giotsas, and D. Hutchison (May 2023). "On Improving the Accuracy of Internet Infrastructure Mapping". In: *IEEE Access, vol 11*. Insititute Of Electrical and Electronic Engineers, pp. 59935–59953. DOI: 10.1109/ACCESS.2023.3281333

The following publications have been generated while developing this thesis, and to an extent have guided the thesis into what it has become:

H. Alshaer, N. Uniyal, K. Katsaros, K. Antonakoglou, S. Simpson, H. Abumarshoud, H. Falaki, P. McCherry, C. Rotsos, T. Mahmoodi, R. Nejabati, D. Kaleshi, D. Hutchison, H. Haas, and D. Simeonidou (2020). "The UK Programmable Fixed and Mobile Internet Infrastructure: Overview, Capabilities and Use Cases Deployment". In: *IEEE Access, Vol 8*. Insititute Of Electrical and Electronic Engineers, pp. 175398–175411. DOI: 10.1109/ACCESS.2020.3020894

S. Simpson, A. Farshad, P. McCherry, A. Magzoub, W. Fantom, C. Rotsos, N. Race, and D. Hutchison (Nov. 2019). "DataPlane Broker: Open WAN control for multi-site service orchestration". In: *2019 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, pp. 1–6. DOI: 10.1109/NFV-SDN47374.2019.9040084

C. Rotsos, A. Marnerides, A. Magzoub, A. Jindal, P. McCherry, M. Bor, J. Vidler, and D. Hutchison (Nov. 2020). "Ukko: Resilient DRES management for Ancillary Services using 5G service orchestration". In: *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 1–6. DOI: 10.1109/SmartGridComm47815.2020.9302980

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

## 1.1 Background

A flourishing economy depends on a resilient Internet, which is crucial for the advancement of a country, forming a backbone for modern society to progress and prosper, supporting businesses, fuelling innovation, and connecting communities. There is much more to network resilience than simply preventing network failures or disaster recovery. A resilient network acknowledges the inevitability of incidents, prioritising robust system design, rapid restoration of services, and pre-emptive planning to mitigate the impact of outages.

The Internet was originally designed to be capable of rerouting traffic packets around outages, but in recent years, the growing use of Content Delivery Networks (CDNs), increased centralisation, and the trend toward Internet flattening are arguably affecting the intended resilience of Internet traffic delivery (Gill, Arlitt, Li, and Mahanti, 2008) (Merrill and Narechania, 2023).

Indeed, Merrill and Narechania express concerns about the current trends of departing from the older decentralised Internet tier 1,2 and 3 models to a more flattened and concentrated Internet where CDNs establish their own proprietary networks directly with Internet service providers (Merrill and Narechania, 2023). Their results suggest that this flattening reflects a change in the topology of the

Internet, and argue that it is difficult to overstate the significance of this change for the structure of the Internet. They argue that this change means that the Internet has effectively shifted away from data packet transmission across the tier 1 providers and towards a network of CDNs. Although this new model has helped deliver better and more secure services to a wide range of countries around the world and enabled new applications such as streaming audio and video, it has not come without a cost. There are now central points of failure on the network that resist scrutiny and oversight. These CDNs are effectively black boxes that frustrate efforts to improve reliability.

Recent global outages such as Fastly (Duffy, 2021), (Rockwell, 2021), (Medina, 2021), the Akamai outage (Akamai, 2021) and more recently the Australian national outage arguably caused by Optus (Reuters, 2023), are becoming worryingly common.

The Internet is a vast and complex mesh of Autonomous Systems (ASs) with few tools and limited capabilities to measure the topology, leaving many aspects of the Internet structure as an opaque cloud. To understand, analyse and resolve Internet routing issues, researchers and network operators need to be able to view the routing topology. However, protocols have been designed in such a way as to mask many details of Internet operation (Merrill and Narechania, 2023). This lack of public visibility into the Internet's structure affects the effectiveness of risk assessments and disaster planning, which can have detrimental effects on security and reliability.

Current tools such as Traceroute to infer geolocation have been shown to have many limitations and challenges (Willinger and Roughan, 2013), whilst the use of BGP (Border Gateway Protocol) to map the Internet infrastructure results in many geolocation ambiguities (Winter, R. Padmanabhan, King, and Dainotti, 2019), (Giotsas, Smaragdakis, Huffaker, Luckie, and Claffy, 2015). These challenges highlight the need for not only new tools and methodologies to improve IP geolocation accuracy but also innovative approaches to enhance the utility of existing tools including Traceroute. By refining how Traceroute data is collected, interpreted, and cross-referenced with complementary datasets such as DNS names, latency

measurements, and external infrastructure databases, it is possible to address some of its inherent limitations. Combining traditional tools with modern analytical techniques, such as machine learning or multi-path measurements, can extract more granular and reliable geolocation data, thereby improving the precision and robustness of Internet infrastructure mapping

**Roadmap** Section 1.1 highlights the importance of a resilient Internet, the impact of centralization and Internet flattening, and the limitations of current tools like traceroute and BGP. Section 1.2 explores the need for transparency, better routing security, and geographic awareness to mitigate risks and improve network resilience. Section 1.3 summarizes the structure and focus of each chapter, linking them to the research objectives. Section 1.4 defines the goals of creating fine-grained Internet maps, improving IP geolocation, and enhancing IPv6-based routing. Section 1.5 introduces key innovations, including advanced geolocation methods and a novel IPv6 extension header for geopolitical routing. Section 1.6 explains the origins of the research material.

## 1.2 Motivations for Internet Cartography

A report by the UK National Cyber Security Centre 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 (NCSC, 2021). It is therefore paramount to develop appropriate methods and practices to make BGP more secure, thus maintaining the integrity of the routing system which relies almost exclusively on it. The European Network and Information Security Agency (ENISA) concluded in a 2015 report that the current lack of structural transparency is the biggest obstacle to addressing the inherent vulnerabilities and architectural shortcomings of the Internet routing system (ENISA, 2015).

The ability to perceive, comprehend, and project the status and activities within

cyberspace is referred to as Cyber Situational Awareness (CSA). CSA enables informed decision-making to address security risks and vulnerabilities. Furthermore, it extends beyond monitoring network activity to understanding the geographical, logical, and operational dynamics that influence cyber resilience and security.

Knowledge of the geographical locations of the Internet infrastructure is a necessary requirement for CSA. It can allow us to understand and mitigate risks related to topological vulnerabilities and design more resilient networks and routing policies (Motamedi, Rejaie, and Willinger, 2015). For example, the ability to predict what would happen if a colocation facility or Internet Exchange (IXP) fails can inform better fallback policies and more efficient resource allocation. Enhancing security and resilience in the routing system is critical for ensuring uninterrupted Internet service delivery, especially as the threat landscape continues to evolve.

To develop appropriate prediction techniques, we must measure the relevant routing paths and infer the interconnection points traversed. Analysis of these paths can provide clues about connectivity changes that will prevent choke points, single point failures, or serious performance degradation due 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. Building a resilient Internet infrastructure will also mitigate the risks associated with geopolitical tensions, natural disasters, and cyberattacks.

Critics may contend that transparency reduces the time adversaries spend identifying attack vectors and allows them to exploit 'low-hanging fruit' more readily. However, transparency is a crucial enabler for proactive identification and resolution of vulnerabilities, reducing the overall attack surface. It fosters collaboration among security professionals, leading to faster and more effective solutions to emerging threats and reinforcing the broader security ecosystem.

By highlighting weaknesses, organisations can prioritise and strengthen their defences, ensuring that security measures are robust and well-tested rather than relying on obscurity, which creates a false sense of security. Historical data breaches

demonstrate that determined attackers will exploit vulnerabilities regardless. Transparency equips defenders to close gaps before they are widely exploited and builds trust with stakeholders by showing a commitment to security and accountability—an essential trait in today's interconnected digital environment.

While transparency may inform attackers, its benefits far outweigh the risks. The greatest danger lies in relying solely on obscurity, which often leads to overlooked vulnerabilities and inadequate defences. A balanced approach, integrating transparency with strong and well-designed security measures, delivers a more effective and sustainable solution. Furthermore, embedding resilience as a foundational design principle ensures that the Internet can adapt and recover rapidly from attacks and even failures, maintaining continuity of services under highly adverse conditions.

## 1.3 Thesis Outline

This thesis is organised as follows:

Chapter 1 provides a background on the current state of IP geolocation, the motivations, and the objectives of this thesis. Chapter 2 reviews the current methods and literature, while also providing the terminology and preliminaries of Internet cartography. In Chapter 3, we introduce and evaluate current methods and tools. Chapter 4 uses the tools and methods described in Chapter 3 to introduce and evaluate a new idea of using Internet Exchange Points as Vantage Points to get closer to populated areas, thus reducing overall error margins. Chapter 5 uses the methods described in Chapter 4 which creates Vantage Points from IXPs along with many other new ideas and methods to finally create fine-grained multilayer maps of the Internet infrastructure.

Chapter 6 introduces an advanced IPv6 network infrastructure mapping tool, designed to enhance and expand the capabilities of existing mapping tools such as traceroute.

Chapter 7 introduces an innovative IPv6 extension header that incorporates geopolitical awareness into network routing, which will provide network administrators and end users with some control over the route their data packets take across the Internet.

Chapter 8 provides a conclusion, limitations and a look ahead to possible future work in this important area.

## 1.4 Objectives and Scope of the Thesis

The first section presents the current techniques in IP Geolocation and builds upon that research by developing new methods in IP infrastructure mapping with the ultimate goal of creating fine-grained multilayer maps of the Internet infrastructure that are currently lacking. The aims of this work are as follows.

*In Chapters 1-5*

- To extend DNS-based geolocation from city-level to facility-level and address shortcomings of the state-of-the-art tools with respect to their limited geographical coverage.
- 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 bidirectional interfaces.
- Construct a data set 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 including London.
- To evaluate the inferences and estimate its success using a carefully curated dataset obtained by two of the largest London IXPs.

*In Chapter 6*

- Introduce a new tool for mapping IPv6 network infrastructure, with the objective of surpassing current methodologies in accuracy, comprehensiveness, and utility.

- Detail the necessity for such a tool, discussing enhancements over existing technologies, outlining the proposed changes to network protocols and infrastructure,
- Demonstrate the tool's potential benefits for network analysis and troubleshooting.
- Attempt to offer a more detailed visualisation of the structure of the Internet, addressing both technical implementations and the broader implications for network research and administration.

*In Chapter 7*

- Design a new IPv6 extension header that incorporates geopolitical information in the form of country code bitfields.
- Evaluate technical feasibility and, if possible, to assess the practicality of implementing and deploying the new extension header in real-world network environments, considering current router capabilities and infrastructure
- Analyze security and privacy implications, critically examining the security and privacy concerns that may arise from the use of the proposed extension header, especially in the context of international data transmission and legal compliance.
- Study the impact of geopolitical routing, investigating the potential impacts and benefits of incorporating geopolitical considerations into IPv6 routing, from both technical and policy perspectives.

*In Chapter 8*

The thesis ends with the conclusions that can be drawn and points toward the next steps that need to be taken.

## 1.5 Contributions

Using existing state-of-the-art tools, new methods and procedures have been created to discover and develop finer-grained maps of the Internet infrastructure. These new methods can facilitate a better understanding of the resilience of the Internet infrastructure and allow for the prioritisation of robust system design, rapid restoration of services, as well as providing vital data for pre-emptive planning.

A new IPv6 extension header has been conceptualised and presented, which aims to incorporate a geopolitical dimension into each data packet, allowing administrators and end-users to dynamically adjust network paths based on country codes of transit networks. This addresses a growing need for data controllers and processors to comply with the data protection laws of each country.

A new IPv6 extension header has been conceptualised and presented, which aims to incorporate a geopolitical dimension into each data packet, allowing administrators and end-users to dynamically adjust network paths based on country codes of transit networks. This addresses a growing need for data controllers and processors to comply with the data protection laws of each country. Additionally, a new IPv6 network infrastructure mapping tool has been conceptualised, representing a significant step forward in addressing the limitations of current Internet mapping methodologies. It offers a more precise and comprehensive mapping of the Internet topology, including geolocation data and peer connections of network nodes. The proposed changes to the node information protocol, along with suggested modifications to router kernels and installation procedures, underscore a holistic approach to improving the accuracy and utility of network mapping.

While the focus remains on the design and theoretical development of these innovations, the proposed tools pave the way for the creation of more sophisticated applications

## 1.6 Acknowledgment of Prior Work

This section explains the origin of material for the chapters in this thesis. Chapters 2, 3 and 5 are based on material published in 'On Improving the Accuracy of Internet Infrastructure Mapping' (McCherry, Giotsas, and Hutchison, 2023) in which McCherry conducted the bulk of the research and technical work in the publication. The rest of the chapters were composed for the benefit of the thesis by McCherry.

# Chapter 2

# Related Work

## 2.1 Introduction

This chapter introduces the terminology and preliminaries of mapping the Internet infrastructure and reviews the relevant literature on Internet cartography. **Roadmap** Section 2.2 reviews passive, active, and hybrid geolocation techniques, noting the strengths and limitations of using static databases, latency-based measurements, and multi-source data integration for mapping network infrastructure. Section 2.3 identifies limitations in traditional methods, including lack of ground-truth data, traceroute inaccuracies, and challenges with infrastructure diversity and circuitous routing. Section 2.4 argues for refining geolocation methods to address these challenges, highlighting the potential of hybrid approaches to achieve higher accuracy, particularly at the facility level.

## 2.2 Traditional IP Geolocation Methods

Mapping network infrastructure has been a foundational aspect of understanding and optimising the Internet's functionality, with various tools and methodologies developed over the years to explore and document this complex landscape. The early tools and techniques involved tools such as Ping, which was designed to

measure the reachability of hosts across an IP network. Ping works by sending ICMP "echo request" packets to the target host and listening for "echo response" replies (Mühlbauer, Feldmann, Maennel, Roughan, and Uhlig, 2006). Ping has been instrumental in diagnosing network connectivity issues. Traceroute was developed in 1988 and allows the mapping of the path of packets through an IP network in transit to their destination. The traceroute program incrementally sets the Time To Live (TTL) values of subsequent ICMP packets and then observes the node where the packet is dropped, which is revealed through ICMP REPLY "time exceeded" messages. The REPLY messages contain the IP address where the packet was dropped, and traceroute measures the time delay between sending the REQUEST packet and receiving the REPLY packet. Furthermore, ICMP messages have been used in various network diagnostic tools to map network infrastructure by identifying active hosts and potential routing issues. However, the use of Ping and Traceroute to map the Internet infrastructure has been shown to produce many small to large errors over the years (Motamedi, Rejaie, and Willinger, 2015), (Willinger and Roughan, 2013), (Oliveira, Pei, Willinger, B. Zhang, and L. Zhang, 2010). This has not deterred researchers from attempting to improve IP geolocation using a host of various methods based on these tools.

IP-based geolocation maps an IP address to the geographical 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.

### 2.2.1 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 (Davis, Vixie, Goodwin, and Dickinson, 1996) that are designed to hold the geographical coordinates of the IP address host.

However, they are rarely created by administrators (Graham-Cumming, 2014). Davis et al cite that there are security reasons for this such as LOC record information could be used to plan a penetration of physical security. This highlights that location records can be very sensitive data. DNS LOC records are publicly available and there is no verification that the location associated with a domain is accurate.

Another source of passive geolocation data is the WHOIS protocol (Daigle, 2004), which stores information on the owners of Internet resources, including IP addresses. Among this information is often the address of the organisation 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) RFC8805 (Kline, Duleba, Szamonek, Moder, and Kumari, 2020). Despite the technical potential of Geofeeds, network operators and ISPs rarely adopt them because they perceive the risks, costs, and efforts outweigh the benefits. Many ISPs and cloud providers dynamically reassign IP addresses based on demand and manually updating and maintaining accurate Geofeed files is seen as an additional burden. All of this means that geolocation data can quickly become out of data causing inaccuracies in the Geofeed data of these IP addresses. Finally, several commercial geolocation services that use proprietary methods provide location data to subscribers, such as Maxmind (Maxmind, 2024), IP2Location (IP2location, 2024), and Neustar (Neustar, 2024). However, past research on the accuracy of these databases shows that commercial databases can be highly inaccurate, especially for router and infrastructure IPs (Gharaibeh, Shah, Huffaker, H. Zhang, Ensafi, and C., 2017) (Poese, Uhlig, Kaafar,

Donnet, and Gueye, 2011) (Shavitt and Zilberman, 2011). For example, research carried out in 2017 concluded that whilst the accuracy of IP location was high at country level, regional accuracy was worse with the best databases giving an accuracy of 50%. The best city-level accuracy was 30% and there was no data for facility-level accuracy (Komosny, Voznak, and Rehman, 2017).

A technique used by many commercial IP geolocation companies is to build a database of mappings between geolocation and IP addresses over time. This data can come from many sources such as end users, the results of traceroute projects, DNS lookups, and other sources, however, geolocation databases are notoriously difficult to keep up to date, and their accuracy depends on the source and age of the data they reference (Poese, Uhlig, Kaafar, Donnet, and Gueye, 2011).

### 2.2.2 Active IP Geolocation

Active IP geolocation is based on network-level latency measurements between a node with a well-known location (landmark) and the IP address that must be geolocated. Assuming that the Speed of the Internet (SoI) is known, the latency can then be translated to the distance from the landmark (V. Padmanabhan and Subramanian, 2001). Because of the high overheads in maintenance that passive geolocation methods require and the inherent ambiguities that can entail from historically out-of-date data, active geolocation tends to be more accurate than passive geolocation. However, active, IP geolocation incurs a much higher measurement overhead, and it is hard to scale to geolocation of millions of IP addresses. Additionally, the SoI is not fixed, but depends on the transmission medium and the network conditions. Geoping is one of the earliest active geolocation techniques introduced in 2001 by Padmanabhan and Subramanian. Geoping measures the latency between multiple landmarks and creates a latency vector for each landmark. It then measures the latency from all the landmarks to the target IP and geolocates it to the landmark with the most similar latency vector.

In 2006, 'Constraint-Based Geolocation' (CBG) was proposed as an improvement

of Geoping. CBG also employs measurements from multiple landmarks but 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. Multilateration uses multiple landmarks in the region of a target to estimate the target is somewhere within a circle centred at the landmark and with a radius corresponding to the one way trip time of a traceroute from the landmark to the target. If any circles overlap then the possible region where the target is located is reduced to this overlapping region (Gueye, Ziviani, Crovella, and Fdida, 2006).

In 2022 a new algorithm based on router error training was proposed (Zu, Z. Luo, and F. Zhang, 2022), which requires an exhaustive mapping of the Metropolitan Area Network (MAN) of the city where the target IP should be located to infer its street address location. Although this technique achieves high accuracy, it is limited to cities with a suitably large number of measurement Vantage Points.

### 2.2.3 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, which would provide a useful tool for detecting poor routing structures and understanding why damaging routing events occurred (NCSC, 2021). A method in which a combination of data sources could be used, such as crowd-sourcing, reverse DNS records, tagged naming schemes, Round Trip Time (RTT) delay-distance models, and Internet exchange points, was proposed. The following methods attempt to use one or more of these data sources.

In "Topology-Based Geolocation" (TBG), the authors 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 (Katz-Bassett, John, Krishnamurthy, Wetherall, Anderson, and Chawathe, 2006). 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 is a model-based active geolocation service and uses a probabilistic approach to derive a generic model of the relationship between network delay and geographic distance rather than using a predetermined SoI value or separate calibration data for each point of reference (Laki, Mátray, Hága, Sebok, Csaba, and Vatta, 2011). This delay distance model was then used to geolocate an IP address. The authors of "Towards street-level client-independent IP geolocation" refined the granularity of CBG to achieve street-level geolocation (Wang, Burgener, Flores, Kuzmanovic, and Huang, 2011). To this end, they mine web-based geolocation hints for locally hosted web servers to significantly expand the list of passive landmarks. They tried to leverage the observation that "many entities host their web services locally", but since then the trend of cloud-hosted services and resource centralisation certainly inhibit the applicability of their technique.

The developers of Octant claim that it is a comprehensive framework for the geolocalization of Internet hosts and that it considers the locations of intermediate routers as landmarks to geolocate the target (Wong, Stoyanov, and Sirer, 2007). 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 minimisation constraint satisfaction problem. Although Octant achieves better accuracy than CBG, the authors noted that extracting useful positive and negative information is a challenge. In contrast to CBG and TBG, the authors allowed for circuitous routes and used 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.

RIPE Atlas 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 (RIPE, 2015). One of the Atlas 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 more than 10 ms are discounted. The probe with the minimum latency to the target IP is then selected, and the distance is calculated using the signal transmission speed through the optical fiber 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.

New approaches involve the use of the Border Gateway Protocol (BGP), which is the protocol used to make routing decisions between Autonomous Systems (AS) on the Internet. BGP was not designed as a mapping tool, but data derived from BGP tables by researchers has been crucial to understanding the topology of the Internet at a macro level (Giotsas, Smaragdakis, Huffaker, Luckie, and Claffy, 2015). Projects such as RouteViews (RouteViews, 2024) and RIPE NCC (RIPE, 2015) have aggregated BGP information to analyse and visualise Internet paths and interconnections between ISPs. Giotsas et al. developed a method known as Constrained Facility Search (CFS), which combines data from various sources such as Internet Exchange websites, PeeringDB, and traceroute measurements to infer the connection facility of a specific IP address (Giotsas, Smaragdakis, Huffaker, Luckie, and Claffy, 2015). Using this method, they were able to locate 71% of the router interfaces to a specific facility.

A method has been developed called Hints-Based Geolocation (HLOC), which extracts geohints from router DNS names (Scheitle, Gasser, Sattler, and Carle,

2017), similar to Octant (Scheitle, Gasser, Sattler, and Carle, 2017). It then validates these hints by selecting several RIPE Atlas probes based on the extracted geohints 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 9 ms over a maximum distance of 900 km to accommodate packet buffering, processing, and scheduling delays. If 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.

Motamedi et al. 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 (Motamedi, Yeganeh, Chandrasekaran, Rejaie, Maggs, and Willinger, 2019).

DNS names do not accurately map geolocations without improved lookup tables, prompting the proposal to use Looking Glass servers as additional landmarks (Livadariu, Dreibholz, Al-Selwi, Bryhni, Lysne, Bjørnstad, and Elmokashf, 2020). Livadariu et al. also investigated the accuracy of RIPE 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.

A significantly improved DNS to geolocation lookup was demonstrated by compiling an extensive list of regular expressions (Luckie, Huffaker, Marder, Bischof, Fletcher, and Claffy, 2021). Luckie et al. reported that their solution reads the DNS records that network operators embed into router interface hostnames and can then accurately and comprehensively extract geographic data. They evaluate their results

against ground truth data sets and report a much improved success rate compared to other solutions such as HLOC. Machine Learning was applied to the task of learning DNS names and their locations, demonstrating significant improvements over previous academic baselines and proving to be both complementary and competitive with commercial databases (Dan, Parikh, and Davison, 2021a). The authors state that the benefits of this approach are that the solution does not require any human annotation or intervention and given enough training data it can automatically learn rules that were previously handcrafted making the approach easier to scale world-wide.

In further research, Dan et al. 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 make the assumption all IPs in that prefix are also in that location (Dan, Parikh, and Davison, 2021b). Other researchers argue that such an approach can result in large inaccuracies and can come at the cost of precision especially when IP ranges are not contiguous geographically (Motamedi, Rejaie, and Willinger, 2015), (Luckie, Huffaker, Marder, Bischof, Fletcher, and Claffy, 2021). The combination of IP interpolation with other methods such as using traceroutes and ground truth datasets for validation, may mitigate some of these criticisms.

In-band Operations, Administration, and Maintenance (IOAM) is a network measurement and monitoring technology. IOAM is a relatively new approach that embeds data collection information directly into data packets as they traverse the network (Iurman and Donnet, 2020). This is a method that allows for detailed tracking and measurement of packet flows and offers granular insights into network performance and topology. IOAM enables devices to sample service traffic in real time at high speed, adds IOAM information (metadata, including the device ID, inbound and outbound interfaces, and timestamp) to the sampled data, and can proactively send the sampled data to an analyser for analysis. In this way, the network running status is detected and monitored in real time. IOAM is

gaining traction and is supported by various network equipment manufacturers and standards bodies, but its deployment across the Internet is not uniform, and it is more commonly found in specific networks that prioritise detailed performance monitoring and operational visibility.

Over the years, several large-scale projects have been initiated to create comprehensive mappings of the Internet's infrastructure, each focusing on different aspects of the network to address various research and operational challenges.

The Centre for Applied Internet Data Analysis (CAIDA) has played a pivotal role in Internet research by conducting extensive studies and developing mapping initiatives. CAIDA employs a variety of tools, including Traceroute and BGP (Border Gateway Protocol), to analyze the Internet's topology, connectivity, and performance issues. Their work includes generating large-scale datasets, visualizations, and insights that help researchers and policymakers understand the dynamics of Internet infrastructure and its evolution. CAIDA's focus extends beyond mapping, addressing broader challenges such as security, scalability, and the resilience of the Internet ecosystem (CAIDA, 2024).

Internet Atlas, another significant project, concentrates on mapping the physical infrastructure of the Internet. This includes documenting undersea cables, terrestrial fiber optic networks, data centers, and other critical components that underpin global connectivity. By focusing on the physical layer, Internet Atlas provides valuable insights into the geographic dependencies of Internet infrastructure, helping researchers and governments evaluate risks related to natural disasters, geopolitical tensions, or other disruptions. Its maps and analyses are crucial for enhancing the understanding of the Internet's physical mechanisms and identifying vulnerabilities (Berkeley, 2024).

Measurement Lab (M-Lab) takes a different approach by emphasizing Internet performance transparency. This global collaboration of researchers collects open data on various Internet performance metrics, including network speed, latency, and reliability. By making these datasets freely available, M-Lab enables researchers,

policymakers, and industry stakeholders to study the health and performance of the Internet. M-Lab's efforts support the goal of increasing accountability and transparency in Internet infrastructure, empowering users with data to advocate for better services and fostering an open, equitable Internet (Mlab, 2024).

Each of these projects contributes a unique perspective and set of tools to the broader understanding of Internet infrastructure. Together, they help build a more comprehensive picture of the Internet's topology, performance, and physical foundation, enabling improved resilience, security, and accessibility in the face of evolving challenges.

## 2.3 Challenges

Several previous studies have identified numerous issues with techniques that attempt to map the Internet topology and geolocate IP addresses using traceroutes. (Motamedi, Rejaie, and Willinger, 2015) (Willinger and Roughan, 2013), which is summarised in this section.

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. The scarcity of ground-truth datasets, as noted by researchers, forces reliance on incomplete or coarse-grained abstractions of Internet topology (Motamedi, Yeganeh, Chandrasekaran, Rejaie, Maggs, and Willinger, 2019). These abstractions miss many details of interconnections and render them largely irrelevant to real-world Internet engineering problems (Oliveira, Pei, Willinger, B. Zhang, and L. Zhang, 2010), and many of the findings based on simplistic models are controversial or misleading, due to the incompleteness and inaccuracies of the maps produced (Willinger and Roughan, 2013).

Layer 2 clouds are largely opaque to tools that use 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, it was observed that a significant fraction of proxy clients were located several hundred to thousands of kilometers from the location of the proxies (V. Padmanabhan and Subramanian, 2001). Network delay measurements are oblivious to this and incorrectly return the location of the proxy server.

Traceroute Round Trip Times (RTT) includes both application-layer and network-layer delays, and if a measurement device is overloaded or underresourced, then the RTT times may be inflated. This is a problem that RIPE Atlas probes may encounter (especially older versions) (Holterbach, Pelsser, Bush, and Vanbever, 2015).

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 example, Figure 2.1 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. The Speed of Light (SoL) constraint is a geolocation sense check that uses the fact that network packets cannot travel faster than the speed of light in fiber-optic cables (typically approx $0.66c$). This constraint helps rule out unrealistic paths where the measured RTT implies an impossible travel distance. If we make the assumption that 0.66 of the speed of light is approximately 200 km/ms then the one-way trip distance can be assumed to be RTT multiplied by 100 km.

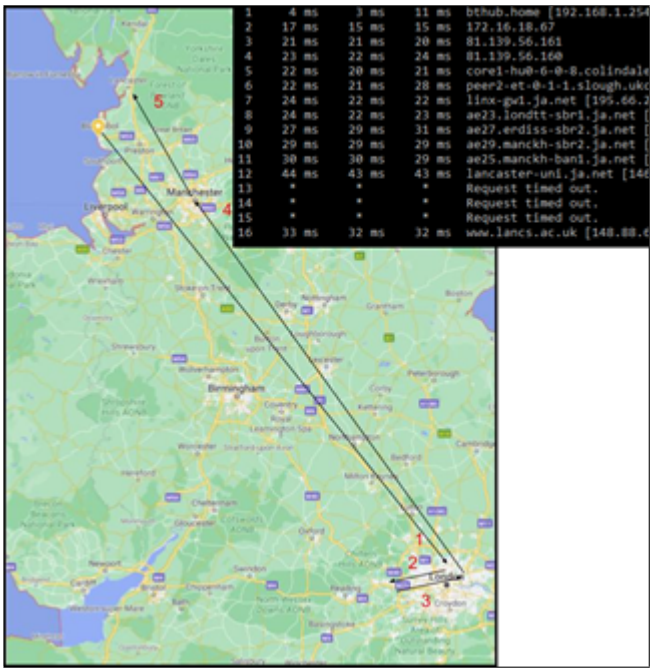


Figure 2.1: Example of circuitous route from Blackpool to Lancaster

$$\begin{aligned} &\text{Distance to Target} = Time \times Speed \\ &\text{Distance to Target} = 32 \text{ ms} \times (200 \text{ km/2}) \\ &\text{Distance to Target} = 32 \text{ ms} \times 100 \text{ km/ms} \\ &\text{Distance to Target} = 3200 \text{ km} \end{aligned} \tag{2.1}$$

The calculated distance to target is therefore 3200 Kilometers. However, the actual distance from Blackpool to Lancaster is 40 kilometers and dividing this by 200 km/ms 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. The actual time it took to make the one way trip is RTT/2 which is 32 ms/2 or 16 ms. 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 9 ms to account for these issues.

It should be noted that the estimated packet speed in fibre of 200 km/ms and the distance to target calculation of 100 km multiplied by RTT are used throughout this thesis.

Another complication of traceroute-based measurements is the diversity of infrastructure in different regions of the world, leading to different delay coefficients (Candela, Gregori, Luconi, and Vecchio, 2019). 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.

## 2.4 Summary

Although active IP geolocation can provide real-time updates and requires no administrative upkeep, many active IP geolocation solutions employ active measurements, notably traceroute, to discover network interfaces and topology. However, the traceroute tool was designed primarily for troubleshooting and its use in network discovery is not what it was designed for. IP geolocation remains an essential but complex field with many unresolved challenges. Existing techniques, including passive (database-driven), active (latency-based), and hybrid methods, offer varying degrees of accuracy but struggle at facility-level resolution.

Problems such as circuitous routes, different router configurations, ICMP filtering, route congestion, and technologies such as microwave links, SDN, ATM, and MPLS clouds can have varying effects on delay-based geolocation techniques. IP interpolation is commonly used to fill gaps, its assumptions about IP block contiguity often lead to inaccuracies. Passive methods suffer from out-of-date or completely incorrect information; therefore, both active and passive methods appear to have their own strengths and weaknesses. BGP-based geolocation provides logical topology insights but lacks physical accuracy, as BGP paths do not necessarily correlate with real-world geography. Geofeeds, discussed in RFC 8805 (Kline, Duleba, Szamonek, Moder, and Kumari, 2020), present a potential solution, allowing ISPs to self-publish IP geolocation data, yet participation remains low due to privacy concerns and lack of incentives.

To depart from oversimplified models, it has been argued that it is necessary to identify the geolocation of Internet infrastructure, which would provide a useful tool for detecting poor routing structures and understanding why damaging routing events occurred (Aben, 2013). Aben proposed a method in which a combination of data sources could be used, such as crowd-sourcing, reverse DNS records, tagged naming schemes, RTT delay-distance models, and Internet Exchange Points. Emerging techniques such as machine learning-based geolocation and crowdsourced ISP data offer promising directions but require wider adoption and standardization.

Overall, achieving high-accuracy geolocation, particularly at the facility level, remains a work in progress, necessitating new hybrid approaches that integrate network measurements, AI, and real-time data validation. A hybrid mix of active and passive techniques has the potential to alleviate these weaknesses and offer the most accurate IP geolocation solutions.

The key contribution of the methods developed in Chapters 4 and 5 is that they will extend DNS-based geolocation from the city level to the facility level and address shortcomings of state-of-the-art geolocation techniques with respect to their limited geographical coverage. We showcase the applicability of our work by geolocating over a thousand IP addresses at the level of colocation facilities. Although the data set 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 metropolitan Internet hubs.

# Chapter 3

# Introduction of Geolocation Routing Tools

## 3.1 Introduction

This chapter begins by introducing the tools used by researchers and analysts in the IP Geolocation arena and the current methods used and built upon by the tools developed in Chapters 4 and 5 of this thesis. It then examines and evaluates these current methods to ensure that there is a firm basis for the new tools developed later in this thesis. **Roadmap** Section 3.2 explains the four distinct levels of Internet mapping. Section 3.3 describes Paris Traceroute and how it differs from that of normal traceroute. Section 3.4 highlights the role of IXPs as critical nodes in the Internet ecosystem, facilitating data exchange and enhancing network reliability. It explores their physical location, importance in reducing latency, and the method for mapping interconnection facilities in the UK. Section 3.5 introduces the Peeringdb website which provides a freely available resource on Internet connections, Section 3.6 introduces OpenStreetMap as an open-source geographical database used extensively in this thesis for visualizing geolocation data. Section 3.7 introduces the IXPDB website which provides comprehensive world wide data on Internet Exchange Points. Section 3.8 provides an overview

of the RIPE Atlas platform, its methodologies like the Single-Radius engine, and its significant role in geolocating IP infrastructure. Section 3.9 examines the CAIDA ITDK dataset for router geolocation and compares its accuracy with verified ground-truth datasets, identifying areas for improvement. Section 3.10 evaluates the performance of Constraint Based Geolocation (CBG) using RIPE Atlas probes and identifies challenges like insufficient overlapping vantage points and latency issues that reduce geolocation accuracy. Section 3.11 discusses the steps of RIPE's Single-Radius method for geolocation and highlights its accuracy in specific regions but acknowledges its limitations in poorly connected areas. Section 3.12 explores the Shortest Ping method as an alternative to RIPE's Single-Radius engine, noting its simplicity and comparable accuracy but also its dependency on available landmarks. Section 3.13 summarizes findings, critiques traditional methods like CBG and Single-Radius, and sets the stage for the proposed solution in the next chapter, leveraging IXPs as additional landmarks to enhance IP geolocation accuracy.

## 3.2 Resolutions in Geographical Mapping

Researchers outline four distinct levels of resolution for mapping the Internet topology, each providing a different perspective on the network's structure (Motamedi, Rejaie, and Willinger, 2015):

- Interface Level: This is the most granular level of Internet topology. It represents the individual IP interfaces of routers and end-hosts. Each IP address corresponds to a specific interface, providing the highest level of detail but lacking aggregation.

- Router Level: At this level, the topology aggregates multiple IP interfaces into a single entity, representing the physical or logical router they belong to. This level simplifies analysis by summarizing all the IP addresses of a single router into a single identifier.

- AS (Autonomous System) Level: An Autonomous System (AS) refers to a network or collection of networks managed by a single administrative entity with independent connections to multiple networks. The AS level provides a high-level view of the Internet's structure, focusing on the relationships between large administrative domains.

- PoP (Point of Presence) Level: A Point of Presence (PoP) is a collection of routers owned by an AS in a specific geographic location, also known as an interconnection facility. This level strikes a balance between granularity and abstraction, making it ideal for geographical mapping of network infrastructure.

These resolutions provide different vantage points for studying and analyzing the Internet's topology. Motamedi et al consider the Point of Presence (PoP) as the ideal resolution for the geographical mapping of network infrastructure as it captures the infrastructure's regional characteristics while maintaining an appropriate level of abstraction. This thesis proposes a new method to map the Internet infrastructure using this PoP layer, starting with the discovery of the geolocation of each interconnection facility in the UK.

## 3.3 Traceroute and Paris Traceroute

Traceroute is a network diagnostic tool that maps the route packets take from a source to a destination by sending packets with incrementally increasing TTL (Time-to-Live) values. The time taken for the packet to reach its destination and return to the source is known as Round Trip Time (RTT). It is commonly used in latency measurements and IP geolocation techniques to estimate physical distances between network nodes. Traceroute helps in identifying routing paths and detecting latency issues. However, traditional traceroute has limitations when dealing with networks that use load balancing. Load balancing is a technique used by network devices to distribute traffic across multiple paths to optimize performance and redundancy.

In per-flow load balancing traditional traceroute sends probes with varying header fields, causing probes to take different paths in load-balanced networks. This creates inconsistent or incorrect path representations. Additionally traditional traceroute can erroneously report paths that do not exist due to the variation in the flow identifiers of its probes.

Paris Traceroute is an advanced tool designed to address these issues, providing more accurate and consistent results. It mitigates these issues by ensuring that all probe packets belong to the same "flow" as defined by load balancers. It achieves this by carefully crafting probe packets so that their header fields remain consistent.

The fields that must remain the same for a packet to be considered part of the same flow are:

- Source and Destination IP Address.
- Protocol Type.
- Source and Destination port numbers.
- Other fields in the IP Header.

Paris traceroute ensures changes are not made to fields used for load balancing. Instead, it makes changes to header fields that are still within the first 28 octets, such as, in the case of TCP probes, the sequence number, and in UDP probes, it is the checksum field. This requires manipulating the payload to yield the desired checksum, as packets with an incorrect checksum are liable to be discarded. In ICMP probes, it is a combination of the ICMP identifier and the sequence number. Paris Traceroute sets the value of the ICMP identifier and sequence number to keep constant the header checksum of all probes to a destination.

Packets with the same flow are then sent over the same link and by eliminating false path diversity, Paris traceroute provides a more accurate depiction of the network topology. Furthermore Paris traceroute is used by default by the RIPE Atlas platform.

## 3.4 Internet Exchanges

Internet Exchange Points (IXP) are key physical infrastructures in the Internet ecosystem. They allow Internet Service Providers (ISPs), content delivery networks (CDNs), and other large enterprise networks to interconnect directly, rather than through third-party networks. Direct interconnections serve multiple purposes in reducing costs, improving the latency of data exchange between networks, and improving bandwidth capacity whilst increasing redundancy and reliability of the Internet infrastructure. Data centres typically host IXPs, offering a central location in which to store their networking equipment, such as switches and routers that facilitate the exchange of Internet traffic. Internet Exchange Points can be viewed as the main crossroads of a country's Internet communications and excel at keeping data local to that country.

Internet Exchanges consist of a set of switches through which participating Internet service providers, transit providers, and content delivery networks (CDN) exchange data. They are housed in colocation facilities, are generally located close to large populations, and are therefore essential to the Internet network infrastructure. IXPs connect the facilities where they interconnect at layer 2 of the OSI model; therefore, data entering an IXPs network at one facility can traverse the IXP network and exit at any other facility where the IXP interconnects. Since Layer 2 traffic does not require IP-level routing for data transmission, it cannot be easily tracked within the IXP network. The absence of Layer 3 (IP-level) routing information within the Layer 2 domain means that data packets cannot be geolocated until they leave the Layer 2 network at a gateway router. This limitation can obscure the geographic path of data traversing an IXP, complicating efforts to accurately map the Internet's infrastructure. Addressing these challenges requires innovative approaches, such as leveraging additional data sources or combining Layer 2 physical maps with Layer 3 measurements, to enhance the precision of IP geolocation in networks heavily reliant on IXPs. Additionally, to complicate matters further, there are two methods by which an organisation may wish to connect to an IXP, direct peering, or remote

peering. Direct peering requires an organisation to have physical presence at a colocation facility where the IXP also has a presence, while remote peering allows an organisation to peer with the IXP using one of the IXP's partners, generally over layer 2 Multiprotocol Label Switching (MPLS) clouds. MPLS is a network routing technique that directs data packets based on labels rather than traditional IP addresses (Rosen, Viswanathan, and Callon, 2001). MPLS improves traffic engineering, reduces latency, and enables remote peering by allowing traffic to be routed through virtual Layer 2 networks without needing direct physical connections. Should a network wish to exchange traffic through an IXP, it is usually the case that a peering agreement is set up between two network operators. This ensures that the two network operators agree to exchange traffic between their networks without charging each other, under mutually agreed-upon terms. It enables direct interconnection between ISPs, CDNs, or other networks to optimize performance and reduce costs. However, the terms of these agreements can vary depending on the volume of traffic exchanged and the type of data. If traffic volumes are significantly unbalanced, then other compensation terms may be agreed. Once the agreements are in place, the networks are physically connected to the IXP switches. This connection can be made through a direct physical link if they have their own network equipment on site at the IXP location, or it can be connected using a virtual connection and facilitated by a third party. Once physical connections are established, the networks can start exchanging traffic. Should a user request data from one network from a service on another network, the data can now be directly routed through the IXP, which bypasses transit over other, possibly distant, networks, reducing latency and providing a better experience for the user. The Border Gateway Protocol (BGP) is used to route and manage traffic flow. IP address ranges are announced by networks using BGP to other networks, which allows for the efficient routing of data to its destination.

## 3.5 Peeringdb Website

PeeringDB is a publicly accessible, user-maintained database that provides details on Internet exchanges, network facilities, and AS interconnections (PeeringDB, 2024). It is widely used for geolocation research, network engineering, and interconnection planning. It is shown in Figure 3.1.

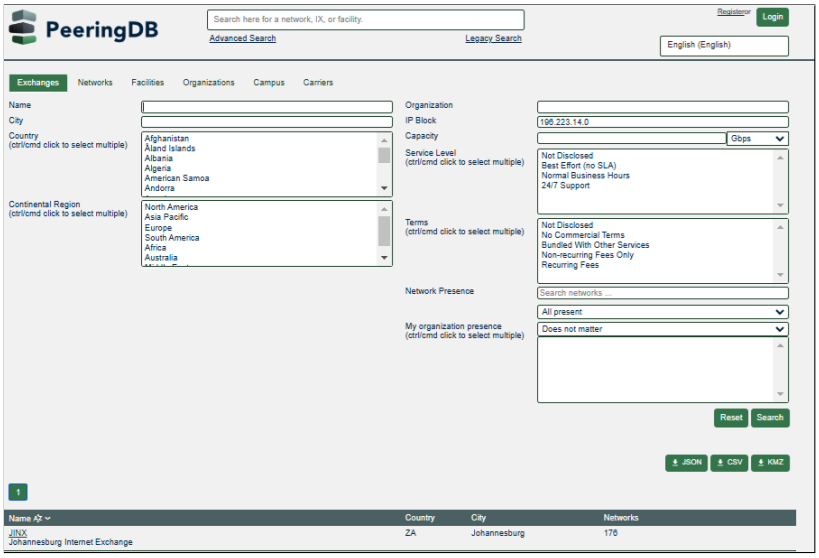


Figure 3.1: PeeringDB Website

## 3.6 OpenStreet Map

OpenStreetMap (OSM) is an open-source project that creates a freely editable geographical database in which tags can be created to provide information about elements (Openstreetmap, 2024), as shown in Figure 3.2. It is used extensively in this thesis to portray geolocation data.

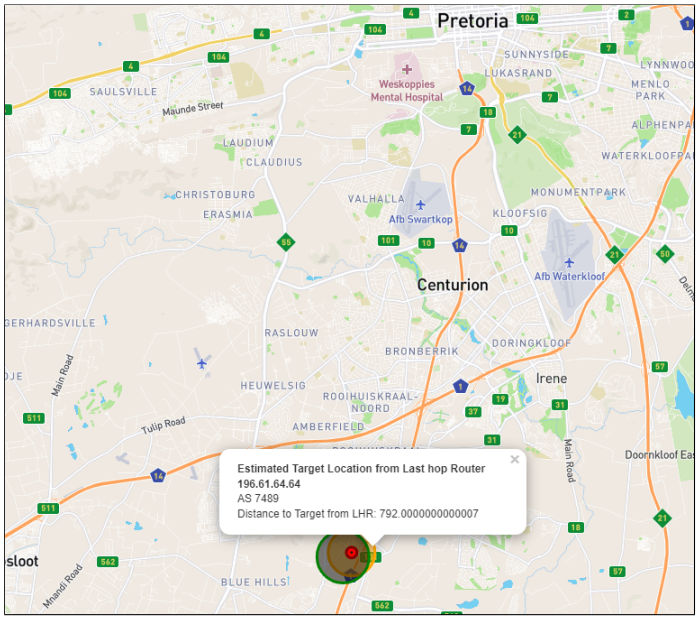


Figure 3.2: Example Element Tagging using PeeringDB and OpenStreetMap

## 3.7 IXPDB website

In Chapter 5 we make use of the Internet Exchange Point Database (IXPDB) which is a publicly accessible, dynamic repository that provides comprehensive data on Internet Exchange Points (IXPs) worldwide. It aggregates and updates information directly from IXP management systems, ensuring an accurate and current representation of the global interconnection landscape.

The IXPDB encompasses detailed insights into IXP architecture, membership, hardware configurations, capacity, and utilisation. This extensive dataset supports diverse use cases, including commercial analysis, network automation, and historical

benchmarking. The data can be explored through an intuitive website interface or accessed in detail via a free API, facilitating seamless integration into research and practical applications.

## 3.8 RIPE Atlas Platform

RIPE Atlas has been mentioned already, but in this chapter we discuss this platform in more detail as its use in this thesis has been of particular importance.

In 2010 the Réseaux IP Européen's Network Coordination Centre (RIPE NCC) began developing the Atlas platform, which collects information on Internet connectivity and reachability through thousands of measurement devices around the world to gain a better understanding of the state of the IP layer of the Internet in real time (RIPE, 2015).

IP-based geolocation is the mapping of an IP address to the geographic location of a real-world connected device to the Internet. IP geolocation can involve mapping the device's IP address to latitude and longitude, country, region (city), or facility (premises) (Gueye, Ziviani, Crovella, and Fdida, 2006). Capturing an accurate view of the Internet topology can allow researchers and network engineers to design and evaluate new protocols and services or to analyse the vulnerability of network infrastructures (Motamedi, Rejaie, and Willinger, 2015).

A geolocation landmark, also known as Vantage Point (VP), is a network node with a known and verifiable geographic location, such as a RIPE Atlas probe, a server in a well-known data center, or an IXP facility. These landmarks serve as reference points for geolocation algorithms. One of the first IP Geolocation approaches developed was the Shortest Ping method which chooses the landmark with the shortest Round Trip Time (RTT) to the target IP address as being the approximate geographical location of the target IP (V. Padmanabhan and Subramanian, 2001). Since then many researchers have attempted to improve the accuracy of IP geolocation (Katz-Bassett, John, Krishnamurthy, Wetherall,

Anderson, and Chawathe, 2006) (Gueye, Ziviani, Crovella, and Fdida, 2006) (Scheitle, Gasser, Sattler, and Carle, 2017) (Wong, Stoyanov, and Sirer, 2007).

RIPE Atlas developers designed their solution for geolocating the core Internet infrastructure, IPmap's Single-Radius Engine, to incorporate the Shortest Ping method (Du, Candela, Huffaker, Snoeren, and Claffy, 2020). This requires a number of landmarks that are geographically close to the target IP address (Du, Candela, Huffaker, Snoeren, and Claffy, 2020),(Katz-Bassett, John, Krishnamurthy, Wetherall, Anderson, and Chawathe, 2006),(Gueye, Ziviani, Crovella, and Fdida, 2006),(V. Padmanabhan and Subramanian, 2001). Therefore, the more landmarks, or 'probes' in IPmap's case, that are deployed in a specific region, the greater the chance that one is geographically close to the target of interest and therefore the greater the accuracy. RIPE's IPmap uses more than 11,000 probes worldwide and currently provides the largest number of landmarks for research, test, and troubleshooting purposes (RIPE, 2015).

In 2015 Giotsas et al. (Giotsas, Smaragdakis, Huffaker, Luckie, and Claffy, 2015) developed a method they named Constrained Facility Search (CFS) which infers the physical interconnection facility where an interconnection occurs. This is a hybrid method of geolocation relying on the passive collection of records from user maintained databases and Network Operating Centre websites, as well as active traceroute measurements carried out by the RIPE Atlas platform, LG servers, IPLANE project and CAIDAs ARK platform. They also supplement this data with BGP queries from Looking Glass servers which gave them information about the peering router such as ASN and IP information. The authors claim that the accuracy of this method outperforms heuristics based on naming schemes and other IP geolocation methods.

## 3.9 Internet Topology Data Kit

A primary concern of IP geolocation is the lack of ground-truth data against which new methods can be accurately tested. However, there are some databases that claim to hold up-to-date and accurate information.

The Centre for Applied Internet Data Analysis (CAIDA) maintains a large database of calculated router IP geolocations. This database is a part of CAIDA's Internet Topology Data Kit (ITDK). The ITDK contains data about connectivity and routing gathered from a large cross-section of the global Internet and is useful for studying the topology of the Internet at the router-level, amongst other uses. However, no dataset is without its limitations. Despite its reputation, even CAIDA data is not immune to inaccuracies or gaps. For example, in this research, discrepancies were observed when comparing CAIDA geolocation results to those obtained using alternative methodologies. These inconsistencies highlight the inherent challenges of relying on a single dataset as a definitive ground truth. The observed fallibility of CAIDA data underscores the importance of treating it as a baseline rather than an absolute authority. While it provides a valuable starting point, its reliability can be enhanced when combined with complementary datasets, such as RIPE Atlas measurements, or through the application of cross-validation techniques. Moreover, the iterative refinement of methods that leverage CAIDA data can help identify and correct potential errors, ultimately improving its utility for research.

ITDK is generated using various methods, combining sources from various websites and databases across the Internet (CAIDA, 2024). The source dataset used in testing in this chapter is based upon the 2021-03 dataset which combines 3 router geolocation datasets from MaxMind, Holistic Orthography of Internet Hostname Observations (HOIHO) and an internally generated IXP database that has been collated from sources such as BGP Looking Glass servers, PeeringDB website, and Packet Clearing House (PCH).

Methods developed later in this thesis use geolocation data available directly from

two of the largest Internet Exchanges in the UK, London Internet Exchange (LINX) and London Access Point (LONAP), to assess the accuracy of the CAIDA dataset. LINX and LONAP, as operational entities, are uniquely positioned to provide precise geolocation data for their routers, given their direct responsibility for configuring and maintaining these devices. The data reflects the physical placement of their infrastructure, which ensures a high level of reliability.

However, while their operational records are highly credible, it is essential to recognize potential limitations. Errors in documentation, infrastructure updates, or the limited scope of their operational domain could introduce inaccuracies. To address these concerns, this study employs cross-validation techniques, comparing LINX and LONAP data with independent geolocation methodologies where possible. By leveraging the trusted data from these IXPs while maintaining a critical perspective, this approach ensures a robust evaluation of the CAIDA dataset's accuracy

The precision of the CAIDA dataset was tested using this LINX and LONAP data, whereby, 83 million UK-based IP addresses from the CAIDA dataset were downloaded along with 1.7 million interface/node names and their geolocations. Each of these IP addresses was linked to the location of that node name as indicated by the CAIDA data set. The geolocation of the CAIDA dataset was then compared with that of the LINX/LONAP dataset, where an IP address appeared in each of the datasets, as shown in Figure 3.3.

It should be noted that whilst the LINX/LONAP dataset is accurate to facility location level, the CAIDA dataset is accurate only to city level. So, whilst CAIDA may show an IP as being in London, the LINX/LONAP dataset can give the actual name and address of the facility where that IP address is hosted.

The LINX/LONAP combined dataset contains 1838 IP addresses with a known geolocation; of these, 1237 were also contained within the CAIDA dataset and therefore 601 IP addresses were unknown to CAIDA. This shows that the CAIDA dataset does not have access to the fullest possible set of IXP data and combining

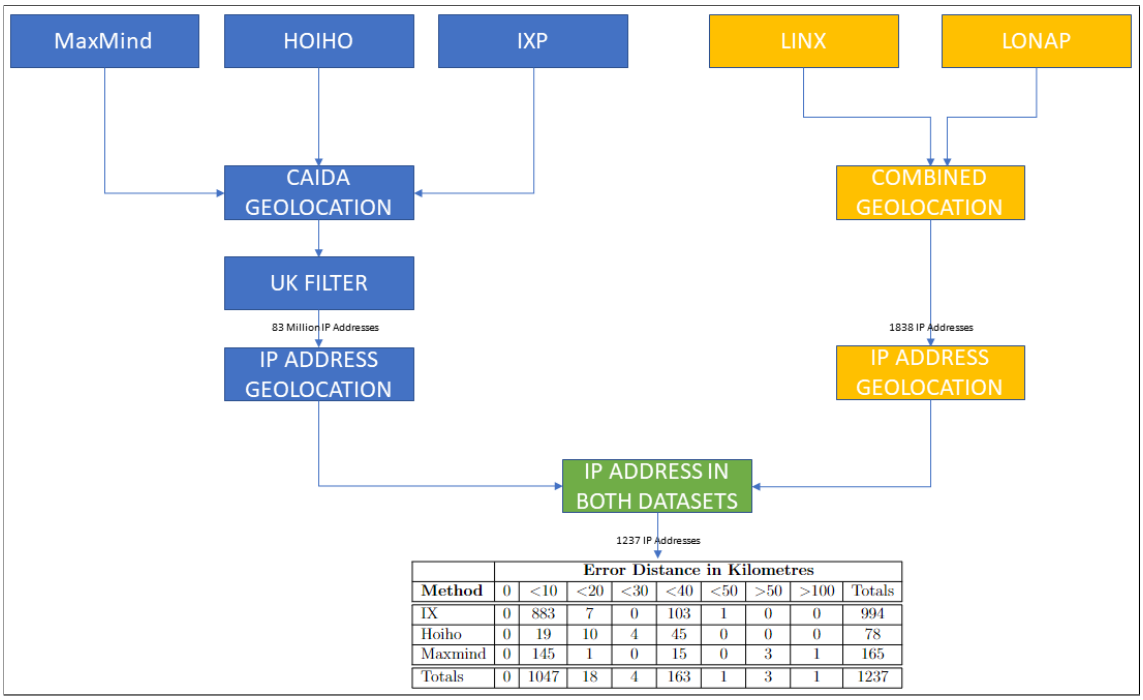


| | Error Distance in Kilometres | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **Method** | 0 | <10 | <20 | <30 | <40 | <50 | >50 | >100 | Totals |
| IX | 0 | 883 | 7 | 0 | 103 | 1 | 0 | 0 | 994 |
| Hoiho | 0 | 19 | 10 | 4 | 45 | 0 | 0 | 0 | 78 |
| Maxmind | 0 | 145 | 1 | 0 | 15 | 0 | 3 | 1 | 165 |
| Totals | 0 | 1047 | 18 | 4 | 163 | 1 | 3 | 1 | 1237 |

Figure 3.3: Comparing CAIDA Geolocation with Ground Truth Data

these website sources, such as LINX, LONAP, EQUINIX etc., within the CAIDA dataset would immediately improve the number of known IP addresses whilst also improving the CAIDA dataset.

| | **Error Distance in Kilometres** | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| **Method** | 0 | <10 | <20 | <30 | <40 | <50 | >50 | >100 | Totals |
| IXP | 0 | 883 | 7 | 0 | 103 | 1 | 0 | 0 | 994 |
| Hoiho | 0 | 19 | 10 | 4 | 45 | 0 | 0 | 0 | 78 |
| Maxmind | 0 | 145 | 1 | 0 | 15 | 0 | 3 | 1 | 165 |
| Totals | 0 | 1047 | 18 | 4 | 163 | 1 | 3 | 1 | 1237 |

Table 3.1: Comparing CAIDAs three Source Datasets for Geolocation Accuracy

As can be seen in Table 3.1, all 3 methods were only able to geolocate to the resolution of the city (less than 10 km), hence why none of them managed to locate the IP addresses accurately. This is the closest accuracy that the the various methods used by CAIDA were designed for. The methods are only valid down to city level where each method selects a centre point for the city of London, which

varies depending on method but all are within 10 km of the actual location. The less-than 10 kilometres error distance column is therefore the most populated and is mainly due to the number of accurately calculated city locations. For example, each of the 3 methods locates the centre of London at slightly different coordinates, CAIDA's in-house IXP method locates any London based IP addresses as being at Charing Cross in London, whilst, HOIHO locates London based IP addresses at Heathrow Airport in London and Maxmind locate London based IP addresses at Westminster in London. Although these are all close to the actual facilities, they do represent errors in distance that the LINX/LONAP combined dataset is able to geolocate more accurately. It can also be seen that in the 'less than 40 Km' column the errors are high. This is because all three CAIDA methods were unable to distinguish between the city of Slough and the city of London. All three methods geolocate any IP addresses located in Slough to the individually estimated centres in London; hence there is an error distance of approximately 35 kilometres which is the distance from Slough to London. It is puzzling why CAIDA's in-house IXP method should do this, as its use of data from sources such as PeeringDB and Packet Clearing House provides more accurate information, and this method would be expected to be able to distinguish between Slough and London. One final note is that whilst the Maxmind database maps these IP addresses to the Westminster geocoordinates, it fails to indicate any city at all. Therefore, it should be noted that a search for 'London' within the MaxMind dataset would fail to find these IP addresses.

## 3.10 Constraint Based Geolocation

In order to test Constraint Based Geolocation (CBG) on the UK infrastructure, the RIPE Atlas platform has been employed to create 900+ (32x31) traceroutes between each of 32 RIPE anchors, whose IP address and geolocation are known. As there is fairly good RIPE coverage in parts of the UK it was expected that CBG would perform well; however, this was not the case. For example, the RIPE Atlas probe

6562 shown in figure 3.4 (yellow circle) is located in Whitechapel, London, and it has many RIPE Atlas probes (red circles), which can be used as Vantage Points (VP) close by, so one would assume this would provide a good opportunity for CBG's triangulation algorithm.

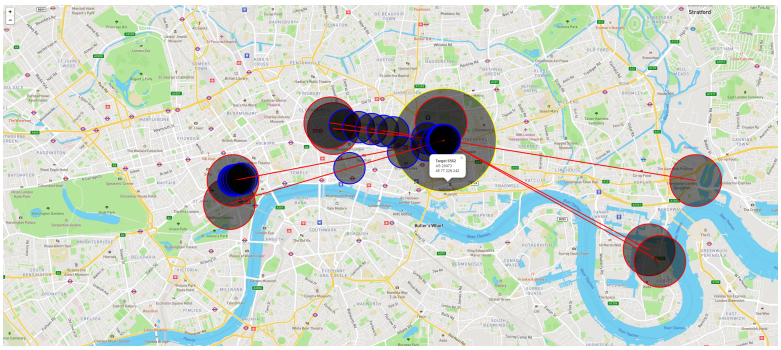

Figure 3.4: RIPE Atlas Probe Distribution close to target

As can be seen in figure 3.5, CBG's method is to draw a greater circle around each Vantage Point using the formulae derived at (2.1)

It is then expected that these greater circles will intersect, forming a smaller target area which would then create a more accurate geolocation result. However, this was not the case; not only was triangulation completely ineffective since the larger circles do not even intersect, but the estimated possible location area was 125 km wide (smallest green circle) as seen in Figure 3.5, which is much worse than expected. In contrast, RIPE's IPMAP single-radius engine locates the target to the nearest city of Bethnal Green, which is approximately 3 km radius; however, this may be somewhat biased in favour of IPMAP as the target is one of RIPE'S own known anchors, which provides RIPE with its actual location. The cause of this poor CBG result is the latency across the Internet and is depicted in Figure 3.4 where the red lines represent packet speeds of less than 0.2 x speed of light (0.66 x speed of light is expected). According to RIPE the anchors themselves should be connected to high-speed links and therefore issues such as last-mile latency should not prove a factor, so one has to wonder where these delays are originating from.

There are various types of delay that can affect the RTT values. Transmission

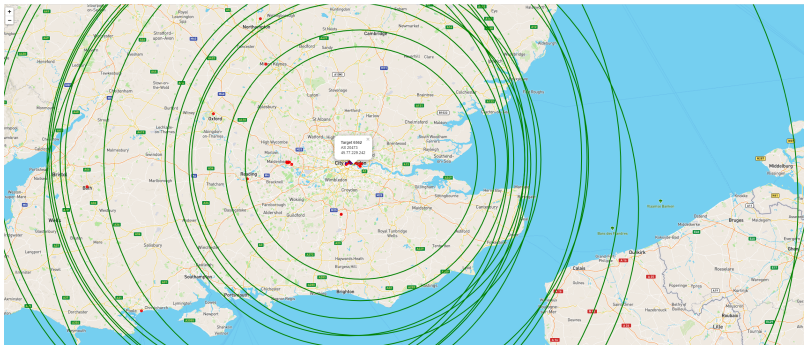

Figure 3.5: Using Constraint Based Geolocation to Geolocate a Target RIPE Probe

delays are caused by the originating device, which in this case is the RIPE Vantage Point anchor; could these be underspecified? Processing delays are caused by the ISP or various networks en route which need to process the packet header to decide on its route, propagation delays are caused by the type of medium a packet travels through, and, from a highly connected device to another highly connected device as is the case with RIPE anchors, one would expect the best possible medium. Finally, queueing delays are caused by the amount of time a packet waits to be processed on the target device, which again is a RIPE anchor, and suspicion would again have to fall on the specification of the device.

Table 3.2 below shows the outcome of using CBG to geo-locate each of the rest of the anchors with similar unsatisfactory results ranging from target location accuracy of between 65 km and 4000 km. The constrained column details whether CBG was able to reduce the location area with overlapping greater circles. In all cases where this happened, the constrained area only reduced the overall area by a negligible amount.

| Probe | IP Add | Constrained | Location Accuracy in km |
|---|---|---|---|
| 6087 | 5.57.16.65 | no | 125 |
| 6182 | 141.170.19.12 | yes | 450 |
| 6214 | 178.237.173.220 | no | 150 |
| 6382 | 90.223.193.3 | no | 80 |
| 6405 | 37.10.44.14 | yes | 275 |
| 6423 | 176.74.17.75 | yes | 200 |
| 6446 | 185.40.232.202 | yes | 165 |
| 6451 | 107.162.220.5 | no | 75 |
| 6471 | 52.56.61.239 | yes | 175 |
| 6499 | 153.92.43.251 | no | 400 |
| 6501 | 37.143.141.141 | no | 750 |
| 6512 | 80.82.241.134 | yes | 650 |
| 6515 | 153.92.43.249 | yes | 450 |
| 6516 | 153.92.43.250 | yes | 450 |
| 6519 | 46.227.202.97 | no | 400 |
| 6532 | 109.232.177.220 | no | 2750 |
| 6552 | 185.57.191.228 | no | 825 |
| 6559 | 185.97.160.7 | no | 65 |
| 6562 | 45.77.229.242 | no | 125 |
| 6609 | 185.184.236.30 | no | 650 |
| 6647 | 5.62.127.14 | no | 650 |
| 6670 | 92.223.59.77 | no | 125 |
| 6674 | 194.50.88.164 | no | 4000 |
| 6695 | 194.81.236.229 | yes | 265 |
| 6699 | 35.234.152.175 | yes | 265 |
| 6716 | 82.148.224.6 | no | 215 |

| Probe | IP Add | Constrained | Location Accuracy in km |
|---|---|---|---|
| 6738 | 185.232.117.201 | yes | 165 |
| 6843 | 86.188.235.234 | no | 2500 |
| 6879 | 90.223.193.1 | no | 115 |
| 6892 | 164.39.242.17 | no | 165 |
| 6945 | 156.154.80.254 | no | 65 |
| 6948 | 193.57.144.24 | no | 140 |
| 6971 | 46.101.90.215 | no | 65 |
| 7021 | 185.194.168.88 | no | 515 |

Table 3.2: CBG Accuracy using RIPE Probes

In order to discover exactly where this latency is occurring, a request for the installation of an additional RIPE anchor was made and accepted by RIPE. Lancaster University's Information Systems Services (ISS) agreed to host this RIPE anchor after looking into the relevant technical and security issues. More installations at other Universities such as Edinburgh and Bristol would allow for an end-to-end latency test where every aspect of a traceroute between two anchors can be thoroughly investigated. This would also have the bonus of increasing RIPE coverage across the UK and provide Universities with RIPE credits, which can be used by future researchers to create measurements on the RIPE platform. However, obtaining relevant permissions to install equipment at each university, time and cost limitations plus RIPE had a long waiting list to supply their anchors at the time meant this was not feasible in the required time frame.

## 3.11 RIPE's Single-Radius Method

The RIPE single-radius method uses 4 steps to locate target IP addresses (Du, Candela, Huffaker, Snoeren, and Claffy, 2020).

1. Map the target IP to the AS, announcing its containing prefix using RIPE Routing Information Service (RIS) BGP data. RIS is a collection of BGP routing data collected from multiple vantage points across the Internet. It is widely used for studying routing behavior, detecting anomalies, and analyzing Internet topology. Then it finds a set of RIPE Atlas probes topologically close to the target IP. Schedule a ping measurement from the selected probes. Return an estimated measurement duration to the user.

2. Collect all resulting RTTs and discard those above 10ms as timings above this could be caused by various reasons such as congestion, routing inefficiencies, high latency links (e.g., satelite links), packet loss, and retransmission due to poor network quality, administrative throttling, device load, Quality of Service policies, security configurations, outdated or faulty equipment, traffic shaping, response times and finally physical distance. Convert remaining RTTs to one-way latencies, RTT /2. This filtering assumes that geolocation using distant probes (e.g. on another continent) is not effective due to physical distance.

3. Select probe $p$ with minimum latency and convert it to distance $d$ using a distance delay coefficient of $2/3c$.

4. Use the location of $p$ as the centre of circle $C$ with radius $d$. Select 100 closest cities to $p$ using the RIPE Worlds database, based on the shortest straigh line distance between $p$ and the city. Select only cities inside circle $C$, hence lower latencies yield fewer cities. Rank cities by population size and return the highest-ranked one to the user.

Taking Atlas's obfuscation policy into account, Atlas Probe 16430 is located within a 1 km radius of the city of Johannesburg in South Africa. It has an IP address

of 154.126.223.204. Using this probe as a sample target and using the formulae derived at (2.1) yields a result that returns Johannesburg as the nearest city and an RTT value of 7.095553 ms. this would equate to a greater circle around Johannesburg with an error radius of 702 km.

This greater circle would encompass most of South Africa and obviously cannot be used for accurate geolocation due to this huge target radius. However, using the RTT value to calculate the greater circle error radius for the single-radius engine may be unfair. Previous researchers suggest that the average perimeter of a metropolitan area is approximately 40 km (Du, Candela, Huffaker, Snoeren, and Claffy, 2020), (Gharaibeh, Shah, Huffaker, H. Zhang, Ensafi, and C., 2017) or 50 km (Giotsas, Nomikos, Kotronis, Sermpezis, Gigis, Manassakis, Dietzel, Konstantaras, and Dimitropoulos, 2021). This places the target IP address somewhere in the Johannesburg metropolitan area.

## 3.12 Shortest Ping Method

The Shortest Ping method could also use the RIPE Atlas platform. A researcher would manually group a selection of probes that are within the estimated target area and carry out the first 3 steps of RIPE's Single-Radius method. Once the results are returned, The probe with the smallest RTT value can be selected, placing the probe at the center of a larger circle with a radius equal to RTT/2, as illustrated in Figure 3.6. Each ping from a probe defines a circle, with the target located within the circle's boundary.

Using the same sample target of 154.126.223.204 and the other 66 probes as sources to ping the target, this second method results in the probe 1000492 having the smallest round-trip time. Probe 1000492 is located in Elandsfontein in South Africa. It has a one-way trip time of 0.236335 ms, which is equivalent to a target radius of 47 km, as shown in Figure 3.6. The red circle depicts this probe's calculated target area, and the green circles show the discounted probe target area, which did

not have the shortest ping.

Multilateration is a technique used in geolocation to determine an unknown position based on distance measurements from multiple known reference points. In IP geolocation, multilateration is used by combining latency-based estimates from multiple vantage points to pinpoint a target's location. If any of the other probes greater circles had intersected the red greater circle, we could add the Constraint Based Geolocation (CBG) method using multilateration (Gueye, Ziviani, Crovella, and Fdida, 2006) to further reduce the target radius. Unfortunately, in this case, no other probes intersect, thus multilateration cannot be used because the probe with the shortest ping is in the hop paths of all other chosen probes. This appears to be a very common scenario where the last few hops of each traceroute use the same path, thus the CBG method cannot be used.

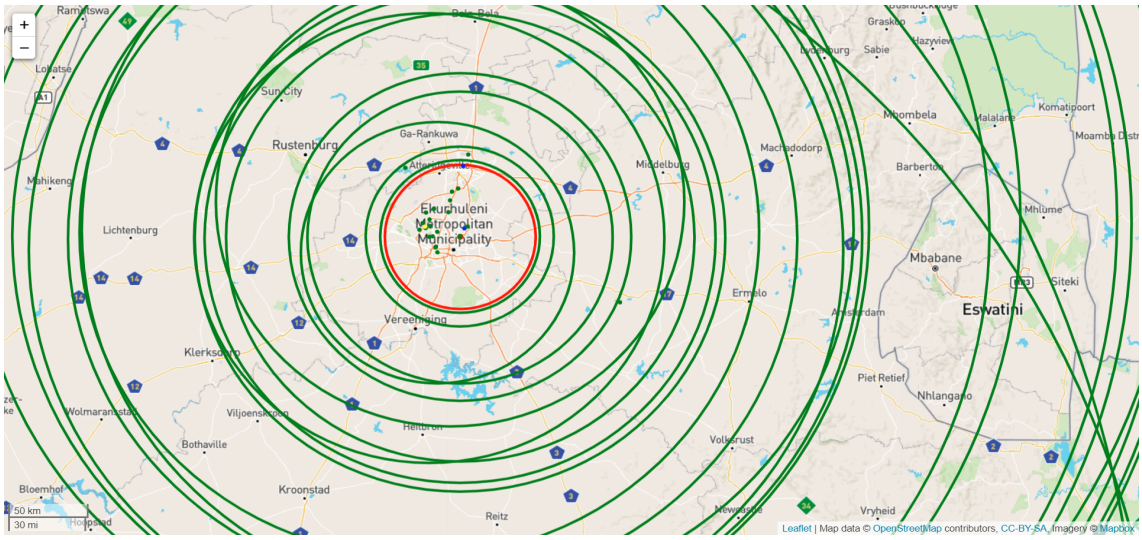


Figure 3.6: Geo-locating IP address 154.126.223.204 Using Shortest Ping

This method offers similar accuracy to that of the Single-Radius engine, 47 km target radius vs 40 km target error radius, over the single-radius method. It should be noted that if CBG were an option, that is, if there is more than one route to the target in any measurement, then the accuracy of the shortest ping method would improve.

## 3.13 Summary

RIPE Atlas developers designed their solution for geolocating core Internet infrastructure, viz. IPMAPs Single-Radius Engine, to incorporate the Shortest Ping method (Du, Candela, Huffaker, Snoeren, and Claffy, 2020). This requires a number of landmarks that are geographically close to the target IP address (Katz-Bassett, John, Krishnamurthy, Wetherall, Anderson, and Chawathe, 2006), (Gueye, Ziviani, Crovella, and Fdida, 2006), (V. Padmanabhan and Subramanian, 2001). Therefore, the more landmarks, or probes in IPMAP's case, that are deployed in a specific region, the greater the chance that one is geographically close to the target of interest and therefore the greater the accuracy. RIPE's IPMAP uses more than 11,000 probes worldwide and currently provides the largest number of landmarks for research, test, and troubleshooting purposes (RIPE, 2015).

In 2020 Du et al. evaluated the accuracy, coverage, and consistency of RIPE's IPMAP Single-Radius engine (Du, Candela, Huffaker, Snoeren, and Claffy, 2020). They recognised that their findings had a geographical bias of ground truth due to the fact that most IP addresses from their ground-truth dataset were located in western europe and the contiguous US. They determined that Single-Radius accuracy appears to differ by region and realised that their results may not hold in all regions.

Ding et al. agree with Du et al. and furthermore observe that many classical delay-based IP geolocation algorithms (Katz-Bassett, John, Krishnamurthy, Wetherall, Anderson, and Chawathe, 2006) (Gueye, Ziviani, Crovella, and Fdida, 2006) (Scheitle, Gasser, Sattler, and Carle, 2017) (Wong, Stoyanov, and Sirer, 2007) are well-suited for richly connected networks, such as those in the United States and Western Europe (Ding, X. Luo, Dengpan, and Liu, 2017).

However, Ding et al. conclude that delay-distance correlation algorithms are seriously affected where regions are poorly connected, but also believe that the delay distance correlation of some subnetworks within those regions may be better than the overall regional delay distance correlation (Ding, X. Luo, Dengpan, and

Liu, 2017).

Du et al. suggest a potential improvement to the Single-Radius method such as using a multilateration engine which would use results from multiple probes to improve results in regions where Atlas node deployment is sparse. Shichang et al. reason that to obtain more accurate results, a multilateration method such as constraint-based geolocation (CBG) should select probing hosts and landmarks that are in the same richly connected subnetwork with the target host (Ding, X. Luo, Dengpan, and Liu, 2017).

For Constraint-Based Geolocation to be useful and multilateration/triangulation to be a success, multiple IP-to-Geographical landmarks must be available; however, if RIPE node deployment is sparse as suggested by Du et al. then CBG is unlikely to be successful due to a lack of landmarks. The problem is how to increase the number of known landmarks that are close enough to target addresses and thereby improve the possibility of utilsing the CBG method. The method proposed in the next chapter uses Internet Exchange Points (IXP) and interconnect facilities to add additional landmarks that may be closer to target IP addresses to increase IP geolocation accuracy. The main influencing factors for accuracy when using these facilities is the ability to detect which of the IXP facilities involved in the IXP are being used in the route to the target and how far the IXP is from the target. Furthermore, if two or more IXP facilities are involved in the target measurements, then the use of CBG may further increase the accuracy.

# Chapter 4

# Internet Exchanges as Additional Landmarks

## 4.1 Introduction

The additional use of Internet Exchange Points (IXP) as landmarks, also known as Vantage Points, as described in Chapter 3, is an ideal solution as their geographical locations and IP address ranges are publicly known. However, there is a problem with the use of IXPs as landmarks, although much research has been carried out, the task of geolocating the physical interfaces of IXPs is still a challenging problem (Motamedi, Yeganeh, Chandrasekaran, Rejaie, Maggs, and Willinger, 2019) for the following reasons:

- Layer 2 Nature of IXPs: IXPs primarily operate at Layer 2 of the OSI model. As a result, traffic routing and switching happen without detailed Layer 3 information like IP headers, making it difficult to directly associate physical interfaces with geolocation data.

- Sparse Publicly Available Data: While some IXPs provide information about their physical locations and member connections, much of the detailed geolocation data at the interface level is proprietary or unavailable for public

analysis. This lack of transparency complicates efforts to geolocate interfaces accurately.

- Dynamic Nature of Infrastructure: IXP infrastructures often evolve with changing membership, upgrades, and new interconnection agreements. The dynamic nature of these configurations means that static geolocation data may quickly become outdated.

- Physical and Logical Separation: In some cases, the physical interfaces of an IXP may not correspond directly to its logical topology. For example, an IXP may have facilities in multiple cities, and data traversing the IXP could appear to originate or terminate in different physical locations than expected.

- Existing Tools Have Limitations: Motamedi et al. highlight that tools like traceroute struggle to provide granular insights into Layer 2 networks. Layer 2, also known as the Data Link Layer in the Open Systems Interconnection (OSI) model, is responsible for switching and forwarding data frames within a local network segment (Braden, 1989). Unlike Layer 3 (Network Layer), which uses IP addresses for routing, Layer 2 relies on MAC addresses, making it difficult to track geographical movement of data packets within IXPs. This limitation underscores the broader challenge of associating physical interfaces with their actual geolocation.

- Need for Ground-Truth Data: Accurate geolocation of IXP interfaces often requires validation against ground-truth data provided by the IXPs themselves. However, reliance on such data introduces biases and inconsistencies, as it may not always be verified independently.

**Roadmap** Section 4.2 introduces the concept of using Internet Exchange Points (IXPs) as landmarks to improve geolocation accuracy. The section discusses the challenges of geolocating IXPs and the benefits they provide, such as reduced latency and improved geolocation precision. Section 4.3 highlights the difficulties

in geolocating IP addresses through multi-facility IXPs due to their dispersed geographical nature. It also discusses methods to infer packet paths using Border Gateway Protocol (BGP) community values and ASN facility data. Section 4.4 presents a Python-based method to identify the entry and exit facilities of packets in multi-facility IXPs. The method utilizes traceroutes, PeeringDB, and RIPE whois to improve geolocation accuracy by reducing RTT and error radius. Section 4.5 acknowledges the limitations of the method, such as challenges posed by remote peering and cases where multiple facilities might share common ASNs, leading to larger error radii. Section 4.6 demonstrates the effectiveness of the proposed method through a case study using South African Internet infrastructure. It also validates accuracy by comparing estimated locations with RIPE's obfuscated probe locations. Section 4.7 details the application of the method to South Africa's simpler infrastructure, providing comprehensive traceroute measurements and establishing the accuracy of the geolocation approach in both single and multi-facility scenarios. Section 4.8 extends the testing to the UK's infrastructure, using PeeringDB, OpenStreetMap, and RIPE Atlas to map IXPs and their associated facilities. the section also demonstrates the method's scalability and adaptability to complex networks. Section 4.9 summarizes the innovative use of IXPs as geolocation landmarks, highlighting the method's ability to enhance geolocation accuracy and reduce errors. It also discusses how this approach contributes to a more precise mapping of network topologies and acknowledges that newer IXP gateway databases make parts of the method redundant.

## 4.2 Internet Exchange Points

In order to use Internet Exchange Points (IXP) as landmarks, a list of a country's Internet Exchanges, facilities, and their geographical locations is required. A list of ASs and the facilities that host each AS is also required, all of which for the majority can be accessed at PeeringDB. PeeringDB is a freely available user-maintained

database of networks, facilities and Internet Exchanges that provides comprehensive details on address, geolocation, and other useful information (PeeringDB, 2024). PeeringDB identifies IXPs, facilities, and networks by its own identification numbers, and these identification numbers are used in the following chapters.

IXPs are located near population centres and are connected to retain the maximum amount of local traffic within their own country (Scheitle, Gasser, Sattler, and Carle, 2017). IXPs are also allocated their own IP address ranges, so any traceroutes that pass through these points can be geolocated to the IXP facilities' published geographical addresses. Round Trip Times (RTT) can then be calculated from the IXP to the target instead of the source to the target, reducing packet trip distance, thus reducing errors caused by latency, link congestion, circuitous routes, infrastructure diversity, and buffering. A further benefit of using IXPs as landmarks is that RIPE probes as well as RIPE anchors can be used as sources to further increase the number of Vantage Points; this is because the IXPs are closer to the target than the source, the geolocation coordinates are known, and the RTT will be much reduced. This technique could be used in poorly connected or richly connected networks such as the UK to further improve the accuracy of IP geolocation, and this will help with risk assessment and mitigation.

## 4.3 Multi Facility IXPs

As mentioned above, Motamedi et al. pointed out a problem with the use of IXPs as landmarks which is their use of multiple facilities that have different geographical locations that are connected at layer 2. This makes it much more difficult to geolocate the location of the required IP address. A packet may route through a specific IXP network with an IP address assigned to that IXP, but in many cases that IP address can be in one or more different physical locations. Mi2 is a method developed to infer interconnections within a single colocation facility by analyzing traceroutes and BGP data (Motamedi, Rejaie, and Willinger, 2015). While it effectively maps AS-level

connectivity within a facility, it does not solve the problem of mapping IXPs that span multiple geographically dispersed locations. Since many IXPs operate across multiple data centers, Mi2 lacks the ability to distinguish which specific facility a packet traverses. Additionally, Mi2 primarily relies on Layer 3 (IP-level) data, making it ineffective for tracking traffic within Layer 2 switching domains that are common in IXPs. This thesis builds on the strengths of Mi2 by incorporating ASN-to-facility mapping and resolving multi-facility ambiguities to improve geolocation accuracy.

Researchers also use Border Gateway Protocol Community (BGP) values to infer many hidden peer-to-peer links across IXPs (Giotsas, Smaragdakis, Huffaker, Luckie, and Claffy, 2015). The Border Gateway Protocol (BGP) is the routing protocol used to exchange reachability information between Autonomous Systems (ASes) on the Internet. It is critical for determining the best paths for data packets but does not inherently consider geographical constraints. Giotsas et al compiled a dictionary of 109 BGP community values which are used to annotate AS ingress points offering a valuable source of validation data. Tools such as Bdrmap (Luckie, Dhamdhere, Huffaker, and Clark, 2016) and MAP-IT (Marder and Smith, 2016) attempt to infer border interfaces; however, none of these methods has been designed to geolocate, or "pin", the interface to a geographical location.

Once the packet reaches an IXP it is very difficult to know which facility it reached and which facility it came from (Motamedi, Rejaie, and Willinger, 2015). In some cases, an IXP has only one facility, so this makes it easy as in the case of Figure 4.1 which shows the estimated target location (green circle) and the actual target location (red circle) to be very accurate (within 720 metres).

However, when multiple facilities are involved, the geolocation tool must incorporate sense checks as a fundamental step to identify the facility through which the packet is entering or exiting. IXPs operate at Layer 2 (Ethernet switching), meaning packets can traverse multiple facilities without generating new IP hops. This makes it impossible to determine the actual path just from traceroutes,

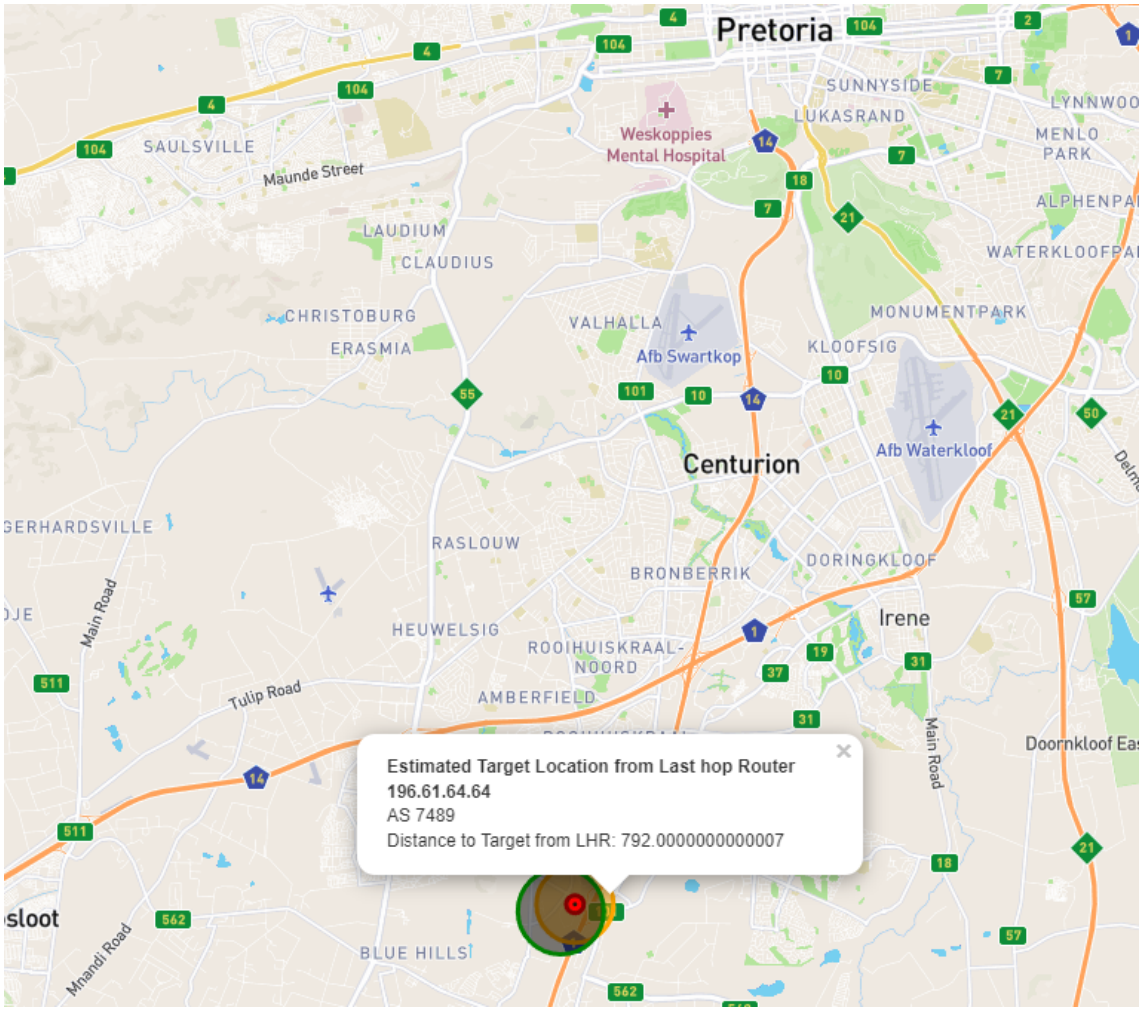


Figure 4.1: Geolocating an IXP with a single facility

requiring additional geographical and routing logic to deduce where a packet enters and exits. Outliers in geolocation occur when a calculated or inferred geolocation result deviates significantly from expected physical or routing behavior. this can be due to an incorrect facility selection due to Multi-Facility IXPs, for example, An ASN may announce routes at an IXP in London while actually forwarding traffic from Manchester, leading to geolocation errors. Another example may be if A packet from Manchester to London may unexpectedly route through an IXP in Leeds, creating a misleading geolocation result. Network congestion or buffering can inflate RTT values, making a target appear further than it is. Also, Many databases

(e.g., MaxMind, CAIDA) rely on static mappings, which do not update dynamically as networks evolve. Packets may also take different paths in each direction due to asymmetric routing and remote peering allows an ASN to appear at an IXP without having a physical presence there.

These checks are crucial for ruling out facilities that the IXP uses that cannot possibly be a part of the traceroute, such as when a facility is located in the wrong country/region/city or when a Speed of Light (SoL) sense check indicates that the facility is too far from the target to provide a logical or feasible distance. RTT measurements play an integral role in this process, as they help eliminate facilities that exceed the maximum allowable speed and distance limitations. At each facility, it is proposed to identify the upstream and downstream ASNs by performing a database lookup of the IP addresses associated with traceroute entries and exits. By cross-referencing this data, the shared facility between the IXP and the ASN can be pinpointed. This approach aims to reveal the precise path a traceroute takes through an IXP, providing greater clarity into its traversal. BGP communities or BGP AS-PATH information may also add more detail or validation (Giotsas, Smaragdakis, Huffaker, Luckie, and Claffy, 2015).

## 4.4 Resolving the issue of Multi Facility IXPs

Knowing the entry and exit routes that a packet takes in an IXP provides additional Vantage Points (VPs), as we can access the geolocation of those facilities from PeeringDB. A Python program was developed to provide a method that could detect which facilities a packet enters and exits a multi facility IXP. This method has the following steps:

- Create a traceroute from source to target using RIPE atlas probes
- Detect which IXP the packet traverses using a list of prefixes used by UK IXPs gathered from PeeringDB.

- Using RIPE whois, identify the ASN that the packet travels to after leaving the IXP.
- From PeeringDB find the common facility for the IXP and the ASN.
- Calculate the RTT value from facility to target as the result of :

$$RTT(target) - RTT(IXPcommonfacility) \tag{4.1}$$

- Calculate the distance from facility to target using the result from equation 4.1, above, and the equation previously provided at equation 2.1.

By identifying the exiting interface, its geolocation can serve as a Vantage Point, enabling a reduction in the effective RTT to the target and subsequently minimizing the overall error radius.

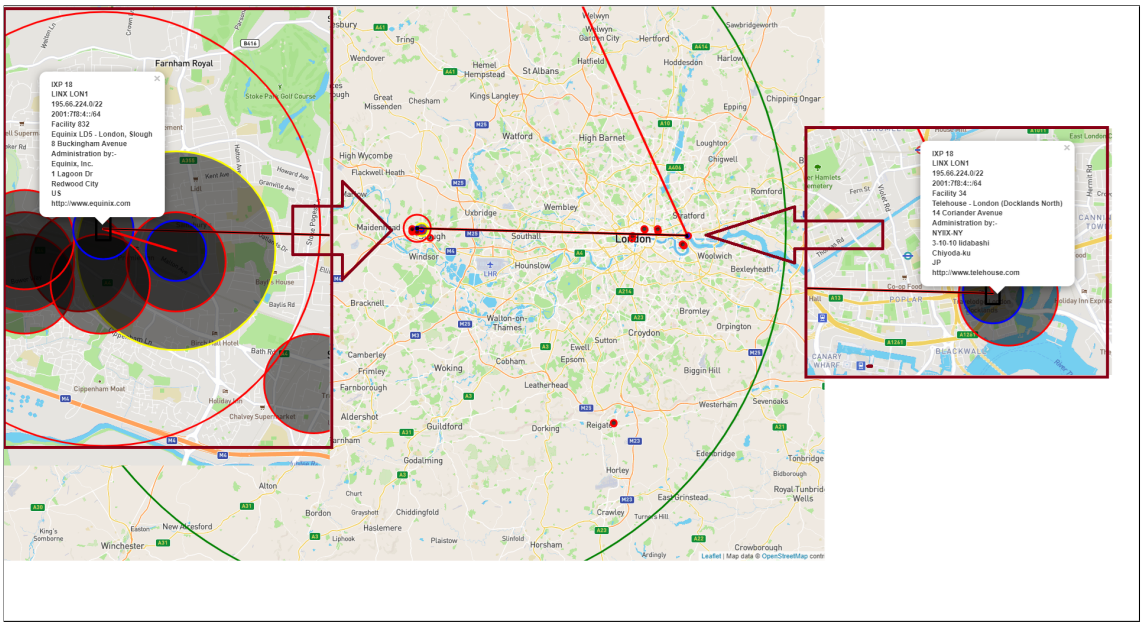


Figure 4.2: The large red circle shows the error distance of this method compared against the error radius of Constraint Based Geolocation (CBG) denoted by the green circle. The smaller red circles are the location of other Atlas probes which can be ignored and the blue circles denote traceroute hops

In Figure 4.2 we have a target RIPE Atlas probe 6695 which we know the approximate location of and is denoted by the yellow circle when taking into account RIPE's 1 km obfuscation policy. This probe has an IP address of 194.81.236.229. In this test, we have carried out a traceroute from Manchester towards the target probe. We can see that it routes via the LINX Lon1 Internet exchange and forwards to Facility 832 in Slough. If we did not know the actual location of the target probe this would be the last known Vantage point and therefore the RTT time divided by 2 between Facility 832 and the target provides the maximum possible radius of a circle that the target probe could be located. This is the large red circle in the diagram. We can see that the actual location denoted by the yellow circle and the estimated location denoted by the red circle are very similar.

This is an improvement in accuracy compared to other active IP geolocation methods, and it became evident that a paper based on this method could be produced, see (McCherry, Giotsas, and Hutchison, 2023). This method can also allow for the discovery of the ingress facility by discovering the ASN of the router before arriving at the IXP facility, and this could be used to help with secondary information on the IXP such as bandwidth allocations.

## 4.5 Limitations

However, it became apparent that there was a problem with this method, facilities using remote peering to connect to IXPs will not have a common facility with that of the IXP that traceroutes pass through, and this presented a problem. Remote Peering is a method where an AS connects to an IXP without having a physical presence at the IXP's facility. Instead, the AS connects via an intermediary provider that extends Layer 2 connectivity over long distances. This can obscure the actual physical location of a peering session, making geolocation more challenging. If remote peering occurs before entering a facility, then there is no problem because we are only interested in the exiting facility as in Figure 4.3.

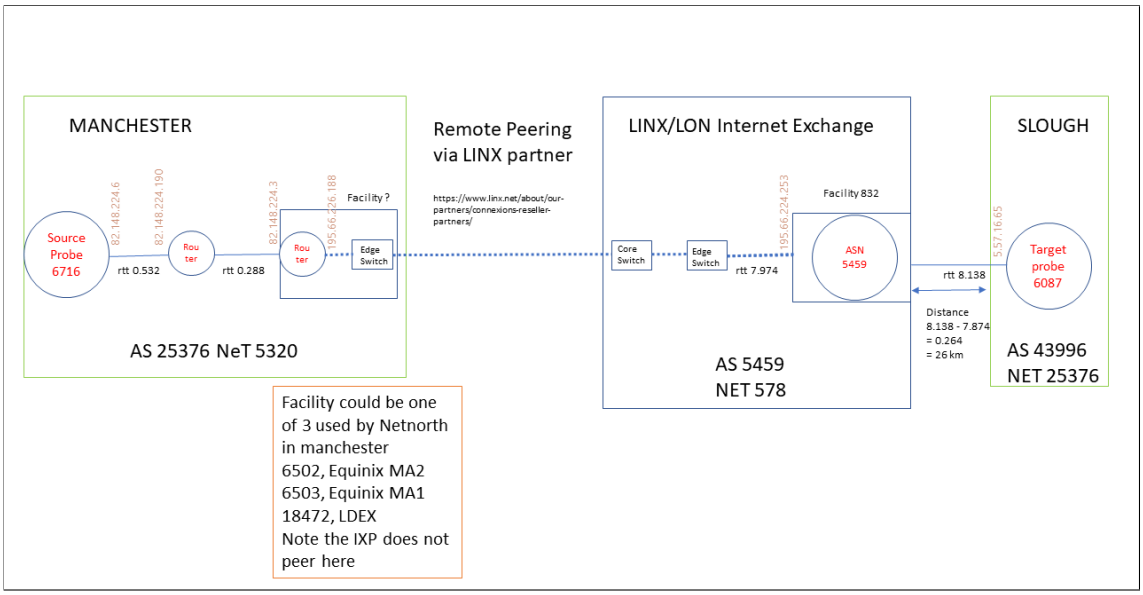


Figure 4.3: Remote Peering before entering an IXP network

If remote peering occurs between the IXP and the target, then we cannot be sure from which of the IXP facilities a traceroute will exit. We can infer the ASN of the next hop after the IXP which will provide us only with the facilities that a packet may traverse: see Figure 4.4. This would impose a minor limitation on this new method.

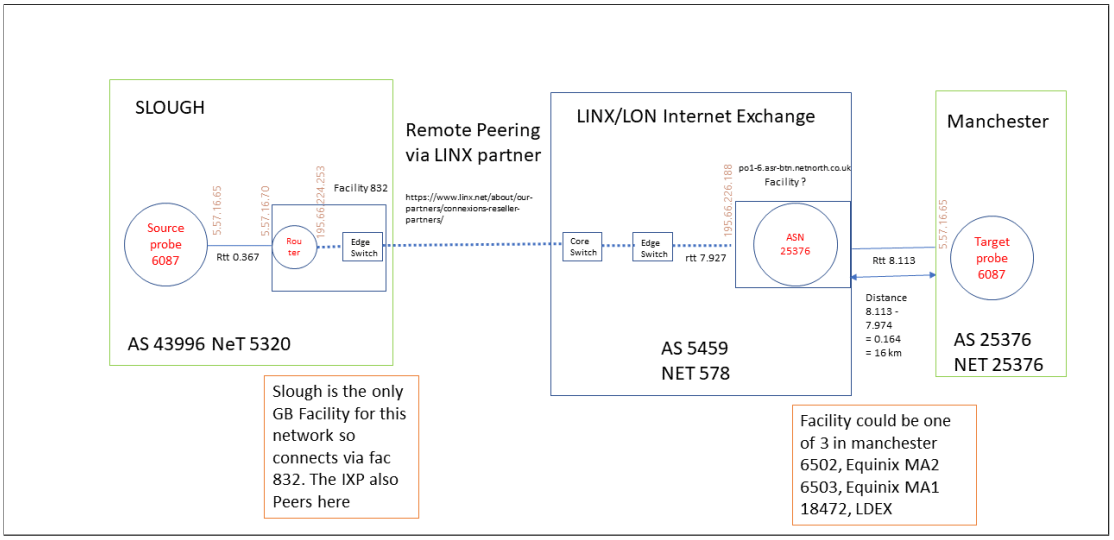


Figure 4.4: Remote Peering after leaving an IXP network

In this example, we are using a target probe whose geolocation is already known, but if this was not the case, the error radius would be the combination of the calculated distance from each of the three possible exiting facilities as shown in Figure 4.5. The error radius is not significantly increased in this diagram but the 3 facilities are not a large distance apart. Some IXP's facilities can be over hundreds of miles which would obviously increase the error radius by a significant amount.

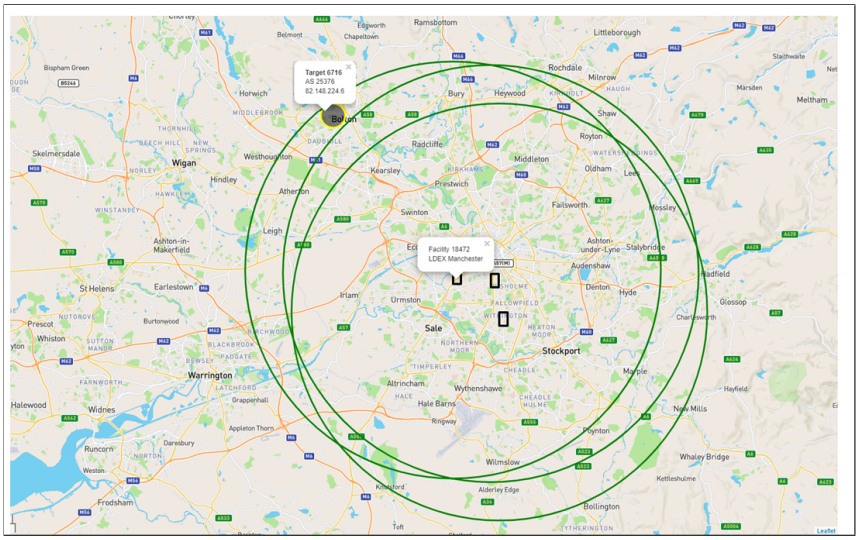


Figure 4.5: Three possible exiting facilities leads to a larger error radius consisting of all three greater circles combined

## 4.6 Proof of Concept

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 at many locations, such as the Packet Clearing House website (PCH) (PCH, 2024), the IXPDB website (IXPDB, 2024), and the PeeringDB website (PeeringDB, 2024). Among these directories, PeeringDB has the most comprehensive list. Simple data extraction can be performed because PeeringDB is a freely available network database that contains a well-updated list of IXPs (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. PeeringDB is incomplete, as it excludes data from certain IXPs, a limitation highlighted in previous research (Klöti, Ager, Kotronis, Nomikos, and Dimitropoulos, 2016). This causes additional failures in the code to recognise the geographical location of the IP addresses registered with those IXPs.

## 4.7 South Africa

In order to test the method, South Africa was chosen because of its simpler Internet infrastructure. In 2008, due to the lack of a copper wire backbone, data transmission across the African continent was difficult. Only three submarine fibre optic cables connected the entire continent to the global Internet, of which two were located in North Africa (Ngari and Petrack, 2024). Since then the addition of more submarine cables has vastly improved global Internet connectivity and, due to advances in fibre-optic technology, Africa has a chance to leapfrog over the older copper infrastructure found in first-world countries to design and build a continent wide modern fibre-optic backbone. Figure 4.6, provides a general visualization of Africa's fibre infrastructure and population density, offering additional context but not forming the primary basis of this analysis. It should be noted that this study focuses on South Africa's

connectivity and does not attempt to make inferences about the broader state of connectivity across the African continent.

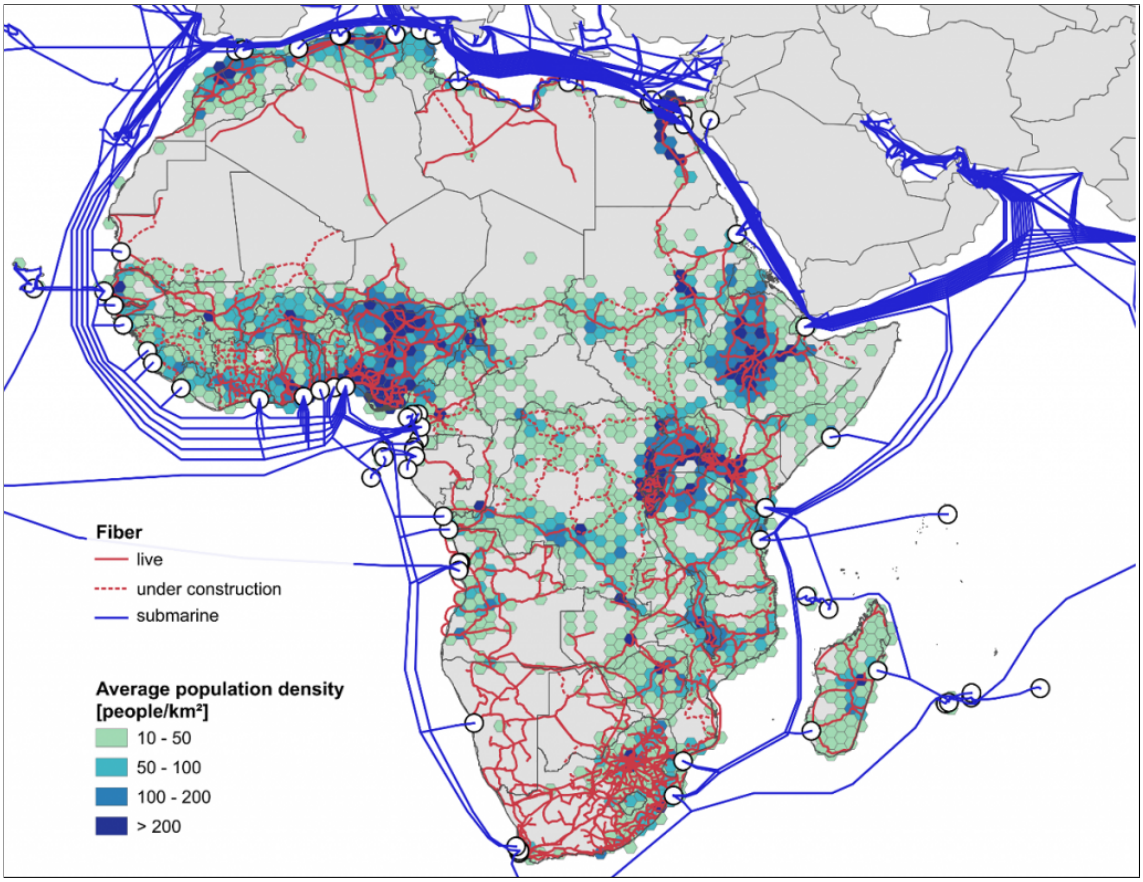


Figure 4.6: Visualization of fibre infrastructure and population density in Africa. Reproduced from (Ngari and Petrack, 2024)

South Africa currently has 67 active RIPE Atlas Nodes and, to establish proof of method, each probe is used as a source and a target to create a matrix of 67 x 66 traceroutes across South Africa with each probe acting as a target for the other 66 source probes. This creates 4422 traceroutes in total with more than 40,000 hops across South Africa.

For each traceroute, whilst we focus on IXP infrastructure, we can ignore all the hops up to the hop containing an IP address that is recognised as belonging to an Internet Exchange Point (IXP). We then need to establish the AS to which this IP address connects and the facilities where this AS is hosted.

### 4.7.1 Method Overview

Step 1 of the proposed new geolocation method involves querying the PeeringDB REST API to obtain a list of a country's 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 facilities Internet Exchanges have deployed their switching equipment to build an OSI layer 2 map of the IXP network infrastructure. When available, 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 and using probes that are in the same country as the target to reduce errors. Traceroutes are created in both directions, to and from each probe, creating a mesh of thousands or even hundreds of thousands of 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 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 Internet Exchange websites, to locate the position of the router where these hops are interfacing, considering the previously discovered facility and IXP information.

The fundamental difference between the single facility example and the multi facility example explained in this chapter lies in the complexity of determining the exact location where a packet exits the IXP network. In the single facility example it is straightforward to infer the geolocation of any IP address traversing the IXP as there is only one geograpically located facility (i.e., Landmark). In the Multi facility

example the exit point is ambiguous and the packet may exit at one of three facilities. If sanity checks are unable to eliminate any of these facilities then the possible area in which the target is located will be a combination of all three regions. These two examples are discussed in detail in the next sections.

### 4.7.2 Single Facility Example

Single-facility IXPs are relatively less common than multi-facility IXPs, especially in regions with well-established Internet infrastructure. However, their prevalence can vary depending on the region and the stage of development of the Internet ecosystem. In developing regions, single-facility IXPs may be more common due to limited resources, lower traffic volumes, and the state of the local Internet ecosystem (InternetSociety, 2021). Single-facility IXPs are often the starting point for interconnection in these areas, providing foundational infrastructure for local traffic exchange. These IXPs might suffice to meet the needs of local Internet service providers (ISPs), content delivery networks (CDNs), and enterprises. These IXPs can effectively serve as regional or niche hubs (InternetSociety, 2021).

Due to privacy, probe locations are obfuscated for all RIPE Atlas users. The location that is provided by the API will be within 80 to 400 metres of the actual location submitted to the RIPE Atlas platform.

Taking into account Atlas's obfuscation policy, Atlas Probe 1000237 is located within 80 to 400 metres of Samrand Business Park, 37 km north of Johannesburg. It has an IP address of 196.61.64.64 and is selected as the target probe. Sixty-six traceroute measurements are created from each of the other probes to this target. The measurements in which we are interested are:

- Shortest RTT times from IXP to target.
- Shortest RTT time from IXP to Last Hop Router (LHR)
- Shortest RTT time from target to Last Hop Router (LHR)

Figure 4.7 depicts the four traceroutes with the shortest trip times and includes the IXP Ingress and Egress, the last hop Routers and the IXPs.

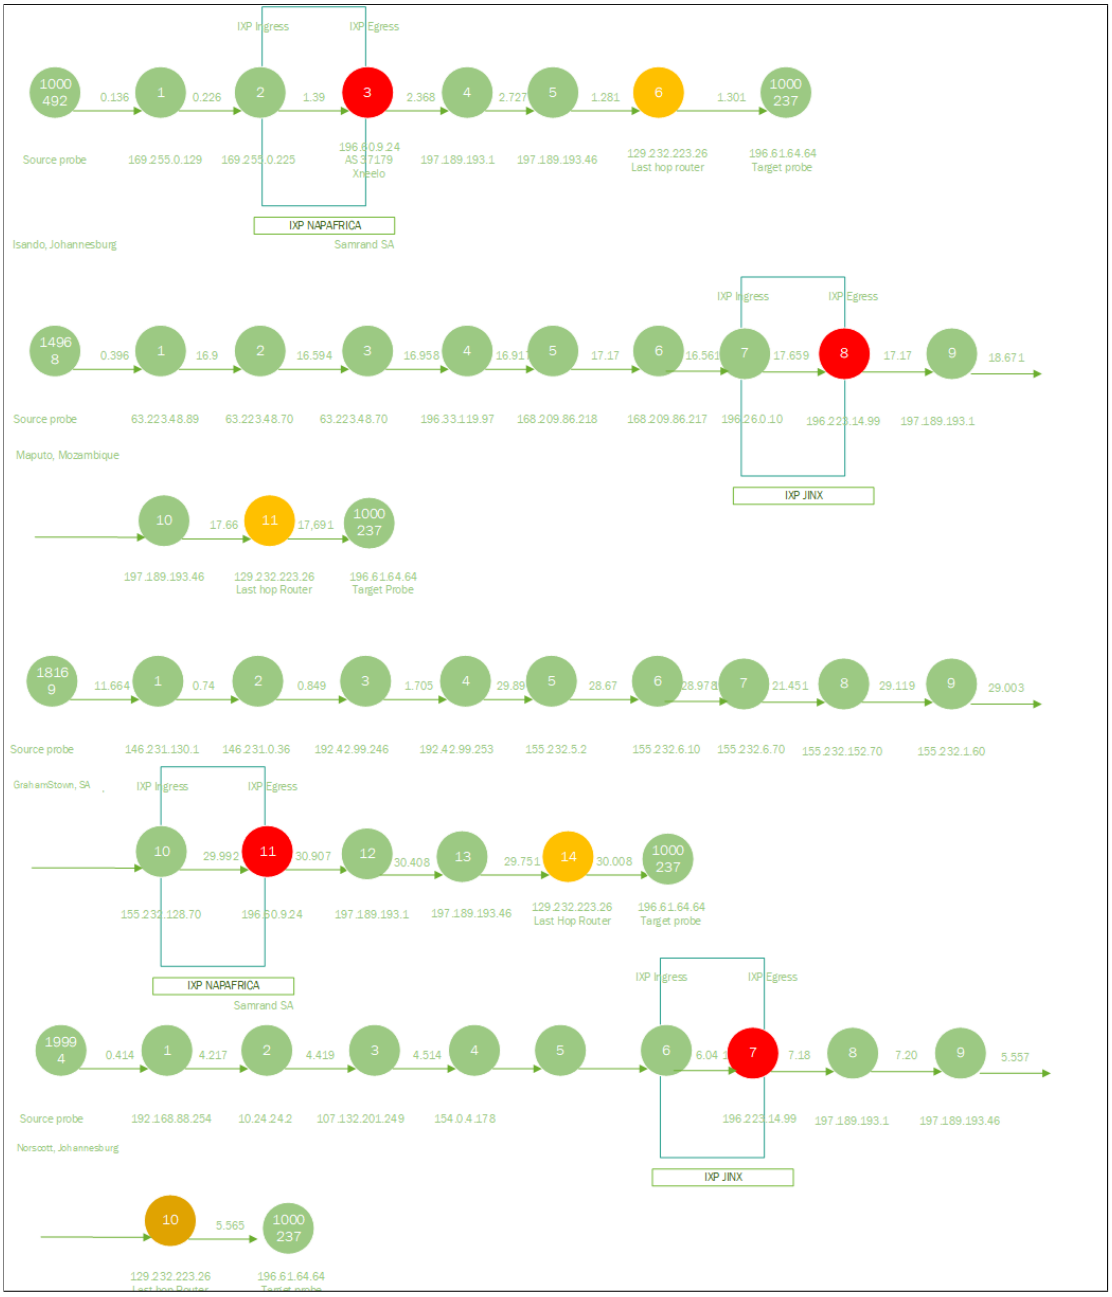


Figure 4.7: Four example traceroutes to Target 196.61.64.64, Probe 1000237

Table 4.1 shows the routes taken at each hop, as packets pass through the IXPs and Table 4.2 shows the results we can take from these traceroutes.

| Source Probe | hop1 ip | hop1 rtt | hop2 ip | hop2 rtt | hop3 ip | hop3 rtt |
|---|---|---|---|---|---|---|
| **1000492** | 169.255.0.129 | 0.136 | 169.255.0.225 | 0.226 | 196.60.9.24 | 1.39 |
| **14968** | 63.223.6.89 | 0.396 | 63.223.48.70 | 16.9 | 63.223.48.70 | 16.594 |
| **18169** | 146.231.130.1 | 11.664 | 146.231.0.36 | 0.74 | 192.42.99.246 | 0.849 |
| **19994** | 192.168.88.254 | 0.414 | 10.24.24.2 | 4.217 | 102.132.201.249 | 4.419 |
| | | | | | | |
| **Source Probe** | **hop4 ip** | **hop4 rtt** | **hop5 ip** | **hop5 rtt** | **hop6 ip** | **hop6 rtt** |
| **1000492** | 197.189.193.1 | 2.368 | 197.189.193.46 | 2.727 | 129.232.223.26 | 1.281 |
| **14968** | 196.33.119.97 | 16.958 | 168.209.86.218 | 16.917 | 168.209.86.217 | 17.17 |
| **18169** | 192.42.99.253 | 1.705 | 155.232.5.2 | 28.89 | 155.232.6.10 | 28.67 |
| **19994** | 154.0.4.178 | 4.514 | unknown | unknown | unknown | unknown |
| | | | | | | |
| **Source Probe** | **hop7 ip** | **hop7 rtt** | **hop8 ip** | **hop8 rtt** | **hop9 ip** | **hop9 rtt** |
| **1000492** | 196.61.64.64 | 1.301 | | | | |
| **14968** | 196.26.0.10 | 16.561 | 196.223.14.99 | 17.659 | 197.189.193.1 | 17.17 |
| **18169** | 155.232.6.70 | 28.978 | 155.232.152.70 | 21.451 | 155.232.1.60 | 29.119 |
| **19994** | 196.223.14.99 | 6.041 | 197.189.193.1 | 7.18 | 197.189.193.46 | 7.202 |
| | | | | | | |
| **Source Probe** | **hop10 ip** | **hop10 rtt** | **hop11 ip** | **hop11 rtt** | **hop12 ip** | **hop12 rtt** |
| **1000492** | | | | | | |
| **14968** | 197.189.193.46 | 18.671 | 129.232.223.26 | 17.66 | 196.61.64.64 | 17.691 |
| **18169** | 155.232.128.70 | 29.003 | 196.60.9.24 | 29.992 | 197.189.193.1 | 30.907 |
| **19994** | 129.232.223.26 | 5.557 | 196.61.64.64 | 5.565 | | |
| | | | | | | |
| **Source Probe** | **hop13 ip** | **hop13 rtt** | **hop14 ip** | **hop14 rtt** | **hop15 ip** | **hop15 rtt** |
| **1000492** | | | | | | |
| **14968** | | | | | | |
| **18169** | 197.189.193.46 | 30.408 | 129.232.223.26 | 29.751 | 196.61.64.64 | 30.008 |
| **19994** | | | | | | |

Table 4.1: Fastest Time to Target 196.61.64.64 Probe 1000237, Red = IXP, Orange = Last Hop Router

| Source Probe | i to t | i to lr | t to lr | s to lr | s to t |
|---|---|---|---|---|---|
| **1000492** | -0.89 | -0.109 | 0.02 | 1.281 | 1.301 |
| **14968** | 0.032 | 0.001 | 0.031 | 17.66 | 17.691 |
| **18169** | 0.016 | 0.247 | 0.257 | 29.751 | 30.008 |
| **19994** | -0.476 | -0.484 | 0.008 | 5.557 | 5.565 |

Table 4.2: Fastest Time to Target results, Blue = Quickest Time

The four measurements in table 4.1 each show their respective hop's IP address and the RTT time to that hop. Most importantly they show the IXP that is traversed. The five columns in table 4.2 displays the IXP to target RTT time (i to t), the IXP to last hop router RTT time (i to lr), the target to last hop router RTT time (t to lr), the source probe to last hop router RTT time (s to lr), and finally the source to target RTT time (s to t). Each of these columns has the shortest relative trip time coloured blue.

From the 66 traceroute measurements we find that source probe 18169 has the shortest time from IXP to target of 0.016 milliseconds, as highlighted by the blue box in the "i to t" column, and hop 11 uses the IP address 196.60.9.24 which according to Peerindb belongs to NAPAFRICA Internet Exchange based in Johannesburg with a PeeringDB assigned IXLAN identification of 592. This is connected to a PeeringDB-assigned network identification of 9791 (see Figure 4.8).

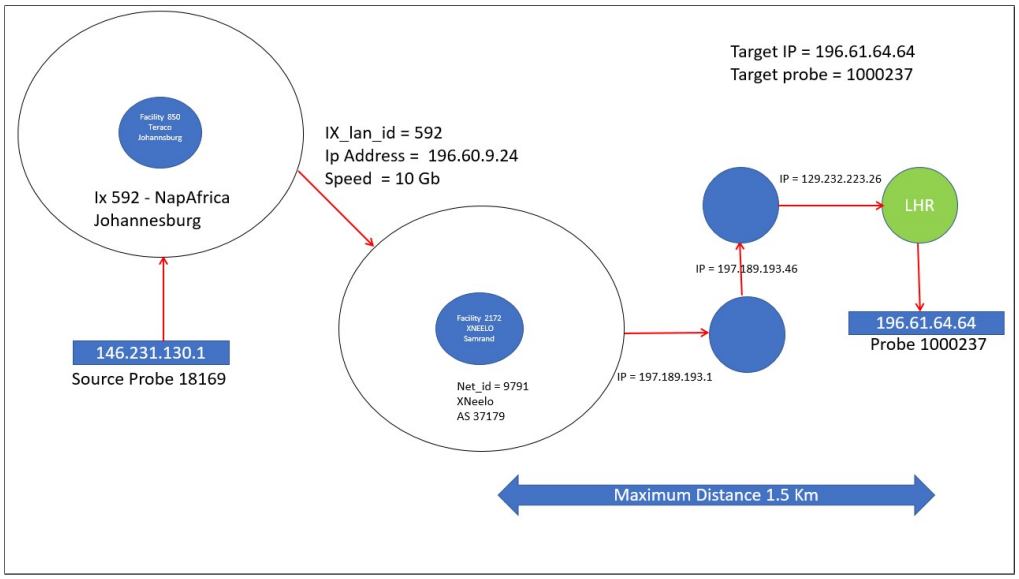


Figure 4.8: Single Facility IP Geolocation

Network 9791 belongs to a hosting company named Xneelo who owns AS37179 and we can see from PeeringDB that this network/AS is hosted at only one facility which is in Samrand, South Africa, with a latitude of -25.927622 and a longitude of 28.140755.

Using the 'Distance to Target' calculation (2.1), it is determined that the target is no more than 1.5 km from the Samrand Facility and is likely to be located at the same site, considering packet processing and latency factors. To evaluate the process, we compared this estimated location with the actual location of the probes as provided by the RIPE Atlas platform. Considering RIPE's 400-meter probe obfuscation policy, the estimation aligns closely with the actual location, as shown in Figure 4.9. While this result demonstrates the method's potential to achieve more precise IP geolocation than current state-of-the-art techniques, such as Single-Radius or shortest-ping methods, it is important to note that this conclusion is based on a single case study. Further testing across diverse scenarios and geographic regions would be needed to comprehensively validate the robustness and general applicability of this method.

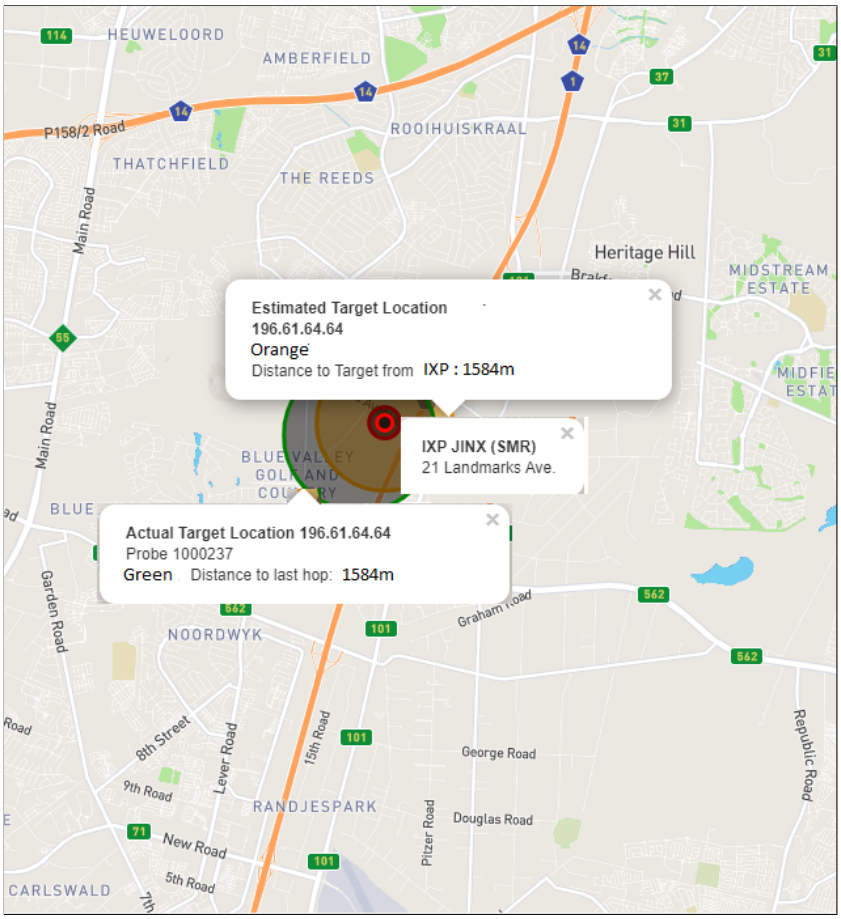


Figure 4.9: Estimation Vs actual location Orange = Estimated, Green = Actual (Taking into account RIPEs 400 m obfuscation policy

### 4.7.3 Multi Facility Example

The fundamental difference between the single-facility example in Section 4.7.2 and the multi-facility example in this section lies in the complexity of determining the exact location where a packet exits the IXP network. In the single-facility example the IXP exists at only one physical location and since there is only one facility, it is straight forward to infer the geolocation of any packet traversing the IXP. However, in a multi-facility example the IXP spans multiple geographically dispersed facilities. Therefore, a packet can enter at one facility and exit at another, making IP geolocation more difficult. Additionally, The error radius can increase significantly since multiple possible exit locations exist. Also, sense checks are needed to rule out physically impossible paths, but uncertainty still remains.

Probe 22221 is situated in Cape Town, has an IP address of 196.40.111.174, and is selected as the target probe; traceroute measurements are created from each of the other probes to this target. The measurements in which we are interested are:

- Shortest RTT times from IXP to target.
- Shortest RTT time from IXP to Last Hop Router (LHR)
- Shortest RTT time from target to Last Hop Router (LHR)

Figure 4.10 below shows the three shortest measurements. The three measurements each show their respective hop's IP address and the RTT time to that hop. The final five columns display the IXP to the target RTT time (i to t), the IXP to last hop router RTT time (i to lr), the target to last hop router RTT time (t to lr), the source probe to last hop router RTT time (s to lr), and finally the source to target RTT time (s to t). Each of these columns has the shortest respective RTT time coloured green.

From the 66 traceroute measurements we find that source probe 1000707 has the shortest time from IXP to the target of 0.023 milliseconds, and hop 9 uses the IP address 196.223.22.98 which belongs to the Cape Town Internet Exchange (CINX.)

PeeringDB assigns this IXP an identification of 344. CINX uses three facilities in Cape Town, and this IP address peers with a PeeringDB assigned identification of 9791 which is Xneelo PTY known as AS37153. Xneelo only has one facility, which is in Samrand, Johannesburg, with a peering assigned facility identification of 2172. This facility is more than 1200 km away, and a simple sense check would warn that it is impossible for a packet to travel that distance in 0.023 milliseconds. Xneelo must have some presence locally that is not documented, most likely a remote peering connection as discussed in Section 4.5

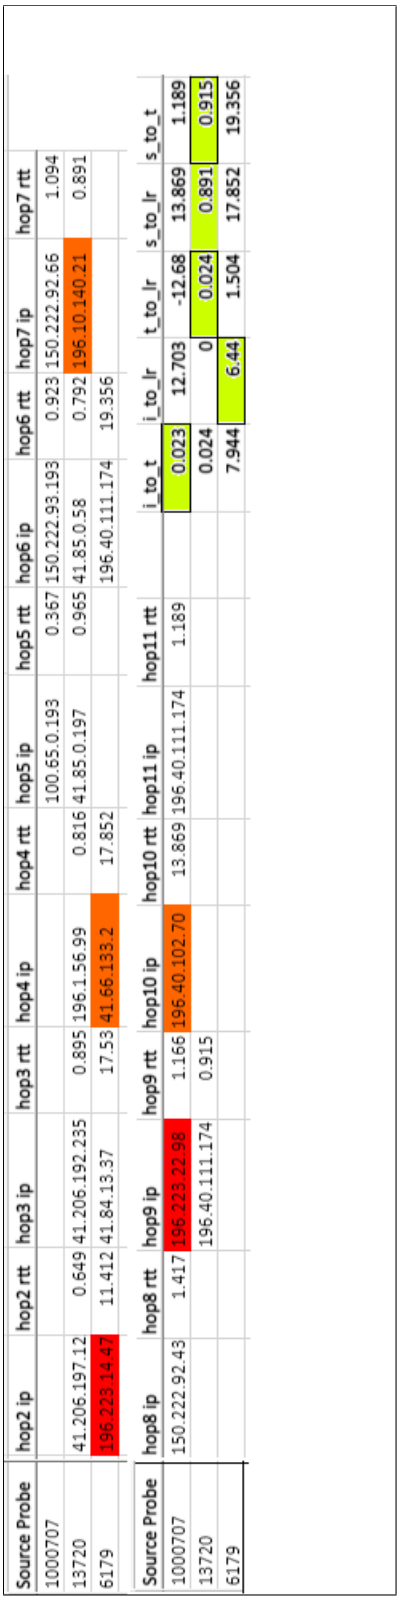

| Source Probe | hop2 ip | hop2 rtt | hop3 ip | hop3 rtt | hop4 ip | hop4 rtt | hop5 ip | hop5 rtt | hop6 ip | hop6 rtt | hop7 ip | hop7 rtt |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1000707 | | | | | | | 100.65.0.193 | 0.367 | 150.222.93.193 | 0.923 | 150.222.92.66 | 1.094 |
| 13720 | 41.206.197.12 | 0.649 | 41.206.192.235 | 0.895 | 196.1.56.99 | 0.816 | 41.85.0.197 | 0.965 | 41.85.0.58 | 0.792 | 196.10.140.21 | 0.891 |
| 6179 | 196.223.14.47 | 11.412 | 41.84.13.37 | 17.53 | 41.66.133.2 | 17.852 | | | 196.40.111.174 | 19.356 | | |

| Source Probe | hop8 ip | hop8 rtt | hop9 ip | hop9 rtt | hop10 ip | hop10 rtt | hop11 ip | hop11 rtt | i_to_t | i_to_lr | t_to_lr | s_to_lr | s_to_t |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1000707 | 150.222.92.43 | 1.417 | 196.223.22.98 | 1.166 | 196.40.102.70 | 13.869 | 196.40.111.174 | 1.189 | 0.023 | 12.703 | -12.68 | 13.869 | 1.189 |
| 13720 | | | 196.40.111.174 | 0.915 | | | | | 0.024 | 0 | 0.024 | 0.891 | 0.915 |
| 6179 | | | | | | | | | 7.944 | 6.44 | 1.504 | 17.852 | 19.356 |

Figure 4.10: Shortest time to Target 196.1.58.54, Probe 13720
Red: IXP, Orange: Last Hop Router, Green: Shortest RTT

The packet's RTT time at the IXP was 1.166 ms, and the packet's RTT time at the target was 1.189 ms. Therefore a packet took approximately (1.189 ms - 1.166 ms)/2 to travel the distance from IXP to reach its destination, which therefore took 0.0115 milliseconds and, referring to our formulae at (2.1) this equates to a maximum distance of 2.277 km however, we do not know from which of the 3 CINX facilities the packet transited. Therefore, the possible target area is a combined area of 2.277 km around each of the 3 facilities (orange circles) as shown in Figure 4.11. The actual target location is shown as a green circle, taking into account RIPE's obfuscation policy. The red dots are the locations of the three CINX facilities. It can be seen that the actual location of the target is well enclosed within one of the three estimated circles.

An additional point to note is the RTT time for hop 10, which is over 13 ms. It could be of interest to Xneelo why this latency is occurring at this hop. A DNS lookup, using NSlookup, of this IP address gives the URL value 'core-access-switch1-vlan1001.cpt.host-h.net.'; presumably, the subdomain "CPT" is short for Cape Town. Perhaps it is caused by traffic engineering, i.e. low priority for packets with that specific routers destination, or ICMP packets, but it may be that the packet is travelling all the way to Xneelo's facility in Johannesburg for no essential reason. However, this hop does not seem to affect packets in transit, as the final hop RTT is back to within a reasonable latency. Resolving small issues such as this could result in improved network latency and greater network efficiency.

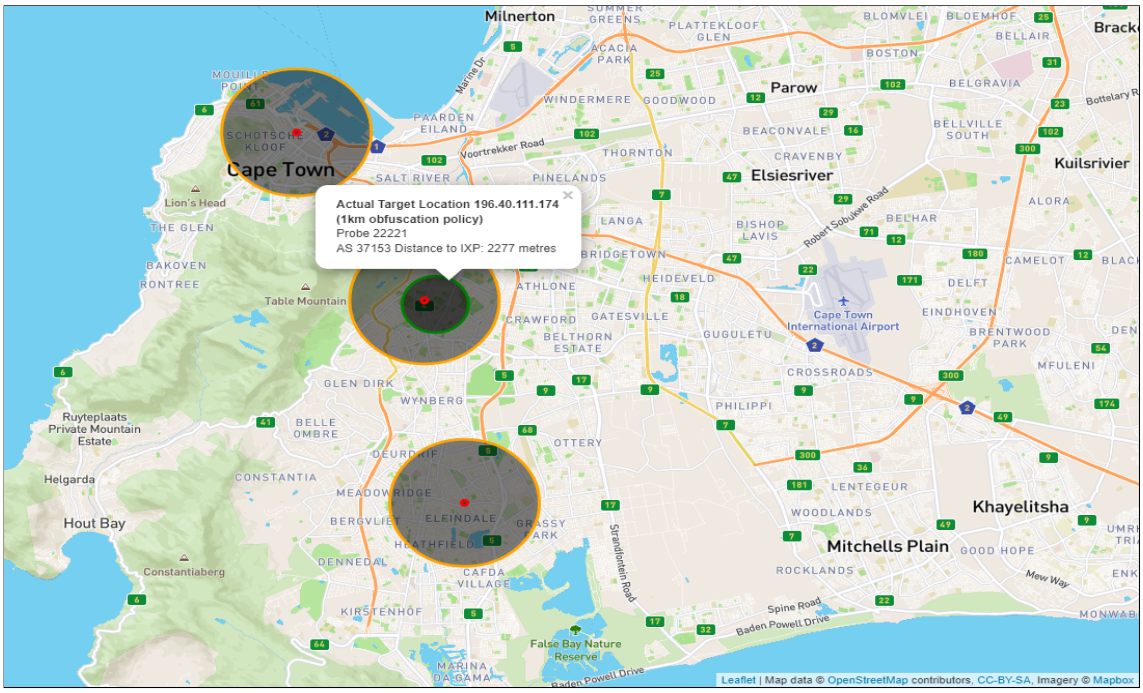


Figure 4.11: Multi Facility Example

## 4.8 Testing the use of IXPs on the UK Infrastructure

A list of UK-based facilities was extracted from PeeringDB along with the geographical coordinates of each facility. Where facility records have no geolocation information available, the facility's address is entered into Nominatim (nominatum, 2024), which is a tool to search OSM data by name and address (geocoding) and generate 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 16th February 2023), only sixteen facilities had to be manually geolocated. Figure 4.12 shows the number of facilities (black rectangles) geolocated in the London area using PeeringDB, Nominatim, and OSM.

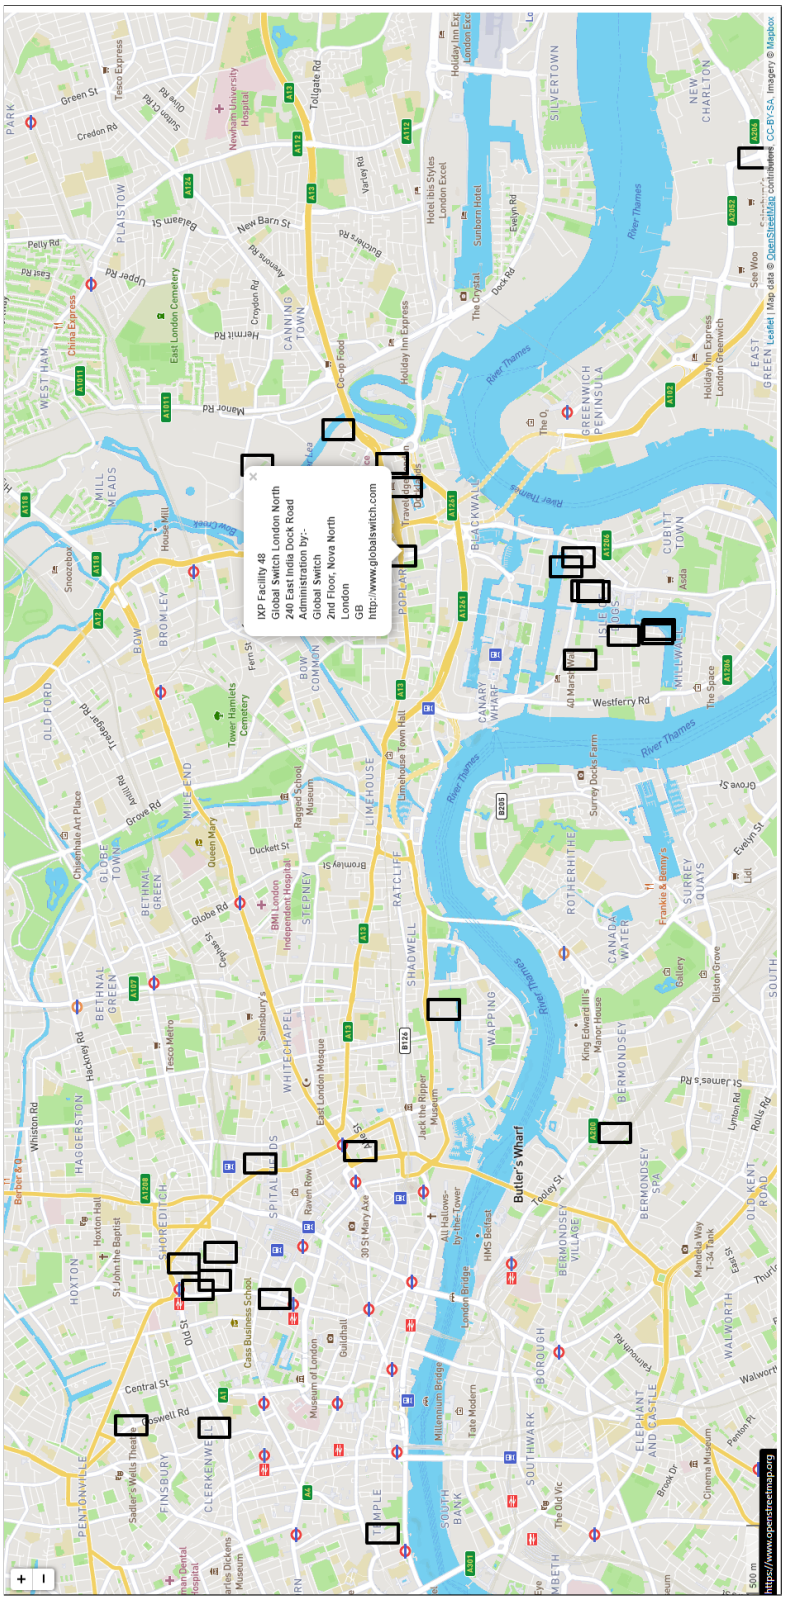


Figure 4.12: Geo-Mapping Interconnection Facilities using PeeringDB and OpenStreetMap

According to PeeringDB, there are currently 26 Internet Exchanges in the London area, although two are listed with no connected networks. Packet Clearing House (PCH) lists 15 active IXPs, whereas the IXPDB website lists nine IXPs with connected networks. To map an Internet Exchange, PeeringDB is queried to discover the facilities at which each IXP publicly interconnects, and these layer 2 networks are mapped to OpenStreetMap, as shown in Figure 4.13. Public information is not available to map the actual physical cables, so, while point-to-point connections are depicted in this Figure, the exact nature of the network topology is uncertain. It is possible that the network could involve various forms of mesh configuration, where some or many points may be interconnected. However, the principle of the Layer 2 logical network remains consistent regardless of the details of the physical connection.

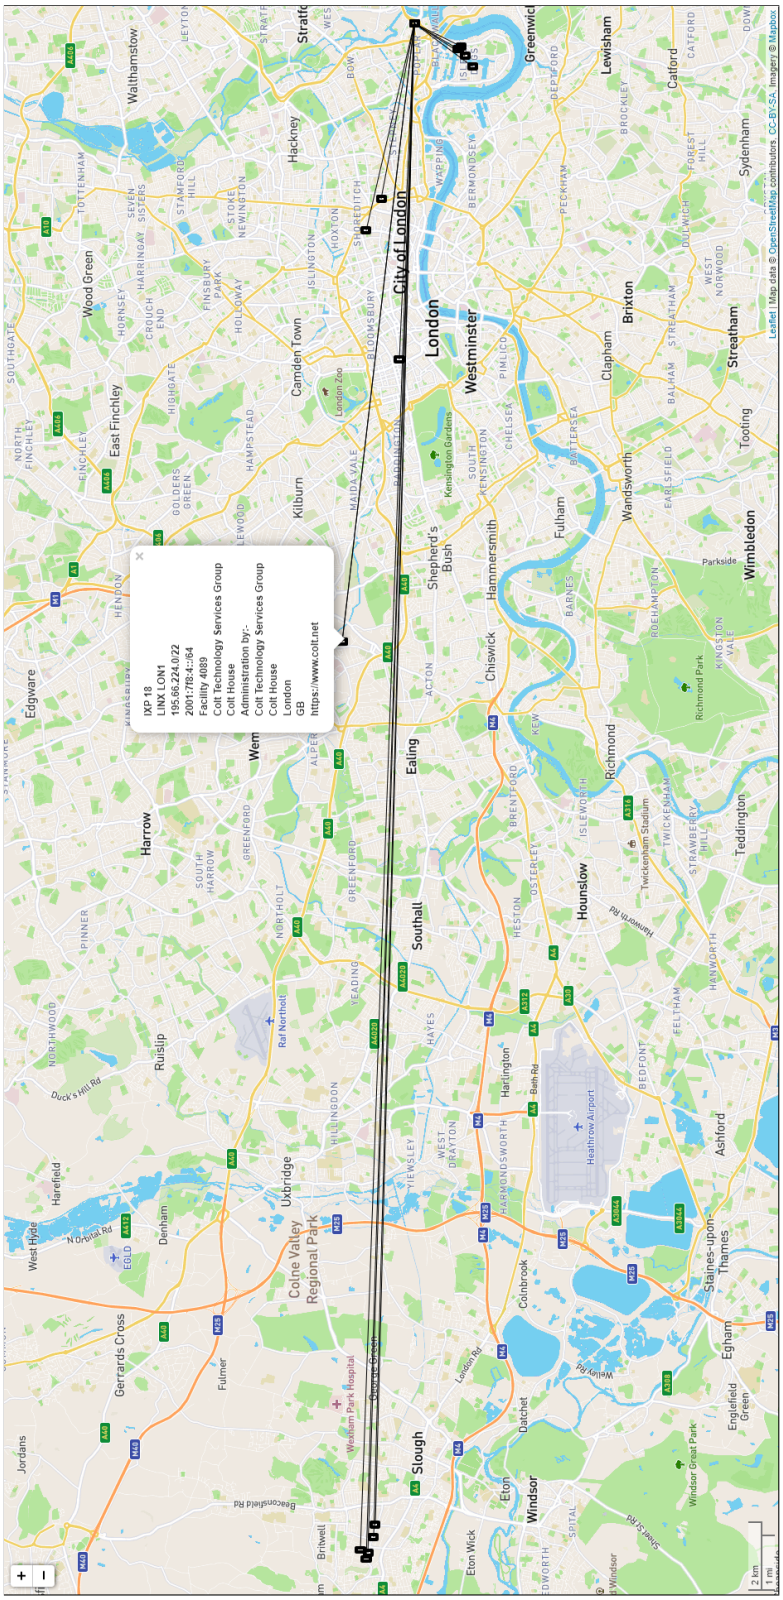


Figure 4.13: LINX LON1 IXP showing its Public Peering Facilities and its Layer 2 Network

Once the public peering points from PeeringDB have been mapped, we can refer to the IX website to collect any additional IXP 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. In addition, many ports are marked with a port type 'Connexions', which are LINX's reseller partners, and provide information on clients who 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 IXPs 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 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 (CAIDA, 2024) and Looking Glass (LG) servers which are publicly accessible network diagnostic tools provided by ISPs, IXPs, and network operators. They allow users to query real-time BGP routing and traceroute data from different network vantage points, aiding in geolocation research and 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 avoids this problem by adapting the header fields of the probe packet in a manner that allows all probes to follow the same path; per-flow load balancing is an option (Augustin, Cuvellier, Orgogozo, Viger, Friedman, Latapy, Magnien, and Teixeira, 2006). The

RIPE Atlas platform uses the Paris Traceroute as default. Although it cannot enumerate paths in all situations, it has been shown to perform considerably better than the classic traceroute.

### 4.8.1 Discovering Target Locations using IXPs as Landmarks Example

First of all a traceroute must be created from a source RIPE Atlas probe towards a target IP address. At each hop towards the target, a test is carried out to see if the IP address matches known IXP IP addresses. If the test proves successful, we can then take note of the IP address and ASN of the entry point and compare its peering facilities with the IXP's peering facilities. If we find a single facility that is shared between the entry ASN and the IXP then we can be sure that this facility is the correct one that is used by the route. The exit facility is discovered in a similar fashion, whereby the IP address and ASN of the hop leaving the IXP is noted and its peering facilities are compared with the IXP peering facilities. If the exit ASN and IXP share only a single facility, we can be assured that the facility is the correct exit facility. For example, Probe 6182 is located in Leeds and has an IP address of 141.170.19.12. Probe 6087 is located in Slough and has an IP address of 141.170.19.12. A traceroute from the Leeds probe 6182 routes via the following hops as shown in Figure 4.14 and discovers the exact facility used by an IXP that the traceroute passes through allowing the creation of a landmark which is closer to our target and therefore reduces the error distance.

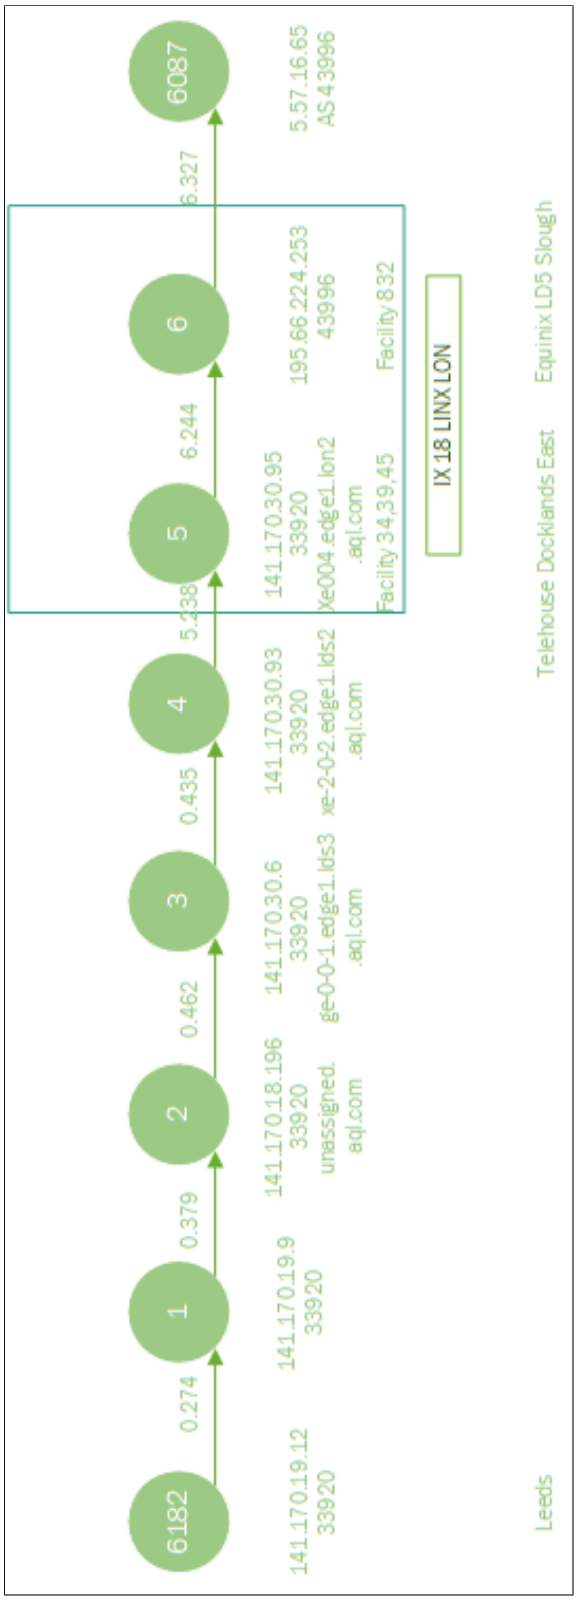


Figure 4.14: Discovering an IXP's exit facility

Hop 5 is the IXP entry hop; its IP address is owned by AS33920, and its peering facilities are the following PeeringDB assigned identifications [39, 45, 896, 76, 2384, 34].

Hop 6 is IXP 18 and its IP address is hosted by Linx LON whose peering facilities according to PeeringDB are [34, 39, 40, 43, 45, 46, 79, 399, 534, 832, 2262, 835, 4404, 4360, 4089, 3152, 6535, 3399].

Hop 7 is the IXP exit hop; its IP address is owned by AS43996, and its peering facilities are [832, 63, 1, 705, 225, 58].

Comparing entry facilities with IXP facilities results in 3 facilities that are possible candidates [39, 45, 34]; however, all of these are owned by Telehouse in London and are in a tight geographical area in London Docklands, so for this purpose any of the three can be chosen as the entry facility. It should be noted that in many cases only one facility is shared between the IXP and the entry ASN making this process much more accurate because the shared facility becomes the de facto meeting point for the ASN and the IXP. It should also be noted that in some cases two or more geographically distant ASN's are shared between the IXP and the entry ASN making the process far more difficult. This introduces ambiguity because it is unclear which facility is being used for the specific data path observed in the traceroute. However, if traceroutes are sent from multiple locations towards the target then the number of shared facilities may be reduced. Additional sense checks, such as the speed of light over a distance, may also eliminate certain facilities.

We carry out the same process for the exiting ASN's (43996) facilities which are [832, 63, 1, 705, 225, 58] and discover that the only facility that is shared between the IXP and the exiting ASN is 832, which is Equinix LD5 located in Slough and therefore this must be the exiting facility as shown in Figure 4.15. If it were the case that we were unable to narrow down the number of facilities, for example, if all 6 of the facilities listed above hosted ASN 43996 then we would have to draw possible geolocation areas around all 6 IXPs as discussed in Section 4.7.3

It should also be noted that in this case we have not discovered the exact facility

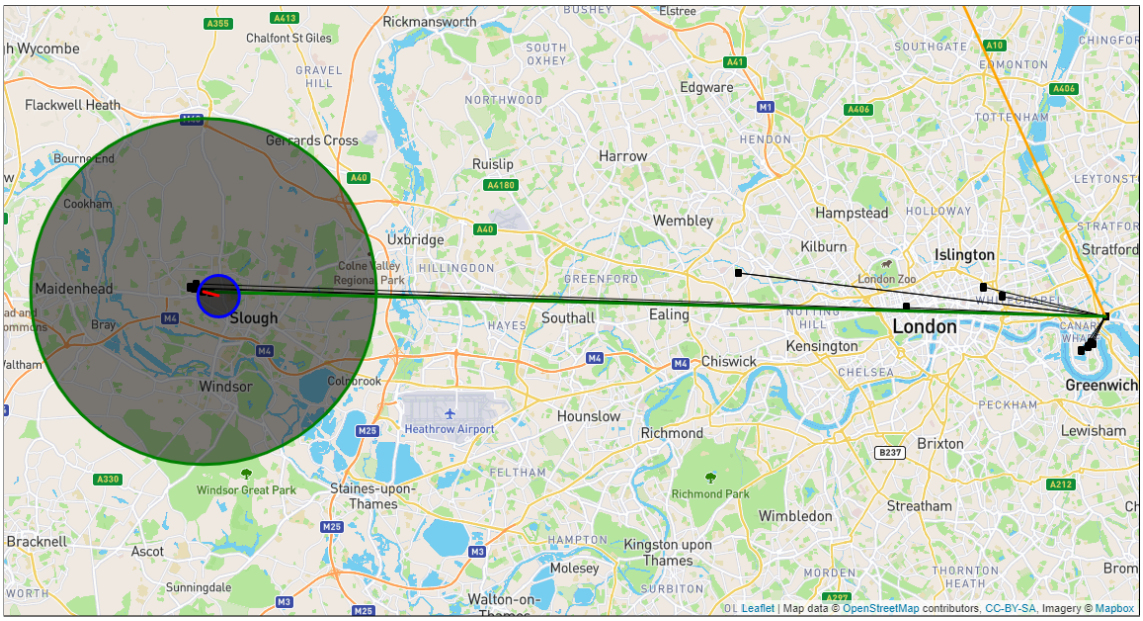


Figure 4.15: Using an IXP as a Vantage Point

where the packet enters the IXP; the packet could enter the IXP at any of the three Telehouse facilities based in the Docklands in London and shown as black rectangles in Figure 4.15. However, this is not important as we are only interested in geolocating the target, and for that, we are using the exiting IXP as the nearest Vantage Point. The orange line intersecting with the Telehouse facilities denotes the packet path, and as a visual cue regarding estimated speed, the orange colour shows that the speed of the packet is travelling slightly less than optimal. After entering the IXP LINX LON1, the packet then travels through the IXP layer 2 network towards the exiting facility; this is shown as a green line and denotes the packet is now estimated to be travelling faster at close to the maximum packet speed in a fibre medium.

Upon discovery of the exiting facility, we now have a Vantage Point that should be far closer than any existing VPs. Closer Vantage Points reduce geolocation errors that are caused by latency factors such as processing, queueing, congestion, and transmission delays and are multiplied by the number of intermediate routers that are crossed.

To discover the approximate location of the target IP address, we can now

subtract the RTT value of the exiting IXP hop from the RTT value of the target hop: $6.327 - 6.244 = 0.083$ ms. This equates to the time it can take for a packet to travel from the IXP exit facility to the target and return. We can multiply this by the commonly used speed of a packet in a fibre medium which is 0.66 x speed of light which equals approximately 200 km per millisecond, and hence 200 *.083 = 16.6 km. This figure is for the round trip and, in fact, we only need the one-way distance so we can half this figure to give an approximate target distance from the facility of 8.3 km. The blue circle shows the actual target location according to RIPE Atlas taking into account its 1 km obfuscation policy, whilst the green circle shows the calculated estimated location area. At this point, it can be seen that the packet takes a considerable amount of time to travel the actual distance to the target, and the visual cue for this is that the line is now red. This also accounts for the relatively large target area.

## 4.9 Summary

This chapter explores the innovative use of Internet Exchange Points (IXPs) as Vantage Points to improve the accuracy of IP geolocation and reduce errors associated with traceroute methods. Using the strategic positioning of IXPs, which are typically located near major population centres, the study proposes that these facilities can serve as closer reference points, minimising issues such as congestion, latency, and circuitous routing paths that often skew geolocation data.

The chapter introduces a methodology for identifying and utilising IXPs as key geolocation markers. It discusses resolving the challenges associated with multi-facility setups, ensuring that geolocation can benefit from the presence of these IXPs without being confounded by their complexity.

Proofs of concept are provided using data from South Africa and the United Kingdom, demonstrating the effectiveness of employing IXPs as geolocation landmarks. By reducing the geographical distance between the source and the target

of the traceroute and avoiding unnecessary detours through less direct routes, the approach promises enhanced accuracy and reliability in the analysis of network paths.

Ultimately, the chapter underscores the potential of IXPs to act as pivotal landmarks in geolocation efforts, contributing to a more precise mapping of network topologies and improving the Quality of Service for end users.

The use of IXPs as Vantage Points is shown to provide very accurate IP Geolocation when the exiting facility can be determined. The method of comparing the ASN of the IXP and the ASN of the next-hop IP address to determine the correct facility is used in Chapter 4, but is effectively made redundant by the introduction of IXP gateway databases (which various Internet Exchanges began making available to the public towards the end of the research carried out for this thesis). This new method for discovering the location of IXP facilities is introduced in Chapter 5.

## 4.10 Conclusions

IXPs are pivotal in understanding Internet topology because they serve as central hubs for data exchange between networks. Identifying the physical facilities where IXPs are located provides valuable data points for geolocation whilst the UK's IXP landscape, characterized by well-documented facilities like LINX and LONAP, offers an excellent testing ground due to the high density of IXPs and their detailed public data.

Traceroute data reveals the paths packets take through networks, but identifying the exact entry and exit points within multi-facility IXPs remains challenging. Layer 2 networks used within IXPs obscure the packet paths between entry and exit points, hindering precise geolocation although single-facility IXPs are easier to geolocate since the data exchange occurs in a single physical location, reducing ambiguity.

Using data directly from LINX and LONAP as ground truth proved invaluable, as these IXPs physically configure routers and maintain precise geolocation data.

However, reliance on ground truth data assumes the accuracy and timeliness of these records, which may not always align with dynamic network changes.

The proposed methods in this chapter shows promise in achieving greater accuracy compared to traditional geolocation techniques like shortest-ping or Single-Radius methods. Combining IXP-specific data with additional sense checks (e.g., Speed of Light RTT constraints) can refine the accuracy of geolocation further.

Combining traceroute analysis, IXP database records, and supplementary tools enhances the reliability of geolocation methods. A holistic approach that accounts for both technical (e.g., RTT) and infrastructure-level (e.g., IXP database) data provides a clearer picture of network topology.

By testing IXPs as geolocation anchors within the UK, this chapter underscores the potential while highlighting the need for refinement, validation, and expansion to broader contexts.

# Chapter 5

# Improving Internet Infrastructure Mapping

## 5.1 Introduction

Chapter 4 establishes why IXPs can serve as reliable geolocation landmarks, while Chapter 5 builds upon this concept by actively using IXPs to improve Internet infrastructure mapping through a systematic methodology. This chapter builds on the previous methods put forth in Chapter 4 for using Internet Exchange Points (IXPs) as Vantage Points to create fine-grained multilayer maps of the Internet structure. It transitions from theory to application, demonstrating how IXPs can be leveraged to enhance geolocation accuracy and infrastructure mapping. It develops a methodology that integrates IXP-based geolocation into a systematic mapping approach. Finally, it validates the approach through real-world testing.

The new procedure will use a similar combination of methods, databases, and tools described in Chapter 4, however, data from the IXPDB website will now be referenced in step 4. This will replace the earlier method of finding a common facility between the IXP ASN and the outgoing ASN, as described in Section 4.3, with the aim of further improving the mapping of the Internet infrastructure to a facility level. The steps are repeated below for clarity however step 4 now includes the

interrogation of the IXPDB website and removes the comparison of ASN's method.

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 IXP 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 information that will be useful for future research. To map these intermediate hops, a software tool was created, which reads the data from the traceroute measurements created in Step 3, and queries various sources, such as PeeringDB, DNS, the Internet Exchange websites, and now the IXPDB website to locate the position of the router where these hops are interfacing, considering the previously discovered facility and IXP information.

Chapter 2.3 details the difficulties in using delay-distance models and suggests that the use of Round Trip Times (RTTs) can lead to highly misleading results. This chapter develops a new procedure that combines state-of-the-art methods to avoid many of the fundamental problems in Internet topology mapping while creating finer-grained Internet maps than those currently available. The procedure is tested on the UK infrastructure by conducting a series of tests using distributed measurement points provided by the RIPE Atlas platform.

**Roadmap** In Section 5.2, we detail the objectives of this chapter. Section 5.3 discusses the use of the RIPE Atlas platform and the usage metrics. In Section 5.4, we test the new procedure, and in Section 5.5, the procedure is automated by creating rules that can be used by a new software tool that is developed, and the results are provided. Section 5.6 discusses many of the points raised in this chapter. Section 5.7 raises the spectre of the impending global shift towards IPv6 and now pivots the thesis towards IPv6 geolocation.

## 5.2 Objectives

The objectives of this chapter 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.
- Introduce a new technique to create constraints in DNS geohints inference. Although 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 any differences in the forward and reverse measurements.
- 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 including London.

- To evaluate the inferences and estimate its success using a carefully curated dataset obtained by two of the largest London IXPs.

## 5.3 RIPE Atlas Measurements

The use of the RIPE Atlas platform requires credits which can be spent to request User-Defined Measurements (UDMs) from the platform. RIPE Atlas probe hosts earn these credits for the time their probes remain connected and for the number of measurement results they generate. This is designed to serve as a means of measuring the level of contribution to and consumption of resources in the RIPE Atlas system.

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 was used to calculate the RIPE Atlas credit cost of a user-defined measurement:

$$\text{Traceroute credit cost} = 10 \times N \times (\frac{S}{1500} + 1) \tag{5.1}$$

*Where: N = Number of packets per traceroute (default is 3), S = packet size (default is 40)*

Thus, the total approximate cost is $1190 \times 40 \times \approx 1 = 47600$ credits.
It should be noted that 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.

## 5.4 The Method in Action

Where possible, each hop on a traceroute was assigned to a facility using a combination of DNS lookup, Internet Exchange website information, and PeeringDB information. The combined information creates new Vantage Points on the way to destinations; these will be invaluable in further research, especially when there is a dearth of ground-truth data, as many researchers have recognised.

However, the use of these Vantage points does come with a warning. IP-to-geolocation mappings can change over time due to dynamic factors such as network reconfigurations, ISP mergers, or changes in infrastructure ownership. However, the value of the discovered vantage points lies in their utility at the time of their documentation. These geolocation pairs represent a snapshot of the Internet's topology and infrastructure during the study period. As such, they provide valuable ground truth data for benchmarking geolocation methodologies and serve as a historical reference point for evaluating changes in the Internet's infrastructure over time.

Future researchers can leverage these vantage points to:

- Validate and refine new geolocation techniques.
- Analyze temporal trends in IP geolocation data by comparing past and present mappings.
- Gain insights into the stability or volatility of specific regions or networks in the Internet topology.

To mitigate the risk of relying on outdated data, future work should incorporate continuous validation of these vantage points through active measurement and cross-referencing with updated datasets. While some geolocation pairs may become obsolete, the methodology and framework established in this research remain robust and adaptable for identifying new vantage points as the Internet evolves.

### 5.4.1 Example of Network Mapping

An example of this logical network mapping is shown in Figure 5.1, where a traceroute is first carried out from probe 6515 towards probe 6087.

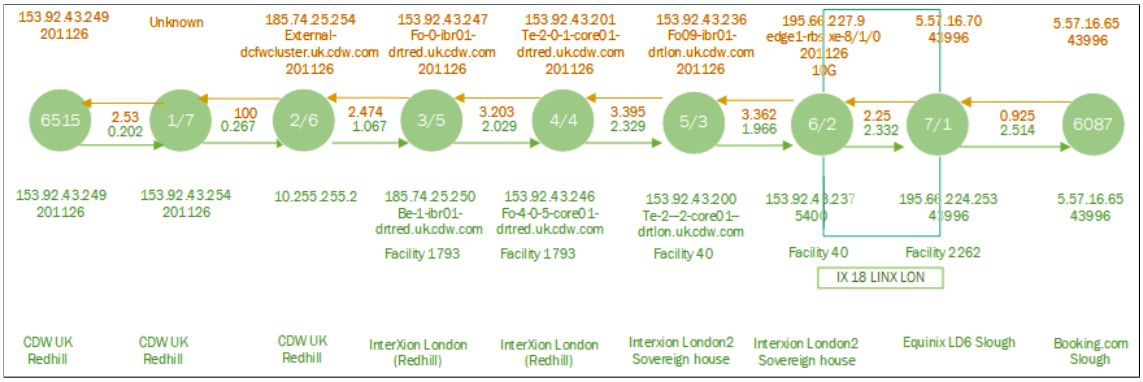


Figure 5.1: Forward and reverse traceroute measurements between two probes demonstrating facility and IXP mapping.

The IP address of each hop corresponds to an ingress port interface on a router located in a specific location. The hop times for this measurement are listed in Table 5.1.

| Hop | RTT (m/s) | IP Address | DNS |
|---|---|---|---|
| 1 | 0.202 | 153.92.43.254 | |
| 2 | 0.267 | 10.255.255.2 | |
| 3 | 1.067 | 185.74.25.250 | Be-1-ibr01-drtred.ukcdw.com |
| 4 | 2.0153 | 192.43.246 | Fo-4-0-5-core01-drtred.ukcdw.com |
| 5 | 2.329 | 153.92.43.200 | Te2–2-core01-drtlon.ukcdw.com |
| 6 | 1.966 | 153.92.43.237 | |
| 7 | 2.332 | 195.66.224.253 | |
| 8 | 2.514 | 5.57.16.65 | |

Table 5.1: Traceroute Measurement Table

A DNS look-up for each IP address on the route can provide useful information regarding the location of the port/router combination. For example, in Figure 5.1, we find that the first hop has no assigned DNS name and is probably the default gateway of the probe owner (CDW UK). According to the source probe details the

IP address is in the same subnet. The RTT value of 0.202 ms provides a target location area of 20 km although some of this could be caused by networking delays. 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 gives a further target location area of 20 km.

In 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 in 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. This is not just simply a case of using RTT values but combining them with DNS records and database records to obtain a fuller picture. Additionally, the DNS data at this hop would suggest that the previous 2 hops were also at Redhill. The hop 4 DNS name shows that we are still in Redhill, but the 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 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 that hop 6 is still located at the same facility as hop 5 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 assigned prefix range of LINX LON1 IX 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), organisation, 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 assume 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 5.1. The traceroute timings are listed in Table 5.2.

| Hop | RTT (m/s) | IP Address | DNS |
|---|---|---|---|
| 1 | 0.925 | 5.57.16.70 | |
| 2 | 2.25 | 195.66.227.9 | edge1-rbsxe-8/1/0 |
| 3 | 3.362 | 153.92.43.236 | Fo09-ibr01-drtlon.ukcdw.com |
| 4 | 3.395 | 153.92.43.201 | Te-2-0-1-core01-drtred.ukcdw.com |
| 5 | 3.203 | 153.92.43.247 | Fo-0-ibr01-drtlon.ukcdw.com |
| 6 | 2.474 | 185.74.25.254 | External-dcfwcluster.uk.cdw.com |
| 7 | unknown | 185.74.25.254 | |
| 8 | 2.53 | 153.92.43.249 | |

Table 5.2: Reverse Traceroute Measurement Table

Hop 1 provides little information on 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. By 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 10 Gb 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 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 travelled to the Redhill 1793 Interxion facility, adding another VP to the table. The DNS name of Hop 5 shows that the packet is still in Redhill but has now moved to an edge router, adding another VP to our VP table, as shown in Table 5.3.

Hop 6 has the DNS name of 'external-dcfw-cluster.uk. cdw.com', which does not provide any clues regarding its location. The IP address of hop 7 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 5.2. 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.

In Figure 5.2, colour coding is used only as a visual cue to denote the approximate transmission speeds 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. However, it should be noted that each router may prioritise ICMP packets differently depending on their target, and timings can also suffer from packet forwarding decisions, circuitous routes, different router configurations,

| IP Address | Facility | Longitude | Latitude | DNS | Port | Speed |
|---|---|---|---|---|---|---|
| 185.74.25.250 | 1793 | 51.2476 | -0.1571 | be-1-ibr01-drt-red.uk.cdw.com. | none | Unk |
| 153.92.43.246 | 1793 | 51.2476 | -0.1571 | fo-4-0-5-core01-drt-red.uk.cdw.com. | none | Unk |
| 153.92.43.200 | 40 | 51.4998 | -0.0107 | te-2-0-1-core01-drt-lon.uk.cdw.com | none | Unk |
| 153.92.43.237 | 40 | 51.4998 | -0.0107 | fo-0-0-0-20-ibr01-drt-lon.uk.cdw.com. | none | Unk |
| 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-drt-lon.uk.cdw.com. | none | Unk |
| 153.92.43.201 | 1793 | 51.2476 | -0.1571 | te-2-0-1-core01-drt-red.uk.cdw.com. | none | Unk |
| 153.92.43.247 | 1793 | 51.2476 | -0.1571 | fo-0-0-0-20-ibr01-drt-red.uk.cdw.com | none | Unk |
| 185.74.25.254 | 1793 | 51.2476 | -0.1571 | external-dcfw-cluster.uk.cdw.com | none | Unk |

Table 5.3: Vantage Points Table

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 due 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 5.2, 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 additional 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.

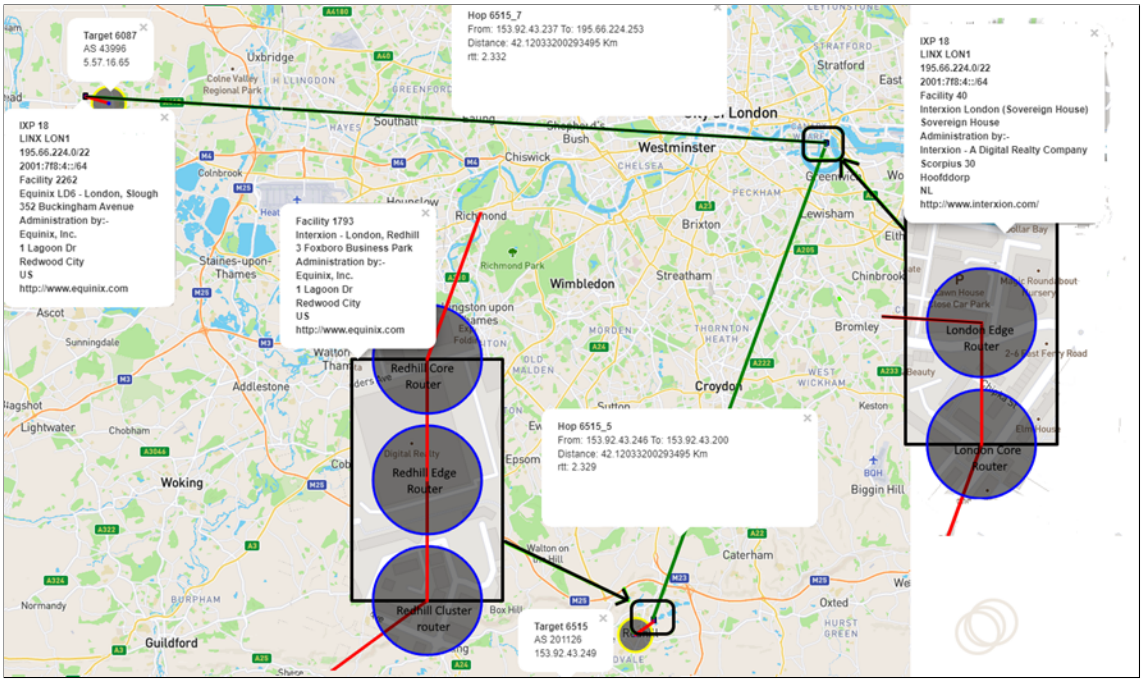


Figure 5.2: Example of UK Infrastructure Mapping incorporating LINX LON1 Internet Exchange Point.

The reverse measurement from RIPE Probe 6087 to RIPE Probe 6515, shown in Figure 5.1, follows the same path in this case, but it is highly likely that

many measurements will follow alternative reverse paths. In fact, 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 and 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.

To improve confidence in the results, additional validation of the IP address location within the VP Table could involve contacting the various facilities or AS organisations to confirm the location. 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 at 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. The location of Hop 2 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 through a private network, perhaps a VPN or a point-to-point connection. The IP address of hop 3 is 185.74.25.250, while the egress IP address of hop 2 is 185.74.25.254, indicating that it is on the same LAN or possibly on a point-to-point link. The RTT difference in timing between hop 2 (0.267 ms) and hop 3 (1.067 ms) 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 one 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 5.4, 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 the facility or the company confirmed the validation. 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.

| IP Address | Facility | Longitude | Latitude | Conf | DNS | Port | Speed |
|---|---|---|---|---|---|---|---|
| 185.74.25.250 | 1793 | 51.2476 | -0.1571 | 2 | be-1-ibr01-drt-red.uk.cdw.com. | none | Unk |
| 153.92.43.246 | 1793 | 51.2476 | -0.1571 | 2 | fo-4-0-5-core01-drt-red.uk.cdw.com. | none | Unk |
| 153.92.43.200 | 40 | 51.4998 | -0.0107 | 2 | te-2-0-1-core01-drt-lon.uk.cdw.com | none | Unk |
| 153.92.43.237 | 40 | 51.4998 | -0.0107 | 2 | fo-0-0-0-20-ibr01-drt-lon.uk.cdw.com. | none | Unk |
| 195.66.224.253 | 2262 | 51.5243 | -0.6380 | 2 | None | edge5-eq4 xe-3/0/3 | 10G |
| 195.66.227.9 | 40 | 51.4998 | -0.0107 | 2 | None | edge1-rbs xe-8/1/0 | 10G |
| 153.92.43.236 | 40 | 51.4998 | -0.0107 | 2 | fo-4-0-5-core01-drt-lon.uk.cdw.com. | none | Unk |
| 153.92.43.201 | 1793 | 51.2476 | -0.1571 | 2 | te-2-0-1-core01-drt-red.uk.cdw.com. | none | Unk |
| 153.92.43.247 | 1793 | 51.2476 | -0.1571 | 2 | fo-0-0-0-20-ibr01-drt-red.uk.cdw.com | none | Unk |
| 185.74.25.254 | 1793 | 51.2476 | -0.1571 | 3 | external-dcfw-cluster.uk.cdw.com | none | Unk |

Table 5.4: Vantage Points Table with Confidence Column

## 5.5 Automating the Method

Trialling of the automated geolocation methods is done within the same geographical area of London and South England that was used in the earlier manual geolocation efforts. This decision serves two important purposes.

First, it allows for a direct and meaningful comparison between the manual and automated methods by ensuring consistency in the test region and dataset. By focusing on the same geographical area, we can evaluate the effectiveness and accuracy of the automated methods relative to the manually obtained results, thereby providing a clear measure of their performance.

Second, the choice aligns with the focus of the original sponsorship proposal, which emphasizes research within the UK. Concentrating on the UK not only fulfills the requirements of the sponsorship but also enables the research to address region-specific characteristics of the Internet infrastructure, such as the density and configuration of IXPs and colocation facilities, which are critical to understanding and improving geolocation techniques.

By combining these considerations, the selection of London and South England as the trial area provides both a practical and strategic framework for evaluating and advancing the proposed automated geolocation methods.

In this section, we describe the automation of the 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 on a traceroute are first passed through a filter that tests the hop results against five assumptions (or rules).

### 5.5.1 Gateway Router Location - Rule 1

#### 5.5.1.1 First Hop

If this is the first hop, it is assumed that the first router encountered will probably be the gateway router for the source probe. This may or may not be in the same location as that of the probe. A gateway is a network node used in telecommunications that

connects two networks. Gateways serve as an entry and exit point for a network, as all data must pass through or communicate with the gateway before being routed. In most IP-based networks, the only traffic that does not go through at least one gateway is traffic flowing among nodes on the same local area network segment. In this method, we class the gateway router as being the first router encountered and therefore only need to apply this rule to the first hop. Rule 1 applies only to hop 1 and performs an initial sanity check.

#### 5.5.1.2 Sanity Checks

To provide a sanity check, a test is conducted to discover whether the RTT to this router is less than 1 ms; if less than 1 ms, then further tests are now applied. This is the only use of the delay-distance model. 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 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 performed. 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.

#### 5.5.1.3 Multiple or No Facilities

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 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 the IP address of this interface is performed. Any resulting DNS address is again filtered 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 United Kingdom town or city names where facilities are known to be located.

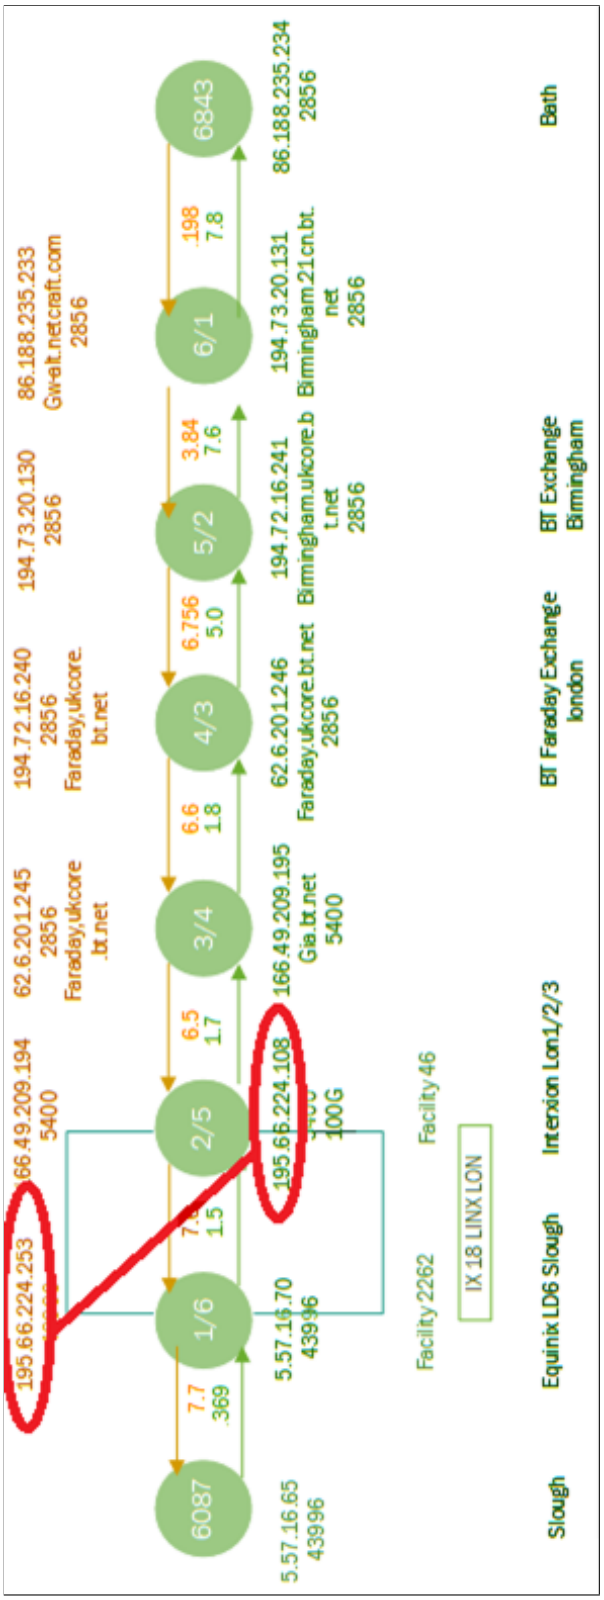


Figure 5.3: Forward and reverse traceroutes from IP address 5.57.16.65 to 86.188.235.234, showing outgoing and incoming IP addresses in the same prefix range.

#### 5.5.1.4 Example

Figure 5.3 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 has an RTT of 0.369 ms, 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.

#### 5.5.1.5 Database Searches

A search of the LINX IXP membership database shows that the 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 5.3 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 geocoordinates of Equinix Facility LD4 can be used in our Vantage Points table.

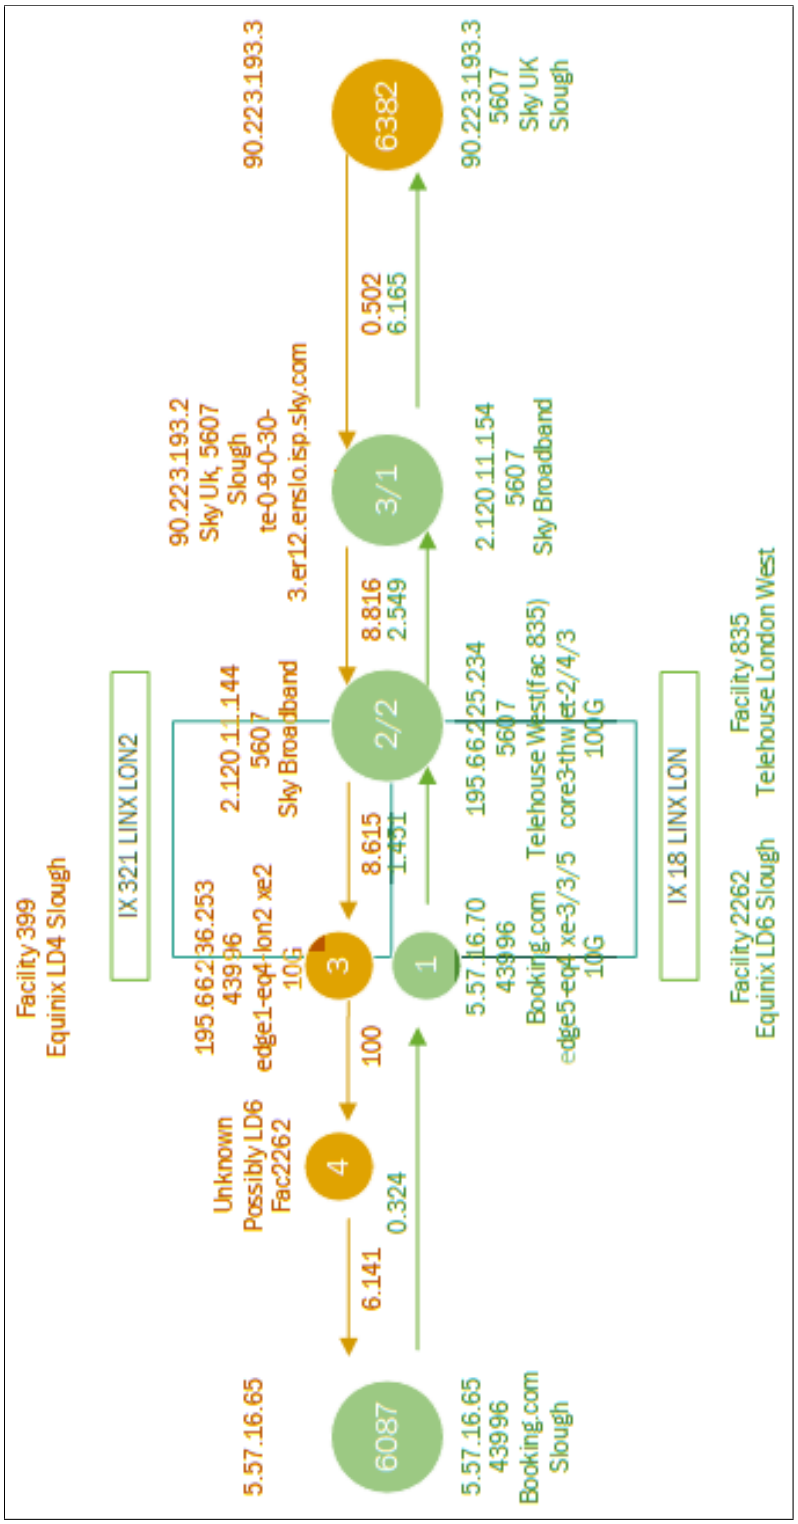


Figure 5.4: Forward and reverse traceroutes from IP address 5.57.16.65 to 90.223.193.3, showing incoming penultimate IP address blocked and outgoing IP addresses at 195.66.225.234. Also showing asynchronous forward and reverse routes using different Internet Exchanges on the forward and return legs.

#### 5.5.1.6 Further Verifications

We need to compare this traceroute with that shown in Figure 5.4, 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 determine 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 (IXP 18 LINX London) and reverse routes (IXP 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. 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 5.5.

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 0.198 ms 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 city or town 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.

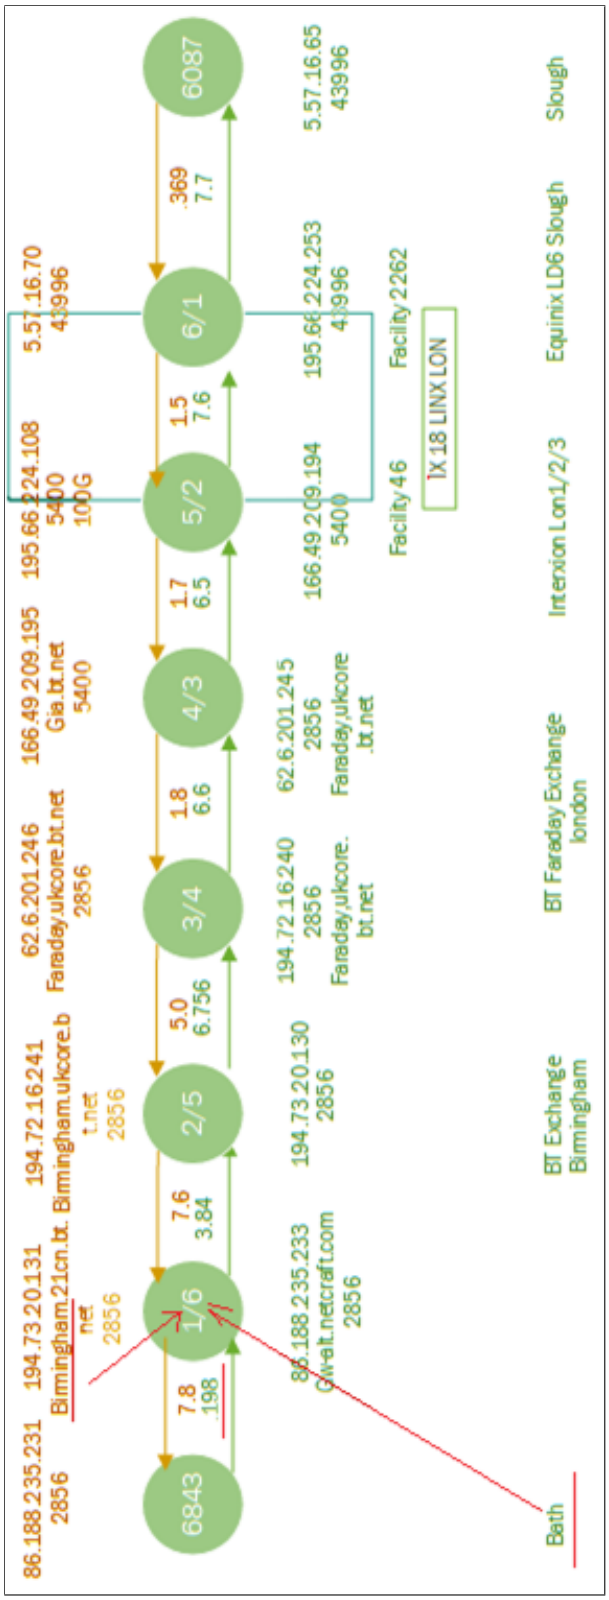


Figure 5.5: Forward and Reverse traceroutes from IP address 86.188.235.231 to 5.57.16.65, showing false positive in finding IP address location.

### 5.5.2 Private IP address - Rule 2

All IPv4 addresses can be divided into two main groups: global (public, external), which are those used on the Internet, and private (local, internal), which are those used in LANs and can also be used in VPNs, 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. 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 provides 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 manner.

### 5.5.3 Target IP address - 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. if something other than a RIPE Atlas probe has been used as a target then a reverse DNS lookup can be carried out to discover any further information as to the location and perhaps provide some verification of the results returned from the RTT values.

### 5.5.4 IXP Test - Rule 4

This rule determines whether the IP address of a hop is registered on an Internet Exchange. A test is conducted on each hop's IP address to determine if it falls within any IX-registered prefix. The initial design would then look at the previous hop to discover the ASN entering this IXP and compare it with a list of IXPs and their ASN peers from PeeringDB in an attempt to discover the sole facility where

the IXP and previous ASN are peers. Therefore, it would provide an entry facility. The same method is then applied to the existing facility using the successive hop ASN. However, finding the entry facility in this manner is unnecessary, as the IXP entry facility is not as important as the exit facility, which can be used more easily 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 process in which it looked at the successive hop to discover the ASN exiting this IXP and compared it with a list of IXPs and their ASN peers in an attempt to discover the only facility in which the IXP 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 (IXPDB, 2024). The IXPDB database also integrates the data from third-party sources. The website provides a comprehensive and corroborated view of the global landscape of interconnections. The combined data can be viewed, analysed, and exported through a web-based interface or API.

It is this IXPDB database that is now initially interrogated in an attempt to discover whether the IP address of a hop is registered on an Internet Exchange. It was discovered that the IXPDB database is not as comprehensive as advertised; some IXPs mark their IP to facility information as private, so no information is uploaded to IXPDB. One of these IXPs 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 method previously discussed of comparing the successive ASN with the known peers of each Internet Exchange. This is described in more detail in Chapter 4.

### 5.5.5 DNS Lookup - Rule 5

If no other rule is applicable, then a DNS lookup is made on the IP address of this hop, and any resulting DNS address is passed through a series of REGEX search strings to extract the city name or town 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 IP address is successfully located. 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 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 location of the IP address; therefore, we are unable to add this IP address to our Vantage Point table.

Referring back to our previous example in Figure 5.1, 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 location of this particular hop. If we examine the reverse traceroute, we can see that the reverse hop provides us with a DNS address 'External-dcfwcluster.uk.cdw.com'.

Combining this with our rules regarding sanity checks and incoming/outgoing IP addresses on the same subnet range as explained in Rule 1, Section 5.5.1.2, 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 traceroute is indeed the same router, and therefore we can geolocate this 10.255.255.2 and the reverse 185.74.25.254 IP address to Redhill.

### 5.5.6 Automated Results

To demonstrate the 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, Each UK probe is used as a target in turn with the rest of the UK probes acting as sources. As RIPE Atlas creates each measurement towards a target it assigns a measurement number which must be recorded in order to read the results at a later time see Figure 5.6. Additional data is recorded such as the IP address of each probe and its geographical coordinates which are store within the RIPE ATLAS platform.

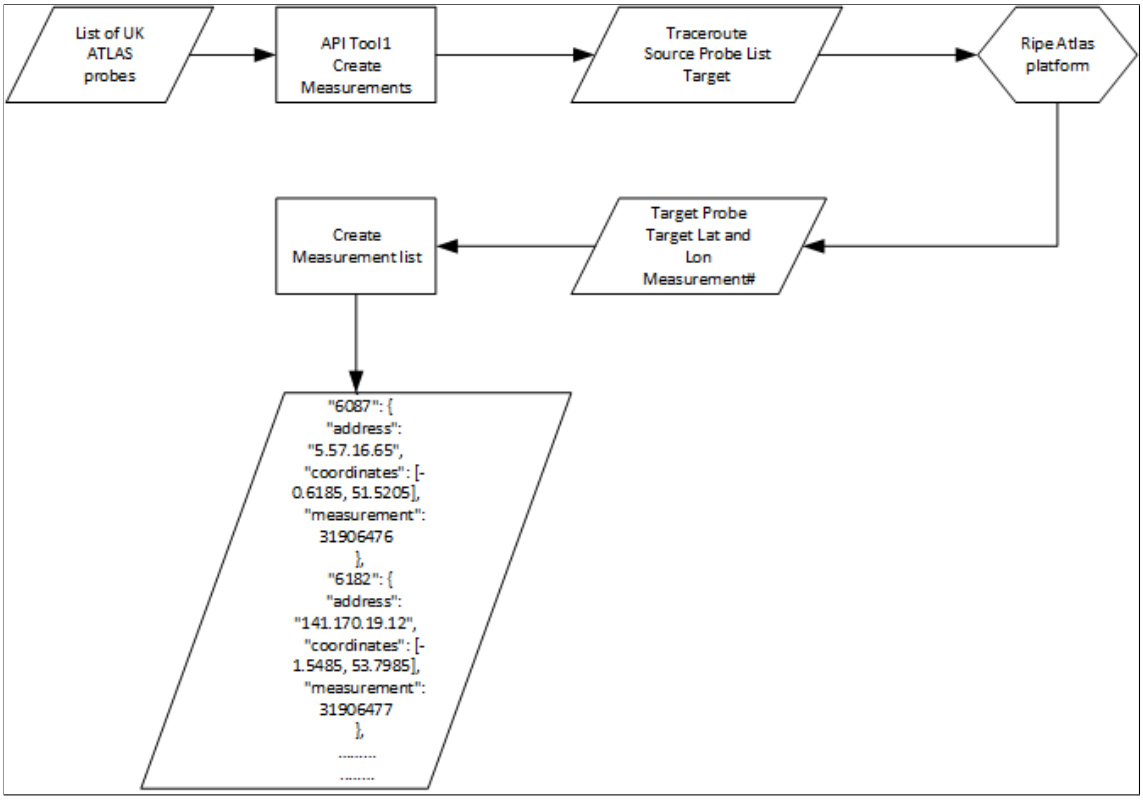


Figure 5.6: Measurement Creation API Tool showing the creation of RIPE ATLAS measurements and recording the measurement IDs for later use

Once RIPE ATLAS completes the requested measurements, the Read-Measurements Script reads the measurements from it and creates a local JSON file, see Figure 5.7

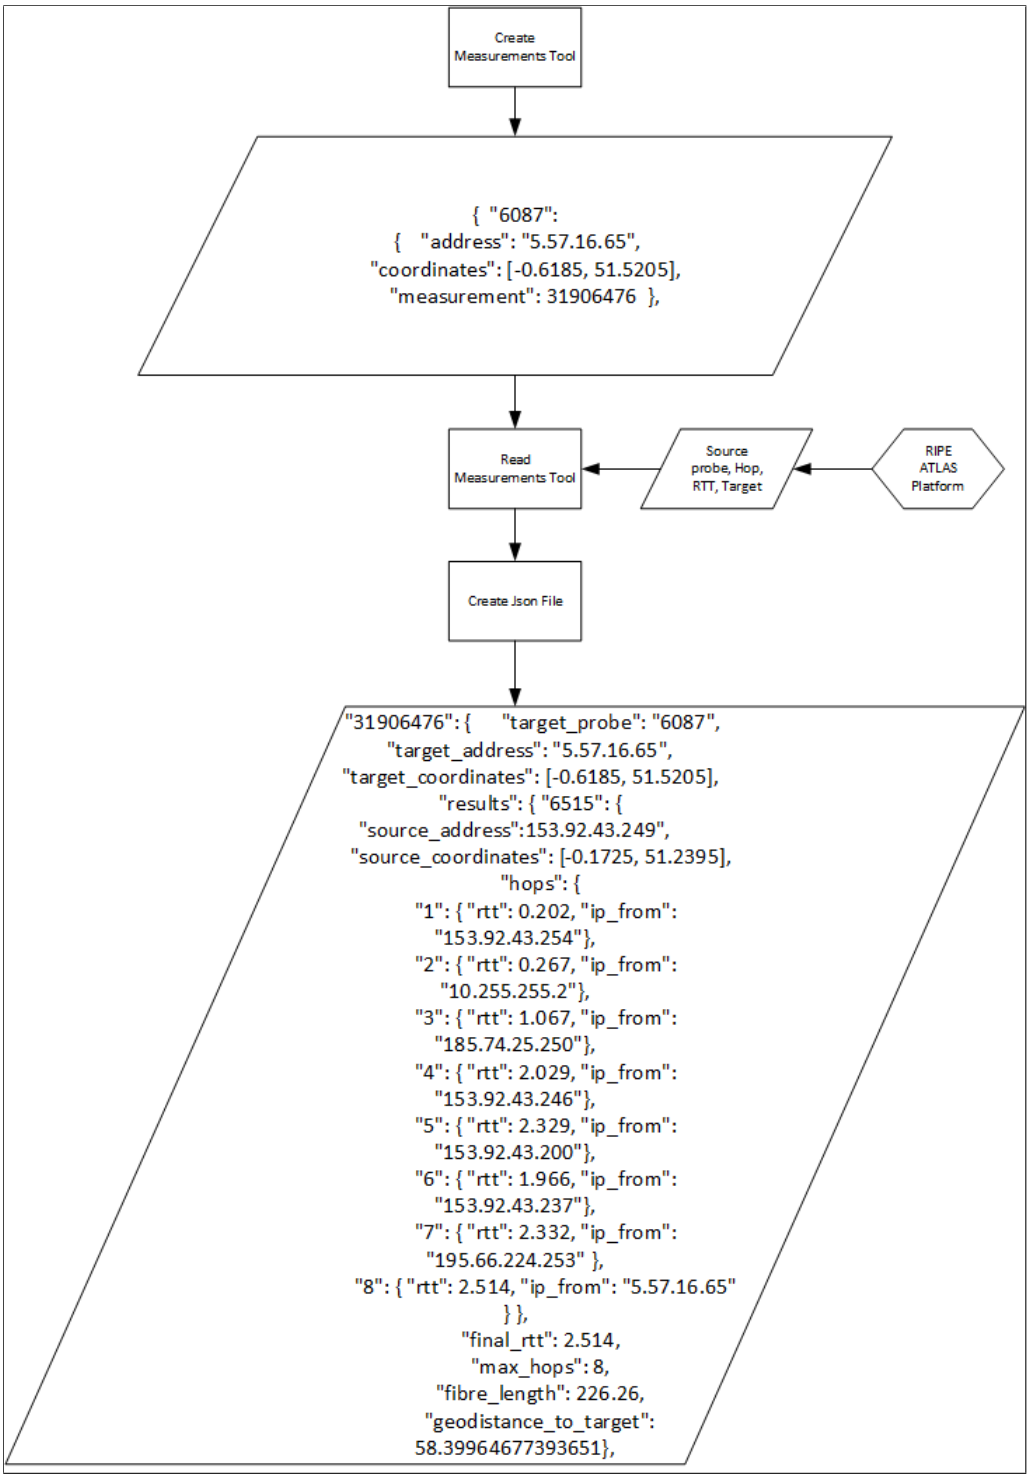


Figure 5.7: Read-Measurements Script showing the reading in of RIPE ATLAS measurements using the previous list of measurement numbers and recording each measurements routing information.

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 output. The fourth optional tool maps these results to an OpenStreetmap, as shown in Figure 5.2.

The results of applying the described rules and methods through an automated process are presented in Table 5.5. This table also shows the effect of the 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 IXP datasets become available, the 'Common_Fac' method will not need to be used.

| **Rules and Methods** | **Before Add LINX Data** | | **After Add LINX Data** | | **After Add LONAP Data** | |
|---|---|---|---|---|---|---|
| | Failures | Successes | Failures | Successes | Failures | Successes |
| 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 | 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 5.5: Rules and Methods Success and Failure Table

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.

The rules described in 5.5 use these methods as described below.

1. The Gateway Router Location Rule 1 uses the following methods:

    - Reverse DNS
    - REGEX
    - Reverse Traceroute
    - Reverse DNS of the IP address returned from the Reverse Traceroute
    - Regex of the DNS returned from the Reverse Traceroute.

2. The Private IP address - Rule 2 only carries out a test to see if the IP address lies within a private IP address range and does not use any of the methods.

3. The Target IP address - Rule 3 uses the following methods:

    - Reverse DNS

4. The IXP Test - Rule 4 uses the following methods:

    - Common Facility later made mostly redundant by Facility to IP Table

5. The DNS Lookup - Rule 5 applies to all hops that have not fallen into one of the rules above. It is similar to Rule 1 without the initial Delay-Distance test and uses the following methods:

    - Reverse DNS
    - REGEX
    - Reverse Traceroute
    - Reverse DNS of the IP address returned from the Reverse Traceroute.
    - Regex of the DNS returned from the Reverse Traceroute.

It should be noted that many of the IP addresses were tested multiple times using different traceroutes; therefore, a much larger number of successes and failures occurred compared with the discrete number of IP addresses. The first two columns of Table 5.5 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 IXP 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 1,190 traceroutes used in the test, 1,047 individually discrete IP addresses were discovered, 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 prevents them from being 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. By analysing the individual contribution of each geolocation rule described in the previous section, we arrive at the following observations:

- 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 5.6 summarises these statistics according to geolocation rules. If we rule out the 26 private IP addresses from our formulae because private addresses only provide 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 success rate of 33.6%.

| | |
|---|---|
| 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 because they were anycast) |
| TOTAL | 372 (3 excluded due to being anycast) |

Table 5.6: Rules Total Successful Geolocations

Luckie et al. downloaded 1.39 million IP addresses from the CAIDA ITDK 2020/2021 datasets. It was discovered that 220,000 had geohints and from these 220,000 they geolocated 183,000 which works out at 13% of the original 1.39 million. Although our initial dataset contains only 1047 discrete IP addresses, the location of 369 of them represents 33.6%, which is a significantly higher percentage, admittedly a smaller, more dense dataset.

## 5.6 Discussion

Relatively little research has been done on geolocating IP addresses to the facility level. It is claimed that researchers were able to geo-locate 13% of IP addresses at the city level when comparing their data set against the CAIDA ITDK data set (Luckie, Dhamdhere, Huffaker, and Clark, 2016). When trying to compare our data set with the CAIDA ITDK datasets, it was only possible to compare city locations because the CAIDA data set is based on city-level resolution rather than facility level. Although the CAIDA dataset has 1.7 million UK nodes, our dataset has 492 overlapping IP addresses for comparison. Of the 1,047 distinct IPs in our dataset, 288 were geolocated at the city level by CAIDA and 369 at a facility level by our method; this includes 89 new IPs that 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.

The solution proposed in Chapter 5 uses a confidence level mechanism to provide some idea of the precision of the methods explained in this thesis. Of the 1190 traceroutes employed in the test, 1021 individually discrete IP addresses were found, of which 343 were geolocated using the procedure and methods described. This gives a success rate of 33.6% in geolocating the 1021 IP addresses to a confidence level of 3 or higher where a 1 is fully confident, a 2 is probable, a 3 is likely and a 4 is a "best guess", as discussed in Section 5.4.1

A REGEX filter is one of the main components of this solution. First, the

DNS name of each IP address is discovered, and then it is entered into a REGEX script in an attempt to discover the town or city where it is located. Although several researchers have already worked on this issue, such as Luckie et al. (Luckie, Huffaker, Marder, Bischof, Fletcher, and Claffy, 2021) and Dan et al. (Dan, Parikh, and Davison, 2021a), a new solution has been created and reported here, specifically designed for the UK Internet infrastructure. The limited number of UK towns 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; an investigation of the 3000+ failures of this REGEX technique would lead to more comprehensive results. However, many of these failures will be due to The machine learning techniques of Dan et al. and the learning of geographic naming conventions from Luckie et al. could also significantly improve these results.

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. When none of the previous rules apply, Rule 5 uses a set procedure to try to narrow down the IP address location:

1. Attempt to discover the DNS address of this IP.
2. The DNS address is processed 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 other DNS hints, sanity checks, ASN lookups, and RTT values.
4. If no facilities were found, perform a reverse traceroute, if available, and subject any reverse DNS to the same REGEX filter.
5. Failing all of this, carry out a Common Facility comparison to determine the facilities at which the previous ASN and current ASN interconnect.

If none of these methods proved successful, we were unable to locate the IP address.

In Step 2, DNS parsing uses a regular expression script similar to that developed by Luckie et al. (Luckie, Huffaker, Marder, Bischof, Fletcher, and Claffy, 2021) and Dan et al. (Dan, Parikh, and Davison, 2021a). 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 knowledge of the local infrastructure. 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 exchange 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 located 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.londen12.uk.bb.gin.ntt.net'. Faelix has identified facilities with names, such as AEBI, its full DNS name is '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 challenges described in Chapter 2.3, there are still limitations that need to be addressed. ICMP echo packets are often treated as second-class by routers and target hosts. This means that ICMP echo requests and responses may have a lower priority than traffic, which is considered more important. The end result indicates that the round 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 used, where the IP address and geolocation are already known to be used as sources and targets to initiate 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 more than 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 data set of more than one million VPs from which future research on IP geolocation can be based. This figure was achieved by starting with the fact that there are over 600 RIPE Atlas probes in the UK and in a fully connected mesh network, each probe acts as a source and sends traceroutes to all other probes as targets. The number of traceroutes required to achieve this full mesh is calculated using the formula :

$$\textit{Number of Traceroutes} = \textit{Number of Probes} \times (\textit{Number of Probes} - 1)$$

Substituting 600 Probes:

$$\textit{Traceroutes} = 600 \times (600 - 1) = 600 \times 599 = 359,400$$

On average, each traceroute discovers 1–10 unique IP address/geocoordinate pairs (referred to as vantage points or VPs).

Using this range, the total number of VPs discovered would be :

$$\textit{Total VPs} = \textit{Number of Traceroutes} \times \textit{Average VPs per Traceroute}$$

Using the lower and upper bounds for discovery:

$\textit{Total VPs (minimum)} = 359,400 \times 1 = 359,400$ $\textit{Total VPs (maximum)} = 359,400 \times 10 = 3,594,000$

Taking a conservative estimate of approximately 1 million VPs based on the mid-range discovery rate (e.g., an average of 3 VPs per traceroute):

$$\textit{Estimated VPs} = 359,400 \times 3 \simeq 1,078,200$$

In addition, it should be noted that the IXPDB dataset is seemingly 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.

### 5.6.1 OpenstreetMap

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, routing of data, geographical data. OpenStreetMap and other products can now convey information that can then be rendered in 3D, and it would be useful to investigate the efficacy of these products compared to this 2D view.

### 5.6.2 High Latency issues

In Section 3.6 we discovered a high-latency issue with the RIPE probes in and around the London area. In order to discover exactly where this latency is occurring it would be useful to be fully in control of the source and target probes at each end of a traceroute to fully analyse where the delays are occurring. The installation of additional RIPE anchors at other universities such as London, Edinburgh, and Bristol would allow end-to-end latency tests where every aspect of a traceroute between two RIPE Atlas probes can be thoroughly investigated.

### 5.6.3 IPv4 Geolocation Future work

Future work could involve creating the same traceroute measurements over extended periods and using different routes, adding alternative hops as backup paths, or finding completely new paths, allowing new infrastructure details to be realised and more 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 which involved finding the common peers of an IXP 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.

## 5.7 Transition to IPv6 and Future Considerations

The new IPv4 geolocation methods and tools significantly enhance the granularity and precision of Internet infrastructure mapping. These advances address many existing challenges and offer promising solutions for IPv4-based networks. However, the Internet is now rapidly transitioning to IPv6, driven by the need for a larger address space and improved functionality.

Given the limitations of IPv4 and the impending global shift towards IPv6, it is imperative to explore how our geolocation techniques can be adapted and enhanced for IPv6. IPv6 not only provides a vast address space, but also introduces new features and capabilities that can be leveraged to improve geolocation and routing security. The thesis will now focus on IPv6 and the improvements that this can offer in IP Geolocation.

# Chapter 6

# IPv6 - The Future for IP Geolocation

## 6.1 Introduction

Accurate maps of the Internet infrastructure would allow network engineers and operators to improve and optimise the allocation of network resources such as backup routes, routers, proxies, replica servers, and data centres. Detailed and complete maps of the Internet's topology, annotated with the geographic locations of network equipment, could help with the study of a wide range of security-related problems and protocols. Locating the source of malicious traffic or assessing the vulnerability of the Internet to blackouts or attacks on parts of its physical infrastructure would inform network planners of the best possible recovery solutions.

Over the years, Traceroute, Ping, and the Border Gateway Protocol (BGP) have been used by many researchers to discover the location of routers and hosts on the Internet. Traceroute and Ping are primarily diagnostic tools used to measure the route path and latency to a destination IP address, whilst BGP is a routing protocol designed to manage how packets are commercially routed across the Internet, none of which were designed for determining geographical locations. In fact, the uses of these tools are described as merely 'engineering hacks' proposed by researchers to

collect information about the Internet topology (Motamedi, Rejaie, and Willinger, 2015).

The accuracy of IP-based geolocation can be affected by many factors such as circuitous routes, dynamic IP address allocation, the use of proxy servers or VPNs, and layer 2 clouds such as MPLS and ATM that are generally opaque to a traceroute.

**Roadmap** Section 6.2 sets the aim of creating a superior IPv6 mapping tool, focusing on a detailed visualisation and security implications. Section 6.3 explains the necessity for a new geolocation tool. Section 6.4 gives some background on Node Information Queries. In Section 6.5, the document proposes essential changes to the Node Information Protocol, router kernels, and installation processes to enable this advanced mapping capability, enhancing data accuracy and security. Section 6.6 describes the operational design of the tool, including iterative queries for comprehensive network mapping. In Section 6.7 the challenges of introducing a new tool are investigated along with the barriers to deployment and potential solutions. Section 6.8 outlines measures to safeguard the tool's use, addressing potential risks and configuration options. Section 6.9 suggests various methods of testing the proposed solution. Section 6.10 carries out a technical evaluation, an operational evaluation, a security and privacy evaluation, and an economic and regulatory evaluation. Section 6.11 provides a feasibility and impact analysis study. Section 6.12 looks at alternative approaches to this issue. Chapter 6.13 provides a short summary and Chapter 6.14 starts a discussion on segment routing and some ideas to improve security which could be used with these new methods

## 6.2 Objectives and Scope of this Chapter

The primary objectivee of this chapter is to introduce a new tool for mapping IPv6 network infrastructure, with the aim of surpassing current methodologies in accuracy, comprehensiveness, and utility. Other objectives include detailing the necessity for such a tool, discussing enhancements over existing technologies, outlining the proposed changes to network protocols and infrastructure, and demonstrating the tool's potential benefits for network analysis and troubleshooting.

The work described here is intended to offer a more detailed visualisation of the Internet's structure, addressing both technical implementations and the broader implications for network research and administration. This chapter also attempts to address many of the related security issues; however, these range in complexity from simple to highly complex (Motamedi, Rejaie, and Willinger, 2015) as explained in section 6.7.

## 6.3 Necessity for a new purpose-built IP Mapping tool

The mapping of the network infrastructure has long been essential for optimising and understanding the functionality of the Internet. Traditional IPv4 geolocation methods typically involve databases that map IP addresses to geographical locations using data from various sources, such as traceroute results, DNS lookups, and user reports. However, many of these methods often rely on delay-distance formulas to estimate the geographic location of IP addresses based on network latency measurements. Such formulas have been shown to be seriously inaccurate due to various factors, including network congestion, routing asymmetries, and circuitous paths that introduce errors in latency-based estimates.

Despite their utility, these methods have shown limitations and inaccuracies in mapping Internet infrastructure, producing errors ranging from small inaccuracies to significant deviations. Studies have demonstrated these shortcomings, highlighting the need for more accurate and reliable mapping techniques. Additionally, traditional IP geolocation methods, which rely on databases that map IP addresses to geographical locations, often suffer from outdated or inaccurate data, further complicating precise infrastructure mapping.

The emergence of IPv6 offers a promising avenue to improve the accuracy of geolocation. IPv6 provides enhanced capabilities over IPv4, allowing for the introduction of a more precise mapping of the Internet infrastructure. The greater functionality of IPv6 enables new geolocation methodologies that can overcome the limitations of previous tools by providing more detailed and accurate network topology information. Leveraging IPv6's capabilities, a new geolocation tool can provide a more accurate and up-to-date view of Internet topology, ensuring better routing decisions and improved network performance. Such an innovative approach could represent a significant advancement in geolocation technology, addressing previous limitations, and offering a robust solution for the mapping of the Internet

infrastructure.

## 6.4 Background on Node Information Queries

This chapter proposes a brand new method for discovering the network topology by designing a new network mapping tool that is based on the IPv6 Node Information Queries protocol described in RFC 4620 (Crawford and Haberman, 2006). The Internet Engineering Task Force (IETF) developed IPv6, which was intended to replace IPv4 and would deal with the long-anticipated problem of IPv4 address exhaustion. In December 1998, IPv6 became a draft standard for the IETF. The KAME project was launched in 1998 and concentrated on the research and development of IPv6 technologies, succeeding in the global standardisation of basic IPv6 specifications and establishing the framework required for the commercial marketing of IPv6 technologies (Kudou, Suzuki, Hagino, Yamamoto, Shima, Uehara, Wakikawa, Mitsuya, Momose, Jinmei, and S., 2006). The project included the development of a protocol for asking an IPv6 node to provide certain network information, such as its hostname, IP addresses, or fully qualified domain name. In addition, a direct query mechanism for other information has been found to be useful in serverless environments and for debugging (Crawford and Haberman, 2006). This chapter emphasises the need for a more accurate and resilient mapping tool given the inadequacies of Traceroute, Ping, and BGP to capture the full complexity of the Internet topology.

## 6.5 Proposed Changes to Node Information Queries

### 6.5.1 Overview

This thesis proposes an extension to IPv6 Node Information queries as specified in RFC 4620 (Crawford and Haberman, 2006). Currently, the mechanism can be used to learn the addresses and names of nodes and is also useful when there are no global

routing or DNS name services available.

According to RFC 4620 A 'Node Information Query' (NI Query) message is sent by a querier node to a responder node in an ICMPv6 packet addressed to the queried address. Currently, the query can contain a Subject Address or a Subject Name. The responder sends a 'Node Information Reply' (NI Reply) to the querier, containing information associated with the node at the Queried Address. Both NI Query and NI Reply have the same format (see Figure 6.1) and are carried in ICMPv6 packets.

| Offsets | Octet | 0 | | | | | | | | 1 | | | | | | | | 2 | | | | | | | | 3 | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Octet | Bit | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
| 0 | 0 | Type = 139/140 | | | | | | | | Code | | | | | | | | Checksum | | | | | | | | | | | | | | | |
| 4 | 32 | Qtype | | | | | | | | | | | | | | | | Flags | | | | | | | | | | | | | | | |
| 8 | 64 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 12 | 96 | Nonce | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 16 | 128 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 20 | 160 | Data | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 24 | 192 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 28 | 224 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 32 | 256 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Figure 6.1: Node Information Messages

The Type field is an 8-bit identifier and is an NI Query when set to 139 and an NI Reply when set to 140.

When the Type field is set to NI Query, the code field can currently contain the following options:

- 0 - Indicates that the Data field contains an IPv6 address that is the subject of this query.
- 1 - Indicates that the data field contains a name that is the subject of this query, or is empty, as in the case of a NOOP.
- 2 - Indicates that the data field contains an IPv4 address that is the subject of this query.

When the Type field is set to NI Reply, the code field can currently contain the following options:

- 0 - Indicates a successful response. The Response Data field may or may not be empty.
- 1 - Indicates that the responder refuses to supply the answer. The Response Data field will be empty.
- 2 - Indicates that the QType of the query is unknown to the responder. The Response Data field will be empty.

The Checksum field stores an IPv6 checksum.

The QType field is a 16 bit field that designates the type of information requested in an NI Query or supplied in a NI Reply. Currently, there are five values of QType specified.

- 0 - No Operation
- 1 - unused
- 2 - Node Name
- 3 - Node Addresses
- 4 - IPv4 Addresses.

The Flags field are QType specific flags that may be defined for certain queries and replies.

The nonce field is an opaque 64-bit field to help avoid spoofing and/or to aid in matching replies with queries. Its value in a query is chosen by the querier. Its value in a reply is always copied from the corresponding request by the responder. The nonce must be a random or good pseudo-random value to foil spoofed replies.

The data field in an NI query will currently contain the subject's address or name and contains the information specified by the QType in a reply. It is important to note that the length of the data may be inferred from the IPv6 header's payload length field as detailed in RFC8200, the length of the fixed portion of the NI packet,

and the lengths of the ICMPv6 header and intervening extension headers (S. Deering and R. Hinden, 2017).

### 6.5.2 Concept and Structure

To achieve the goals set out in this document, changes to the Node Information Query and Reply protocols, router's kernel, and equipment installation procedures are required.

#### 6.5.2.1 Node Information Protocol Changes

The QType field currently uses the options shown in Figure 6.2.

| Qtype | Qtype Name |
|---|---|
| 0 | NOOP |
| 1 | unused |
| 2 | Node name |
| 3 | Node Addresses |
| 4 | IPv4 Addresses |

Figure 6.2: QType options

It is proposed to make changes to the NI Query format by adding additional options to the QType field where:

- 5 - Indicates that the data field contains a list of the IP addresses of the node's peers.
- 6 - Indicates that the data field contains the node's geolocation.

#### 6.5.2.2 Router Kernel Changes

The proposed options will require some changes to a router's kernel. To provide the data for option 5 the routing table needs to be interrogated to provide the peer's IP addresses for each interface. Additionally, a new Access Control List (ACL) will need to be added, which has the ability to test and allow for Node Information

Queries, and for additional security, can also test for the source originator which may be from a list of trusted IP addresses.

#### 6.5.2.3 Router Installation Procedural Changes

To provide the data for QType option 6 it is proposed that upon initial installation of a router or middlebox, the geolocation data of the installed equipment is entered. For increased security, the actual location can be obfuscated by entering slightly offset data to protect privacy. An obfuscation offset similar to this is used in the RIPE ATLAS platform.

## 6.6 Network Mapping Tool

### 6.6.1 Introduction

A tool similar to traceroute could be developed using a node's peer IP addresses which would be returned using one of the new options described above from a Node Information Query. Altogether, this tool would make three queries to a given node:

1. Using one of the existing QType options (QType = 3 Node Addresses) Query the node for its set of interface IP addresses.

2. Using one of the proposed new QType options (QType = 6 Node geolocation) query the node for its geolocation.

3. Using one of the proposed new QType options (QType = 5 Node Peers) query the node for its peer IP addresses.

### 6.6.2 Network Mapping Example

Figure 6.3 shows an example of a user interrogating a node responder and receiving geolocation data, interface IPs and peer IPs which can be saved to a JSON file. The tool can then iterate over the adjoining nodes using the information in this JSON

file addresses to collect further information gradually, building up a comprehensive geographical network map. The JSON File shown in the diagram is returned from the Responder. The Json file contains the IP addresses of the responders peers, each of these peers can now be used as responders to gather further information and so on.

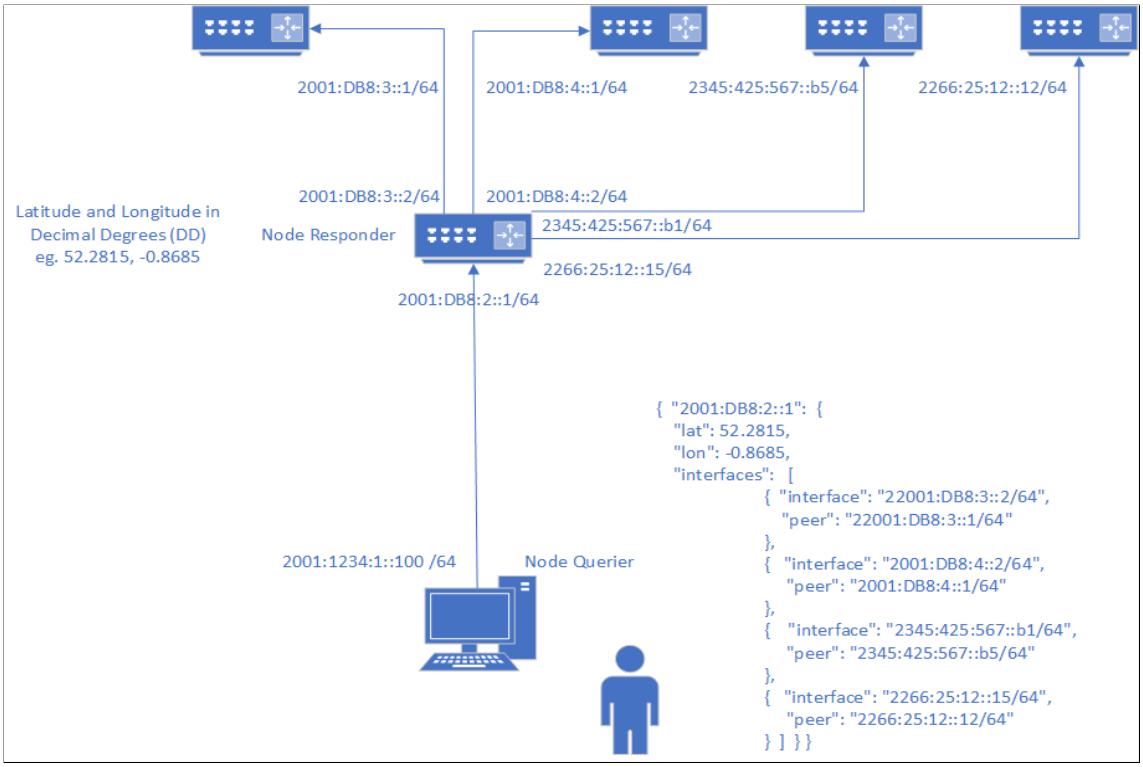


Figure 6.3: Tool Router Iteration

## 6.7 Barriers To Deployment

### 6.7.1 Overview

Mapping the Internet's infrastructure involves navigating technical, logistical, and ethical challenges such as improving the accuracy and completeness of geographically locating network infrastructure. New technologies such as dynamic routing, load balancing, proxies, and the use of content delivery networks (CDNs) obscure the true path data packets take, creating erroneous conclusions and making it difficult to map the network accurately. Privacy and security considerations must be taken into account when collecting and publishing detailed network topology information and must be balanced against the risk of exposing vulnerabilities to malicious actors.

Attempts to map the network infrastructure reflect the evolving complexity of the Internet and the continuous development of tools and methodologies to better understand it. Basic utilities like ping and traceroute to advanced protocols like BGP and IOAM have provided valuable insights into the Internet's structure and performance, even as the landscape continues to change. However, these tools are not designed to provide a complete and accurate map of the Internet infrastructure. A complete solution to this mapping problem has eluded researchers to date, and there are many complex barriers to deployment of this solution, technical, operational, economic, regulatory and geopolitical barriers to name just a few.

### 6.7.2 Technical Barriers

Many ISPs and organizations still use IPv4 or dual-stack implementations, limiting the widespread effectiveness of an IPv6-only solution. This is further discussed in Section 7.14.

The proposed QType extensions (QType = 5 for peer IPs, QType = 6 for geolocation) are not currently defined in RFC 4620, meaning no existing routers, operating systems, or network devices support them. This means that major router vendors (Cisco, Juniper, Huawei) would need to modify their software stacks and

firmware to accommodate these changes. We would need to work with major router vendors (e.g., Cisco, Juniper) to propose these changes through IETF working groups.

If every query for network geolocation, peer addresses, and interface information is processed in real-time, this could introduce significant CPU load and increase packet processing times on routers. Large-scale querying (e.g., thousands of queries per second across an IXP) could result in congestion and resource exhaustion.

#### 6.7.2.1 Potential Solutions

- **Unsupported IPv6 solution:** Implement the tool in dual-stack environments, allowing for fallbacks to conventional methods when IPv6 is unavailable.
- **Standardization Effort:** Submit an Internet-Draft to the IETF for review, outlining the benefits and security considerations of the extension. Also, we could engage with the IPv6 Maintenance Working Group (6man) within the IETF.
- **Router Vendor Collaboration:** Work with major vendors (Cisco, Juniper, MikroTik, Arista) to incorporate the feature in future firmware updates. Develop open-source implementations (e.g., Linux-based routers, FRRouting, Bird) to demonstrate feasibility.
- **Backwards Compatibility with Existing Systems:** If routers do not support the new extensions, they should still respond to basic Node Information Queries, ensuring a graceful degradation of functionality.
- **Rate Limiting and Query Throttling:** Implement rate limiting at the router level to prevent excessive queries or only allow authenticated and trusted queries (e.g., network operators, IXPs, security researchers) such as that described in Section 6.14.2.1.

- **Caching and Query Optimization:** Nodes could cache recent query responses to reduce redundant processing.

### 6.7.3 Operational Barriers

Network operators are generally risk-averse, preferring proven technologies over experimental ones and therefore ISPs, cloud providers, and IXPs who would need to support the proposed geolocation-aware routing mechanisms will need to be convinced of the benefits of this solution. The proposed method would require standardisation within the IETF to ensure interoperability across different networks. Other proposed solutions such as Geofeeds as proposed in RFC8805 are rarely used which suggests a general reluctance to adopt new location-based protocols. Many networks use BGP-based traffic engineering, which conflicts with per-packet geolocation-based routing. Operators would need to modify routing policies significantly, which could be expensive and time-consuming.

#### 6.7.3.1 Potential Solutions

- **Risk Averse Nature:** Demonstrate real-world benefits, such as improved traffic engineering, better DDoS mitigation, and enhanced compliance with local data laws.
- **Incremental Deployment Strategy:** Instead of requiring a full network overhaul, implement the solution as an overlay service.
- **Operators rely on BGP for business-driven routing:** Develop a hybrid model by combining both BGP and Geolocation
- **Need for IETF Standardization:** Early engagement with IETF, BGP integration, alignment with RPKI.

### 6.7.4 Economic Barriers

ISPs and backbone providers may hesitate to invest in new routing capabilities without clear financial incentives whilst deploying the proposed method at scale could require new router hardware, software updates, and staff training Unlike BGP, which serves direct business interests such as transit agreements and peering negotiations. Geopolitical-aware routing provides no direct revenue stream and therefore ISPs may not see an immediate business case to justify implementation. Some network operators and IXPs may not want to expose their infrastructure layout due to competitive or security reasons.

#### 6.7.4.1 Potential Solutions

- **Leverage Security and Compliance as Selling Points:** Highlight improved regulatory compliance for data sovereignty laws can reduce imposed financial penalties.
- **Pilot Programs with Select IXPs and ISPs:** Partner with small-to-medium-sized ISPs and IXPs for limited-scope testing.

### 6.7.5 Regulatory and Geopolitical Barriers

Data sovereignty laws vary by country, making geopolitical-aware routing highly complex. Some countries mandate local data storage and routing policies, which might conflict with user-defined routing preferences. Governments could exploit the technology to enforce censorship by restricting access to certain regions. The ability to route around warzones or sanctioned countries could be seen as a violation of existing regulations. Different jurisdictions have different approaches to data routing and privacy, making it difficult to enforce a unified standard. Some networks may refuse to honor geopolitical routing requests, rendering the system only partially effective. Revealing geolocation data at the router level may conflict with data privacy laws such as GDPR, whilst some network operators and IXPs may not

want to expose their infrastructure layout due to competitive or security reasons.

#### 6.7.5.1 Potential Solutions

- **Obfuscation and Granular Control Over Data Exposure:** Allow nodes to return approximate locations rather than exact coordinate and provide network operators with the ability to disable or restrict.
- **User Consent and Policy-Based Exposure:** Develop opt-in models where organizations choose what data to expose and allow different exposure levels such as city level or building level data.
- **Secure Data Handling and Compliance Audits:** Ensure that geolocation data handling complies with GDPR and industry best practices. Implement audit logs to track who queries geolocation data and for what purpose.

### 6.7.6 Security Barriers

The introduction of new query types could create new attack vectors, such as DDoS attacks, unauthorized data leaks, and man-in-the-middle (MITM) exploitation. Attackers could forge NI Queries to probe network infrastructure and discover sensitive routing details. ICMPv6 packets can be exploited for network discovery or resource exhaustion in DDOS attacks and amplification attacks. Malicious actors could flood routers with Node Information Queries, causing service degradation. Nation-state actors, hackers, or competitors could use the data to map infrastructure vulnerabilities. Attackers could route traffic through specific nodes based on geolocation, enabling interception, censorship, or data manipulation.

#### 6.7.6.1 Potential Solutions

- **Access Control Lists (ACLs) and Authentication:** Queries should only be answered for whitelisted IPs (e.g., within an ISP, an IXP, or a trusted security organization).

- **IPv6 Router Advertisement (RA) security mechanisms:** IPv6 includes features like RA Guard which protects from rogue routers (Gont, 2014). Also SEND that helps protect against rogue router advertisements and neighbor discovery attack (Arkko, Kempf, Zill, and Nikander, 2005).
- **Encryption and Authentication for Responses:** Encrypt responses to prevent interception or modification by attackers and require authentication (e.g., cryptographic signatures, TLS, or DNSSEC-like verification) for both queries and responses.

## 6.8 Security and Privacy Concerns

The following are legitimate uses of the extended options that the software tool adds to the ICMP Extended Echo functionality:

- To determine the geolocation of an IP address.
- to determine the peers to which a router connects.

However, malicious parties may be able to use the tool to obtain additional information such as:

- Interface bandwidth.
- The type of device that supports the interface (e.g., vendor identity).
- The operating system version that the device uses.

Understanding this risk, network operators often establish policies that restrict access to the ICMPv6 Extended Echo functionality. In order to enforce these policies, nodes that support ICMP Extended Echo functionality must support the following configuration options:

- Enable/disable the ICMP Extended Echo functionality. By default, the ICMP Extended Echo functionality is disabled.

- The L Bit decides whether Link-Local or Global addresses are used. Define the L-bit enabled settings. By default, the option to set the L bit is enabled and the option to clear the L bit is disabled.
- Define enabled query types (i.e., by name, index, or by address); by default, all query types are disabled.
- For each enabled query type, define the prefixes from which ICMP Extended Echo Request messages are permitted.
- For each interface, determine whether the ICMP Echo Request messages are accepted.

When a node receives an ICMP Extended Echo Request message that it is not configured to support, it must silently discard the message. In order to protect local resources, implementations should rate-limit incoming ICMP Extended Echo Request messages.

In order to foil spoofed reports, the nonce must be a random or good pseudo-random value. Where multiple independent processes are used to send NI queries, the nonce value may be used to deliver replies to the correct process. Each process must check the received nonce and ignore extraneous responses (Crawford and Haberman, 2006). The responder must return a NI Reply with ICMPv6 Code = 2 and no Reply Data if the QType is unknown. The responder should rate-limit such replies as it would ICMPv6 error replies and IP Security (IPsec) (Mühlbauer, Feldmann, Maennel, Roughan, and Uhlig, 2006) should be used where true communication security is required (Conta, S. Deering, and Gupta, 2006).

## 6.9 Testing and Evaluation

To assess the feasibility, effectiveness, and impact of the proposed IPv6 Node Information Query extension, a combination of simulation, controlled deployment, and real-world testing is necessary.

- **Simulation-Based Testing:** Tools such as NS-3 (Nsnam, 2025), or OMNeT++ (Omnet++, 2025) to model the protocol in a virtual IPv6 network. These could also simulate query and response interactions to evaluate latency, processing overhead, and network load. They would allow us to model large-scale deployments to assess query propagation, response delays, and congestion impacts.

- **Testbed Deployments:** Controlled environments like PlanetLab (University, 2025) or RIPE Atlas (RIPE, 2015) testbeds. Collaboration with IXPs and ISPs: Work with regional IXPs (e.g., LINX, LONAP) to assess real-world feasibility. Implement in software routers such as BIRD (Bird, 2024) and commercial routers supporting IPv6 extensions.

- **Security and Privacy Audits:** Evaluate resistance to spoofing, denial-of-service (DoS), and unauthorized queries. Ensure queries do not expose sensitive network topology information. Test ACL mechanisms to confirm only authorized queries receive responses.

- **Comparative Performance Analysis:** Compare against existing methods such as traceroute based geolocation, BGP based traffic engineering and commercial geolocation databases. Metrics may include geolocation accuracy, network overhead and scalability under load.

- **Real-World Case Studies:** Deployment in an IXP such as LINX or LONAP. Measure how well IXPs can integrate geolocation-aware queries into their network monitoring and routing policies. Examine how it impacts routing

decisions and traffic engineering in a real-world multi-AS environment. Deploy within a single ISP's infrastructure to test intra-network and inter-AS effects.

- **User Acceptance and Policy Review:** Gather industry perspectives on practical deployment concerns using feedback surveys. Submit an Internet-Draft to gauge community acceptance and interoperability requirements. Ensure compliance with GDPR and national security regulations regarding geolocation data usage.

## 6.10 Critical Evaluation

This solution introduces an innovative approach to addressing the challenges of network topology discovery. However, like any new protocol modification, it presents both significant advantages and notable challenges when deployed in real-world environments.

### 6.10.1 Technical Evaluation

The proposed extension to IPv6 Node Information Queries (Crawford and Haberman, 2006) introduces a direct query mechanism for geolocation and peer discovery, significantly improving the accuracy of Internet infrastructure mapping.

#### 6.10.1.1 Advantages

- **Greater Accuracy than Traditional Methods:** Unlike latency-based geolocation, which suffers from circuitous routing and congestion, the proposed method provides direct, ground-truth location data. It eliminates error-prone inference techniques such as the delay/distance mechanisms used by many geolocation methods.

- **Enhances Existing Network Mapping Tools:** It complements traceroute, BGP, Looking Glass, and RIPE Atlas by integrating explicit geolocation

information. the method provides peer discovery, allowing for more complete topology visualization.

- **Scalable Querying Mechanism:** Uses ICMPv6 Node Information Queries, which are lightweight and designed for extensibility. The iterative querying model allows nodes to be mapped gradually, preventing network overload.
- **IPv6-Ready and Future-Proof:** Designed for IPv6, aligning with the long-term Internet migration away from IPv4. Enhances Internet resilience by improving topology awareness.

#### 6.10.1.2 Disadvantages

- **Requires Router and Kernel Updates:** The QType field must be extended to support new queries, requiring IPv6 router firmware upgrades. Many ISPs and network operators may not prioritize these updates.
- **Potential Scalability Issues:** Large-scale querying could create high ICMPv6 traffic, leading to network overhead. Nodes could be overloaded if too many geolocation requests are made.
- **Lack of Interoperability Without Standardization:** This solution would need to be adopted as an IETF standard to ensure compatibility across different network architectures. Without standardization, different vendor implementations could create inconsistencies.

### 6.10.2 Operational Evaluation

Network operators are risk-averse and typically hesitant to adopt experimental technologies unless there is a clear operational benefit. The success of this solution depends on widespread ISP and IXP buy-in.

#### 6.10.2.1 Advantages

- **Improves Infrastructure Visibility for ISPs and IXPs:** Provides operators with a real-time network topology map, improving routing efficiency. Could assist in troubleshooting latency issues and misconfigurations.
- **Facilitates Traffic Engineering and Network Optimization:** Knowing the exact geolocation of peers allows better routing decisions. Could be used for load balancing and optimized transit selection.
- **Supports Data Sovereignty and Compliance Efforts:** Assists with regulatory compliance by ensuring that data stays within specified geographic regions. Enables geopolitical routing policies to prevent traffic from traversing restricted jurisdictions.

#### 6.10.2.2 Disadvantages

- **Resistance from ISPs and IXPs:** Operators prefer proven technologies over experimental ones. BGP based routing policies conflict with per-packet geolocation aware routing, requiring significant modifications.
- **Requires Major Changes to Network Configuration Policies:** Existing routing is optimized for cost and efficiency, not geopolitical constraints. Operators would need to redefine routing policies, requiring additional training and network adjustments.
- **Lack of Adoption History:** Other location-based protocols such as Geofeeds (Kline, Duleba, Szamonek, Moder, and Kumari, 2020) have not seen widespread adoption, suggesting hesitancy towards geolocation-based routing.

### 6.10.3 Security and Privacy Evaluation

The solution offers new capabilities, however, it also introduces potential security vulnerabilities and privacy concerns. Unauthorized access to geolocation data could

lead to network reconnaissance attacks.

#### 6.10.3.1 Advantages

- **Enhances Security Monitoring and Attack Detection:** Network administrators can validate geolocation data to detect misrouted or hijacked prefixes. Helps in identifying rogue ASNs or malicious BGP behavior.
- **Allows Controlled Access via Access Control Lists (ACLs):** Queries can be restricted to trusted entities, reducing the risk of unauthorized data exposure. Implementing ICMP rate-limiting can mitigate potential DDoS abuse.
- **Can Support Privacy-Enhanced Geolocation:** Nodes can intentionally obfuscate geolocation by providing an approximate location. This balances transparency with security, preventing precise tracking of network infrastructure.

#### 6.10.3.2 Disadvantages

- **Risk of Network Reconnaissance by Attackers:** Malicious actors could map out ISP infrastructure using Node Information Queries. Could assist in state-sponsored cyber operations by revealing critical routing paths.
- **Exposes Network Topology Information:** Some ISPs intentionally obscure their infrastructure for competitive or security reasons. The proposed protocol could force unwanted transparency, creating new attack vectors.
- **Could Violate ISP and Government Privacy Policies:** Some governments and organizations enforce strict privacy controls on network geolocation. Geolocation-aware routing could conflict with existing ISP policies, leading to regulatory resistance.

### 6.10.4 Economic and Regulatory Evaluation

For ISPs and backbone providers, any new protocol must align with economic incentives to be adopted at scale. The proposed solution lacks direct financial benefits, making adoption unlikely without regulatory or competitive pressures.

#### 6.10.4.1 Advantages

- **Regulatory Compliance as a Market Differentiator:** ISPs that adopt this solution can advertise compliance with data sovereignty laws such as GDPR. The solution could be marketed as a security-enhancing feature to government agencies and enterprise clients.
- **Potential Cost Savings from Network Optimization:** Better infrastructure visibility can reduce latency issues and misconfigurations, leading to lower transit costs. It could enable more efficient peering strategies, improving network economics.
- **Pilot Deployment Can Target Select IXPs and ISPs:** Smaller regional ISPs and mid-sized IXPs could be targeted for initial adoption. Early adopters can test feasibility before industry-wide rollout.

#### 6.10.4.2 Disadvantages

- **No Immediate Revenue Model for ISPs:** Unlike BGP, which directly impacts transit costs and peering agreements, geolocation-aware routing does not provide an immediate business incentive.
- **Deployment Costs Are Significant:** The solution requires new router firmware, training for network engineers, and operational policy updates.
- **Regulatory Uncertainty:** Some governments might oppose geolocation-based routing due to national security concerns. Conflicting data localization laws could create compliance risks for multinational ISPs.

## 6.11 Feasibility and Impact Analysis

While the solution could enhance network resilience, security monitoring, and regulatory compliance, it also raises privacy concerns and could be exploited for network reconnaissance. The need for router kernel modifications, IETF standardization, and interoperability testing further complicates deployment. Targeted implementations in research networks, security-focused enterprises, or regulatory-driven environments may be more viable before widespread adoption can be achieved.

The proposed IPv6 Node Information Query extension is technically feasible, leveraging existing protocols to provide more accurate Internet infrastructure mapping. However, its real-world deployment faces significant operational, economic, and security barriers. Network operators are risk-averse and require strong incentives to adopt new geolocation-aware routing mechanisms, especially given the lack of direct financial benefits compared to existing BGP-based traffic engineering.

The proposed solution offers a technically innovative approach to Internet infrastructure mapping; however, its practical deployment faces significant challenges. Without stronger incentives for ISPs, clearer regulatory frameworks, and standardized implementation through the IETF, widespread adoption is unlikely in the near future. However, targeted deployment in smaller IXPs, research networks, or government-regulated infrastructure could serve as a stepping stone for broader industry adoption.

Ultimately, while the geolocation-aware IPv6 Node Information Query extension is a viable tool for enhanced Internet mapping, it must overcome technical, operational, security, and economic hurdles before becoming a practical standard for global use. Future research should explore hybrid approaches, where existing geolocation techniques are enhanced with selective, opt-in query-based geolocation, striking a balance between accuracy, privacy, and feasibility.

## 6.12 Alternative Solutions

Anycast Routing routes traffic to the nearest available node (Abley and Lindqvist, 2006), optimizing performance without requiring explicit geolocation data however, it does not provide exact infrastructure mapping. Anycast Routing is further discussed in Section 7.12.

Encrypted Routing Metadata allows geolocation queries but encrypts responses so only authorized entities can decrypt them; however, this adds computational overhead (Reddy, Wing, and M., 2023).

A Private Peer-to-Peer Geolocation Exchange is where nodes share geolocation data only with direct peers (OECD, 2024), limiting public exposure; however, this can reduce overall visibility of network topology.

## 6.13 Summary

It is widely believed within the industry that the use of the ICMP protocol can add a security issue to the entire solution; however, ICMP is a suite of protocols designed to facilitate communication and error reporting between network devices. Its use in network diagnostics, error reporting, and management tasks are key functions and the idea that potential skilled attackers cannot gain information about a network if ICMP is disabled provides no real security benefits while actively impeding essential network functionalities and is a counterproductive practice (Bartels, 2024).

## 6.14 Discussion

### 6.14.1 Segment Routing

Building on the capabilities discussed in Section 6.6.2 Segment Routing emerges as a powerful tool to enhance the precision and control of packet routing within the global Internet.

Source routing is a routing methodology that allows the sender to either partially or fully determine the route a packet will take through the network. This is in contrast to traditional routing, where routing decisions are made incrementally at each router/node along the path. This allows for easier troubleshooting and allows a host to 'know' all possible paths to the destination.

Segment routing is a type of source routing that is being developed by the IETF (Filsfils, Nainar, Pignataro, Cardona, and Francois, 2015) and is specified in RFC 8402. In a segment-routed network, an ingress router may prepend a header to packets that contain a list of segments, which are instructions that are executed on subsequent routers in the network. These instructions may be forwarding instructions, such as an instruction to forward a packet to a specific destination or interface. With segment routing, the network no longer needs to maintain a per-application and per-flow state. Instead, it obeys the forwarding instructions provided in the packet.

Segment Routing can operate with an IPv6 data plane, and integrates with the rich multi-service capabilities of MPLS (Multiprotocol Label Switching), including Layer 3 VPN (L3VPN), Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), and Ethernet VPN (EVPN). This capability would ensure that the user has a suitable path available before even sending the data.

Segment Routing could make use of the JSON file described in Sections 6.6 and 7.14.3 by using it as the input to a tool which would be capable of predetermining a route through a network to avoid those countries that a user deems undesirable for their data to transit.

Using the information returned from the node information query, network administrators can establish predetermined paths for data transmission. This capability allows for the prepending of route instructions directly to each packet header, ensuring that packets follow the exact path intended, adhering to both geopolitical and performance requirements.

For example, the JSON file returned from a NODE Information Query as

described in Section 6.6 and augmented with the additional country field as described in Section 7.14.3 would be analysed to discover a transit path through the network. All interfaces on each router which contained a country code in a user's 'No Transit' country code list would be removed from the data packets possible routing path. Once a valid path is found through the Internet to the destination, this route would prepend the route instructions within each packet header, which all segment routers on the transit path would obey.

Figure 6.4 shows a valid route (green line) from the UK to the USA if a user had designated Russia and China as 'No Transit' countries.

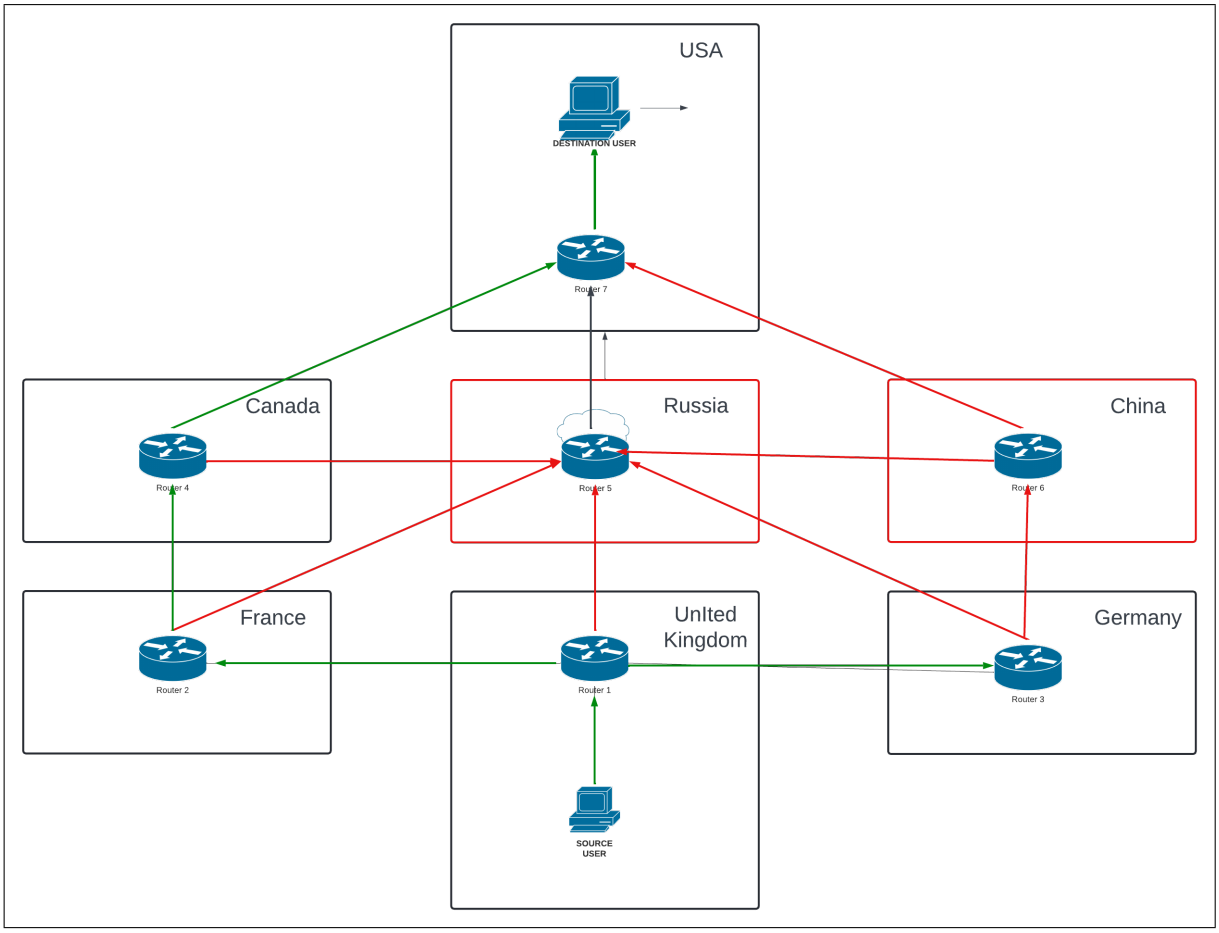


Figure 6.4: Example of Segment Routing by Transit Country

Segment routing is particularly valuable when precise control over the path a

packet takes is necessary. This includes scenarios where geopolitical restrictions must be enforced, specific network performance goals are required, or redundancy and fault tolerance need to be ensured. By embedding route instructions within the packet header, segment routing eliminates the need for per-hop state maintenance in intermediate routers, streamlining the routing process and reducing complexity. This approach is essential in scenarios where predefined routing paths must be strictly followed to meet legal, security, or operational requirements. Segment routing ensures that the data adheres to the intended route, leveraging the detailed network information gathered through the node information query.

### 6.14.2 Security

The IPv6 solutions proposed in Chapters 6 and 7 introduce advanced capabilities for routing and network management. However, these enhancements also expose the infrastructure to significant security risks. Unauthorised access, data interception, and malicious exploitation of the network could compromise the integrity and confidentiality of the global infrastructure. The following sections of this chapter offer some solutions to the issues posed.

#### 6.14.2.1 Using an Authentication and Authorisation Server

Additional security could be added by introducing a trusted authentication and authorisation server. ICMPv6 NI Query functionality allows the subject address to be different from the queried address. Therefore, all queries could go through a specific server as shown in Figure 6.5 that authenticates the queries by applying strict security measures (1) before passing them to routers and middle boxes (2). It will then forward the "Echo Replies" back to the originator of the query (3). This will allow network administrators to confine requests for this information to specific queries originating from the IP addresses of these trusted authentication servers.

This server ensures that only authorised and authenticated users can access sensitive routing and network information. It protects network integrity by prevent-

ing unauthorised users from querying or manipulating network infrastructure and protects against potential security breaches. Furthermore, it ensures that sensitive network information is only accessible to trusted entities, protecting against data leaks. The server would also mitigate the risk of attackers gaining access to routing information, which could be exploited for malicious purposes and limits access to legitimate users, preventing the abuse of network resources and protecting the infrastructure from unnecessary load. This layer of security is crucial for maintaining the integrity and confidentiality of the network.

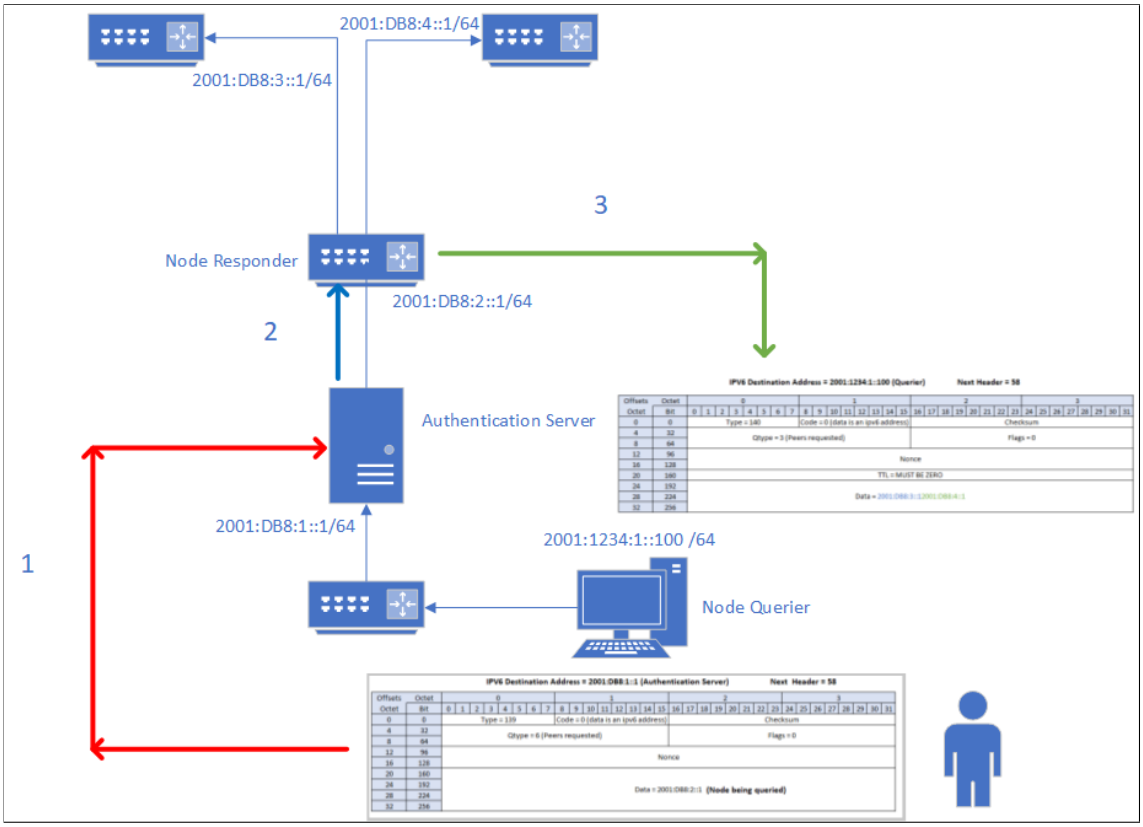


Figure 6.5: IPv6 Node queries and Replies via an Authentication Server

### 6.14.2.2 Alternative Method using TLSA

However, as already pointed out, this solution would use ICMPv6 replies, which are known to be disabled in many routers to ensure that ICMP flooding does not take

place and, therefore, it is likely that the ICMP requests and replies may be blocked in transit in some cases.

An alternative method could involve the use of HTTPS requests to download the JSON file. However, this would require each router to provide a web-based interface. This method could use the TLS authentication record (TLSA), which associates a TLS server certificate or public key with the domain name where the record is found. A TLSA record stores the fingerprint of a TLS/SSL certificate in the DNS of the domain where the router resides. This DNS record provides an additional layer of validation and verification for TLS connections, ensuring that users can authenticate to the router to which they are connecting. However, it does not provide authorisation, and either an authorisation server would still be required or each router could provide authorisation although this would not scale well.

However, even with the introduction of a suitable authorisation method, this solution is still not without its problems. TLSA also requires DNSSEC to be enabled along with a web server service on each router. These services would use valuable CPU cycles, and this solution may be rejected by router manufacturers.

#### 6.14.2.3 Privacy Enhancing Technologies (PET)

Hiding sensitive data from malicious parties requires advanced methods using the latest information hiding techniques (IHT) and privacy enhancement technologies (PET). PETs are a wide range of technologies that allow organisations to collect, share, and use data while mitigating the privacy risks that arise from these activities (ICO, 2024). PETs can involve the random injection of data to make it look like noise or synthetic data which is generated from real data using machine learning and provides similar results to the real data. The Information Comissioners Office provides a list of PETs that may be applicable to this solution. Further research into these technologies may be useful to avoid providing a potential misuse of the information that the methods and tools described in this thesis can provide.

#### 6.14.2.4 Summary

This section explores how the proposed IPv6 geolocation tool can be enhanced using Segment Routing, which enables the sender to define a packet's exact path through the network. By leveraging the JSON data returned from Node Information Queries, network administrators can precompute safe and efficient routes that avoid undesirable regions embedding these instructions into packet headers. This allows for fine-grained control, improved security, and compliance with data sovereignty laws.

The section also addresses security enhancements. It proposes the use of a central authentication and authorization server to filter and verify all Node Information Queries, ensuring only trusted entities access sensitive data. As an alternative to ICMPv6 (which may be blocked), the use of HTTPS with TLSA records and DNSSEC is discussed, though it may be resource-intensive for routers. Lastly, Privacy Enhancing Technologies (PETs) are recommended for further research to protect geolocation data from misuse by obfuscating exact locations.

# Chapter 7

# Geopolitically Aware Routing

## 7.1 Introduction

IPv6 introduced a significant evolution in the area of Internet protocols, which resolved many of the issues with the limitations of IPv4. Clearly, IPv6 provides an improved framework for the future of the Internet. One of the improvements is the concept of Extension Headers (EHs); these are designed to improve and enhance the protocol's flexibility and functionality.

This chapter looks at the development of a novel IPv6 extension header, tailored to incorporate geopolitical awareness into network routing. The new extension header aims to incorporate a geopolitical dimension into each data packet, optionally allowing network paths to be dynamically adjusted based on country codes of transit networks. This addresses a growing need for data controllers and processors to comply with the data protection laws of each country, respecting the geopolitical sensitivities which are inherent in global data transmission. The ethical, security and policy considerations surrounding the implementation of this new extension header will be examined, offering a comprehensive view of how it intersects with societal and regulatory concerns. Although the focus of this chapter is on IPv6, the methods and tools developed for IPv4 geolocation provide a solid foundation for our IPv6 advancements. The transition from IPv4 to IPv6 geolocation is not a divergence,

but rather an evolution that builds on our previous work. The techniques discussed in this chapter will complement and enhance our existing IPv4 tools, providing a comprehensive suite of geolocation and routing solutions for both current and future Internet architectures.

In summary, this chapter will strive to combine advanced protocol design and practical, legal, and ethical considerations in an increasingly complex global Internet infrastructure.

**Roadmap** Section 7.2 reiterates the objectives of this chapter. In Section 7.3 we justify the need for a geo-political extension header and briefly outline IPv6's evolution from IPv4, focusing on the role and functionality of existing headers. Section 7.4 provides more detail about IPv6 extension headers, In Section 7.5, we present the design of the new extension header: the conceptual framework, technical specifications, and the integration strategy with IPv6. In Section 7.6, we discuss methods for router configuration and country code administration. Section 7.7 considers technical, security, privacy, and operational challenges, along with mitigation strategies such as P4 programming, caching, and alternative solutions like SRv6. In Section 7.8 we discuss the security aspects and privacy considerations of the EH header. In Section 7.9 we discuss what would be necessary to create a test environment to simulate the Internet environment using the new IPv6 header. Section 7.10 provides a critical analysis and evaluation of the solution. In Section 7.11, we address relevant technical feasibility issues and analyse router hardware enhancements necessary for practical deployment. Section 7.12 investigates possible alternative methods of achieving the goal of geopolitical routing and Section 7.13 offers a summary of the solution. Section 7.14 Finishes with a discussion on Extension Header recognition and fault finding amongst other topics.

## 7.2 Objectives and Scope of this Chapter

At this point it is appropriate to summarise the objectives and revisit the scope of this key chapter, as follows.

The main objective is to conceptualise and detail the design of a novel IPv6 extension header that incorporates geopolitical information in the form of country code bitfields.

If possible it will assess the practicality of implementing and deploying the new extension header in real-world network environments, considering the current router capabilities and infrastructure. Another objective is to critically examine the security and privacy concerns that may arise from the use of the proposed extension header, especially in the context of international data transmission and legal compliance. It will study the impact of Geopolitical routing, investigating the potential impacts and benefits of incorporating geopolitical considerations into IPv6 routing, both from technical and policy perspectives.

The chapter focuses on the technical aspects of designing the new extension header, including its structure, intended functionality, and integration with the IPv6 protocol. It will explore methods for administering country codes at the router level and it will carry out an evaluation of how the new extension header might affect network performance, including processing overhead, bandwidth implications, and compatibility with existing network hardware and protocols. Also, it will carry out an examination of the security risks and compliance challenges associated with the use of the extension header, including data privacy and cross-border data flow concerns.

It will look at the practical and ethical considerations, discussing the practical deployment challenges and ethical implications of geopolitical-aware routing, including the potential for censorship or regional isolation.

## 7.3 Necessity for a Geo-politically Aware Extension Header

When it comes to the transfer of personal data outside the European Union (EU) and the European Economic Area (EEA), the United Kingdom Data Protection Act 2018 (UK-Government, 2023) (DPA), which is the UK's implementation of the EU's General Data Protection Regulation (EU-GDPR), has rules regarding the countries through which data packets can transit. Personal data can be freely transferred within the EU and EEA countries, as they all maintain the same level of data protection. Under the EU-GDPR, the UK has been considered a third country since Brexit. However, the UK has been granted adequacy, which means that personal data can flow from the EU/EEA to the UK without additional safeguards. Likewise, the Data Protection Act allows for data transfers to the EU/EEA as it considers them adequate.

However, stricter regulations are attached to the transfer of personal data outside these regions. According to the UK Data Protection Act 2018, such transfers can only be made to countries that provide an adequate level of data protection, as determined by the UK Government, or under certain conditions such as:

- Adequacy Decisions: The UK Government can determine that a non-EU country offers an adequate level of data protection. This allows for the easier transfer of data to these countries.
- Standard Contractual Clauses (SCCs): These are legal contracts drawn up between the data sender and the recipient in the non-EU country, ensuring the protection of personal data.
- Binding Corporate Rules (BCRs): These are internal rules adopted by multinational companies to allow transfers within the same corporate group to entities in countries without adequacy decisions.

- Specific Derogations: In absence of an adequacy decision or SCCs/BCRs, transfers can be made under specific conditions such as explicit consent from the individual whose data is being transferred, or for the performance of a contract.

A landmark ruling by the Court of Justice of the European Union (CJEU) on 16 July 2020, known as Schrems II, was made in the case of Data Protection Commissioner v. Facebook Ireland and Maximillian Schrems (Case C-311/18), and the decision has significant implications for international data transfers between the European Union (EU) and countries outside the EU, particularly the United States. The CJEU invalidated the EU-US Privacy Shield framework, which was a mechanism used by companies to legally transfer personal data from the EU to the United States (Tracol, 2020).

The court found that the US surveillance laws did not provide an adequate level of protection for personal data, as required by the General Data Protection Regulation (GDPR). Specifically, the court was concerned that US law did not afford EU citizens sufficient rights to challenge US government access to their data. Secondly, the court imposed stricter requirements on the use of SCC's where companies must now assess data protection laws in the destination country to ensure that they provide an adequate level of protection to that in the EU and implement additional safeguards if this is not the case. Thirdly, organisations are required to perform thorough data transfer impact assessments (DTIAs) to evaluate whether the legal framework of the importing country meets the EU's standards for data protection and, if this is not the case, they must take additional steps to protect the data. This has led to significant challenges for companies that transfer data internationally especially those that rely on US based services.

As of March 21, 2022, new mechanisms for international data transfers have been introduced in the UK, including the International Data Transfer Agreement (IDTA) (ICO, 2022a) and the UK Addendum to the EU Standard Contractual Clauses (SCC) (ICO, 2022b). These tools are necessary for ensuring that any personal data

transferred outside the UK is protected by appropriate safeguards, particularly when transferring data to countries that do not have an "adequacy decision" from the UK government.

In essence, while the core principles of the UK Data Protection Act 2018 remain relevant, the mechanisms and tools for international data transfers have evolved, and organisations must now ensure compliance with the latest requirements, such as using the IDTA or the UK Addendum to the SCCs for cross-border data transfers (Kapko, 2024).

The new Standard Contractual Clauses (SCCs) introduced several important changes compared to the old SCCs, reflecting the evolving data protection landscape, especially post-Brexit. The new SCCs are modular and designed to cover different data transfer scenarios, including transfers between controllers, processors, and subprocessors. This is a significant change from the old SCCs, which were more rigid and less adaptable to different contractual relationships. The new SCCs require organisations to perform detailed data transfer impact assessments. This ensures that data protection levels in the destination country are adequate and in line with GDPR standards. Data importers must notify data exporters and data subjects if they receive government access requests and assess the legality of the orders, potentially challenging them if necessary. Controllers and processors under the new SCCs are required to demonstrate compliance with data protection principles, reflecting the emphasis of GDPR on accountability. This includes keeping detailed records of processing activities and implementing appropriate technical and organisational measures (Lawbite, 2022).

The UK Addendum to the International Data Transfer Agreement (IDTA) provides a way for UK-based organisations to continue using the EU Standard Contractual Clauses (SCCs) for international data transfers while ensuring compliance with UK-specific data protection laws. The UK addendum modifies the EU SCCs by replacing EU-specific legal references with UK equivalents, making them applicable under UK law post-Brexit and allowing organisations that operate across

both the EU and UK to use a single set of SCCs for data transfers by adding the UK addendum simplifying compliance and reducing administrative burden. Businesses had until March 2022 to update existing contracts that were based on old SCCs. After this date, all such contracts must comply with the IDTA or include the UK Addendum (ICO, 2022a).

The route that data packets take can often involve countries that do not meet these EU and UK GDPR standards, and organisations must ensure that the appropriate safeguards are in place to protect the data during its transit. There is also an emphasis on accountability and transparency, which requires data controllers and processors to take responsibility for ensuring that any personal data is protected throughout its journey, regardless of the countries it transits through. This has led to increased scrutiny and changes in the way organisations manage data flows, often requiring more direct and secure data transfer routes.

Beyond the legal requirements for data transfer, there are several other compelling motivations for developing a geopolitically aware IPv6 extension header.

- Security concerns: as cyber threats become more sophisticated, it is important to ensure that data does not traverse through regions with lax cybersecurity standards or where data interception is more likely is crucial. A geopolitically aware extension header can help mitigate these risks by enforcing secure routing paths.

- Compliance with international regulations: beyond the UK's legal framework, various countries have specific regulations regarding data sovereignty and privacy. A geopolitically aware header could help organisations comply with these diverse regulations by ensuring that data transits only through compliant jurisdictions.

- Performance optimization: routing data through geopolitically stable regions can improve network performance and reduce the likelihood of data loss or corruption due to political instability, censorship, or network disruptions in

certain regions.

- Ethical considerations: some organisations may wish to avoid routing data through regions with known human rights violations or where data might be used for unethical purposes. A geopolitically aware routing option allows organisations to align their data transmission practices with their ethical standards.

- Resilience and redundancy: in the event of geopolitical tensions or conflicts, having the ability to dynamically adjust routing paths to avoid affected regions can enhance the resilience and redundancy of the Internet infrastructure, ensuring continuous and reliable data transmission.

- The Border Gateway Protocol: is a powerful protocol for managing inter-domain routing but it has inherent limitations when it comes to enforcing precise routing paths based on geopolitical considerations. Whilst BGP can be configured to influence routing decisions through policies and AS path preferences, it lacks the capability to explicitly and reliably route traffic away from specific undesirable countries, this is because BGP focuses on path selection based on network performance and policy, rather than the granular control over geographic routing. The solution allows for selective routing of individual packets in real time by end users and not network administrators, enabling tailored decisions. For example, certain packets (e.g., humanitarian aid data) could be routed through sensitive regions like warzones, while others (e.g., less critical traffic) could avoid those areas altogether.The proposed IPv6 extension header solution addresses these limitations by providing a mechanism that embeds geopolitical routing requirements directly into the packet headers.

These motivations, combined with legal considerations, provide a strong case for the development and implementation of a geopolitically aware IPv6 extension header. However, critics may argue that the decisions about which countries that

the routing of data should avoid is best made at the government or organisation level but there are several purposes and benefits of providing this flexibility:

- Providing users and network administrators the ability to define routing policies allows for a tailored approach that aligns with specific operational, legal, or security priorities that might not be fully addressed by a blanket government or organisational policy.
- Empowering users to enforce data sovereignty over their personal or organisational data enhances trust and ensures they can protect their interests in ways that blanket policies may not sufficiently address.
- Providing flexibility at the user or administrator level ensures rapid and dynamic adaptation to emerging threats or geopolitical changes, which may not always be addressed promptly by static, centralized policies.
- Allowing network administrators to set routing preferences provides autonomy and control, ensuring that organisations can meet their internal standards while still adhering to overarching governmental or regulatory frameworks
- Enabling users and administrators to decide routing paths allows for high-priority or sensitive use cases, ensuring critical data can still reach necessary destinations even when blanket rerouting policies may otherwise disrupt essential services.
- Allowing users to refine routing preferences ensures that broad restrictions do not unintentionally disrupt legitimate activities or create inefficiencies for specific use cases.
- User-defined routing fosters transparency and accountability by enabling individuals and organisations to take ownership of their data paths, ensuring alignment with their unique priorities and values.

- Allowing flexibility in routing decisions encourages innovation in routing strategies, helping to optimize network performance and discover novel approaches to secure and efficient data transmission.

## 7.4 Background on IPv6 and Extension Headers

On November 17, 1994. RFC 1883, Internet Protocol, Version 6 (IPv6) Specification, was drafted by the Internet Engineering Steering Group, and the proposed standard resulting from this was published in 1995. The core set of IPv6 protocols became an IETF draft standard on August 10, 1998. One of the main opportunities of IPv6 is that extension headers are a fundamental part of the design, and these provide a flexible method in which we can extend IPv6 functionality to include new ideas that were not conceived during its initial design.

The rapid expansion of the Internet has exhausted its 4 billion IPv4 addresses, and the introduction of IPv6 marks a significant evolution in Internet technology. IPv6 addresses this limited address space and also introduces a more efficient and flexible protocol design. The implementation of extension headers improves the functionality and scalability of IPv6 and is a key feature among these advancements.

Some applications and upper-layer protocols assume that a packet is unmodified in transit, except for a few well-defined fields such as the TTL field, this also includes protocols that define their own integrity-protection mechanisms such as checksum fields (Aben, 2013). However, Network Address Translation devices and other middle boxes can also modify the contents of packets and the IPsec architecture (Mühlbauer, Feldmann, Maennel, Roughan, and Uhlig, 2006) added security to the IP model, although transport-mode IPsec is not currently widely used over the Internet. That is not to say that IPsec cannot be used with this solution, the original header is not encrypted in transport mode, thus allowing routers to read and make use of the geographical header.

IPv4 has a rigid format, whilst IPv6 adopts a more modular approach using a fixed-size base header with optional extension headers. This more streamlined approach creates an easy-to-use standard packet structure while allowing for additional features, ensuring efficient packet processing. Extension Headers are critical for customising packet handling for specific functions, such as security, routing, and fragmentation.

The Hop-by-Hop (HBH) Options header is used to carry optional information. In the first version of the IPv6 specification, all nodes along a packet's delivery path were required to process Hop-By-Hop options (B. Hinden and S. E. Deering, 1998), which proved to be impractical due to:

- Current High Speed Routers were unable to process the hop-by-hop options at wire speed.
- Packets containing Hop-by-Hop options would often be sent to the "slow path", degrading performance for possibly important traffic.
- A Denial of Service attack on the router could be created by exploiting a mechanism that forces external packets to the routers "slow path" which means that the "slow path" is at risk of being flooded.
- Packets could contain multiple Hop-by-Hop options, making the previous issues worse by increasing the complexity required to process them.

There are many issues with the use of extension headers, particularly those that need to be processed by nodes along a packet delivery path (R. Hinden and Fairhurst, 2020). The IPv6 specification was updated and published in 2017 and changed the procedures for nodes encountering Hop-By-Hop extensions (S. Deering and R. Hinden, 2017). However, these changes allowed routers to only examine and process Hop-By-Hop headers if configured to do so. This change meant that routers complied with the IPv6 specification even when they did not process the Hop-By-Hop header, which does not fix the problems highlighted above, merely circumventing them.

APNIC Labs measured the drop rate of IPv6 packets that contain an HBH extension header consisting of an 8 octet padding option (Huston, 2022a). They added this option to the TCP data streams and then checked to see if the receiver acknowledged receipt of this packet in the TCP sequence number flow. They made 5000 measurements per day for 65 days in 2022 from a set of IPv6 servers to a

collection of IPv6 client hosts, they found that the average drop rate of the HBH option was 92% although APNIC's measurement cannot discern between network drop and host drop. Additionally, some manual tests have shown that many IPv6 implementations discard incoming packets with unexpected HBH extension headers.

A draft RFC makes several recommendations, including that IPv6 nodes 'MUST only process a Hop-by-Hop Options header if it can be done in the fast path of the router' (R. Hinden and Fairhurst, 2020) as shown in Figure 7.1.

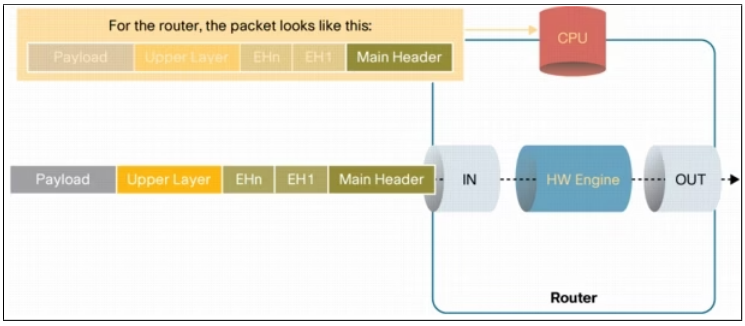


Figure 7.1: Forwarding IPv6 Packets with Extension Headers other than Hop-by-Hop in the absence of ACLs via the 'Fast Path' (Cisco, 2006)

Secondly, only one option should be contained within a Hop-By-Hop extensions header to simplify processing. However, all vendor equipment has been designed to forward IPv6 traffic that contains a Hop-by-Hop Extension Header to go through the slow forwarding path (Cisco, 2006); see Figure 7.2.

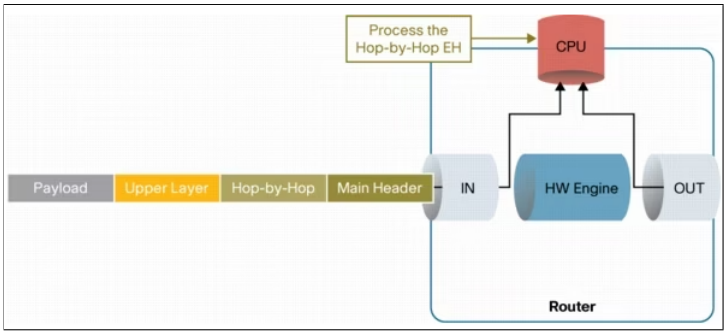


Figure 7.2: Forwarding IPv6 Packets with the Hop-by-Hop Extension Header via the 'Slow Path' (Cisco, 2006)

Therefore, currently 'Fast Path' forwarding is not feasible in the case of any

packet containing a Hop-By-Hop extension header. This will obviously add some limitations to the use of a geopolitical extension header which uses a Hop-By-Hop Extension Header. While pragmatic solutions like BGP are effective for general traffic management, they are not equipped to handle the nuanced requirements of secure and compliant data transit. The proposed header is not a replacement but an enhancement, addressing use cases where security and compliance outweigh speed.

## 7.5 Proposed IPv6 Extension Header Design

### 7.5.1 Overview

The Earth is currently comprised of 195 recognized sovereign states (UN, 2024). The proposed solution requires that each country require a bit set to zero or a bit set to one to allow or disallow transit. It is proposed that a 32 byte field will cover this, plus provision for additional support for the formation of new countries. IPv6 extension headers, such as the Routing Header or Destination Options Header, often vary in size but generally aim to be as small as possible to minimise processing overhead and maintain efficiency. For example, the basic size of a Routing Header can vary, but common implementations often involve relatively small sizes, typically in the range of 8-16 bytes, depending on the number of included addresses. The Fragment Header is smaller, often just 8 bytes, as it only needs to carry essential fragmentation information. However, it is not unusual for extension headers to vary in size depending on the specific use case. A 32 byte field is feasible within the IPv6 standard, as the protocol is designed to accommodate variable-length headers. The key consideration is to ensure that the added processing load and potential impact on network performance are justified by the functionality provided by the new header.

This concept requires modification of the router firmware and its table structure. A country field will need to be added to each interface within the IFIndex (Interface Index: see Section 7.6 for a more complete description). This table will be populated

manually during initial installation. Scalability will need to be investigated especially considering the dynamic nature of network topologies and international routing paths. This will also add a small amount of processing at each router. The impact on network performance would need to be carefully evaluated. There are security and privacy implications, potential interception, or misrouting of packets are just some of the issues faced. Reliance on ICMP error responses adds complexity. Where the route is disrupted, a router will have to check if an alternative "Safe Path" exists. Routers will return an ICMP error message if they detect that there is no safe alternative route.

### 7.5.2 Concept and Structure

It is proposed that any data packet travelling across the Internet that is required to avoid transiting specific countries should include an IPv6 Hop-By-Hop (HBH) Extension Header (EH) that informs all transited routers of the countries that the packet should not be routed through. It is expected that a network administrator will enable this "Safe Path" option at his gateway router so that any IP communication exiting their network will have the extension header injected into each packet. Alternatively, individual users/developers may wish to select the "Safe Path" option within their application to ensure that only the data from their particular application follows a "Safe Path". A "Safe Path" option can be provided within an application by the developer to allow a user to select it. Additionally, they would be able to select the countries through which they do not wish this data to travel.

Routers will have been modified to tag each of their interfaces with the destination countries; see Section 7.6. When a packet is processed at each router, the Hop-By-Hop extension header will be read, and the router will match, using 'OR' logic, the bits which are set to one in the Extension Header against those that are set in the outgoing interface country field within the IFIndex table. If, for example, an extension header has the first 3 bits of the 'Do Not Transit' field set. This would be compared with a routers' interfaces proposed new IFIndex table 'countries' field,

and if any of the first three bits are set in this field, the packet will not be allowed to transit any interfaces where a countries "Do Not Transit" bit is set; this is where the 'OR' logic is applied, see Figure 7.3. The router must either look for a backup route or send an ICMPv6 "Destination Unreachable" error (type = 1) back to the source.

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| HBH Header | 1 | 1 | 1 | 0 | 0 | 0 | ..... |
| Interface E0/0 'Countries' | 1 | 0 | 0 | 0 | 0 | 0 | ..... |
| result | 1 | OR 0 | OR 0 | OR 0 | OR 0 | OR 0 | = 1 |

Figure 7.3: Example of Logic used between a HBH and a Router Interfaces Countries Field

### 7.5.3 Crafting the Safe Path Extension Header

IPv6 allows nodes along a packet path to optionally process a Hop-by-Hop header (S. Deering and R. Hinden, 2017).

The Hop-by-Hop Options header can carry a variable number of "options" that are encoded by the type length value (TLV) as shown in Figure 7.4. However, the proposed geopolitical header will contain just a single option.

| Offsets | Octet | 0 | | | | | | | | 1 | | | | | | | | 2 | | | | | | | | 3 | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Octet | Bit | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
| 0 | 0 | Option Type = 100NNNNN | | | | | | | | Option length = 00100000 | | | | | | | | | | | | | | | | | | | | | | | |
| 4 | 32 | 'Do Not Transit' Field | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 8 | 64 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 12 | 96 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 16 | 128 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 20 | 160 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 24 | 192 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 28 | 224 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| 32 | 256 | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

Figure 7.4: IPv6 with Hop-By-Hop Extensions header and Geo-political Field

The Option Type field is an 8 bit identifier where the two highest-order bits signify the action that must be taken if the processing IPv6 node does not recognise the option type (S. Deering and R. Hinden, 2017):

1. 00 - skip over this option and continue processing the header.

2. 01 - discard the packet.

3. 10 - discard the packet and, regardless of whether or not the packet's Destination Address was a multicast address, send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognised Option Type.

4. 11 - discard the packet and, only if the packet's destination address was not a multicast address, send an ICMP Parameter Problem, Code 2, message to the packet's source address, pointing to the unrecognised option Type.

The third-highest-order bit of the Option Type specifies whether or not the Option Data of that option can change en-route to the packet's final destination:

- 0 - Option Data does not change en route

- 1 - Option Data may change en route

In the case of our geopolitical extension header, the data will not change and therefore this bit will be set to 0 as shown in Figure 7.4.

The final 5 bits of the Option Type field (denoted by 'NNNNN' in Figure 7.4) is the option type assigned by the Internet Assigned Number Authority (IANA). This will be added once we have registered our extension header with IANA.

The option length field will be set to 32 (octets) to provide adequate space for the data.

### 7.5.4 Country Code Bitfields

The 32 byte data field will consist of 195 bits (and 61 spare bits) with a bit set to one for each relevant country on the user's "do not transit" list.

The United Nations (UN) recognises 193 countries and 2 "Observers States", the Holy See/Vatican City and Palestine, which are self-ruling territories but not

full-fledged countries bringing to a total of 195 countries (UN, 2024). It was hoped to use an international standard such as the ISO 3166 standard – Codes created and maintained by the International Organisation for Standardisation (ISO) for the representation of names of countries and their subdivisions or the telephone codes used by the International Telecommunication Union (ITU) in ITU-T standards E.123 and E.164; however, these numbers do not lend themselves easily to this solution. The ISO numbers end at 894 and the ITU numbers end at 999 which means that a minimum of 894 or 999 bits would have to be reserved. This equates to 112 bytes, which is significantly more than the 64 bytes of extension header data that cisco's hardware accelerated platforms are designed to handle (Cisco, 2006). Cisco's selected size is considered sufficient to handle the most common chains of EH currently used with various IPv6 traffic types and services. In Cisco routers, when the size of the EH chain exceeds the resources allocated in hardware and the upper-layer protocol filters are applied, the IPv6 traffic will be software switched by the Line-Card CPU.

Instead, we will use an alphabetical coding standard in which each country is assigned a code in alphabetical order. A 32 bit field allows for 256 countries, of which only 195 are currently used at this time, leaving 61 bits spare for the creation of new countries. Countries that cease to exist can be left in place and not used or replaced in router firmware updates by new countries. The list and the bits assigned to the countries are shown in Figure 7.5.

Each code is the reference to the relevant bit within the field, for example, Bulgaria is code 26 so this would be bit 26 of the countries field in the IFIndex table discussed in section 7.6. It will also be the reference to the relevant bit in the Countries field of the Hop-By-Hop extension header. However, as the first 16 bits are taken up by Option Type and Option Length fields, the Countries field begins at bit 16 so Bulgaria is referenced by its code of 26 plus an offset of 15 which is equal to bit 41 of the Hop-By-Hop extension header. Continuing with this example, if Bulgaria is a country that should not be transited, then the country field would

| Code | Country name | Code | Country name | Code | Country name | Code | Country name | Code | Country name | Code | Country name | Code | Country name | Code | Country name |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 1 | Afghanistan | 26 | Bulgaria | 51 | East Timor | 76 | Iceland | 101 | Lithuania | 126 | New Zealand | 151 | Sao Tome and Principe | 176 | Tonga |
| 2 | Albania | 27 | Burkina Faso | 52 | Ecuador | 77 | India | 102 | Luxembourg | 127 | Nicaragua | 152 | Saudi Arabia | 177 | Trinidad and Tobago |
| 3 | Algeria | 28 | Burundi | 53 | Egypt | 78 | Indonesia | 103 | Macedonia | 128 | Niger | 153 | Senegal | 178 | Tunisia |
| 4 | Andorra | 29 | Cambodia | 54 | El Salvador | 79 | Iran | 104 | Madagascar | 129 | Nigeria | 154 | Serbia | 179 | Turkey |
| 5 | Angola | 30 | Cameroon | 55 | Equatorial Guinea | 80 | Iraq | 105 | Malawi | 130 | Norway | 155 | Seychelles | 180 | Turkmenistan |
| 6 | Antigua and Barbuda | 31 | Canada | 56 | Eritrea | 81 | Ireland | 106 | Malaysia | 131 | Oman | 156 | Sierra Leone | 181 | Tuvalu |
| 7 | Argentina | 32 | Cape Verde | 57 | Estonia | 82 | Israel | 107 | Maldives | 132 | Pakistan | 157 | Singapore | 182 | Uganda |
| 8 | Armenia | 33 | Central African Republic | 58 | Ethiopia | 83 | Italy | 108 | Mali | 133 | Palau | 158 | Slovakia | 183 | Ukraine |
| 9 | Australia | 34 | Chad | 59 | Fiji | 84 | Jamaica | 109 | Malta | 134 | Palastine | 159 | Slovenia | 184 | United Arab Emirates |
| 10 | Austria | 35 | Chile | 60 | Finland | 85 | Japan | 110 | Marshall Islands | 135 | Panama | 160 | Solomon Islands | 185 | United Kingdom |
| 11 | Azerbaijan | 36 | China | 61 | France | 86 | Jordan | 111 | Mauritania | 136 | Papua New Guinea | 161 | Somalia | 186 | United States of America |
| 12 | The Bahamas | 37 | Colombia | 62 | Gabon | 87 | Kazakhstan | 112 | Mauritius | 137 | Paraguay | 162 | South Africa | 187 | Uruguay |
| 13 | Bahrain | 38 | Comoros | 63 | The Gambia | 88 | Kenya | 113 | Mexico | 138 | Peru | 163 | South Sudan | 188 | Uzbekistan |
| 14 | Bangladesh | 39 | Republic of the Congo | 64 | Georgia | 89 | Kiribati | 114 | Micronesia, Federated States of | 139 | Philippines | 164 | Spain | 189 | Vanuatu |
| 15 | Barbados | 40 | Democratic Republic of the Congo | 65 | Germany | 90 | Korea, North | 115 | Moldova | 140 | Poland | 165 | Sri Lanka | 190 | Vatican |
| 16 | Belarus | 41 | Costa Rica | 66 | Ghana | 91 | Korea, South | 116 | Monaco | 141 | Portugal | 166 | Sudan | 191 | Venezuela |
| 17 | Belgium | 42 | Cote d'Ivoire | 67 | Greece | 92 | Kuwait | 117 | Mongolia | 142 | Qatar | 167 | Suriname | 192 | Vietnam |
| 18 | Belize | 43 | Croatia | 68 | Grenada | 93 | Kyrgyzstan | 118 | Montenegro | 143 | Romania | 168 | Swaziland | 193 | Yemen |
| 19 | Benin | 44 | Cuba | 69 | Guatemala | 94 | Laos | 119 | Morocco | 144 | Russia | 169 | Sweden | 194 | Zambia |
| 20 | Bhutan | 45 | Cyprus | 70 | Guinea | 95 | Latvia | 120 | Mozambique | 145 | Rwanda | 170 | Switzerland | 195 | Zimbabwe |
| 21 | Bolivia | 46 | Czech Republic | 71 | Guinea-Bissau | 96 | Lebanon | 121 | Myanmar (Burma) | 146 | Saint Kitts and Nevis | 171 | Syria | 196 | |
| 22 | Bosnia and Herzegovina | 47 | Denmark | 72 | Guyana | 97 | Lesotho | 122 | Namibia | 147 | Saint Lucia | 172 | Tajikistan | 197 | |
| 23 | Botswana | 48 | Djibouti | 73 | Haiti | 98 | Liberia | 123 | Nauru | 148 | Saint Vincent and the Grenadines | 173 | Tanzania | 198 | |
| 24 | Brazil | 49 | Dominica | 74 | Honduras | 99 | Libya | 124 | Nepal | 149 | Samoa | 174 | Thailand | 199 | |
| 25 | Brunei | 50 | Dominican Republic | 75 | Hungary | 100 | Liechtenstein | 125 | Netherlands | 150 | San Marino | 175 | Togo | 200 | |

Figure 7.5: Proposed Country Codes

be coded as shown in Figure 7.6. This method allows for the selection of multiple undesirable countries by simply setting the desired bits.

| Offsets | Octet | 0 | | | | | | | | 1 | | | | | | | | 2 | | | | | | | | 3 | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Octet | Bit | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 | 29 | 30 | 31 |
| 0 | 0 | Option Type = 100NNNNN | | | | | | | | Option length = 00100000 | | | | | | | | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 4 | 32 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 1 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 8 | 64 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 12 | 96 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 16 | 128 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 20 | 160 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 24 | 192 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 28 | 224 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| 32 | 256 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 | Bit 41 = 26 + 15 = Bulgaria = "Do Not Transit" | | | | | | | | | | | | | | | |

Figure 7.6: Example of setting Country Code

## 7.6 Router Configuration

The Interface Index (IFIndex) value is one of the most commonly used identifiers in SNMP-based network management applications. It is a unique identification number

associated with a physical or logical interface. RFC1213 (MIB2) (McCloghrie and Rose, 1991) defines an initial IFIndex as follows:

> "Each interface is identified by a unique value of the IFIndex object, and the description of the IFIndex constrains its value as follows: Its value ranges between 1 and the value of IFNumber. The value for each interface must remain constant at least from one re-initialisation of the entity's network management system to the next re-initialisation."

The latest Cisco IOS releases add support for an IFIndex value that can persist across reboots, which proves particularly useful not only for our solution but for other features such as:

- SNMP: monitoring the interfaces counters
- Netflow: reporting of the interface IFIndex
- RMON: events/alarms based on specific interfaces
- EXPRESSION/EVENT MIB: creation of a new MIB variable based on interface counters.

The format of the IFIndex table would therefore be updated with a countries field which would denote the countries that each interface connects to as shown in Figure 7.7

| | Name | Type | Description |
|---|---|---|---|
| | size | INTEGER32 | the size of this row |
| | ifIndex | INTEGER32 | The interfaces ifIndex |
| | enablePersistence | INTEGER32 | 1 if persistence is enabled |
| | ifDescr | OCTECT STRING | The Interface Description |
| *Proposed New Field | countries | ARRAY[INTEGER32] | 32 bytes where a bit set to 1 = destination country(ies) |

Figure 7.7: Proposed change to the IFIndex table

### 7.6.1 Router Behaviour

Routers that are not configured to support Hop-by-Hop options are not expected to examine or process the contents of this option as per RFC8200.

Routers that support Hop-by-Hop Options but are not configured to support this Option should read the two highest-order bits of the Extension Header Option Type Field to determine whether to skip over this header or discard the packet and send an ICMP error message in accordance with RFC8200.

Routers that support this geopolitical extension header will initially follow the rules laid out in Section 7.5.3

To improve scalability and performance, several modifications to router processing are recommended:

- **Firmware Updates for Efficient EH Processing:** Ensuring EHs are processed in hardware rather than software.
- **IFIndex Table Optimization:** Reducing lookup overhead by caching country mappings for frequent destinations.
- **Policy-Based Routing Approaches:** Allowing organizations to configure routing policies based on regulatory compliance requirements.

### 7.6.2 Interface Country Tagging

The tagging of each interface to the relevant countries could be carried out upon the initial installation of a router. It is at this time that an engineer creates the connections to its peers, and choosing the destination countries of each of the routers interfaces should not add overly to the workload. Alternatively, a Node Information Query as described in Chapter 6 could be sent to a router's upstream routers to enquire about the upstream routers' country of location. This may be a better method as it could be done automatically without any manual intervention on the part of the installation engineer.

### 7.6.3 Integration with IPv6 architecture

The first step of this procedure would require the router installation engineer to configure the country or countries to which each of the router interfaces is connected. This information will be stored within the proposed additional countries field inside the IFIndex table as discussed in Section 7.6.

Secondly, should a company consider it part of their policy, there are two alternative methods of tagging outgoing data packets:

1. A network administrator or engineer could enable an option on a host which would inject the geopolitical header containing the countries he does not wish any packets to transit into all outgoing packets. The host could be the gateway or outgoing router for an entire organisation or just simply each individual PC.

2. Alternatively, this could be left to individual users and provided as an option at the application layer where a user can choose to add the geopolitical header to an outgoing file.

Thirdly, as this packet traverses the network, routers will perform an additional routing action, perhaps using an Access Control List (ACL), which will be to check their IFIndex table for the interface that the packet will be sent out. If this proves to be a country in the do not transit list then the router will check for alternative paths, and if this proves fruitless will either discard the packet, return an ICMP Destination unreachable error message or ignore the header and forward the packet regardless.

This new method incorporates the following procedure:

1. A router's firmware will be updated to change the IFIndex table to include a 32 byte field for each interface.

2. During initial installation an engineer will set corresponding bits within this field dependent on the country that each of the routers interfaces connects to. The engineer will also configure an Access Control list (ACL).

3. On receipt of a packet containing this header the ACL will check the country field against its IFIndex to ensure that the next hop destination country is not on the "Do Not Transit" list.

4. If the interface's destination country is on the "Do Not Transit" list then the router may check for an alternative path and if unsuccessful send an ICMP Destination Unreachable message back to the source. If the router does not check for an alternative path then it will also send a ICMP Destination Unreachable message back to the source.

5. If the router is successful in finding a valid interface it forwards the packet.

### 7.6.4 Policy-Based Routing Approaches

The Hop-By-Hop type field shown in Figure 7.4 allows for differences in company policies. This will allow the network administrator or user to be able to decide whether a router which does not recognise this header should continue routing the packet or discard the packet and send an error message back. Whilst the second option, discard and send error message may be the default setting, it may be a company's policy to add the extension header to all outgoing data as "Best Effort", but allowing for data to continue to be transmitted should it come across routers which do not support this header.

## 7.7 Barriers to Deployment

### 7.7.1 Technical Barriers

Geopolitical-aware routing would require updates to routing firmware and hardware to process the new header type without significantly degrading performance. Vendors such as Cisco disable or deprioritise hop-by-hop headers for security and efficiency reasons. The solution requires real-time checking of packets against geolocation constraints which adds overhead at high speeds in large scale networks with millions

of packets per second. This increases processing demands on routers which could lead to network bottlenecks and reduced performance. This could also provide an attack vector for DDOS attacks.

#### 7.7.1.1 Potential Solutions

- **Hybrid Approach with BGP Integration:** Instead of relying solely on per-packet processing, aligning BGP route policies with the extension header could reduce unnecessary processing overhead.
- **Caching Mechanism for IFIndex Table Lookups:** Frequently accessed routes can be stored in a short-term cache to avoid repetitive slow-path processing (Kasnavi, Berube, Gaudet, and Amaral, 2008).
- **Adoption of Segment Routing over IPv6 (SRv6):** Instead of Hop-by-Hop EHs, SRv6 could define predetermined "safe paths" without per-hop intervention (Filsfils, Nainar, Pignataro, Cardona, and Francois, 2015).
- **Fast-Path Processing through P4 Programming:** P4-based programmable routers, discussed in Section 7.9, can analyze and enforce the geopolitical routing policy in hardware, avoiding slow-path inefficiencies (Bosshart, Daly, Gibb, Izzard, McKeown, Rexford, Schlesinger, Talayaco, Vahdat, Verghese, and Walker, 2014).

### 7.7.2 Operational Barriers

Many current routers, especially those in large networks, do not process IPv6 extension headers efficiently, as they are not supported in fast-path processing. Indirect impacts on communications such as increased latency, inconsistent router performance, and disruptions in path stability can degrade the accuracy and reliability of geolocation efforts; however this is true of all Hop by Hop headers but a solution to this is outside of the scope of this thesis.

#### 7.7.2.1 Proposed Optimizations

- **Edge-Based Processing:** Instead of processing EHs at every transit router, handling can be limited to edge routers at organizational boundaries.
- **Router Firmware Enhancements:** Upgrading IPv6 forwarding pipelines to support fast-path processing of geopolitical EHs (Kasnavi, Berube, Gaudet, and Amaral, 2008).
- **Adoption of Alternative Routing Mechanisms:** Exploring alternatives such as Anycast routing and AI-based dynamic route selection (discussed in Section 7.12).

Many firewalls and middleboxes will examine and process an entire IPv6 packet before it makes a decision to either forward or discard a packet, due to differences in extension header formatting this process can be slow and clumsy. Some firewalls will drop any packet they cannot recognise, and several widely used firewalls do not recognise all the extension headers that have been standardised since the IPv6 protocol was introduced (Carpenter and S. Jiang, 2013).

### 7.7.3 Security Barriers

Increasing header overhead can make packets larger than the Maximum Transmission Unit (MTU), leading to fragmentation and potential performance degradation (S. Deering and R. Hinden, 2017). Attackers may exploit this by sending oversized packets to disrupt services, for example, fragmentation-based attacks. Many firewalls and security devices are not configured to properly parse or filter IPv6 extension headers (Huston, 2022b). Attackers may exploit this lack of filtering to bypass security checks by embedding malicious payloads within headers. IPv6 allows for multiple extension headers in a single packet (S. Deering and R. Hinden, 2017). Attackers can craft packets with excessive or malformed headers, consuming processing resources on routers and firewalls and leading to Denial-of-Service

(DoS) attacks. Fragmentation can lead to reassembly vulnerabilities. Attackers can send fragmented packets to bypass intrusion detection systems (IDS) or exploit weaknesses in packet reassembly logic to execute DoS attacks. IPv6 extension headers can be spoofed to manipulate the flow of packets (Davies, Mohacsi, and Beijnum, 2007). Attackers can redirect traffic through malicious nodes or unintended paths, potentially capturing sensitive data or disrupting service. Hop By Hop extension headers require processing by every router along the path. Attackers can craft packets with complex header chains, leading to excessive resource consumption as routers parse each header. This can result in cascading failures in the network infrastructure. Extension headers can be used to embed data in ways that evade detection (Blumberg, Pihelgas, Kont, Maennel, and Vaarandi, 2016). Attackers can use these covert channels for exfiltration of sensitive data or command-and-control communication. Many middleboxes (e.g., NATs, firewalls, intrusion prevention systems) struggle to handle IPv6 extension headers (Huston, 2022b). Attackers can exploit these incompatibilities to bypass security mechanisms, potentially accessing protected networks or injecting malicious traffic. When extension headers are used in encrypted traffic (e.g., IPsec), their handling can become ambiguous. Misconfiguration or lack of standardisation can lead to traffic being mishandled, potentially exposing sensitive data or creating attack vectors

#### 7.7.3.1 Security Enhancements

- **Integration with IP Authentication Header:** Authenticating packets carrying the geopolitical extension header to prevent spoofing.
- **P4-based Security Filtering:** Programmable P4 switches can analyze and reject malformed or suspicious EHs (Bosshart, Daly, Gibb, Izzard, McKeown, Rexford, Schlesinger, Talayaco, Vahdat, Verghese, and Walker, 2014).
- **Real-Time Anomaly Detection:** AI-based security systems can monitor traffic patterns for irregular geopolitical header usage and prevent abuse.

## 7.8 Security and Privacy Considerations

It is common for routers to ignore the Hop-by-Hop option header or to drop packets containing a Hop-by-Hop options header. Routers implementing IPv6 accordingly only examine and process the Hop-by-Hop options header if explicitly configured to do so. This emphasises the need to ensure that the data being transmitted across networks remains protected against unauthorised access, tampering, and potential misuse. The extension header will include sensitive geopolitical routing information; it must be designed with robust security measures to prevent exploitation and to maintain the confidentiality and integrity of transmitted data. This is crucial to maintaining trust in the network infrastructure and ensuring compliance with various legal and regulatory frameworks that govern data protection.

IPv6 packets incorporating the Geopolitical Extension Header can be authenticated using the IPAuthentication Header (Kent, 2005a). A node should include an Authentication Header when sending a geopolitical extension header if a security association exists for use with the IP Authentication Header for the destination address. The security associations may have been created through manual configuration or through the operation of some key management protocol. Received Authentication Headers in IPv6 packets containing Geopolitical Extension Headers should be verified for correctness, and packets with incorrect authentication should be ignored and discarded. The importance of securing the extension header against misuse or exploitation by malicious actors must be emphasised, particularly as it could be used to manipulate routing paths for unauthorised or harmful purposes. It highlights the need to implement robust security measures to protect sensitive geopolitical routing information and to ensure that data privacy is maintained during transmission across various jurisdictions. This consideration is crucial to prevent the extension header from becoming a vector for attacks or privacy breaches, ensuring that it contributes positively to network security and compliance with data protection regulations.

It should be possible for the system administrator to configure a node to ignore

any messages including a Geopolitical Extension Header that are not authenticated using either the Authentication Header or Encapsulating Security Payload. It is expected that a node will default to allowing unauthenticated messages.

The information learned through this method should not be trusted for making security-relevant decisions unless other mechanisms beyond the scope of this document are used to authenticate this information. Confidentiality issues are addressed by the IP Authentication Header (Kent, 2005a) and the IP Encapsulating Security Payload (Kent, 2005b).

Detection and mitigation often require a combination of monitoring tools, firewall rules, and changes in network configuration and implementing strict validation of header formats to prevent malformed packets. The key to effectively managing these issues lies in maintaining a secure and updated network infrastructure, using intrusion detection systems, and implementing best practices for network security.

## 7.9 Testing and Evaluation Framework

A successful deployment of the method depends upon several components being implemented and deployed:

- Router support in nodes as described at Section 7.6.
- Support in the sending node and upper layer protocols as described in Section 7.6.3.

The Extended Berkeley Packet Filter (eBPF) is becoming a recognised method for the testing of IPv6 Extension Headers (Iurman, Vyncke, and Donnet, 2023). eBPF can do so much more than just packet filtering, making the acronym eBPF non-sensical, and therefore eBPF is now considered to be just a standalone term (eBPF, 2024). Whilst eBPF is suited to host systems making it easy to inject Extension headers into outgoing and incoming packets, Hop-By-Hop requirements include the need to process the extension header within the transited routers.

eBPF does not currently have this functionality; in fact, it seems impossible to find any tools which allow a router's kernel software to be easily edited to allow for new Hop-By-Hop extension headers. An alternative method of testing could be the use of Programming Protocol-Independent Packet Processors (P4), which is an open-source programming language that lets end users dictate how networking infrastructure operates. P4 controls integrated circuits in network forwarding devices such as switches, routers, and network interface cards. It is similar to OpenFlow in many respects; however, instead of targeting the control plane, P4 is focused on the data plane. A programmable data plane brings with it a number of advantages:

- It is easy to add new features, supporting new protocols (BGP, OSPF, Spanning Tree, etc)
- Remove unused protocols: free up space for it to focus solely on what the user wants it to do.
- Greater visibility: P4 allows users to program in rules to forwarding devices. It can, for example, create a tag for each packet as it passes through a router or switch. Doing so allows network engineers to get a potentially unprecedented level of visibility into the routing paths of packets to determine network latency.

## 7.10 Critical Evaluation

This section provides a critical analysis of how the protocol would function in real-world Internet deployment, highlighting its strengths and weaknesses.

### 7.10.1 Technical Evaluation

#### 7.10.1.1 Advantages

The proposed solution offers granular control over routing paths unlike BGP, which applies broad policies to all traffic, this protocol allows packet-level routing decisions

based on geopolitical constraints. The protocol enables selective compliance with data sovereignty laws without disrupting global traffic flows. The possibility of using P4 programmable switches offers the possibility that the protocol can process geolocation metadata at line rate, minimizing delays. With regards to security, it prevents unintended data transit through adversarial regions, reducing risks of state-sponsored surveillance, censorship, or interception. The protocol ensure compliance with regulatory frameworks like GDPR and the UK Data Protection Act.

#### 7.10.1.2 Disadvantages

There is currently a marked lack of router support for IPv6 extension headers where many routers will process hop-by-hop headers in the slow path, leading to significant performance degradation. The protocol will increase latency when routing around specific geopolitical regions, this may increase path lengths, leading to higher round-trip times (RTTs). If different ISPs implement conflicting geopolitical policies, route flaps and instability could occur.

### 7.10.2 Operational Evaluation

#### 7.10.2.1 Advantages

The protocol offers greater transparency in routing decisions by introducing traceability and allowing users to audit which countries their traffic passes through. It provides more control for enterprises and critical infrastructure offering businesses that handle sensitive data such as financial transactions, government communications, cloud services the ability to define geopolitical routing policies. The protocol helps critical infrastructure providers (power grids, defense networks, emergency services) to prevent traffic from transiting through untrusted territories.

#### 7.10.2.2 Disadvantages

There will be a lot of hesitancy to adopt the protocol by ISPs due to a lack of financial incentives and many ISPs still use IPv4 infrastructure. However, ISPs could sell such services as premium routing services for security-conscious customers.

### 7.10.3 Security and Privacy Evaluation

#### 7.10.3.1 Advantages

By prioritising the Geopolitical extension header over the BGP protocol it would reduce the attack surface for BGP hijacking and prevent BGP route leaks that might redirect traffic through malicious third parties. The protocol ensures that packets follow intended routes without arbitrary re-routing for surveillance.

#### 7.10.3.2 Disadvantages

The protocol has the potential for misuse by some governments for censorship purposes. The header contains explicit geographical routing constraints, potentially exposing an organisation's routing or business policies revealing sensitive corporate strategies, whilst malicious actors could infer national security policies based on observed routing constraints. Packets with a unique geopolitical routing header can be used for traffic fingerprinting, enabling tracking of individual users or organisations. there is also the potential of privacy erosion when a user's traffic includes routing metadata, advertisers, ISPs, or nation-states could track their geopolitical routing preferences. If transit ISPs log the headers, they can build databases of organisations' routing behaviors and governments could subpoena those logs to identify how networks and businesses are avoiding certain jurisdictions. The retention of historical routing preferences could classify as personally identifiable information (PII) under EU data laws.

### 7.10.4 Economic and Regulatory Evaluation

#### 7.10.4.1 Advantages

The protocol helps multinational organisations ensure that personal data does not transit through non-compliant jurisdictions. It allows cloud providers to offer geopolitically-compliant services to customers and could incentivize governments to develop more regional IXPs, reducing dependency on foreign infrastructure.

#### 7.10.4.2 Disadvantages

Hardware upgrades and software modifications may incur significant costs and conflicting regulations such as GDPR and the US CLOUD Act could create routing disputes.

### 7.10.5 Conclusion

The protocol is technically viable, but requires hardware upgrades and cooperation from ISPs. It reduces BGP hijacking risks and ensures compliance with geopolitical routing policies. However, there are performance concerns whereaby IPv6 extension headers may cause slow-path processing and there is a risk of misuse where governments could exploit this for censorship and surveillance. The protocol should be tested in controlled environments (e.g., large IXPs, cloud providers) before full Internet-scale deployment. Collaboration with IETF, ICANN, and ISPs is essential to ensure standardization, economic viability, and real-world feasibility.

## 7.11 Feasibility and Impact Analysis

There are some necessary prerequisites for this kind of approach to be viable:

- All hosts, routers, and all forms of middleware including firewalls, deep packet inspection (DPI) systems, load balancers, and other middleboxes need to support this option.

- IPv6 packets that carry this HBH option are not arbitrarily discarded by devices on the path or by the destination host.
- All forwarding devices recognise this HBH option and will compare the country codes against those of their interfaces before forwarding packets that contain this HBH extension header.
- All IPv6 hosts will make local adjustments to their routing protocol based on the country code information.
- All IPv6 protocol implementations need to support a socket option to allow upper-layer protocols to inject this Hop-By-Hop extension header.

The adoption of this extension header depends on multiple factors:

- **Incentivizing ISPs to Implement EH Processing:** This could be achieved through regulatory mandates like GDPR compliance enforcement.
- **Economic Viability of Router Upgrades:** ISPs and enterprises may be more willing to invest if the technology aligns with compliance requirements or enhances security.
- **Gradual Implementation with Hybrid Models:** Initial deployment in limited use cases such as cloud providers and large enterprises before full Internet-scale adoption.

## 7.12 Alternative Solutions

There are other methods that offer varying degrees of flexibility, efficiency, and security while addressing the core challenge of avoiding specific countries, regions, or networks. Encryption based routing such as Onion Routing, Virtual private Networks (VPNs) and Tor encrypt packets and/or route through randomized relay nodes that only know the next hop. These solutions can provide improved privacy as

no metadata is leaked, packets appear similar regardless of their final destination and encrypted packets prevent deep packet inspection. However, they potentially may be blocked by authoritarian regimes, for example China's Great Firewall blocks the Tor network (Tor, 2025) and furthermore, latency may be increased due to additional hops.

Segment Routing (SRv6) allows administrators to define policy-based routes using a source-routing paradigm, where network operators set rules for traffic paths without embedding sensitive data inside packets (Filfils, Talaulikar, Voyer, Bogdanov, and Mattes, 2022). This enables fine-grained control over traffic flow by defining specific paths for different applications or traffic types. Furthermore, this reduces the complexity of network configuration and management by eliminating the need for extensive control plane calculations. However, as is the case with all new methods, router support for SRv6 is not currently supported by all ISPs and it is less flexible for individual users, as routing policies are centrally defined. the bes use case for Srv6 is where enterprise and government networks that need geopolitical-aware routing without privacy risks.

Anycast routing uses gateways where the same IP address is assigned to multiple geographically distributed locations. Users are then routed to the nearest available gateway, which then makes geopolitically compliant routing decisions on their behalf (Abley and Lindqvist, 2006). For example, a user in Europe needs to send data to to a U.S. service, and instead of manually avoiding specific transit countries, they send the data to an Anycast gateway in London, which automatically routes traffic based on legal/geopolitical rules. The benefits of this method are that users do not need to disclose routing preferences because routing is handled by Anycast nodes. There is no risk of fingerprinting as all users appear to access the same IP. This reduces the attack surface as there is no need for hop-by-hop metadata exposure. However it does require large scale deployment of Anycast nodes and may not guarentee optimal routing performance as traffic may take longer paths increasing latency. Cloud providers and CDN's could make best use of this method to reduce exposing

user metadata.

Decentralized Trust Networks with Smart Contracts allow users to leverage block chains to create secure and censorship-resistant routing agreements whereby users stake digital assets to ensure routes are followed without interference (Pailisse, Manrique, Bonet, Rodriguez-Natal, Maino, and Cabellos, 2019). For example, in a decentralized VPN network, users pay nodes in a cryptocurrency for routing traffic without revealing routing preferences. This method means there is no centralized logging and no possibility of metdata exposure, routes are cryptographically verified which prevents manipulation and user anonymity is preserved as it is the smart contracts that handle routing logic. However, it is widely known that Blockchain services have scalability concerns (Sanka and Cheung, 2021).

AI-Based Dynamic Route Selection dynamically adjusts paths based on real-time geopolitical risks and network performance (latency, congestion, security). AI analyses BGP announcements, detects routing anomalies, and automatically rerouting traffic without user intervention (Aktas, Shayea, Ergen, Saoud, Yahya, and Laura, 2025). The benefits are that users do not expose their routing constraints, network decisions are based on real-time analysis instead of static metadata and as routing is managed by AI this avoids surveillance risks. However there is a high reliance on accurate real time data and it could be expensive for small companies.

| Solutions | Privacy Risks | Performance | Dep. Complexity | Best For |
|---|---|---|---|---|
| Onion Routing (Tor, VPNs) | Very low | Medium | Medium | Users Avoid Censorship |
| Segment Routing (Srv6) | Very low | High | High | ISPs and Enterprise networks |
| Anycast Routing | Very low | High | High | CDNs and Cloud providers |
| Decentralised Trust Networks | Very low | Medium | High | Privacy Advocates and P2P users |
| AI based Dynamic Routing | Very low | High | High | Enterprise and Financial networks |



Table 7.1: Final Alternatives Comparison table

## 7.13 Summary

The proposed IPv6 geopolitically aware Extension Header could play a critical role in addressing the complexities of global data routing. However, it is not a standalone solution, it should be seen as a complementary innovation that builds upon and enhances existing routing protocols. By addressing its limitations and securing adoption, the header has the potential to significantly impact global data routing in a meaningful way. This chapter outlines the technical and legal challenges inherent in implementing such a system while highlighting its potential to enhance both security and privacy in network operations. By integrating geopolitical considerations into routing decisions, this extension header offers a forward-looking solution to ensure compliance with international regulations and protect data integrity across diverse jurisdictions. The chapter underscores the importance of rigorous testing and evaluation to validate the efficacy of this approach, ultimately advocating its adoption as a means to enhance the resilience and security of the global Internet infrastructure. Alternative approaches like SRv6, Anycast routing, and AI-driven dynamic routing could complement or, in some cases, replace this solution. Future research should focus on performance optimization and incentives for widespread adoption.

## 7.14 Discussion

Chapters 6 and 7 look toward the future, introducing new ideas based on IPv6. However, what happens in the event that IPv6 never fully replaces IPv4? The proposal to add IPv6 style Extension Headers to IPv4 would provide an interim solution to this problem (Herbert, 2024). This specification allows the core IPv6 Extension Headers defined in RFC8200 to be used with IPv4 (S. Deering and R. Hinden, 2017). These Extension Headers include Hop-by-Hop Options, Destination Options, Routing Header, and Fragment Header. The Authentication Header (Kent, 2005a) and the Encapsulating Security Payload (Kent, 2005b) are

already usable with IPv4.

This specification is still only an Internet draft and is classed as a "Work in Progress"; therefore, should this fail to be approved for publication, then the solution will still work in a limited fashion much as Traceroute does now. Traceroute only works when routers ICMP replies are enabled; if they are not enabled, then researchers and analysts have to move on to those parts of the Internet that allow ICMP replies. The same will happen for these solutions; if IPv6 is not enabled on a router or the changes to the router kernel have not been implemented to a specific router, then researchers will need to analyse routes where they are enabled.

### 7.14.1 Extension Header Recognition

The steps required to obtain recognition for a new IPv6 Extension Header are significantly difficult. The first step is to engage with the IETF, write a draft proposal, and submit it for peer review. We would then need to get a Hop-By-Hop Options Type assignment from IANA from the "Destination Options and Hop-by-Hop Options" (Section 5) sub-registry of the "Internet Protocol Version 6 (IPv6) Parameters" registry (Bradner and Paxson, 0200). The new Extension Header type assignment would need to be mentioned in the IETF RFC that describes our Extension Header. The next step is even more challenging; we would need to get our newly approved Extension Header implemented across the Internet. We would have to convince network equipment manufacturers to be able to pass our packets and process them as described above.

The new Extension Header would need to be implemented in core Internet routers, service provider networks, subscriber CPE, enterprise networks, and cloud infrastructure. We would also need to get the Extension Header permitted to be forwarded across all firewalls, deep packet inspection (DPI) systems, load balancers, and other middleboxes. We would also need to ensure that other security functions like Intrusion Prevention Systems (IPSs), packet brokers, web proxy services, malware prevention systems, and other filters do not accidentally or intentionally

block IPv6 packets with our new Extension Header. The final step would be to implement the new Extension Header in every IPv6 protocol stack of operating systems such as Apple macOS, iOS, iPadOS, Microsoft Windows, Google (for Android and ChromeOS), Linux and many other Operating Systems.

Without full industry support, parts of the global infrastructure would not recognise the new extension header and this could give rise to issues such as:

- Packet dropping: routers or devices not recognizing the header might drop packets containing the new extension, resulting in failed transmissions and communication breakdowns.
- Routing issues: non-compliant routers might ignore the header, leading to unintended routing paths that do not adhere to the geopolitical restrictions intended by the extension, undermining its purpose.
- Interoperability challenges: inconsistent support across the global infrastructure could lead to significant interoperability challenges, complicating the deployment and effectiveness of the extension header.
- Fallback mechanisms: the infrastructure will need robust fallback mechanisms to handle unrecognised headers, ensuring that packets can still be routed, albeit without the additional controls intended by the extension.

### 7.14.2 Optional Tracking Field

Whilst conceptualizing this header, more thought was put into the possibility of adding a second 32 byte field, which would record the countries of the router interfaces through which the packet had passed. Although this might prove interesting and perhaps useful, the security aspects of providing a writable field were considered to outweigh significantly the possible benefits.

### 7.14.3 IPv6 Geolocation Fault Finding

One major issue that has been identified is a concern about the added complications of debugging any network problems that the proposals in Chapter 7 would create. An analyst would need to differentiate between when packets are lost due to the geopolitical aware routing header proposed in Chapter 7 and when they are lost due to other causes. To resolve this matter, the ideas proposed in Chapter 6 could additionally allow for a router's country code peer connections to be analysed. This would provide sufficient reporting to investigate problems caused by geofencing.

Also, this would require an additional Node Information Protocol change to the Qtype field. A Qtype of 5 would not only return a list of IP addresses of the nodes peers as proposed in Section 7.6 but would also return the contents of the new "countries field" by asking the router to interrogate its IFTable.

The data returned would then consist of 3 pieces of information for each of the router's interfaces as shown below:

```
###### Data in JSON Format ########
{ "2001:DB8:2::1": { # Router IPv6 Address
        "lat": 52.2815, # Router Latitude
        "lon": -0.8685, # Router Longitude
        "interfaces": # Router interfaces
                [{
                "interface": "22001:DB8:3::2/64",
                "peer": "22001:DB8:3::1/64",
                "countries": 32 Bytes
                },{
                "interface": "2001:DB8:4::2/64",
                "peer": "2001:DB8:4::1/64",
                "countries": 32 Bytes
                }, {
                "interface": "2345:425:567::b1/64",
                "peer": "2345:425:567::b5/64",
                "countries": 32 Bytes
                },{
                "interface": "2266:25:12::15/64",
                "peer": "2266:25:12::12/64",
                "countries": 32 Bytes
                }]}}
```

The countries field would be a set of 32 bytes indicating the countries to which a particular interface is connected, as shown previously in Figure 7.5.

### 7.14.4 Testing and Evaluation Framework

The changes discussed in this chapter requires changes to the operating system kernel, which traditionally has always been difficult to attempt due to its central role and high requirement towards stability and security.

#### 7.14.4.1 Enhanced Berkeley Packet Filter (eBPF)

eBPF is a technology that can run sandboxed programs in the operating system kernel within a privileged context. It extends the capabilities of the kernel, without requiring one to change the kernel source code or load new kernel modules. eBPF allows developers to run programs that add additional capabilities to the operating system at runtime compiled with the aid of a Just-In-Time (JIT) compiler and verification engine.

The wide-scale deployment of eBPF suffers with the problem that it is challenging to build applications that are compatible across a wide range of Linux distributions. This is because the eBPF code must be compiled on the target host to make sure that the program is compatible. Also, each host may have a different kernel, and so kernel struct layouts may have changed. So, a system running Linux 5.5 would need a different compilation to that of, for example, a Linux 5.8 system. If testing on different Linux distributions is required, the eBPF community has developed libbpf + CO-RE, which encodes the struct offsets of kernel structs for a given kernel version. CO-RE uses the BPF Type Format (BTF) that enables the BPF program loader to make adjustments to the precompiled eBPF code so that it looks at the right offsets in memory (Nakryiko, 2020).

Although the use of eBPF may be good for testing, it is primarily designed for Linux and has been tightly integrated into the Linux kernel. While eBPF is native to Linux, there have been efforts to bring similar capabilities to other

operating systems, such as Windows. However, the most robust and widely used implementations of eBPF are found in the Linux ecosystem. Routers built on Linux distributions such as BIRD (Bird, 2024) could of course provide a useful test bench, but this implementation would not be suitable for general deployment on the Internet at large.

#### 7.14.4.2 Programming Protocol-independent Packet Processors (P4)

An alternative testing framework that may be better suited to the global Internet is the P4 programming language, which can control packet forwarding planes on network devices (Bosshart, Daly, Gibb, Izzard, McKeown, Rexford, Schlesinger, Talayaco, Vahdat, Verghese, and Walker, 2014). P4 works in conjunction with SDN control protocols like OpenFlow, and its three goals are that it should allow programmers to change the way routers and switches process packets once they are deployed, switches should not be tied to any specific network protocol, and finally programmers are allowed to be able to describe packet processing functionality independently of the specifics of the underlying hardware.

The discussions and methodologies presented in this chapter highlight the potential of advanced programming languages such as P4 and eBPF to transform how we manage and secure the Internet infrastructure. Both P4 and eBPF offer unique advantages and can be strategically employed depending on the specific requirements and constraints of the network environment.

P4, with its high performance and data plane programmability, is particularly suited for real-time geolocation checks in high-speed networks. eBPF, with its flexibility and ease of deployment, provides a powerful alternative to existing Linux-based systems, offering extensive monitoring capabilities.

### 7.14.5 Summary

By integrating these advanced techniques with the geolocation tools and methodologies developed in this thesis, we can significantly improve the accuracy, security,

and resilience of the Internet infrastructure. Future work should focus on further refining these approaches, exploring their combined potential, and addressing any implementation challenges to fully realise their benefits.

# Chapter 8

# Conclusion and Future Work

This thesis is structured around three core contributions: Chapters 1 to 5 examine and enhance IPv4-based geolocation and infrastructure mapping techniques; Chapter 6 introduces a novel IPv6-based tool using Node Information Queries for accurate and secure infrastructure mapping; and Chapter 7 extends this foundation by proposing a mechanism for enforcing geopolitical routing policies through a new IPv6 extension header. The following sections in this final chapter provide conclusions specific to each area, followed by a discussion of broader contributions, limitations, and future directions.

**Roadmap** Section 8.1 provides the conclusions on the IPv4 tools that achieve notable accuracy improvements, but it acknowledges traceroute limitations. Section 8.2 provides conclusions on the IPv6 Node Information Queries for precise, real-time Internet mapping. Section 8.3 provides conclusions on IPv6 extension header for geopolitical routing. It recognizes challenges with Hop-By-Hop headers but highlights future potential. Section 8.4 explains the contributions made to improve geolocation, Internet resilience, and ethical data use. Section 8.5 discusses the limitations and offers some possible future studies. Section 8.6 offers some final thoughts and emphasizes the need for continued research and innovation.

## 8.1 IPv4 Geolocation Conclusion

The purpose of the research presented in Chapters 1 through 5 was to investigate the current methods used by Internet mapping techniques to determine the optimum method to develop fine-grained infrastructure maps. The research then built on these methods by developing tools and techniques that can help create fine-grained infrastructure maps. Specifically, the new method developed in this work uses four newly created Python tools that gather the locations of UK facilities and maps them to OpenStreetMap. The tools also locate all UK IXPs and maps the network structure and interconnection facilities. The tools, listed in Chapter 8.6, then create measurements from the selected RIPE Atlas probes to create a UK infrastructure map by geolocating every hop within each traceroute, where possible.

The objectives were to extend geolocation granularity from city to facility level by integrating various methods with traceroute-derived constraints. The development of custom Python tools and the new method developed using IXPs as geographic anchors demonstrates this effectively. The thesis enhances geolocation accuracy by combining forward and reverse traceroutes with IXP landmark data, creating new constraints on IP location in a more robust way than RTT-only methods. The byproduct collection of IXP related data and traceroute outputs leads to the creation of a reusable, curated dataset of facility level IP to location mappings, and through a detailed case study (viz. London) the thesis shows the method's ability to map interconnection infrastructure at a metropolitan scale. Furthermore, the thesis assesses inference success using ground-truth data from major IXPs, presenting quantifiable improvements over existing approaches.

Chapter 4 defines and develops a new method to use IXPs as geographical landmarks, ironing out the issues and proving the usefulness of IXPs as geographical anchors. Chapter 5 offers new knowledge by combining these landmark-based inferences (IXPs) with curated traceroute constraints and DNS geohints. It shows that it is possible to significantly improve the accuracy of IPv4-based infrastructure maps from coarse city-level to a more fine grained facility-level.

The achievement of these objectives represents a considerable advance over existing geolocation techniques. However, despite notable progress, the work reaffirmed that traceroute-based inference methods remain fundamentally limited due to routing asymmetries, and other opaque network practices. This reinforced the need for new mechanisms beyond legacy IPv4 tools—laying the groundwork for the solutions explored in Chapters 6 and 7.

## 8.2 IPv6 Geolocation Conclusion

The new tool described in Chapter 6 leverages IPv6 Node Information Queries, as detailed in RFC 4620, to extract detailed node data, including geolocations and peer relationships (Crawford and Haberman, 2006). The proposal to add further options to the existing Node Information Queries will allow for a more comprehensive and accurate view of the Internet's structure than tools relying on traceroute, and will avoid the inherent and fundamental issues surrounding traceroute by using the new queries for gathering comprehensive data including node geolocations and peer connections.

The proposed IPv6 network infrastructure mapping tool will exceed current methodologies in precision and completeness. Using IPv6 node information queries, the tool can provide detailed real-time information on the Internet structure, including geolocations and peer relationships of network nodes.

The proposal for the development of this tool achieves the objectives of retrieving more accurate geolocation and peering data from routers by introducing a novel IPv6 mapping tool that leverages Node Information Queries (RFC4620). It does this by proposing protocol-level extensions, router kernel modifications, and deployment strategies to improve the Node Information protocol's utility for mapping. It offers troubleshooting support by providing a platform for enhanced diagnostics, allowing better visibility into network paths and potential vulnerabilities in IPv6 infrastructure. The tool offers a more comprehensive picture of IPv6 topology than

previous methods, showing great potential for network administrators and also for researchers.

This development represents a significant step forward in addressing the limitations of current Internet mapping methodologies and paves the way towards the design of more sophisticated applications. The proposed changes to the node information protocol, along with modifications to router kernels and installation procedures, underscore a holistic approach to improving the accuracy and utility of network mapping. IPv6 Node Information Queries, when extended with custom QTypes, offer a fundamentally new way to query routers directly for their geolocation and peer relationships moving, beyond inference into explicit topology discovery.

The tool's design emphasises the importance of security, implementing measures to safeguard against misuse while enabling legitimate research and network administration activities. With the introduction of extended query options and the ability to obtain detailed geographical and connectivity data, researchers and network operators will be much better equipped to analyse, troubleshoot, and optimise the Internet infrastructure. With these extensions, Internet mapping can become real-time, verifiable, and significantly more granular, provided adoption and security concerns are addressed.

The evaluation of the IPv6 Node Information Query tool, as illustrated in Figure 6.3, demonstrates its potential to accurately enumerate and geolocate routers across a target network by iteratively probing their node information. This capability directly supports the objective outlined in Section 1.4 to develop a new IPv6-based infrastructure mapping tool that offers improved accuracy, comprehensiveness, and utility over existing methods. The tool's proposed ability to expose router interconnections and physical locations that are otherwise opaque to traditional traceroute techniques can enhance situational awareness and inform routing decisions. However, the proposed implementation also reveals several limitations, particularly the dependency on full compliance with the Node Information Protocol and the challenges of adoption across heterogeneous router platforms. While the tool

marks a significant advancement towards fine-grained IPv6 infrastructure mapping, its broader impact on resilience and security remains contingent upon widespread deployment and standardisation across administrative domains.

## 8.3 IPv6 Geopolitical Routing Conclusion

The objective of Chapter 7 was to conceptualise and design a novel IPv6 Hop-by-Hop Extension Header (EH) that incorporates geopolitical information in the form of country code bitfields. This was achieved through the proposal of a new mechanism enabling per-packet routing preferences that reflect geopolitical constraints. The header allows network administrators and users to encode data sovereignty or policy requirements directly into packet headers, offering granular control over international data paths.

Chapter 7 fulfilled its core objectives, as set out in Section 1.4, by going beyond theoretical design to offer a technically sound, security-conscious, and legally-aware proposal. It explored the practical feasibility of implementation via programmable network tools such as GNS3, eBPF, and particularly P4, the latter enabling in-data-plane enforcement of routing policies. This shifts decision-making to the forwarding plane for improved performance and compliance without control-plane latency.

The chapter also tackled security, privacy, and ethical issues, identifying both the risks (e.g., fingerprinting, surveillance, censorship) and the benefits (e.g., BGP hijack mitigation, traffic accountability, and regulatory compliance). Strategies for mitigation included rate-limiting, PETs, and the use of authentication servers to regulate header use.

A comprehensive feasibility analysis assessed deployment barriers across technical, economic, and legal dimensions. These include router vendor support, performance penalties due to slow-path processing, and varying ISP willingness to adopt new protocols. The work acknowledged these issues but also proposed pathways to adoption, including gradual deployment, economic incentives for ISPs,

and collaborations with IETF working groups.

Critically, Chapter 7 established the importance of this technology for:

- enforcing data sovereignty and jurisdiction-aware routing;
- enhancing resilience in geopolitical conflict zones;
- enabling real-time routing decisions by users or enterprises, rather than only by ISPs;
- creating transparent and accountable routing paths, improving trust in digital infrastructure.

Alternative solutions such as SRv6, Anycast gateways, AI-driven routing, and onion-routing protocols were compared in terms of feasibility, privacy, and scalability. This contextualised the proposed IPv6 geopolitical extension header as a complementary enhancement to existing methods, especially for cases requiring explicit policy compliance.

The first two goals of Section 7.2 (conceptualisation and design) are believed to have been met. However, the evaluation proved to be difficult, as explained in Section 7.9. Despite the challenges in implementation and evaluation, this proposal demonstrates a forward-looking approach to improving Internet security and compliance. The solution in Chapter 7 is more targeted and robust in addressing the need to route data away from specific countries based on legal, security, or ethical considerations. It opens new possibilities for data sovereignty enforcement, jurisdiction-aware routing, and traffic steering in conflict zones or under regulatory mandates, transforming how the Internet respects geopolitical boundaries.

However, Geoff Huston, chief scientist at APNIC, is particularly pessimistic about the use of Hop-By-Hop extension headers (Huston, 2022a) where he agrees with Hinden's view that "Hop-By-Hop options are still not practical to be used widely in the Internet and many operational routers are configured to drop all packets containing a Hop-By-Hop option header" (R. Hinden and Fairhurst, 2020).

Strangely, this statement was written by the same author when proposing a Hop-by-Hop extension header to find the maximum MTU size between a source and a destination. However, it is interesting to note that the statement was removed in a later draft (R. Hinden and Fairhurst, 2021). In their white paper, 'IPv6 Extension Headers (EH) Review and Considerations', Cisco argue that EHs are considered a powerful tool in extending IPv6 to adapt to future protocol requirements and service needs. It is expected that other uses will be identified for the existing EHs and that new EHs will be defined (Cisco, 2006). The Internet Engineering Task Force (IETF) continues to provide recommendations on the use of IPv6 hop-by-hop options (Vyncke, Chittimaneni, Kaeo, and Rey, 2021). Huston claims that concerns surrounding the use of IPv6 Extension Headers are valid (Huston, 2022a), but Cisco's endorsement of IPv6 extension headers as powerful tools gives greater credence to the new mapping tool proposed in Chapter 7 as it uses these headers to allow for dynamic routing adjustments based on country codes, ensuring that data paths comply with geopolitical regulations.

In summary, while full deployment remains challenging due to adoption barriers and interoperability concerns, the IPv6 geopolitical routing proposal presented in Chapter 1.4 marks a significant conceptual and technical advancement. It aligns closely with the core thesis goal of integrating legal, political, and infrastructural concerns into the foundational protocols of the Internet. The proposed solution is a forward-looking contribution towards a more sovereign, secure, and ethically-aligned global network.

## 8.4 Thesis Contributions

This thesis makes five significant contributions to the fields of Internet measurement, IP geolocation, and resilient routing architecture:

1. **Chapter 1-5, a new IPv4 mapping technique using IXPs as geolocation anchors**

    - Develops a novel methodology for improving the accuracy of IPv4 infrastructure mapping by using Internet Exchange Points (IXPs) and co-location facilities as geographical anchors.
    - Introduces bespoke Python tools that combined RIPE Atlas probes, OpenStreetMap data, and curated IXP lists to sense-check traceroute data. This resulted in improved geolocation accuracy from 13% to over 33.6%, marking a notable improvement over existing passive or delay-based methods.
    - Creates a new dataset of geolocated UK router IPs was also produced, offering future researchers a valuable reference.

2. **Chapter 6, a proposed extension to IPv6 Node Information Queries**

    - Introduces a new extension to the ICMPv6 Node Information Query protocol (RFC4620), with two additional QTypes:
        - one for returning a node's geolocation (QType = 6);
        - one for returning a node's directly connected peer IPs (QType = 5).
    - These enhancements transform the NI Query protocol into a tool for real-time infrastructure mapping. This is a first-of-its-kind design and represents a conceptual and technical advancement in IPv6-based active measurement.

3. **Chapter 6, Design and Operation of a New IPv6 Infrastructure Mapping Tool**

- Designed a network-mapping tool based on the enhanced NI Query, which is capable of:
  - iterative traversal of networks by querying for interface addresses, peer relationships, and physical locations;
  - generating a dynamic, JSON-based map of the Internet infrastructure with far more accuracy and granularity than existing traceroute-based tools;
- The operational design accounts for scalability, performance, and security, with mitigation strategies for abuse and rate-limiting.

4. **Chapter 7, a novel IPv6 extension header for geopolitical routing**
   - proposes an IPv6 Hop-by-Hop Extension Header that embeds bitfields representing country codes to enable policy-aware routing. This is a new mechanism for:
     - Enforcing jurisdictional compliance
     - Avoiding data transit through disallowed or adversarial regions;
     - Supporting enterprise or governmental data sovereignty strategies.
   - This is a significant contribution to both Internet governance and network engineering, expanding the role of packet headers in routing control.

5. **A security-conscious framework for protocol innovation**
   - Throughout the thesis, particular attention is paid to the risks introduced by the proposed protocols.
   - propose a multi-tiered security model using ACLs, rate-limiting, authentication servers, and privacy obfuscation;
   - Provides detailed discussions of threat models for geolocation and topology disclosure

- Introduces concrete implementation suggestions using P4, TLSA records, and privacy-enhancing technologies (PETs).
- This goes to demonstrate that enhanced network transparency can be achieved without compromising the integrity or security of critical infrastructure.

## 8.5 Limitations and Proposed Future Studies

It is worth reiterating that the European Network and Information Security Agency (ENISA) concluded in a 2015 report that the current lack of structural transparency is the biggest obstacle to addressing the inherent vulnerabilities and architectural shortcomings of the Internet routing system (ENISA, 2015).

The new method described in Chapter 5 increases the accuracy of IP geolocation by a large amount, from 13% to 33.6% using a highly curated set of input data. However, the fundamental issues of using traceroute as discussed in Section 2.3 still apply to this method and, as Motamedi et al. point out, all of these methods are little more than engineering hacks and may produce large discrepancies between what we think the network looks like and what it is in reality (Motamedi, Rejaie, and Willinger, 2015) .

The introduction of an IPv6 network infrastructure mapping tool, as discussed in Chapter 6, is designed to enhance and expand the capabilities of existing mapping tools such as traceroute. It offers a new and comprehensive approach to capture detailed and real-time information on the Internet structure, including IP addresses of router interfaces, their peers, and geolocations. This tool is aimed at providing researchers and network administrators with a sophisticated visualisation of the Internet topology that presents a significant step forward in the precision and comprehensiveness of mapping the Internet infrastructure.

The effectiveness of the tool proposed in Chapter 6 can be compromised in areas with low adoption of IPv6 or outdated infrastructures. Its deployment strategy

should include modular adaptations that can still operate under limited IPv6 use. This approach ensures that even networks at different stages of IPv6 integration can benefit from the tool, promoting a wider adoption of IPv6 through demonstrated efficacy.

In Chapter 7 we acknowledge the drawbacks of our new IPv6 extension header, for example:

- Using ICMP replies to warn of failures in routing can be prone to flooding.
- Debugging of network problems is made harder due to the additional checks on whether failures are due to geofencing.
- Older routers and middleboxes would have to be replaced.

However, ICMP flooding issues have been around for a long time and the problems need to be resolved, noting that ongoing ICMP issues are outside the scope of this thesis. All equipment has to be replaced at some point, and therefore older routers will eventually give way to new equipment. New tools specifically designed for network analysis can include geofencing diagnosis, as discussed in Chapter 8. Hence, all of these problems are not insurmountable given enough time along with the will to succeed.

As we explore innovative methods to improve IPv6 geolocation and routing, leveraging the P4 programming language presents a promising approach. P4 allows for the programming of the data plane, which enables the implementation of geolocation checking directly where packet forwarding occurs. This approach offers several advantages, as follows.

1. Performance and efficiency: by performing geolocation checks in the data plane, we can achieve real-time processing speeds, minimizing latency and improving overall network performance. P4 allows us to embed geolocation functionality directly into the hardware-level forwarding logic (fast path), as discussed in 7.4, eliminating the need to process packets in the slow path, and ensuring real-time performance with minimal latency.

2. Flexibility: P4's programmability allows us to define and update geolocation policies dynamically, adapting to changing geopolitical considerations and network conditions.

3. Security: real-time geolocation checks improve security by ensuring that data packets comply with geopolitical routing policies, reducing the risk of unauthorised data transit.

4. Scalability: the distributed nature of data-plane operations ensures that our geolocation checking mechanism can scale with increasing network traffic and complexity.

Future work should look at the possibilities of developing P4 programs that use IPv6 extension headers to extract and verify geolocation information. These programs could be integrated into network devices, ensuring seamless operation with existing infrastructure. State management techniques will be used to efficiently handle geolocation data, ensuring that data plane operations remain performant. Rigorous tests could be performed to validate the accuracy and security of the P4 programs, ensuring that they meet the desired performance and reliability standards. By integrating geolocation checks directly into the data plane, P4 allows network devices to make routing or geolocation-related decisions without forwarding packets to the control plane or slow path for additional processing. This would enhance the capabilities of the IPv6 geopolitical routing method described in Chapter 7, providing a robust and scalable solution for modern Internet infrastructure. The testing of the new IPv6 Extension Header described in Chapter 7 could be carried out using network simulation software such as GNS3 or Cisco Packet Tracer, which allows the creation of complex network topologies using virtual routers and switches. Quagga can be used to test and experiment with routing decisions and is an open-source routing software suite that provides implementations of various routing protocols. FRRouting (FRR) is an open source routing stack that supports multiple routing protocols. The BIRD Internet Routing Daemon is a lightweight routing

daemon that supports IPv4 and IPv6. It is commonly used for testing and small-scale deployments. However, it seems that none of these packages provides a simple method of introducing and processing a new IPv6 Extension Header, and this would be a useful addition to any of these solutions by providing researchers and engineers with a simple method of testing new ideas.

Experience from testing is an expected input to any decision to progress this specification. Appropriate inputs might include:

- Reports of implementation experience;
- Measurements of the number of paths where the method can be used;
- Measurements showing the benefit realised or the implications of using different options in the Extension Header.

Although the idea of using country codes inside the Hop-By-Hop IPv6 extension header seems a fine approach, current practical experience with Hop-By-Hop headers in the public IPv6 Internet makes the prospects for this header a challenge. The key functionality of IPv6 extension headers is believed still to be a solid foundation for the ongoing growth and diversification of the Internet and demonstrates its adaptability to future Internet needs and challenges. It is nevertheless appropriate and timely to focus on extension header problem resolution, in order for IPv6 to reach its true potential.

The new tools and methods detailed in Chapters 6 and 7 are based on the implementation of IPv6 on the global stage. In 2020, researchers were lamenting the slow deployment of IPv6 and the imminent depletion of IPv4 addresses (Livadariu, Elmokashfi, and Dhamdhere, 2020). However, recently the trend appears to have turned around; for example, Google is showing a 15% increase in IPv6 users since 2020 to 45% in 2024 who use IPv6 to access their site (Google, 2024). The number of BGP entries in the router BGP forwarding table has increased from 80,000 entries in 2020 to more than 200,000 in 2024, as shown in Figure 8.1.

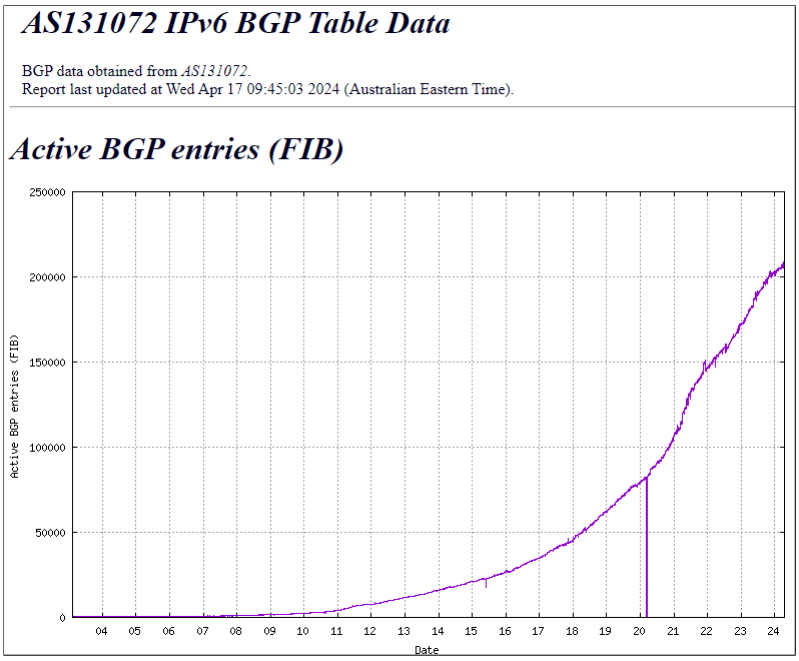


Figure 8.1: Active BGP Entries (FIB) (Huston, 2024)

IPv6 offers so many opportunities for Internet improvement that resolving extension header security issues should be a major focus of researchers and analysts. These conclusions emphasise the need for a strategic focus on the security and functionality enhancements provided by IPv6 extension headers in the development of the IPv6 network infrastructure mapping tool. By resolving these critical security issues and demonstrating the practical benefits of IPv6, the tool would not only advance network management capabilities but also encourage the resolution of long-standing challenges associated with IPv6 deployment globally. This forward-looking approach will help catalyse the broader adoption and optimisation of IPv6 across diverse regions and network setups.

## 8.6 Final Thoughts

The advances offered in this thesis represent significant steps towards a more secure and resilient Internet infrastructure. By leveraging advanced geolocation techniques, innovative IPv6 extension headers, and comprehensive network mapping tools, the

work has laid the groundwork for a future in which IP geolocation is greatly enhanced, and this will lead to more appropriate BGP routing decisions. Continued research and collaboration will be essential to fully realise the potential of these technologies and ensure the ongoing stability of the global Internet infrastructure.

# Appendices

# Appendix A

# Python Scripts

## A.1 Introduction

The software tools and scripts developed for this thesis were written by the author, who is a self-taught programmer and does not possess formal training in software development. While every effort has been made to ensure the accuracy, reliability, and clarity of the code, it is acknowledged that the solutions implemented may not reflect best practices in professional software engineering. The focus throughout has been on achieving functional outcomes that support the research aims of this thesis.

## A.2 Scripts

The following scripts can be found at `https://github.com/Cobber57/thesis`

### A.2.1 Create-Measurements.py

This script automates the creation of traceroute measurements between all active RIPE Atlas anchors located in the UK. It helps build a comprehensive dataset showing how packets travel within the UK Internet infrastructure. The measurements are stored for later use in mapping tools that support fine-grained Internet topology analysis (as explored in Chapters 4 and 5 of the thesis). Creates the initial

trace route measurements between 32 anchors in the UK and adds the measurement info to a file in the measurements folder for use in the read-measurements script.

### A.2.2 create-ixp-list-and-facility-list.py

Creates a list of UK facilities and UK IXPs, uses local synced peeringdb database data to discover geolocation and nominatum as secondary search and finally manual entry as third option.

### A.2.3 Create-uk-facilities-to-networks.py

This script interacts with the PeeringDB API using the ‘peeringdb‘ Python library to extract, enrich, and store network facility data focused on the United Kingdom. It aims to create a geolocated mapping of all networks that peer at UK facilities, including fallback geolocation strategies for facilities without coordinates.

### A.2.4 create-html-facilities-map.py

This script creates an HTML map that visualizes Internet exchange points (IXPs) and colocation facilities in the UK. It processes PeeringDB data, extracts facility and organization details, and generates a Leaflet.js compatible HTML page.

### A.2.5 Read-measurements.py

read-measurements.py reads in the uk-measurments file created by the create-measurements script and accesses the RIPE ATLAS measurements to create a spreadsheet of distances vs RTT times in results/results.xlsx also saves all this infomation to results/targets.json file

### A.2.6 Create-html.py

Create-html.py reads in the file created by the read-measurements script and creates an HTML page showing the geolocations of the probes and creates the map fully

working upto finding the correct IXP entry and exit points

### A.2.7 create-ixp-list-and-facility-list.py

Create a new uk-facilities JSON file peeringdb-test-results/uk-facilities-new.json AND A NEW UK-ips FILE

### A.2.8 check-caida-geolocations.py

Compares IP geolocations against known locations in the LINX and LONAP databases to analyse how accurate CAIDA locations are.

### A.2.9 create-caida-dbs.py

Reads in each node from CAIDA db and inserts them into the ips table but the ips insertions did not work due to size of database so will use the 'data/CAIDA/gb-nodes.json' file to create the ips table instead.

### A.2.10 analyse-results-file.py

Tests for how many records were found without a rDNS lookup and no facility. Tests for how many records were found without a rDNS lookup but had a facility. Tests for How many were succesfully looked up using Rule 5

### A.2.11 myipgeolocator-v19.py

Takes measurements saved from RIPE ATLAS and discovers each hops geographical locations Uses the results/targetsfull.json file created by read-measurments4.py, can use the entire file but to limit how much to use whilst testing, copy and paste just the target and source probe ids into a new file, ie target-6087.json for one target and multiple sources or, target-6087-source-6515.json for just a single source and target.

### A.2.12 myreversedns9.py

Take results and discover the hops locations by using reverse lookup.

# Appendix B

# Data Files

## B.1 Results

### B.1.1 Tables subfolder

Contains the results of all the tests between source and destination probes, filenames are in the format *DESTINATION*.json

### B.1.2 VPtable subfolder

Contains results of discovered Vantage Points (i.e., IXPs) and the facilities they use.

### B.1.3 Results folder JSON files

Contains files which have various results the name of the file depicts just the target or a source to target.

### B.1.4 Peeringdb Results

The following files are in the peeringdb-test-results subfolder.

#### B.1.4.1 uk-facilities-new.json

Lists created from script A.2.2 which lists all UK facilities. Also creates uk-facilities-none.json which is a list of facilities that we could not geolocate.

#### B.1.4.2 uk-ixps.json

A list created from script A.2.2 which lists all UK IXPs and the facilities that they interconnect.

#### B.1.4.3 Uk-Facilities-to-networks-good.json

A list created from script A.2.3 which identifies all the networks interconnecting at each facility. Used in A.2.4 to create a map of the facilities.

#### B.1.4.4 networks-all.json

A listing of all networks sourced from peeringdb

#### B.1.4.5 ipprefixes-all.json

A listing of all IXP and their prefixes sourced from peeringdb.

### B.1.5 Web subfolder

Html files which display the geolocation of source and targets used throughout this thesis.

# Appendix C

# Databases

## C.1 CAIDA Files

Please see https://doi.org/10.21986/CAIDADATA.ARK-ITDK for the caida.db database. Please see https://icloud9.co.uk/phd/thesis for extremely large data files created from the CAIDA database during this thesis:

gb2-nodes.json

gb-nodes.json

midar-iff.nodes

midar-iff.nodes.geo

## C.2 Place Names

Index of Place Names in Great Britain (IPN) 2021 can be downloaded at https://geoportal.statistics.gov.uk/datasets/ons::index-of-place-names-november-2021-in-g about

## C.3 IXP Databases

IX-info subfolder contains LINX and LONAP data downloaded from their resepective sites.

# References


Abley, J. and K. Lindqvist (2006). *Operation of Anycast Services*. https://datatracker.ietf.org/doc/html/rfc4786.

Akamai (2021). *Akamai Edge DNS Causes Massive Outage*. https://constellix.com/news/akamai-edge-dns-causes-massive-outage-2021.

Aktas, F., I. Shayea, M. Ergen, B. Saoud, A. Yahya, and A. Laura (2025). "AI-enabled routing in next generation networks: A survey". In: *Alexandria Engineering Journal*. Vol. 120. Elsevier. DOI: https://doi.org/10.1016/j.aej.2025.01.095.

Alshaer, H., N. Uniyal, K. Katsaros, K. Antonakoglou, S. Simpson, H. Abumarshoud, H. Falaki, P. McCherry, C. Rotsos, T. Mahmoodi, R. Nejabati, D. Kaleshi, D. Hutchison, H. Haas, and D. Simeonidou (2020). "The UK Programmable Fixed and Mobile Internet Infrastructure: Overview, Capabilities and Use Cases Deployment". In: *IEEE Access, Vol 8*. Insititute Of Electrical and Electronic Engineers, pp. 175398–175411. DOI: 10.1109/ACCESS.2020.3020894.

Arkko, J., J. Kempf, B. Zill, and P. Nikander (2005). *SEcure Neighbor Discovery (SEND)*. https://datatracker.ietf.org/doc/html/rfc3971.

Augustin, B., X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira (2006). "Avoiding Traceroute Anomalies with Paris Traceroute". In: *6th ACM SIGCOMM conference on Internet measurement October*, pp. 153–158.

Bartels, R. (2024). *The ICMP Dilemma: Why Blocking it Makes You a Networking Noob*. https://www.linkedin.com/pulse/icmp-dilemma-why-blocking-makes-you-networking-noob-ronald-bartels-ikvnf/.

Berkeley (2024). *Internet Atlas*. https://cltc.berkeley.edu/program/internet-atlas/.

Bird (2024). *Internet Routing Daemon*. https://bird.network.cz/.

Blumberg, B., M. Pihelgas, M. Kont, O. Maennel, and R Vaarandi (2016). "Creating and Detecting IPv6 Transition Mechanism-Based Information Exfiltration Covert Channels". In: *NordSec 2016. Lecture Notes in Computer Science Vol. 10014*. Springer. DOI: https://doi.org/10.1007/978-3-319-47560-8_6.

Bosshart, P., D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayaco, A. Vahdat, G. Verghese, and D. Walker (2014). "P4: Programming Protocol-Independent Packet Processors". In: *Computer Communication Review 2014*. DOI: http://dx.doi.org/10.1145/2656877.2656890.

Braden, R. (1989). *Requirements for Internet Hosts – Communication Layers*. https://www.ietf.org/rfc/rfc1122.txt.

Bradner, S. and V. Paxson (200). *IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers*. https://www.rfc-editor.org/rfc/rfc2780.

CAIDA (2024). *ITDK: Internet Topology Data Ki*. https://doi.org/10.21986/CAIDADATA.ARK-ITDK. DOI: https://doi.org/10.21986/CAIDADATA.ARK-ITDK.

Candela, M., E. Gregori, V. Luconi, and A. Vecchio (2019). "Using RIPE Atlas for Geolocating IP Infrastructure". In: *IEEE Access 7*.

Carpenter, B. and S. S. Jiang (2013). *Transmission and Processing of IPv6 Extension Headers*. https://www.rfc-editor.org/rfc/rfc704.

Cisco (2006). *IPv6 Extension Headers Review and Considerations*. https://www.cisco.com/en/US/technologies/tk648/tk872/technologies_white_paper0900aecd8054d37d.html#wp9000069.

Conta, A., S. Deering, and M. Gupta (2006). *Internet Control Message Protocol*. https://www.rfc-editor.org/rfc/rfc4443.pdf.

Crawford, M. and B. Haberman (2006). *IPv6 Node Information Queries*. https://datatracker.ietf.org/doc/html/rfc4620.

Daigle, L. (2004). *WHOIS Protocol Specification*. https://datatracker.ietf.org/doc/html/rfc3912.

Dan, O., V. Parikh, and B.D. Davison (2021a). "IP Geolocation through Reverse DNS". In: *Internet Technol. 22, 1, Article 17*. DOI: \url{https://doi.org/10.1145/3457611}.

— (2021b). "IP Geolocation Using Traceroute Location Propagation and IP Range Location Interpolation". In: *Companion Proceedings of the Web Conference 2021 (WWW '21 Companion)*.

Davies, E., J. Mohacsi, and I. van Beijnum (2007). *Recommendations for Filtering ICMPv6 Messages in Firewalls - Details security measures for ICMPv6 and potential attack mitigations*. https://datatracker.ietf.org/doc/html/rfc4890.

Davis, C., P. Vixie, T. Goodwin, and I. Dickinson (1996). *A Means for Expressing Location Information in the Domain Name System*. https://datatracker.ietf.org/doc/html/rfc1876.

Deering, S. and R. Hinden (2017). *Internet Protocol, Version 6 (IPv6) Specification*. https://www.rfc-editor.org/rfc/rfc8200.

Ding, S., X. Luo, Y. Dengpan, and F. Liu (Apr. 2017). "Delay-distance correlation study for IP geolocation". In: *Wuhan University Journal of Natural Sciences* 22, pp. 157–164. DOI: 10.1007/s11859-017-1229-2.

Du, B., M. Candela, B. Huffaker, A. Snoeren, and K. Claffy (2020). "RIPE IPmap Active Geolocation: Mechanism and Performance Evaluation". In: *ACM SIGCOMM Computer Communication Review Volume 50 Issue 2*.

Duffy, C. (2021). *Two Obscure Service Providers Briefly Broke the Internet. It Could Happen Again*. https://www.cnn.com/2021/06/09/tech/fastly-cdn-internetrisk/index.html.

eBPF (2024). *Dynamically program the kernel for efficient networking, observability, tracing, and security*. https://ebpf.io/.

ENISA (2015). *Threat Landscape of Internet Infrastructure*. https://www.enisa.europa.eu/publications/iitl.

Filfils, C., K. Talaulikar, D. Voyer, A. Bogdanov, and P. Mattes (2022). *Segment Routing Policy Architecture*. https://datatracker.ietf.org/doc/rfc9256/.

Filsfils, Clarence, Nagendra Kumar Nainar, Carlos Pignataro, Juan Camilo Cardona, and Pierre Francois (2015). "The Segment Routing Architecture". In: *2015 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–6. DOI: 10.1109/GLOCOM.2015.7417124.

Gharaibeh, M., A. Shah, B. Huffaker, H. Zhang, R. Ensafi, and Papadopoulos. C. (2017). "A look at router geolocation in public and commercial databases". In: *Proceedings of the 2017 Internet Measurement Conference*. ACM, pp. 463–469.

Gill, P., M. Arlitt, Z. Li, and A. Mahanti (2008). "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse?" In: *Passive and Active Network Measurement. Lecture Notes in Computer Science, vol 4979*. DOI: \url{https://doi.org/10.1007/978-3-540-79232-1/_1}.

Giotsas, V., G. Nomikos, V. Kotronis, P. Sermpezis, P. Gigis, L. Manassakis, C. Dietzel, S. Konstantaras, and X. Dimitropoulos (2021). "O Peer, Where Art Thou? Uncovering Remote Peering Interconnections at IXPs". In: *IEEE/ACM Transactions on Networking*. Vol. 29, pp. 1–16.

Giotsas, V., G. Smaragdakis, B. Huffaker, M. Luckie, and K. Claffy (2015). "Mapping peering interconnections to a facility". In: *Proceedings of the 11th ACM Conference on emerging networking experiments and technologies*, pp. 1–13.

Gont, F. (2014). *Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard))*. https://datatracker.ietf.org/doc/html/rfc7113.

Google (2024). *IPv6 Statistics*. https://www.google.com/intl/en/ipv6/statistics.html=ipv6-adoption.

UK-Government (2023). *Data protection*. https://www.gov.uk/data-protection.

Graham-Cumming, J. (2014). *The weird and wonderful world of DNS LOC records*. https://blog.cloudflare.com/the-weird-and-wonderful-world-of-dns-loc-records/. Accessed: March 2023.

Gueye, B., A. Ziviani, M. Crovella, and S. Fdida (2006). "Constraint Based Geolocation". In: *Transactions on Networking, Vol. 14, No. 6*. IEEE/ACM.

Herbert, T. (2024). *IPv4 Extension Headers and Flow Label*. https://datatracker.ietf.org/doc/html/draft-herbert-ipv4-eh-03.

Hinden, B. and S. E. Deering (1998). "Internet Protocol, Version 6 (IPv6) Specification". In.

Hinden, R. and G. Fairhurst (2020). *IPv6 Hop-by-Hop Options Processing Procedures*. https://datatracker.ietf.org/doc/html/draft-hinden-6man-hbh-processing-00.

— (2021). *IPv6 Hop-by-Hop Options Processing Procedures*. https://datatracker.ietf.org/doc/html/draft-hinden-6man-hbh-processing-01.

Holterbach, T., C Pelsser, R. Bush, and L. Vanbever (2015). "Quantifying interference between measurements on the RIPE Atlas platform". In: *Proceedings of the 2015 Internet Measurement Conference*.

Huston, G. (2022a). *Hop-by-hop extension headers*. https://blog.apnic.net/2022/04/22/hop-by-hop-extension-headers/.

— (2022b). *IPv6 Extension Headers Revisited*. https:https://blog.apnic.net/2022/10/13/ipv6-extension-headers-revisited/.

— (2024). *Internet Statistics*. https://bgp.potaroo.net/as2.0/bgp-active.html.

ICO (2022a). *Standard Data Protection Clauses to be issued by the Commissioner under S119A(1) Data Protection Act 2018*. https://ico.org.uk/media/

for-organisations/documents/4019538/international-data-transfer-agreement.pdf.

— (2022b). *Standard Data Protection Clauses to be issued by the Commissioner under S119A(1) Data Protection Act 2018*. https://ico.org.uk/media/for-organisations/documents/4019539/international-data-transfer-addendum.pdf.

— (2024). *Privacy Enhancing Technolgies*. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies/what-pets-are-there/introduction/.

InternetSociety (2021). *Effective IXP strategies for the Asia-Pacific*. https://www.internetsociety.org/wp-content/uploads/2021/06/Internet-Peering-in-Asia-Pacific-EN.pdf.

IP2location (2024). *Geolocate IP Address Location using IP2Location*. https://www.ip2location.com/.

Iurman, J. and B. Donnet (2020). "Ipv6 In-Situ Operations, Administration, and Maintenance". In: *Software Impacts Vol 6*. DOI: \url{https://doi.org/10.1016/j.simpa.2020.100036.}.

Iurman, J., E. Vyncke, and B. Donnet (2023). *Using eBPF to inject IPv6 Extension Headers*. https://orbi.uliege.be/bitstream/2268/309796/1/paper.pdf.

IXPDB (2024). *The IXP database*. https://ixpdb.euro-ix.net/en/.

Kapko, M. (2024). *New instruments for data transfer from the UK: UK Extension to the EU-US DPF and IDTA/UK Addendum*. https://legalitgroup.com/en/new-instruments-for-data-transfer-from-the-uk-uk-extension-to-the-eu-us-dpf-and-idta-uk-addendum/.

Kasnavi, S., P. Berube, V. Gaudet, and J. Amaral (2008). "A cache-based internet protocol address lookup architecture". In: *Computer Networks*. Vol. 52. Elsevier. DOI: doi:https://doi.org/10.1016/j.comnet.2007.08.010..

Katz-Bassett, E., J.P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe (2006). "Towards IP geolocation using delay and topology mea-

surements". In: *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pp. 71–84. DOI: \url{https://doi.org/10.1145/1177080.1177090}.

Kent, S. (2005a). *IP Authentication Header*. https://www.rfc-editor.org/rfc/pdfrfc/rfc4302.txt.pdf.

— (2005b). *IP Encapsulating Security Payload (ESP)*. https://www.rfc-editor.org/rfc/pdfrfc/rfc4303.txt.pdf.

Kline, E., K. Duleba, Z. Szamonek, S. Moder, and W. Kumari (2020). *A Format for Self-Published IP Geolocation Feeds*. https://datatracker.ietf.org/doc/html/rfc8805.

Klöti, R., B. Ager, V. Kotronis, G. Nomikos, and X. Dimitropoulos (2016). "A comparative look into public IXP datasets". In: *ACM SIGCOMM Computer Communication Review 46, no. 1.*

Komosny, D., M. Voznak, and S. Rehman (July 2017). "Location Accuracy of Commercial IP Address Geolocation Databases". In: *Journal of Information Technology and Control Vol. 46 / No. 3*. Kaunas University, pp. 333–344. DOI: doi.org/10.5755/j01.itc.46.3.14451.

Kudou, K., S. Suzuki, J. Hagino, K. Yamamoto, K. Shima, K. Uehara, R. Wakikawa, K. Mitsuya, T. Momose, T. Jinmei, and Sakane S. (2006). *The KAME Project*. https://www.kame.net/.

Laki, S., P. Mátray, P. Hága, T. Sebok, I. Csaba, and G. Vatta (2011). "Spotter: A Model Based Active Geolocation Service". In: *Proceedings IEEE INFOCOM*.

Lawbite (2022). *International Data Transfer Agreements (Idtas) Explained*. https://www.lawbite.co.uk/resources/blog/international-data-transfer-agreement-idta.

Livadariu, I., T. Dreibholz, A.A. Al-Selwi, H. Bryhni, O. Lysne, S. Bjørnstad, and A. Elmokashf (2020). "On the Accuracy of Country-Level IP Geolocation". In: *ANRW '20: Proceedings of the Applied Networking Research Workshop*.

Livadariu, I., A. Elmokashfi, and A. Dhamdhere (2020). "An agent-based model of IPv6 adoption". In: *2020 IFIP Networking Conference (Networking)*, pp. 361–369.

Luckie, M., A. Dhamdhere, B. Huffaker, and D. Clark (2016). "bdrmap:Inference of borders between IP networks". In: *Proceedings. ACM IMC*, pp. 381–396.

Luckie, M., B. Huffaker, A. Marder, Z. Bischof, M. Fletcher, and K. Claffy (2021). "Learning to Extract Geographic Information from Internet Router Hostnames". In: *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies*. DOI: `\url{https://doi.org/10.1145/3485983.349486}`.

Marder, A. and J. Smith (2016). "MAP-IT: Multipass accurate passive inferences from traceroute". In: *Proceedings. ACM IMC*, pp. 397–411.

Maxmind (2024). *Detect Online Fraud and Locate Online Visitors*. `https://www.maxmind.com/en/home`.

McCherry, P., V. Giotsas, and D. Hutchison (May 2023). "On Improving the Accuracy of Internet Infrastructure Mapping". In: *IEEE Access, vol 11*. Insititute Of Electrical and Electronic Engineers, pp. 59935–59953. DOI: `10.1109/ACCESS.2023.3281333`.

McCloghrie, K. and M. Rose (1991). *Management Information Base for Network Management of TCP/IP-based internets*. `https://www.rfc-editor.org/rfc/rfc1213`.

Medina, A. (2021). *Inside the Fastly Outage: Analysis and Lessons Learned*. `https://www.thousandeyes.com/blog/inside-the-fastly-outage-analysis-and-lessons-learned`.

Merrill, N. and T. Narechania (2023). "Inside the Internet". In: *Transactions on Networking, 73 Duke Law Journal Online 35*.

Mlab (2024). *Mlab*. `https://www.measurementlab.net/tests/`.

Motamedi, R., R. Rejaie, and W. Willinger (May 2015). "A Survey of Techniques for Internet Topology Discovery". In: *Communications Surveys and Tutorials 17, issue 2*. IEEE.

Motamedi, R., B. Yeganeh, B. Chandrasekaran, R. Rejaie, B. Maggs, and W. Willinger (Oct. 2019). "On Mapping the Interconnections in Today's Internet". In: *Transactions on Networking 27, no. 5*. IEEE/ACM, pp. 2056–70. DOI: doi.org/10.1109/TNET.2019.2940369.

Mühlbauer, W., A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig (2006). "Building an as-topology model that captures route diversity". In: *ACM SIGCOMM CCR, 36(4)*, pp. 195–206.

Nakryiko, A. (2020). *BPF CO-RE (Compile Once – Run Everywhere)*. https://nakryiko.com/posts/bpf-portability-and-co-re/.

NCSC (2021). *Responsible use of the Border Gateway Protocol (BGP) for ISP Internetworking*. https://www.ncsc.gov.uk/files/border-gateway-protocol-technical-paper.pdf. Accessed: 2022-06-12.

Neustar (2024). *Highest Quality IP Decisioning Data*. https://ipintelligence.neustar.biz/portal/.

Ngari, L. and S. A. Petrack (2024). *Internet Infrastructure in Africa*. https://www.empowerafrica.com/internet-infrastructure-in-africa/.

nominatum (2024). *nominatum*. https://www.nominatum.com.

Nsnam (2025). *Network Simulator 3)*. https://www.nsnam.org/.

OECD (2024). *Peer-to-Peer (P2P) Exchange Guidelines*. http://bit.ly/3Rgj5hn.

Oliveira, R., D. Pei, W. Willinger, B. Zhang, and L. Zhang (2010). "The (in) completeness of the observed Internet as-level structure". In: *Transactions on Networking, vol 18, issue 1*. IEEE/ACM, pp. 109–122.

Omnet++ (2025). *Omnet++ Discrete Event Simulator*. https://omnetpp.org/.

Openstreetmap (2024). *Openstreetmap*. https://www.openstreetmap.com.

Padmanabhan, V. and L. Subramanian (2001). "An Investigation of Geographic Mapping Techniques for Internet Hosts". In: *SIGCOMM Computer Communi-*

*cation Review Vol. 31, No. 4*. ACM, pp. 173–185. DOI: https://doi.org/10.1145/964723.383073.

Pailisse, J., J. Manrique, G. Bonet, A. Rodriguez-Natal, F Maino, and A. Cabellos (2019). "Decentralized Trust in the Inter-Domain Routing Infrastructure". In: *IeeeAcess volume 7, pages 166896-166905*. Ieee. DOI: 10.1109/ACCESS.2019.2954096.

PCH (2024). *Packet Clearing House*. https://www.pch.net/.

PeeringDB (2024). *The Interconnection Database*. https://www.peeringdb.com/.

Poese, I., S. Uhlig, M.A. Kaafar, B. Donnet, and B. Gueye (2011). "IP geolocation databases: Unreliable?" In: *SIGCOMM Computer Communication Review 41, Issue 2*. ACM, pp. 53–56.

Reddy, T., D. Wing, and Boucadair M. (2023). *An Approach for Encrypted Transport Protocol Path Explicit Signals*. https://www.ietf.org/archive/id/draft-reddy-tsvwg-explcit-signal-01.html.

Reuters (2023). *Singtel-owned Optus says massive Australia outage was after software upgrade*. https://www.reuters.com/business/media-telecom/singtel-owned-optus-says-massive-australia-outage-was-after-software-upgrade-2023-11-13/.

RIPE (2015). "RIPE Atlas: A Global Internet Measurement Network". In: *The Internet Protocol Journal, Vol 18, No 3*. RIPE.

Aben, E. (2013). "Router Geolocation". In: *Computer Networks 54(9):1490–1501 54(9):1490–1501*.

Rockwell, N. (2021). *Summary of June 8 Outage, FASTLY*. https://www.fastly.com/blog/summary-of-june-8-outage.

Rosen, E., A. Viswanathan, and R. Callon (2001). *Multiprotocol Label Switching Architecture*. https://datatracker.ietf.org/doc/html/rfc3031.

Rotsos, C., A. Marnerides, A. Magzoub, A. Jindal, P. McCherry, M. Bor, J. Vidler, and D. Hutchison (Nov. 2020). "Ukko: Resilient DRES management for Ancillary Services using 5G service orchestration". In: *2020 IEEE International Conference*

*on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, pp. 1–6. DOI: 10.1109/SmartGridComm47815.2020.9302980.

RouteViews (2024). *RouteViews Project*. https://www.routeviews.org/routeviews/.

Sanka, A. and C. Cheung (2021). "A systematic review of blockchain scalability: Issues, solutions, analysis and future research". In: *Journal of Network and Computer Applications*. Vol. 195. ScienceDirect. DOI: https://doi.org/10.1016/j.jnca.2021.103232.

Scheitle, Q., O. Gasser, P. Sattler, and G. Carle (2017). "HLOC: Hints-based geolocation leveraging multiple measurement frameworks". In: *2017 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 1–9. DOI: 10.23919/TMA.2017.8002903.

Shavitt, Y. and N. Zilberman (2011). "A Geolocation Databases Study". In: *Journal on Selected Areas in Communications 29, 10*. IEEE, pp. 2044–2056.

Simpson, S., A. Farshad, P. McCherry, A. Magzoub, W. Fantom, C. Rotsos, N. Race, and D. Hutchison (Nov. 2019). "DataPlane Broker: Open WAN control for multi-site service orchestration". In: *2019 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, pp. 1–6. DOI: 10.1109/NFV-SDN47374.2019.9040084.

Tor (2025). *Censorships*. https://support.torproject.org/censorship/.

Tracol, X. (2020). ""Schrems II": The return of the Privacy Shield"". In: *Computer Law & Security Review, vol. 39, p. 105484*. DOI: https://doi.org/10.1016/j.clsr.2020.105484.

UN (2024). *Member States*. https://www.un.org/en/about-us/member-states.

University, Princetown (2025). *Planetlab*. https://planetlab.cs.princeton.edu/.

Vyncke, E., K. Chittimaneni, M. Kaeo, and E. Rey (2021). *Operational Security Considerations for IPv6*. https://www.rfc-editor.org/rfc/rfc9099.pdf.

Wang, Y., D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang (2011). "Towards street level client-independent ip geolocation". In: *the USENIX Symposium on Networked Systems Design and Implementation (NSDI)*.

Willinger, W. and M. Roughan (2013). "Internet Topology Research Redux". In: *ACM SIGCOMM eBook: Recent Advances in Networking*.

Winter, P., R. Padmanabhan, A. King, and A. Dainotti (2019). "Geo-locating BGP prefixes". In: *Network Traffic Measurement and Analysis Conference (TMA)*, pp. 9–16. DOI: 10.23919/TMA.2019.8784509.

Wong, B., I. Stoyanov, and E.G. Sirer (2007). "Octant: A Comprehensive Framework for the Geolocalization of Internet Hosts". In: *NSDI '07: 4th USENIX Symposium on Networked Systems Design and Implementation*.

Zu, S., Z. Luo, and F. Zhang (2022). "IP-geolocater: a more reliable IP geolocation algorithm based on router error training". In: *Frontiers of Computer Science, 2022, Vol.16 (1), Article no.161504*.