



**The Good, the Bad, and the Ugly
of BGP Routing Modeling:
Confounding Factors, Selective
Announcements and
Location-aware Simulations**

Savvas V. Kastanakis, BSc, MSc
School of Computing and Communications
Lancaster University

A thesis submitted for the degree of
Doctor of Philosophy

January 15, 2025

**The Good, the Bad, and the Ugly of BGP Routing Modeling:
Confounding Factors, Selective Announcements and Location-aware
Simulations**

Savvas V. Kastanakis, BSc, MSc.

School of Computing and Communications, Lancaster University

A thesis submitted for the degree of *Doctor of Philosophy*. January 15, 2025.

Abstract

Over the past quarter-century, substantial research has been devoted to the accurate inference/prediction of AS-level paths on the Internet. Literature, though, has shown that BGP simulations can be highly inaccurate making it hard to take "in-vitro" evaluations at face value. This PhD thesis investigates the complexities of this long standing inference problem by exploring various confounding factors, analyzing the evolution of interdomain routing policies, and examining the potential of geolocation-aware inference models.

As a first step, we identify and quantify the confounding factors to accurate AS-path inference through passive measurements and controlled experiments. From our analysis, two factors are highlighted: (a) the complex routing policies implemented by ASes, and (b) the geolocation agnostic BGP best path selection process.

To delve deeper into the first confounding factor, we explore the evolution of the interdomain routing policies over the last 20 years. We show that while the phenomenon of selective announcements is persistent, not only across time, but also across networks, the influence of AS relationships on path selection has diminished. This underscores the importance of frequent BGP policy inference to keep pace with the evolving landscape of AS connectivity.

As we investigate the second confounding factor, we explore the role of geospatial attributes in geolocation-based routing paradigms such as IP anycast routing. Our analysis shows that 84.06% of anycast ASes rely on selective announcements, a

phenomenon largely influenced by the geolocation-agnostic nature of the BGP best path selection process. Additionally, we demonstrate that anycast ASes follow different routing patterns per geographical regions.

Our work supports the need for more flexible routing models and can aid in the understanding of a variety of interdomain routing applications, such as the measurement of the RPKI adoption, fine-grained interdomain policy learning, interdomain routing verification, privacy-preserving routing and studying routing attacks.

Acknowledgements

This thesis would not have been possible without the support and influence of many individuals.

First and foremost, I extend my heartfelt gratitude to my dear friends Konstantina Fragkouli and Rafail Tsirbas for their unwavering support over the years. I am especially grateful for their hospitality, hosting me in their apartment in Cambridge through both good and challenging times during my Ph.D. studies. Their companionship on drives through valleys, forests and mountains helped me explore both the world around me and the world inside me.

I would also like to express my sincere thanks to my supervisor, Prof. Neeraj Suri, for introducing me to the first-principles thinking and guiding me through both the theoretical and technical aspects of my research. His mentorship has been invaluable, enabling numerous research trips and conference participations, and assisting with the challenging aspects of my Ph.D. The knowledge I gained under Prof. Suri's guidance surpasses the value of any future degrees I might obtain.

I was very fortunate to be guided by my second supervisor, Dr. Vasileios Giotsas. Vasileios is a pioneer in interdomain routing, and I consider myself lucky to have been mentored by one of the foremost researchers in the field. His friendly demeanor, exceptional technical skills, and profound understanding of interdomain routing have saved me considerable time and effort. His guidance has significantly shaped my approach to interdomain routing simulations and scientific research in general.

I am and will always be deeply grateful to my parents for their unconditional and multifaceted love and support, without which achieving my academic goals would have been impossible. Through their own example, they have taught me the importance of determination and compassion in pursuing one's aspirations.

Finally, though it may read unusual, I want to acknowledge and thank myself for persevering through the tough times—when experiments failed, when reality fell short of my expectations, and when the path to earning this Ph.D. seemed like a daunting mountain to climb.

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: 21854

Savvas V. Kastanakis

Publications

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

Savvas Kastanakis, Vasileios Giotsas, and Neeraj Suri (2022). “Understanding the confounding factors of inter-domain routing modeling”. In: *Proceedings of the 22nd ACM Internet Measurement Conference*, pp. 758–759

Savvas Kastanakis, Vasileios Giotsas, Ioana Livadariu, and Neeraj Suri (2023c). “Replication: 20 Years of Inferring Interdomain Routing Policies”. In: *Proceedings of the 2023 ACM on Internet Measurement Conference*, pp. 16–29

Savvas Kastanakis, Vasileios Giotsas, Ioana Livadariu, and Neeraj Suri (2024). “Investigating Location-aware Advertisements in Anycast IP Networks”. In: *Proceedings of the 2024 Applied Networking Research Workshop*, pp. 15–22

Contents

1	Introduction	1
1.1	The Nuts and Bolts of Interdomain Routing	1
1.2	The Significance of Modeling the Interdomain Routing System	3
1.2.1	Predicting and Managing Routing Behavior	3
1.2.2	Performance Optimization	5
1.2.3	Economic Strategy Evaluation	5
1.2.4	Enhancing Security	6
1.2.5	Facilitating Research and Innovation	8
1.3	Challenges	8
1.3.1	Topology Incompleteness	9
1.3.2	Complex AS relationships	9
1.3.3	Simplistic Modeling Abstraction	10
1.3.3.1	Levels of Topology Resolution	10
1.3.3.2	Limitations of Simplistic Abstractions	11
1.4	Contributions	13
1.4.1	Exploring the Confounding Factors of BGP Routing Models	13
1.4.2	Rethinking the Modeling of Interdomain Routing Policies	14
1.4.3	Introducing Geolocation in BGP Simulations	15
1.4.4	Open Source Code and Datasets for Reproducibility	16
1.5	Dissertation Outline	17

2	Background and Related Work	18
2.1	The Internet Routing Policies	19
2.1.1	AS Business Relationships	19
2.1.2	Import Routing Policies	20
2.1.3	Export Routing Policies	21
2.2	Related Work	22
2.3	The Evolution of the Internet Structure	24
2.4	The Incompleteness of the AS Graph	26
2.5	Conclusion	28
3	Confounding Factors	29
3.1	Inference Overview	30
3.1.1	Best-Path Inference Process	30
3.1.2	AS-Relationships Datasets	31
3.1.3	Ground-truth Paths	31
3.1.4	Performance Metrics	32
3.2	Analysis	33
3.2.1	The Path Diversity Problem	34
3.2.1.1	Vanilla Model Accuracy	35
3.2.1.2	Accuracy Without Path Diversity	36
3.2.2	The First-hop Inference Problem	38
3.2.3	The Confounding Factors	40
3.3	Conclusion	42
4	Interdomain Routing Policies	44
4.1	Replication Overview	45
4.1.1	Take-aways from the Original Paper	45
4.1.2	Replication Strategy	46
4.2	Import Policies	49
4.2.1	Route Preference Among Provider, Customer and Peer Routes	49

4.2.2	Consistency of <i>locpref</i> with next-hop	52
4.2.3	Error Introduced by AS relationships	53
4.3	Export Policies	55
4.3.1	Export to Provider	56
4.3.1.1	Inference Algorithm	58
4.3.1.2	Prevalence of <i>SA Prefixes</i>	58
4.3.1.3	Verification of <i>SA Prefixes</i>	61
4.3.1.4	Persistence of SA prefixes	64
4.3.1.5	Causes of SA prefixes	67
4.3.2	Export to Peer	69
4.4	Conclusion	73
5	Location Specific Announcements	74
5.1	Methodology Overview	75
5.2	Selective Announcements in Anycast IP Networks	80
5.2.1	Selective Announcements per AS	80
5.2.2	Selective Announcements per Location	81
5.3	Regionality of Receivers of Selective Announcements	82
5.3.1	Definition of Regionality	82
5.3.2	Regionality Analysis	84
5.4	Conclusion	85
6	Conclusion	87
6.1	Summary of Research Work	88
6.2	Discussion	89
6.3	Future Directions	91
6.4	The Ph. in my Ph.D.	92
	References	94

List of Tables

2.1	Peer Link Statistics of the Internet, 1998-2023, as observed in the CAIDA AS Relationships Graph. The percentage of peer links increased significantly over the last 25 years, indicating the flattening of the AS-topology.	25
3.1	Key-factors that affect AS-path inference.	42
4.1	Characteristics of the ASes used in the import/export policies inference.	48
4.2	List of ASes for which we extracted locpref values along with the percentage of routes that conform to the Gao-Rexford (GR) local preference model.	51
4.3	Typical locpref assignments for 32 ASes which are selected from IRR.	54
4.4	Validation results of AS relationships based on BGP Communities. .	55
4.5	% of SA prefixes and SA origins observed by 21 ASes.	59
4.6	Percentage of SA prefixes verified per AS.	62
4.7	% (#) of multi-homed and single-homed SA origins for AS3257, AS3292, AS3549, AS5511 and AS7018	67
4.8	Causes of SA prefixes	69
4.9	% of peer SA prefixes and % peer SA origins.	71
5.1	Top Anycast Networks based on RUM Uptime.	76

List of Figures

1.1	The Internet is a Network of Networks	2
2.1	Customer Cone (CC) sizes in 2003 (left), and 2023 (right) exhibit similar power-law distributions.	26
3.1	CDF of # distinct AS-paths per AS-pair.	34
3.2	Average inference accuracy	35
3.3	Average accuracy without path diversity.	36
3.4	Frequency per path length.	37
3.5	Average accuracy when first-hop known.	38
3.6	[CAIDA]Exact path match per path length.	39
4.1	Overview of the data used for the import and export routing policies analysis.	47
4.2	The export policies of AS13335, can be observed by its indirect provider AS3257. AS13335 announces prefix p to provider AS3549, but not to AS9498.	56
4.3	CDF of all SA origins in 2023.	60
4.4	Daily and monthly persistence of SA prefixes for AS7018.	63
4.5	Daily and monthly uptimes of SA prefixes for AS7018.	63
4.6	SA prefix ratio over the last 20 years for AS3257, AS3292, AS3549, AS5511 and AS7018.	66

4.7	CDF of union of SA origins over the last 20 years for AS3257, AS3292, AS3549, AS5511 and AS7018.	66
4.8	Causes of a SA prefix	68
4.9	Progressive enumeration of unique SA prefixes by incrementally adding vantage points. The rate of increase does appear to plateau, indicating that our results are a lower-bound estimation of the SA prefix ratio, due to the incompleteness problem of the AS graph.	72
5.1	Tools used and overview of the methodology.	75
5.2	Selective announced prefixes per anycast AS.	78
5.3	Selective Announcement Types.	80
5.4	Regionality levels of the direct neighbors of the top Anycast ASes. In specific regions, big CDNs (e.g., Google, Cloudflare, G-Core, Amazon) rely on regional neighbors to carry their traffic to/from the rest of the Internet. This could be due to regulatory considerations, missing backbone or strategic business interests.	83
5.5	The Anycast AS connects to different ASes in different PoPs. Large PoPs (e.g., PoP1) tend to be connected to both regional and global ASes, while regional PoPs (e.g., PoP2) are typically connected only to regional ASes.	84

Chapter 1

Introduction

1.1 The Nuts and Bolts of Interdomain Routing

At the time of writing this dissertation, advances in five areas: automation, AI, multi-cloud networking, wireless and network security, promise to power the biggest wave of network transformation seen in decades (Cisco, 2021; Cisco, 2024) and the Internet is the common denominator for all these applications to operate. The Internet is a decentralized collection of interconnected component networks/nodes (Autonomous Systems, ASes) (see Fig. 1.1). These networks are composed of end hosts (who are identified by IP addresses) and active forwarding elements (routers) whose role is to direct IP packets as they pass through the network. Routing protocols are used to perform this information propagation. The Internet's routing system is divided into a two-level hierarchy (intra-domain and inter-domain). The single inter-domain routing protocol, i.e., the Border Gateway Protocol or BGP (RFC 4271 Rekhter, T. Li, and Hares, 2006), has provided routing services for the Internet's disparate component networks since the late 1980s. Given the central role of routing in the operation of the Internet, BGP is one of the critical protocols that provide essential coherence to the Internet.

Inter-domain routing does not follow the shortest path principle, but it is based on the economic, performance or security needs of the ASes that constitute it. ASes

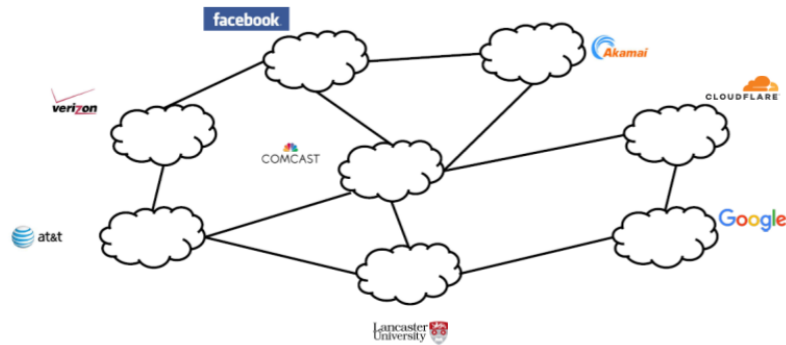


Figure 1.1: The Internet is a Network of Networks

independently define their routing policies (Gao and Rexford, 2001; Huston, 1999) in order to select routes to a certain destination when multiple routes are available (import policies), and to decide to which neighbors to propagate the routes they know (export policies). For instance, the objective of a transit provider may be to maximize its profit, and it may approach this goal through competitive pricing and selective peering. The objective of a content provider, on the other hand, may be to have highly reliable Internet access and minimal transit expenses, and it may pursue these goals through aggressive multihoming and an open peering policy (Dhamdhere and Dovrolis, 2008).

ASes are often unwilling to share proprietary business data such as: the internal network's topology, the list of customers that are buying transit on their networks or their traffic volumes. Routing policies are often protected by non-disclosure agreements, and kept secret as well. However, this opacity of routing policies makes it hard to understand, debug and predict routing decisions. Often to resolve disruptions that occur outside the periphery of an AS requires offline communication among operators, or trial-and-error experimentation. Similarly, predicting the outcomes of topological or policy changes requires to observe the impact of these changes in practice, and calibrate them based on the observed paths. Such practice incurs the risk of outages and network errors. Therefore, the ability to accurately infer the Internet routing policies could significantly improve network operations.

1.2 The Significance of Modeling the Interdomain Routing System

The Interdomain Routing System has been extensively studied over the last two decades, in order to understand, evolve and capture the behavior of the ASes that constitute it but also develop accurate policy models and path prediction capabilities (Dhamdhere and Dovrolis, 2008; Dimitropoulos et al., 2005; Ballani and Francis, 2005; R. V. Oliveira et al., 2008; Mühlbauer, Uhlig, et al., 2007; Gill, Schapira, and Goldberg, 2013; Huston, 1999; Luckie et al., 2013; Mao, Rexford, et al., 2003; Labovitz et al., 2010; Gao, 2001; Gao and Rexford, 2001; Quoitin and Uhlig, 2005; Mühlbauer, Feldmann, et al., 2006; Gill, Schapira, and Goldberg, 2012; Cunha et al., 2016; Tian et al., 2019; Wu et al., 2022; Singh et al., 2021; Sermpezis and Kotronis, 2019; Kastanakis, Giotsas, and Suri, 2022; Kastanakis, Giotsas, Livadariu, et al., 2023c; Kastanakis, Giotsas, Livadariu, et al., 2024).

The necessity to model the interdomain routing system arises from the complex and dynamic nature of the Internet’s architecture. Interdomain routing, involves a multitude of ASes, each of which implements its own routing policies based on economic, performance, and security considerations. These policies and their interactions create a highly intricate and often opaque routing environment, necessitating detailed and accurate modeling to ensure the Internet’s efficient and stable operation (Gao and Rexford, 2001).

1.2.1 Predicting and Managing Routing Behavior

Understanding interdomain routing is crucial for predicting and managing the dynamic nature of global Internet traffic. Changes in one AS can trigger cascading effects throughout the network, leading to shifts in routing paths that might impact distant ASes in unforeseen ways. This interconnectedness means that even minor adjustments in one AS’s policy can cause widespread consequences, potentially leading to performance issues, network congestion, or connectivity problems affecting

many users. When ASes alter their routing policies to enhance traffic management, bolster security, or respond to network congestion, these changes can ripple through the network, occasionally destabilizing large portions of the Internet (Butler et al., 2009; Apostolaki, Zohar, and Vanbever, 2017; Sermpezis, Kotronis, Gigis, et al., 2018; Sermpezis, Kotronis, Arakadakis, et al., 2021).

To address this complexity, accurate interdomain routing models have become indispensable for network researchers and operators. These models allow for the simulation of various policy adjustments, providing a means to analyze potential impacts before they occur. By using these simulations, researchers can predict how different policy scenarios might affect both individual ASes and the global traffic flow, helping to identify possible disruptions or inefficiencies (Gao, 2001; Gao and Rexford, 2001; Gill, Schapira, and Goldberg, 2012; Mao, Rexford, et al., 2003; Mao, L. Qiu, et al., 2005; Dimitropoulos et al., 2005; Cunha et al., 2016).

For network operators, modeling is crucial in anticipating and addressing issues such as traffic bottlenecks, suboptimal routing, or security vulnerabilities that might arise from policy changes. Through detailed modeling, it's possible to test and refine routing strategies, aiming to prevent problems before they impact real-world network performance (Sermpezis, Kotronis, Gigis, et al., 2018; Sermpezis and Kotronis, 2019; Mazloun et al., 2014; Cardona, Francois, and Lucente, 2016; Orsini et al., 2016; Moura et al., 2016; Donnet and Bonaventure, 2008; Z. Li et al., 2018). Given the increasing complexity of the Internet and the demands of modern applications, such proactive approaches are vital.

These models not only shed light on the immediate effects of policy changes but also provide a deeper understanding of the Internet's interconnected nature. As the global Internet evolves and faces new challenges—such as higher traffic volumes, emerging security threats, or shifting economic and political dynamics—modeling becomes a key tool for maintaining network stability. By simulating various scenarios, operators can fine-tune their routing policies to avoid disruptive changes that could lead to widespread network issues (Gill, Schapira, and Goldberg, 2013).

1.2.2 Performance Optimization

The Internet’s performance, from latency and bandwidth utilization to overall user experience, is deeply influenced by routing decisions. Detailed routing models allow operators to understand path characteristics and routing behaviors, enabling them to identify and implement the most efficient routes for data transmission (Ahmed et al., 2017). These models can provide multiple benefits in the overall network performance, e.g., they can help minimize latency, optimize bandwidth utilization, and improve the overall quality of service (QoS) and quality of experience (QoE) (Kastanakis, Sermpezis, Kotronis, Menasché, et al., 2020; Kastanakis, Sermpezis, Kotronis, and Dimitropoulos, 2018; Sermpezis, Kastanakis, et al., 2019). For content providers and ISPs, such performance optimization is essential to stay competitive and meet user expectations for fast and reliable Internet services.

1.2.3 Economic Strategy Evaluation

Economic considerations are undeniably central to how routing decisions are made by different organizations (Norton, 2001; Dhamdhere and Dovrolis, 2008; Bailey, 1997; Gao, 2001). ASes are typically concerned with two primary objectives: a) increasing their revenues or b) reducing their operational costs. This financial focus influences every aspect of how they interact with other networks, and interdomain routing decisions often hinge on these concerns. For example, an AS might have multiple choices when deciding how to route its traffic. It could either rely on a larger transit provider or establish direct peering relationships with other networks (Kastanakis, Giotsas, Livadariu, et al., 2023c). Peering agreements allow networks to exchange traffic directly, often leading to lower costs compared to purchasing transit services. But the decision isn’t always that straightforward. There are multiple variables to consider, such as the cost of setting up and maintaining those connections, potential performance benefits, and how much traffic will flow through these routes (Ahmed et al., 2017).

By employing interdomain routing models, ASes can assess the economic impact

of various routing strategies more systematically. These models provide a way to measure and predict how different routing policies will affect costs and profits. For instance, a large transit provider might want to evaluate whether establishing a peering relationship with another AS will save money in the long term. On the other hand, smaller content providers or enterprises might be looking for ways to reduce their transit fees and optimize their network performance, possibly through strategic multihoming—a process where a network connects to more than one upstream provider to improve redundancy and performance. A real-world example of this is how content delivery networks (CDNs) often seek out multiple peering agreements to cut down on transit costs while ensuring faster delivery times to end-users. Companies like Netflix and Google, which rely heavily on delivering vast amounts of data across the globe, benefit from strategic peering as it allows them to minimize the financial burden of purchasing transit services from third-party providers.

The key advantage of using routing models is that they can ground these decisions in data, rather than intuition or guesswork. These models take into account the complexities of the Internet’s topology, as well as the economic relationships between different ASes. By doing so, they empower network administrators to make more informed decisions that align with their financial objectives.

1.2.4 Enhancing Security

The security of the Internet’s interdomain routing infrastructure has never been more crucial. BGP is a protocol that was not designed with security in mind, leaving it vulnerable to several serious threats, including route hijacking and route leaks. These vulnerabilities allow malicious actors, or even well-intentioned operators who make mistakes, to disrupt Internet traffic, hijack data, or mislead entire networks about where information should be routed (Butler et al., 2009; Huston, Rossi, and Armitage, 2010; Murphy, 2006).

Take the example of route hijacking, where a rogue AS advertises IP prefixes

it doesn't own, tricking other networks into routing traffic through its systems. Similarly, route leaks happen when an AS accidentally or deliberately passes on routes that it shouldn't, potentially exposing traffic to third parties or causing inefficient routing. Both incidents highlight how fragile the Internet's routing infrastructure can be, and how crucial it is to protect it (Lychev, Goldberg, and Schapira, 2013; Goldberg, 2014; Mahajan, Wetherall, and Anderson, 2002; Zhang et al., 2007).

Modeling and simulating BGP behavior has become an essential tool for identifying these weaknesses before they turn into actual problems. By analyzing how BGP announcements propagate across networks and detecting unusual patterns, researchers and network operators can pinpoint where things might go wrong (Sermpezis, Kotronis, Gigis, et al., 2018). This kind of proactive detection is essential. The sooner an anomaly is spotted, the faster it can be addressed, preventing potential outages or security breaches. In essence, these models help anticipate where vulnerabilities lie, enabling more effective mitigation strategies.

There's been a lot of effort over the years to improve BGP's security. One of the most promising advancements is Resource Public Key Infrastructure (RPKI), which verifies the authenticity of IP prefixes (Bush and Austein, 2013). This ensures that only legitimate ASes can announce specific routes, making it much harder for attackers to falsely claim they own parts of the Internet. But despite RPKI's potential, its adoption has been slow. Many networks are hesitant to implement it due to technical complexities or fear of introducing new points of failure. As a result, much of the Internet still operates without these protections, leaving it open to the very risks BGP was meant to solve.

What's needed is a broader approach to securing BGP: not just fixing the problems we see today, but anticipating the threats of tomorrow. This includes encouraging wider adoption of existing solutions like RPKI, but also investing in research to develop new models and tools that can easily integrate into the existing infrastructure.

1.2.5 Facilitating Research and Innovation

For researchers in the field of Internet measurement, interdomain routing models are essential for advancing our understanding and developing new solutions (Lepinski and Sriram, 2017). These models enable the exploration of the Internet’s behavior under various conditions, the testing of novel routing algorithms, and the innovation of approaches to existing challenges (Bush and Austein, 2013). This research is vital for the continuous evolution and improvement of the Internet, driving advancements that enhance efficiency, security, and functionality (Kastanakis, Giotsas, and Suri, 2022; Kastanakis, Giotsas, Livadariu, et al., 2024).

1.3 Challenges

Since its beginning, the Internet topology has experienced fundamental changes in interconnection practices, such as the flattening of the Internet hierarchy (Gill, Arlitt, et al., 2008; Böttger et al., 2018) and the dominance of Content-Distribution Networks (see Section 2.3). This evolution has been noticed for years. Not only stub networks, but also ISPs are using an open peering strategy to peer with more networks (Lodhi, Dhamdhare, and Dovrolis, 2014; Marcos et al., 2018) and there has been noticed a significant performance difference between peering and transit interconnections (Ahmed et al., 2017), which provides one reason of the evolution.

There have been numerous efforts, over the last two decades, to understand the interdomain routing system and develop accurate policy models and path prediction capabilities, nonetheless, the impact of the topological changes on inter-domain policies is still not clear and the state-of-the-art is still unable to accurately infer AS paths due to several open issues (Anwar et al., 2015; Mühlbauer, Feldmann, et al., 2006).

1.3.1 Topology Incompleteness

The topology incompleteness problem remains a fundamental obstacle in interdomain routing modeling. The Internet’s AS-level topology is vast and constantly evolving, leading to incomplete and sometimes inaccurate representations.

Most modeling studies rely on BGP data, Looking Glass servers data and AS relationships. The most significant limitation of the above data collection projects is the large number of missing links, which are divided into two types: *hidden* and *invisible* (R. Oliveira et al., 2009). Hidden links are usually backup upstream links (i.e., customer-to-provider links) that can be observed when the preferred path to a prefix changes, and, invisible links are typically peering links (i.e., peer-to-peer links) which are inherently unobserved due to the limited number of vantage points across the AS graph. Invisible links constitute the majority of missing links and can be located in the periphery of the AS graph (Roughan, Tuke, and Maennel, 2008; R. Oliveira et al., 2009; Khan et al., 2013).

The root cause of this problem is that giant CDNs who originate a large portion of today’s Internet traffic, often operate under a shroud of secrecy regarding their infrastructure details and peering arrangements with other ASes. This, coupled with the complexity introduced by Internet Exchange Points (IXPs), makes it challenging for external observers to map CDNs and their interconnections accurately. To this day, the Internet Measurement community does not have an adequate solution for this issue, hence, our study suffers from the same limitation.

1.3.2 Complex AS relationships

The complexity of AS relationships further complicates interdomain routing modeling. ASes engage in a variety of peering and transit arrangements, driven by economic, performance, and strategic considerations. The traditional categorization of these relationships into simple types like customer-provider and peer-to-peer is insufficient to capture their true complexity:

Hybrid Relationships Many ASes enter into hybrid relationships that exhibit characteristics of both transit and peering arrangements. These relationships may involve specific terms that dictate traffic exchange conditions, pricing, and performance guarantees. The fluid nature of these hybrid arrangements makes it challenging to model them accurately using simplistic categorizations.

Dynamic Policies ASes frequently adjust their routing policies in response to changes in network conditions, business strategies, and external factors such as regulatory changes. These dynamic policies can include changes in route preferences, the introduction of new peering agreements, or adjustments to traffic engineering practices. Capturing the impact of these dynamic policies requires models that can adapt to and accurately reflect ongoing changes.

Multilateral Peering Increasingly, ASes are participating in multilateral peering agreements through IXPs. These agreements allow multiple ASes to exchange traffic more efficiently but add another layer of complexity to the modeling process. IXPs facilitate the rapid establishment and adjustment of peering relationships, which can significantly alter the flow of traffic and the structure of AS paths.

1.3.3 Simplistic Modeling Abstraction

The Internet's layered architecture can be described at various levels of abstraction, each offering a different resolution of the network topology. Understanding these levels is crucial for recognizing the limitations of simplistic models and their impact on the accuracy of routing predictions (Dhamdhere and Dovrolis, 2008; Roughan, Willinger, et al., 2011).

1.3.3.1 Levels of Topology Resolution

The Internet's topology can be conceptualized at multiple levels of resolution, each providing a different perspective on network connectivity:

IP-Level Topology This is the most granular level, focusing on the connectivity

between individual router interfaces. At this level, each IP interface represents a node, and physical connections between interfaces are represented as links. This detailed view captures the precise connectivity between network devices but is often too detailed for large-scale routing analyses.

Router-Level Topology By aggregating interfaces belonging to the same router, the router-level topology simplifies the IP-level view. Routers are grouped into nodes, and connections between these nodes represent links between routers. This abstraction reduces complexity but still retains significant detail about router interconnections.

Point-of-Presence (PoP)-Level Topology This level groups routers into Points-of-Presence (PoPs), which are physical locations where ISPs house their network devices. PoPs serve as access points to the network and aggregate multiple routers. The PoP-level topology offers a balance between detail and manageability, providing a clearer view of network access points and their connectivity.

AS-Level Topology At this highest level of abstraction, nodes represent Autonomous Systems (ASes), and links denote the contractual relationships between them. Each link in the AS graph is a logical construct that aggregates multiple router-level or PoP-level links between ASes. This abstraction simplifies the network into a graph of ASes and their interconnections, but it often overlooks critical details about internal AS structures and precise connectivity.

1.3.3.2 Limitations of Simplistic Abstractions

Simplistic models that rely solely on the AS-level topology, while computationally manageable, present several limitations:

Loss of Internal Structure The AS-level abstraction does not capture the internal structure of ASes or the detailed connectivity within and between ASes (Mühlbauer, Feldmann, et al., 2006; Mühlbauer, Uhlig, et al., 2007). Consequently, models based on AS-level topologies often fail to accurately predict routing paths or address

engineering problems.

Inaccurate Path Predictions Due to the oversimplification of AS relationships and internal routing dynamics, models that use the AS-level abstraction can produce poor path prediction results (Anwar et al., 2015; Kastanakis, Giotsas, and Suri, 2022). For example, routing decisions influenced by specific policies or traffic engineering strategies are not reflected in the AS graph, leading to inaccuracies in predicting how data will traverse the network.

Policy Variability Despite the definition that suggests a single routing policy per AS for inter-AS interactions, real-world ASes often apply different routing policies at various interconnection points (Mühlbauer, Feldmann, et al., 2006). They may announce different prefixes or implement distinct import policies depending on the peering or transit agreements in place. Simplistic models that do not account for this variability cannot accurately represent the diversity of routing behaviors.

To address these limitations, it is beneficial to explore alternative modeling approaches that offer a more nuanced representation of the network (Roughan, Willinger, et al., 2011):

Multi-Graph Models Incorporating multiple types of relationships and policies into the AS graph, such as through multi-graph models, can better capture the diversity of routing policies between ASes. These models allow for a more detailed representation of the complex dynamics governing interdomain routing.

PoP-Level Topology The PoP-level topology represents an intermediate level of granularity that balances detail and manageability. By capturing the aggregation of routers into PoPs, this approach provides a more realistic view of network access points and inter-AS connectivity. However, accurately determining the geographic footprint of ASes and mapping PoP-level relationships can be challenging and typically requires sophisticated techniques like latency measurements and traceroute analysis (Roughan, Willinger, et al., 2011).

1.4 Contributions

This dissertation investigates the confounding factors that affect the accuracy of AS-path inference in interdomain routing and explores methodologies to improve the fidelity of BGP models. The overarching hypothesis is that modern AS-path inference accuracy can be significantly improved by addressing topology incompleteness, incorporating geolocation-aware routing paradigms, and revising the modeling of interdomain routing policies to reflect evolving AS behaviors.

What are the key confounding factors that limit the accuracy of AS-path inference, and how can their impact be quantified? How have interdomain routing policies evolved over the past two decades, and what implications do these changes have for AS-path prediction? Can geolocation-aware models provide more accurate and realistic predictions of AS-paths, especially for applications like anycast routing? By addressing these questions, this dissertation aims to contribute novel insights into the dynamics of BGP modeling and propose practical methodologies to enhance the accuracy of AS-path predictions.

1.4.1 Exploring the Confounding Factors of BGP Routing Models

In this study, we focus on the confounding factors that affect the accuracy of BGP (Border Gateway Protocol) routing models. Utilizing empirical data and sophisticated simulation techniques, we have identified several key factors that significantly influence the precision of AS-path inferences, e.g., the complex routing policies of certain ASes and the geolocation-agnostic BGP best path selection process. By accurately determining the first-hop, we were able to boost the exact-path score from 33.6% to 84.1%, and further incorporating geolocation data improved this accuracy to 94.6%.

Extensive Monte Carlo simulations employing the Gao-Rexford model allowed us to examine the impact of various factors, including missing AS relationships, valley-

free violations, local preference discrepancies, and shortest path violations. Our findings underscore the critical role of geolocation in routing decisions, challenging the traditional location-agnostic models.

This work lays the groundwork for future studies aimed at enhancing the robustness and reliability of inter-domain routing models, thereby contributing to a deeper understanding of Internet routing dynamics.

1.4.2 Rethinking the Modeling of Interdomain Routing Policies

Additionally, we present a fresh perspective on interdomain routing policies, driven by the findings of a replication study of the seminal work by Wang and Gao (Wang and Gao, 2003). This research explores how interdomain routing policies have evolved over the past two decades, particularly focusing on the phenomenon of selectively announced prefixes and the variability in local preference assignments. These two metrics are considered since selective announcements directly impact route visibility, while variability in local preference values reflects differences in routing policies, both of which are critical for understanding AS-level behavior and the evolution of interdomain routing.

The replication study revealed that selective announcements—where Autonomous Systems (ASes) announce their prefixes selectively to manage incoming traffic—remain a significant practice in the Internet ecosystem. Since the original study, the usage of selective announcements has increased by up to 30%, though with considerable variability among different ASes. This persistence and growth highlight the complexity and dynamism of routing behaviors that are often underrepresented in traditional routing models.

Moreover, the study uncovered that the assignment of Local Preference (locpref) values now exhibits greater variability than two decades ago. Locpref, a key BGP attribute, determines route selection when multiple paths are available. This variability suggests that ASes are adopting more nuanced and dynamic routing

strategies, influenced by the economic, performance, and security objectives of individual networks. These findings underscore the necessity for routing models to incorporate more flexible and adaptive policies that can better reflect the real-world practices of ASes.

The longitudinal aspect of the study showed a median increase of more than 20% in selective announcements after 2007, albeit with significant yearly fluctuations. This indicates that while the trend towards selective routing is clear, the exact patterns are influenced by various transient factors, possibly including changes in interconnection agreements, traffic management strategies, and the emergence of new Internet applications.

Importantly, this research emphasizes the need for high-periodicity inference of BGP policies to account for the dynamic nature of AS connectivity and policies. Traditional static models fall short in capturing these rapid changes, leading to inaccuracies in path predictions and routing decisions.

By integrating these insights, our revised modeling approach aims to offer a more accurate representation of interdomain routing. This involves not only considering the static attributes of AS relationships but also the dynamic and often transient nature of routing policies. This enhanced understanding is crucial for improving the reliability and efficiency of Internet routing, ultimately leading to more robust and adaptable network architectures.

1.4.3 Introducing Geolocation in BGP Simulations

Furthermore, we introduce geolocation as a crucial factor in BGP simulations, aiming to bridge the gap between theoretical models and real-world routing dynamics. Traditional BGP models often overlook the geographical aspect of AS-path selection, leading to inaccuracies in path predictions and an incomplete understanding of routing behaviors. By integrating geolocation data into our simulations, we enhance the accuracy and fidelity of these models.

We developed a methodology to geolocate the receivers of selective Anycast

prefix announcements, using tools like Maxmind and PeeringDB to assign country-level characteristics to ASes. This geolocation information allows us to capture regional trends and disparities in Anycast deployment, offering a more granular view of routing policies. Our results demonstrate that geolocation-aware models significantly outperform traditional models, achieving higher accuracy in AS-path predictions.

While BGP itself intentionally abstracts away geographical information to maintain simplicity and scalability, incorporating geolocation data in simulations is appropriate for understanding real-world behaviors. Many routing decisions, such as traffic engineering, anycast deployment, and latency optimization, inherently depend on physical locations. By including geolocation data, we can model the practical considerations that influence routing policies, bridging the gap between BGP’s abstract design and its application in geographically distributed networks.

This research underscores the importance of considering geographical factors in BGP routing simulations. By doing so, we can better understand the influence of physical location on routing decisions and the overall performance of the Internet. Our contributions highlight the potential of geolocation data to enhance the realism and predictive power of BGP models, providing a valuable tool for researchers and network operators to optimize their routing strategies.

1.4.4 Open Source Code and Datasets for Reproducibility

Ensuring the reproducibility of scientific research is a cornerstone of academic integrity and progress. In our research work, we focus on providing open-source code and datasets to facilitate the replication and verification of our results by the broader research community. Our work on BGP routing models and interdomain routing policies is accompanied by a comprehensive repository of code and data, accessible to all interested researchers (Kastanakis, Giotsas, Livadariu, et al., 2023a; Kastanakis, Giotsas, Livadariu, et al., 2023b).

1.5 Dissertation Outline

The rest of the thesis is organized as follows.

Chapter 2: Background and Related Work This chapter lays the groundwork by discussing the essential concepts of BGP and interdomain routing. We review the historical development of BGP path prediction, the role of geolocation in BGP simulations, and contemporary approaches to modeling interdomain routing policies.

Chapter 3: Exploring the Confounding Factors of BGP Routing Models In this chapter, we delve into the confounding factors affecting the accuracy of BGP routing models. Building on our findings from extensive measurements and simulations, we identify key sources of inference errors, such as the first-hop prediction problem and the impact of the geolocation-agnostic BGP best path selection process.

Chapter 4: Longitudinal Study on Interdomain Routing Policies This chapter re-examines the modeling of interdomain routing policies, informed by insights from our replication study of selective announcements and local preference variability. We uncover the widespread use of selective route announcements and the dynamic nature of AS routing strategies. Our findings suggest the need for adaptive policy models that better reflect real-world interdomain routing practices.

Chapter 5: The Role of Geography in Anycast Routing Decisions In this chapter, we demonstrate that selectively advertising anycast prefixes is largely influenced by geography and that anycast ASes follow different routing patterns per geographical regions. These findings suggest the need for routing models that account for the geographical characteristics between and across ASes.

Chapter 6: Conclusions and Future Work In the final chapter, we summarize our findings and propose future research directions. This chapter outlines our vision for advancing the understanding and accuracy of Internet routing modeling.

Chapter 2

Background and Related Work

Chapter 2 builds upon the foundational concepts introduced in Chapter 1 by expanding on the underlying mechanics of the Internet’s routing system. This chapter delves into the essential routing policies that guide the behavior of ASes, particularly focusing on the complex business relationships—customer-provider, peer-to-peer, and sibling-sibling—that influence traffic flow across the global Internet. We will examine how these policies impact both import and export routing decisions, shaping the paths that data packets take as they traverse the Internet. Furthermore, we will review key research studies that have contributed to our understanding of interdomain routing models, shedding light on their strengths and limitations.

By exploring the evolution of the Internet’s structure, especially in terms of the increasing prevalence of Content Delivery Networks (CDNs) and Internet Exchange Points (IXPs), Chapter 2 sets the stage for understanding the challenges posed by the incompleteness of the AS graph. This incomplete topology, coupled with the secretive nature of AS relationships, presents significant hurdles for accurate path prediction. Through this comprehensive review of the literature, we will gain a clearer understanding of the foundational concepts necessary to address the challenges presented in later chapters.

2.1 The Internet Routing Policies

In this section, we present an overview of the AS business relationships and then describe the Internet routing policies.

2.1.1 AS Business Relationships

BGP is a policy-based protocol, therefore, each AS uses the routing policies that best fit its economic, performance, security or traffic engineering goals and there is no need for global coordination among ASes for the Internet to operate (Gao and Rexford, 2001; Huston, 1999).

AS interconnection relies on business agreements that determine financial and technical aspects of their interconnection and traffic exchange. While such business agreements can be arbitrary, they can coarsely be categorized in three types of business relationships (Gao, 2001): (1) In a *customer-to-provider* (c2p) relationship, a customer AS pays a better-connected provider AS to transit its traffic to the rest of the Internet. (2) In a *peer-to-peer* (p2p) relationship, two ASes agree free bilateral traffic exchange between their networks and the networks of their customers. (3) A *sibling-to-sibling* (s2s) relationship expresses the connection between two ASes under the same administrative entity, typically as a result of mergers and acquisitions. Siblings usually do not impose routing restrictions on each other. An AS that has only a single transit provider is called *single-homed*. Often ASes prefer to have multiple providers (*multi-homed*) for resilience and traffic engineering purposes. A few ASes that can access all the rest of the ASes only through customer or peering links do not require transit providers. Those ASes are called *transit-free* and together they form a fully-connected mesh of ASes called the *Tier-1* clique. The business relationships among ASes (AS relationships) may be protected by non-disclosure agreements, so they are often kept secret.

AS relationships impact both how an AS advertises its routes to its neighbors (export BGP policy), and how it selects which route to use when it has multiple

routes available for the same IP destination (import BGP policy). Researchers and engineers have developed algorithms to infer routing policies in the form of AS relationships to study the Internet routing system, with many of those algorithms claiming an accuracy of over 98% (Luckie et al., 2013; Y. Jin et al., 2019; Z. Jin et al., 2020).

2.1.2 Import Routing Policies

A BGP router may receive multiple routes for the same destination IP prefix from different AS neighbors. The router uses the BGP selection process to determine the single most preferable route. This selection process is comprised of the following steps (Cisco, 2023; Juniper, 2023). The process goes to the next step only if the previous step does not result in a single best path.

1. Routes with the highest local preference (*locpref*) value. *locpref* is a non-transitive numerical BGP attribute that denotes the preference of a certain route. Higher *locpref* values imply higher preference for a given AS path.
2. Routes with the shortest AS Path length.
3. Routes with the lowest origin type. Paths that are locally originated (IGP) are preferred over externally originated paths (EGP).
4. Routes with the lowest Multi-Exit Discriminator (MED) value.
5. Routes learned from eBGP over those from iBGP.
6. Routes with the lowest IGP cost to the border router.
7. Oldest routes.
8. Routes with the smallest router ID.

This complexity in the BGP decision process, makes it also challenging for researchers to model it and for operators to predict the impact of their policies.

Some operators switch off some of these steps due to complexity (Gill, Schapira, and Goldberg, 2013).

As shown in the above steps, *locpref* is the highest-priority metric in deciding which route to use. While *locpref* values are arbitrary, ASes generally assign the highest *locpref* values to routes learned from customer ASes, since customer traffic generates revenue, and the lowest *locpref* values to routes learned from provider ASes since provider traffic incurs a cost. Gao and Rexford modeled inter-domain routing to find that such ordering of *locpref* values is necessary in order to ensure convergence in the global routing system (Gao and Rexford, 2001). For this reason this *locpref* allocation pattern is also referred to as the Gao-Rexford model.

2.1.3 Export Routing Policies

Once a router selects the best route towards a destination prefix, it can propagate the best route to its neighboring ASes. The configuration of export policies is similar to those of import policies and can be based on prefix or next-hop.

BGP routes are usually exported following the so-called *valley-free rule*, i.e., a customer route can be exported to any neighbour AS, but a route learned from a peer or a provider can only be exported to customers. Hence, an AS path is valley-free if it follows one of the following patterns (Giotsas, Zhou, et al., 2013): (1) $n \times c2p + m \times p2c$; or (2) $n \times c2p + p2p + m \times p2c$; where n and $m \geq 0$. The sibling links can be inserted freely without changing the valley-free property of a path. The valley-free rule aims to prevent an AS from providing free transit either to their providers or peers, since that would result in consuming resources and paying for traffic exchange that does not pertain to its network.

In addition to the valley-free rule, an AS may select to further restrict the propagation of certain routes for traffic engineering purposes. By selectively advertising routes to different neighbors and AS may be able to control the links which will carry traffic for a specific route. Intuitively, single-homed customers cannot selectively advertise routes to their single transit provider, otherwise not

all of their routes will be globally reachable. However, upstream ASes may choose to selectively advertise routes originated by a single-homed AS. Instead, it is more likely for multi-homed ASes to restrict the propagation of specific routes to their providers. Note that an AS can also selectively advertise routes among its peers.

2.2 Related Work

The pioneering and most classic approach was proposed by Lixin Gao and Jennifer Rexford in 2001 (Gao, 2001; Gao and Rexford, 2001), per which, the routing policies of an AS need to follow the Gao-Rexford model (Gao and Rexford, 2001) in order to be *safe* to converge to a stable state under any link or node failure. The Gao-Rexford guidelines are as follows:

- ASes prefer customer routes over peer and provider routes (import policies).
- An AS can export its customer routes to all neighbors, but, its peer and provider routes only to its customer neighbors (export policies).

Nonetheless, network operators can arbitrarily configure their policies, without any coordination with their neighbors, therefore, a number of ASes might not follow the Gao-Rexford model. This has been indeed observed both by Internet measurement studies (Anwar et al., 2015; Mazloun et al., 2014; Giotsas and Zhou, 2012) and reported by network operators in a 2013 survey (Gill, Schapira, and Goldberg, 2013).

The deviations from the Gao-Rexford model can be likely explained by the evolving economic incentives in a changing IP transit and peering market. During the past two decades, the Internet peering strategies have evolved to become more open, diverse and denser (Lodhi, Dhamdhare, and Dovrolis, 2014). As a result, ASes may prefer *peer* over *customer routes* for performance reasons (Ahmed et al., 2017).

In 2005, Quoitin et al. (Quoitin and Uhlig, 2005) described C-BGP, an open source tool especially designed to allow ISPs to experiment with a model of their

network. In 2006, Muhlbauer et al. (Mühlbauer, Feldmann, et al., 2006), focused on Quoitin’s work and formulated the path diversity in Internet routing. I.e., there may exist multiple paths between a source AS to a destination prefix and considering one router per AS is not sufficient to capture path diversity. The same year, Mao et al. (Mao, L. Qiu, et al., 2005), highlighted the importance of the ability to determine the first-hop on AS-level path inferences.

To derive a model that captures path diversity, in 2007, Muhlbauer et al. introduced a new abstraction: next-hop atoms (Mühlbauer, Uhlig, et al., 2007). Next-hop atoms capture the different sets of neighboring ASes an AS uses for its best routes. They showed that a large fraction of next-hop atoms correspond to per-neighbor path choices. A non-negligible fraction of path choices however do not correspond to simple per-neighbor preferences, but hot-potato routing and tie-breaking within the BGP decision process, which are very detailed aspects of Internet routing.

In 2012, Gill et al. (Gill, Schapira, and Goldberg, 2012) focused on improving the scalability and robustness of simulations for analyzing inter-domain routing techniques. The authors developed a novel routing tree algorithm that computes paths between all source-destination pairs in an AS graph. The algorithm exploits the fact that real-world AS graphs are typically very sparse.

Latest works have steered their approaches towards learning based methods. In 2016, Cunha et al. proposed Sibyl (Cunha et al., 2016). Sibyl attempts to predict the unseen paths and assign confidence to them by using supervised machine learning. Sibyl integrates together diverse sets of traceroute vantage points and uses historical measurements to predict which new ones are likely to match a query. Sibyl achieves a 76% accuracy of the queries that it could match if an oracle told it which measurements to issue, yet it is constrained by the slow and biased traceroute techniques.

In 2019, Tian et al. (Tian et al., 2019) were the first to introduce a data-driven learning-based approach to model route decisions of ISPs. However, they

exploited a limited feature-space which cannot capture enough information related to route decisions. In contrast, the latest path inference work (2022), RouteInfer (Wu et al., 2022), develops a learning-based approach to train a route decision model for predicting the route decisions of the ASes without policies, taking into consideration the node, link, and path features related to route decisions in practice. RouteInfer scores an average 81.64% accuracy.

PredictRoute (Singh et al., 2021), proposed in 2021, is a probabilistic system that predicts network paths between hosts on the Internet using historical knowledge of the data and control plane. PredictRoute discovers 4X more AS hops than other well-known strategies used in practice today. PredictRoute’s predictions differ from the measured path by at most 1 hop, 75% of the time.

Finally, given the multiple evidence that both the Gao-Rexford and the valley-free models do not always explain the actual routing policies with high fidelity, Shao and Gao (Shao and Gao, 2021) highlighted the need for developing new inter-domain routing models. Such models would allow: a) more flexible ranking among BGP routes when modeling the import policies, and b) multiple potential paths to be announced when modeling the export policies.

2.3 The Evolution of the Internet Structure

The Internet evolved from an academic research network to a global critical infrastructure that supports much of our social, economic and political activities. During this transition, the Internet topology has undergone multiple phase shifts. Over the past 25 years, the main change has been described as the “flattening” of the inter-domain AS hierarchy (Gill, Arlitt, et al., 2008). Table 2.1 illustrates the significant growth in the percentage of peer links over the last 25 years.

The Internet started as a research network in 1969 and evolved to a commercial network by 1995, with the rise of the World-Wide-Web. In the early 2000s, the conventional wisdom about the Internet ecosystem, was a multi-tiered hierarchy

Table 2.1: Peer Link Statistics of the Internet, 1998-2023, as observed in the CAIDA AS Relationships Graph. The percentage of peer links increased significantly over the last 25 years, indicating the flattening of the AS-topology.

Year	ASes	Links	Peer Links	% of Peer Links
1998	3549	6475	878	13%
2003	15164	35440	7084	19%
2008	28153	79590	25272	31%
2013	44064	143894	58366	40%
2018	60874	300634	178608	59%
2023	75160	494508	341363	69%

of Internet Service Providers (ISPs). A small clique of international ISPs (Tier-1) were connected with peering links to maintain global connectivity. Regional ISPs (Tier-2) were customers of the Tier-1 ASes and residential networks (Tier-3) were customers of the Tier-2 ASes. Stub networks were at the bottom of the hierarchy, as customers of Tier-3 ASes. The traffic was mostly carried through Tier-1 networks, which received revenue from Tier-2/Tier-3 networks.

Over the past decade, the Internet further evolved into a mesh interconnection network with a dense topology (see Table 2.1) due to the rise of Content Providers (CPs) and Content Delivery Networks (CDNs) (Gill, Arlitt, et al., 2008; Dhamdhare and Dovrolis, 2010; Anwar et al., 2015; Klöti et al., 2016; Labovitz et al., 2010). Big Internet players (Google, Facebook, Amazon) deployed their own private Wide Area Networks (WANs) close to the end users (i.e., in the periphery of the AS graph), to have more control over their end-to-end application performance (Wohlfart et al., 2018; Yap et al., 2017; Schlinker et al., 2017; Giotsas, Nomikos, et al., 2020; Gigis et al., 2021). In this flattening topology era, Internet Exchange Points (IXPs) emerged and played a key role in enabling large CDNs to bypass Tier-1 ISPs (Augustin, Krishnamurthy, and Willinger, 2009; Klöti et al., 2016; Kotronis et al.,




Figure 2.1: Customer Cone (CC) sizes in 2003 (left), and 2023 (right) exhibit similar power-law distributions.

2016). Currently, CDNs originate the largest part of the Internet traffic, and IXPs traffic volumes have become similar to those of Tier-1 ASes, hence, a valid question to ask is whether the Internet actually flattened or if the IXPs replaced Tier-1 ASes in the hierarchical model (Böttger et al., 2018).

One metric that reflects the position of an AS in the IP transit market is the customer cone (CC) size, which expresses the number of ASes that a provider AS can access through routes learned from its customers. We use the CAIDA CC dataset (CAIDA, 2023) to plot the CC distribution in 2003 and 2023 in Figure 2.1. Both distributions look similar, nonetheless, we observe that the maximum *CC* size has increased one order of magnitude (rightmost x-axis value). This increase can be possibly explained by two factors: a) the number of ASes advertised in the BGP Default-Free Zone (DFZ) in 2023 has quadrupled since 2003, while the IPv4 address space is in the exhaustion phase (Huston, 2023). Additionally, over the past two decades there has been a trend of consolidation in the IP transit market, which led to fewer but larger transit providers (Society, 2019).

2.4 The Incompleteness of the AS Graph

Our study relies on BGP data and AS relationships. The most significant limitation of the above data collection projects is the large number of missing links, which are divided into two types: *hidden* and *invisible* (R. Oliveira et al., 2009). Hidden links

are usually backup *c2p* links that can be observed when the preferred path to a prefix changes, and, invisible links are typically *p2p* links which are inherently unobserved due to the limited number of vantage points across the AS graph. Invisible links constitute the majority of missing links and can be located in the periphery of the AS graph (Roughan, Tuke, and Maennel, 2008; R. Oliveira et al., 2009; Khan et al., 2013).

The root cause of this problem is that giant CDNs who originate a large portion of today’s Internet traffic, often operate under a shroud of secrecy regarding their infrastructure details and peering arrangements with other ASes. This, coupled with the complexity introduced by IXPs, makes it challenging for external observers to map CDNs and their interconnections accurately.

While additional measurements, such as deploying more vantage points or using diverse datasets, may marginally improve the visibility of AS relationships, they are insufficient to fully resolve the incompleteness issue. First, even with more vantage points, the observation of AS paths remains constrained by the routing decisions of upstream ASes, which can obscure backup or unused links. Many *p2p* links are intentionally not advertised or used selectively, making them inherently invisible to measurement techniques reliant on observed routing announcements, such as BGP data.

Second, traceroute-based measurements are limited by the availability of responsive intermediate routers, which often mask their identities or omit critical details for privacy or security reasons. As a result, even with broader traceroute deployments, key AS links may remain undetected. Furthermore, certain ASes actively conceal their peering arrangements or employ non-standard routing practices, making their interconnections inaccessible to external observation. This is particularly true for major CDNs and cloud providers, which prioritize operational confidentiality.

Finally, the dynamic and distributed nature of the Internet further complicates measurement efforts. Routing policies can change rapidly due to traffic engineering or network failures, making it challenging to capture a complete and stable snapshot

of the AS-level topology. The highly asymmetric and decentralized deployment of Internet infrastructure means that the addition of vantage points or measurements does not guarantee a proportional improvement in visibility, particularly for peripheral ASes or smaller regional networks that are less connected to global measurement infrastructures.

To this day, the Internet Measurement community does not have an adequate solution for this issue, hence, our study suffers from the same limitation. We analyze the impact of the incompleteness problem in our study in Section 4.3.2.

2.5 Conclusion

In Chapter 2, we explored the fundamental building blocks of interdomain routing by examining the policies and business relationships that govern AS behaviors. We traced the evolution of the Internet’s structure, noting the significant changes brought about by the rise of CDNs and IXPs. We also reviewed key contributions from the literature, highlighting both the achievements and the persistent limitations in understanding and modeling interdomain routing.

However, understanding the structural elements of the Internet alone is not enough. In the next chapter, we will delve into specific factors that further complicate routing model accuracy—namely, the confounding factors that can skew or undermine predictions. These factors, such as the path diversity problem and or the geolocation-agnostic BGP best path selection process, introduce new layers of complexity that must be addressed to refine existing models.

Chapter 3

Confounding Factors

Chapter 3 builds upon the insights gained from the examination of routing policies and AS relationships in Chapter 2 by focusing on the specific confounding factors that complicate BGP routing models. While the previous chapter provided a macro-level view of how ASes interact and exchange traffic, this chapter dives into the micro-level challenges that arise when trying to infer accurate AS-paths. One of the key problems identified is the path diversity issue—where multiple possible paths exist between a source and a destination, making accurate predictions difficult. Another critical challenge is the first-hop problem, where determining the first AS-hop in a path is often prone to error, significantly impacting the accuracy of the entire path inference process.

Additionally, we will investigate the impact of the geo-agnostic BGP best path selection process in the accuracy of the contemporary models, which fail to account for the geographical location of ASes when predicting paths. As we break down these confounding factors, we will also evaluate the impact of missing data, incomplete topology information, and the inherent complexities of AS relationships. This chapter sets the stage for tackling these problems head-on, as we begin to rethink how routing models can be improved by addressing these critical issues.

3.1 Inference Overview

In this section, we break inter-domain routing modeling down to its fundamental blocks to better understand the underlying mechanisms that participate in the inference process. Moreover, we pinpoint the weaknesses of these mechanisms, to highlight possible assumptions/limitations that can cause an error in the inference. Specifically, we describe the best-path inference process as well as the datasets we use in our analysis and then we pinpoint the errors of the modeling process based on a defined set of accuracy metrics.

3.1.1 Best-Path Inference Process

In the Gao-Rexford model (Gao, 2001; Gao and Rexford, 2001), ASes form business relationships which translate into engineering constraints on traffic flows within and across these ASes and can be described as: (1) customer-provider, (2) peer-peer, or (3) sibling-sibling. Generally, the inference process is formulated as a function f , which takes as input the $AS_{topology}$ and the $AS_{relationships}$ (which are link annotations on top of the topology), and produces as output the inferred AS-path from a source AS ($srcAS$ or vantage point) to a destination AS ($destAS$ or origin AS):

$$AS_{path} = f(AS_{topology}, AS_{relationships}, srcAS, destAS) \quad (3.1)$$

Additionally, a basic assumption behind this model is the "valley-free" pattern: a reasonable AS should never provide transit between non-customer ASes because it has no financial incentive. Therefore, BGP paths should have zero or more customer-to-provider AS hops, followed by zero or one peer-to-peer hops, followed by zero or more provider-to-customer hops. Putting all that together, one can first infer the AS-level topology (Gao, 2001; Subramanian et al., 2002) and then leverage it to infer the AS-level paths followed from a source to a destination AS (based on the valley-free rule (Gao and Rexford, 2001)). While the above model and its variations have been used widely in the literature, it is well known that it fails to capture

many aspects of the inter-domain routing system (Mühlbauer, Feldmann, et al., 2006; Roughan, Willinger, et al., 2011; Anwar et al., 2015).

In our analysis, we leverage the simulator of Sermpezis and Kotronis (Sermpezis and Kotronis, 2019), which offers a Python implementation of the Gao-Rexford model. This tool takes as input a set of AS relationships to build the AS-level topology, and uses the Gao-Rexford model to predict AS-path(s) from any *srcAS* to any *destAS*.

3.1.2 AS-Relationships Datasets

At the time of conducting this analysis, the most comprehensive and widely used topology dataset is offered by CAIDA (CAIDA, 2023), using the methods described in (Luckie et al., 2013; Giotsas, Zhou, et al., 2013). A recent promising work, ProbLink (Y. Jin et al., 2019), seeks to improve the accuracy of the ASRank algorithm used in (Luckie et al., 2013).

These topology datasets are not static; they are calculated and updated at regular intervals to reflect the constantly evolving nature of the Internet’s AS-level topology. By incorporating recent BGP announcements, traceroute data, and IXP information, these datasets ensure that new relationships, peering arrangements, and other changes in the network are accurately represented.

We do not seek to improve or modify these algorithms, but understand how the best-available relationship datasets can be used to achieve more accurate AS path inferences.

3.1.3 Ground-truth Paths

To conduct a thorough analysis, it is essential to have a reliable ground truth dataset that will serve as a benchmark for evaluating the accuracy of the Gao-Rexford model in inferring AS paths. By comparing the model’s inferences to this dataset, we can rigorously assess its performance and identify any discrepancies or limitations in the model’s ability to accurately predict the AS path on the Internet. We leverage the

BGPStream API (Orsini et al., 2016) to collect the routing tables of all available vantage points from the RIPE RIS (NCC, 2022) and the RouteViews (Oregon, 2022) projects. We use these routing tables to extract the actual paths observed on the 10th of July, 2021, between arbitrary source and destination ASes.

Regarding the limitations of the observed paths, latest works (Anwar et al., 2015; Giotsas and Zhou, 2012; Giotsas, Luckie, et al., 2014) have shown that a significant portion of observed AS-paths do not follow the valley-free rule in practice; we verify that in our analysis (see Section 3.2.3). Moreover, the large path diversity that ASes demonstrate (Mühlbauer, Feldmann, et al., 2006; Mühlbauer, Uhlig, et al., 2007), makes it extremely challenging to infer the single best-path, when there are more than one best-paths on the real Internet (depending on the vantage point location).

Finally, this ground-truth dataset relies on a limited set of route collectors, such as those operated by RIPE RIS and RouteViews, which are heavily concentrated in specific geographic regions and often hosted by large ISPs or IXPs. As a result, the collected data provides an incomplete view of the AS-level topology, particularly in regions with sparse collector coverage. Additionally, ground-truth data only reflects the routing information visible to these collectors, potentially missing AS relationships, alternative paths, or policies that are not advertised or used in the observed routes. This inherent bias means that certain links, such as backup or selective routes, may be underrepresented, limiting the generalizability of findings derived from this data. Acknowledging these constraints is crucial for interpreting results and understanding the potential gaps in the analysis.

3.1.4 Performance Metrics

In the literature, path similarity is mainly measured using the *exact path match* and *path length match* metrics (Mao, L. Qiu, et al., 2005; J. Qiu and Gao, 2006; Mühlbauer, Uhlig, et al., 2007; Madhyastha et al., 2009; Gill, Schapira, and Goldberg, 2012; Singh et al., 2021). In this thesis, we define and use an extended set of metrics to capture the diverse settings of our controlled experiments. The

metrics are described below:

Exact AS-Path Match The ratio of inferred AS-paths that are the same as the observed AS-paths.

Path Length Match The ratio of inferred AS-paths that have the same length as in the observed path.

First-hop Match The ratio of inferred AS-paths that have the same first-hop as in the observed path.

First and Last-hop Match The ratio of inferred AS-paths that have the same first and last hop as in the observed path.

AS-to-ORG Path Match The ratio of inferred AS-to-ORG paths¹ that are the same as the observed AS-to-ORG paths.

AS-to-Rel Path Match The ratio of inferred AS-to-Rel paths² that are the same as the observed AS-to-Rel paths.

AS-to-Rel First-hop Match The ratio of inferred AS-to-Rel paths that have the same first-hop as in the observed AS-to-Rel path.

Jaccard Similarity The ratio of the intersection of the inferred and observed AS-paths over the union of the inferred and observed AS-paths.

3.2 Analysis

Leveraging the ground truth dataset we compiled in the previous step, we pre-process our data by filtering-out paths with cycles (an AS-path should be *loop-free*), removing paths with private and/or unallocated ASNs and finally removing AS-path prepending from every observed path. The produced dataset contains 94,828,639 different AS-paths between 7,485,974 different AS-pairs, i.e., between a *srcAS* and a *destAS*. To infer the AS-paths between the observed AS-pairs we

¹A path of organization IDs (CAIDA, 2022).

²A path of AS-relationships (CAIDA, 2023).

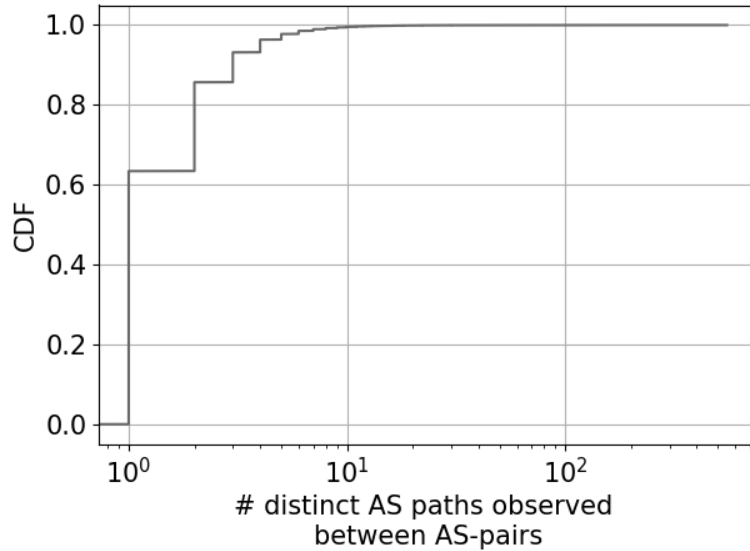


Figure 3.1: CDF of # distinct AS-paths per AS-pair.

conduct Monte-Carlo simulations (Wikipedia, 2022). Specifically, we perform 3 rounds of 10 experiments each, between 1000 random AS-pairs (per experiment) from the ground-truth dataset. For each experiment, we are using as input in the model the state-of-the-art CAIDA and ProbLink AS-topologies (see Equation 3.1), hence, we perform 60.000 simulations in total (3 rounds * 10 experiments/round * 1000 simulations/experiment * 2 AS-topology datasets/simulation).

At first, we study the performance of the model considering the path diversity problem (Mühlbauer, Feldmann, et al., 2006). Following, we explore the performance of the model with and without the first-hop knowledge (Mao, L. Qiu, et al., 2005). Finally, we identify and quantify additional confounding factors that affect the inference accuracy. This analysis sheds light on which variable(s) among the above confounding factors impact the most our AS-path inference accuracy.

3.2.1 The Path Diversity Problem

The Gao-Rexford model produces one single path from a source to a destination, even though there might be multiple options. However, this approach can lead to a




Figure 3.2: Average inference accuracy

heavily random output (Mühlbauer, Feldmann, et al., 2006). We validate that path diversity is an important problem in Fig. 3.1, since we observe that 85% of AS-pairs have at most 2 paths between them. Here, we study the impact of path diversity in our AS-path prediction process.

3.2.1.1 Vanilla Model Accuracy

In Fig. 3.2, we present the accuracy of the model across all performance metrics using both topology datasets (CAIDA and ProbLink). We average the performance scores across all experiments and we can see that the confidence intervals range from 0% to 2%, hence, we claim to have a low degree of uncertainty in our sampling method. Overall, using the CAIDA topology we achieve at least 10% better accuracy across all metrics, compared against the ProbLink topology.

Regarding the main metric of our analysis, *Exact Path Match*, the model scores 33,8% using the CAIDA topology and 21,4% using the ProbLink topology. This is evidence that the limitations addressed in recent work (Mühlbauer, Feldmann, et al., 2006; Anwar et al., 2015) still hold: *the Gao-Rexford model can predict AS-paths*




Figure 3.3: Average accuracy without path diversity.

exactly as they are observed on the real Internet, only 1/3 of the times.

Interestingly, we observe identical results between the *AS-to-ORG Path Match*, and the *Exact Path Match*. This indicates that we should not focus our research trying to fix sibling-sibling relationships, since this would offer diminishing increase in accuracy. Moreover, we see a significant (22,9%) difference between *First-hop Match* and *AS-to-Rel First-hop Match*, which indicates that even though the model can capture the proper AS-relationship in the first hop, unfortunately, it fails to select the proper ASn. I.e., this is an indicator that **(1) tie-breaking in the first-hop can lead to the whole path being inaccurately inferred.**

3.2.1.2 Accuracy Without Path Diversity

To address the path diversity problem, rather than comparing the ground truth paths to a single best inferred path, we compare them to the entire list of candidate best paths (from which the final best path is selected). In other words, we match one inferred path against all the available ground truth paths, and if at least one exact match is found, we consider the inference to be correct. This approach allows us to

account for multiple potential paths and achieve a higher level of inference accuracy when more than one path exists between an AS-pair. By evaluating all candidate paths, we can more accurately assess the model’s performance in handling path diversity and improve its ability to reflect real-world network conditions (Kastanakis, Giotsas, and Suri, 2022).

In Fig. 3.3, we can see that the *Exact Path Match* which can be achieved is 66,7% for CAIDA and 61,3% for ProbLink. Interestingly, the accuracy is 2 times higher than the vanilla model. Hence, the vanilla model can learn the actual best path, but fails to select it in the end, either due to misinferred AS-relationships, or, due to oversimplified path-decision making. Moreover, the *Jaccard Similarity* score

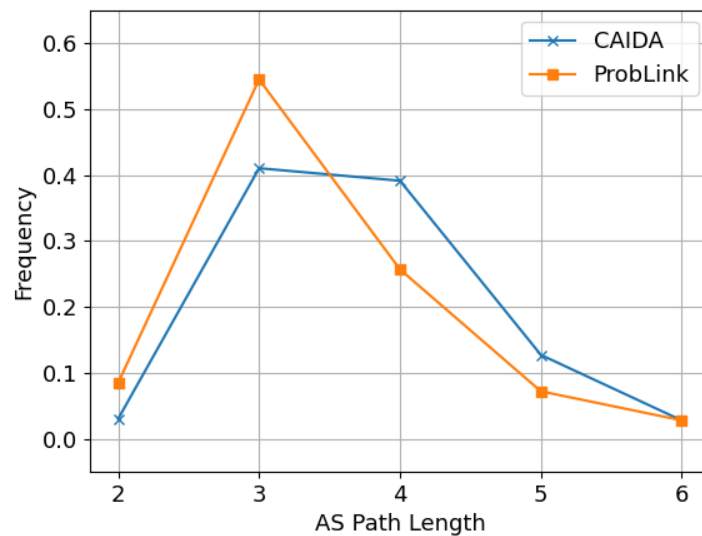


Figure 3.4: Frequency per path length.

is 86,8% for CAIDA and 83,6% for ProbLink. Hence, **(2) the possible "bug" which affects the *Exact Path Match* ratio is one AS in the inferred path**, if we safely assume (see Fig. 3.4) that the most frequent path length is 3 or 4. This suggests that even a minor error, such as one incorrect AS in a relatively short path, can significantly affect the overall accuracy of the model’s path inference.




Figure 3.5: Average accuracy when first-hop known.

3.2.2 The First-hop Inference Problem

In this section of our analysis, we focus on evaluating the efficacy of the model under the assumption that the first-hop in the AS path is already known. The ability to accurately determine the first-hop is a well-established critical factor that significantly influences the overall accuracy of AS-path inference. This, alongside the challenge of inferring AS-relationships, plays a major role in shaping the performance of the model, as highlighted by previous research (Mao, L. Qiu, et al., 2005).

The first-hop inference problem is critical because accurately predicting the initial AS-hop sets the foundation for constructing the rest of the AS path. In this section, we assume the first-hop is correctly predicted. This is a reasonable assumption to make, as many network operators already have detailed knowledge of their first-hop routing decisions, which are often determined by well-documented policies, such as local preference settings or contractual agreements with upstream providers. By making this assumption, we focus on evaluating the accuracy of subsequent hops, which are less directly observable. While this simplification may not account for inaccuracies in first-hop inference in all scenarios, it reflects common

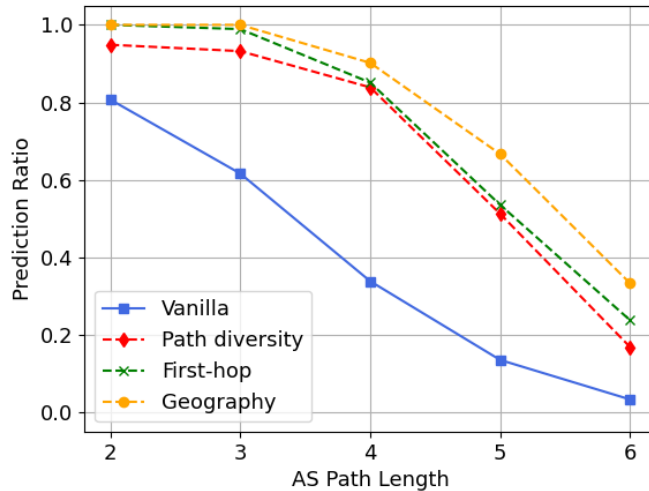


Figure 3.6: [CAIDA]Exact path match per path length.

operational realities and enables meaningful analysis of broader AS-path dynamics.

Towards that goal, we narrow our analysis to only those entries in the ground-truth dataset where the first-hop of the AS-path can be correctly predicted. By doing so, we isolate and examine the model’s capacity to infer the remaining portion of the path, beyond the first-hop. This approach allows us to specifically measure the model’s ability to deal with the complexities of predicting intermediate and end AS-hops, without the additional noise introduced by an incorrect first-hop prediction. Furthermore, concentrating on scenarios where the first-hop is known provides a more controlled environment to assess the model’s efficacy and how well it captures the dynamics of AS-path routing. By studying the problem of inferring the rest of the path accurately, we can provide insights into the broader challenges of AS-path inference and identify areas for improvement. In Fig. 3.5, we plot the inference accuracy of the Gao-Rexford model against the *Exact Path Match*, *AS-to-Rel Path Match*, *Path Length Match* and *Jaccard Similarity* metrics.

We observe that the refined model demonstrates a significant improvement in accuracy, achieving 2.5 times higher accuracy compared to the vanilla model—specifically, 84.1% for the CAIDA dataset and 82.1% for the ProbLink dataset. This substantial

increase is a clear indication that the ability to accurately determine the first-hop has a greater impact on improving inference accuracy than merely addressing the path diversity problem. In particular, **(3) the ability to determine the first-hop leads to a 2.5x increase in accuracy, which proves to be more crucial than addressing path diversity, which only results in a 2x increase.**

This conclusion is further supported by the data presented in Fig. 3.6, which illustrates the model’s accuracy as a function of path length. Across all path lengths, the accuracy is consistently higher when we focus solely on solving the first-hop inference problem, compared to scenarios where only the path diversity issue is addressed. This suggests that the first-hop determination plays a foundational role in the model’s performance, while improvements related to path diversity, although important, provide more incremental gains. Therefore, our findings strongly suggest that prioritizing first-hop accuracy is key to enhancing the overall inference capability of the model, and its impact is more pronounced than efforts aimed solely at handling multiple path options between Autonomous System pairs.

3.2.3 The Confounding Factors

Finally, building on the previous steps of our analysis, we aim to identify and quantify the remaining confounding factors that affect AS-path inference accuracy. Having measured the impact of the path diversity and the first-hop inference problems, we now turn our attention to the remaining sources of error in the inference process. To achieve this, we conduct a study on the specific instances where the model fails to correctly infer the path. By analyzing these ”path misses,” we can better understand the residual challenges and pinpoint additional factors that may be influencing the model’s performance.

We start our analysis by identifying the first *broken link* in the inferred path, i.e., the first AS in the inferred path which is different in the actual path. Then we classify it in one of the following categories:

Missing AS-relationship: In this class, the *broken link* occurs because the AS-

relationships’ dataset does not include a link for the respective ASes in the observed path. Since there is no edge in the simulated AS-topology, the inferred path cannot include the actual link followed in practice.

Valley-free violation: The valley-freeness is a logic outcome of the economic model described by the AS relationships (Giotsas and Zhou, 2012). In this class, the *broken link* occurs because the actual path has a valley, therefore, the Gao-Rexford model by default cannot produce this route or the AS-relationships, which are inferred using the valley-free model, are incorrect.

Local preference violation: The local preference is the most important BGP attribute a router looks at to determine which route towards a certain destination is the “best”. A local-pref violation happens when the model selects a route through a *less preferable neighbor* than the inferred one.

Shortest path violation: The second most important attribute a router looks at is the shortest path attribute. If two paths have the same local preference, then the model selects the shortest path among the best possible options. A shortest path violation occurs when the model selects a *shorter* path than the actual one.

Geo-agnostic path selection: Large content providers are deploying their own wide-area networks, bringing their networks closer to users, and bypassing the BGP selection process on many paths. Due to this flattening of the Internet topology (Chiu et al., 2015; Gill, Arlitt, et al., 2008), it is reasonable to consider the geographical distance between two ASes when inferring AS-paths. Yet, neither the actual BGP best path algorithm (CISCO, 2022) nor the Gao-Rexford model consider geography. Hence, in this occasion, the *broken link* occurs when the model tie-breaks randomly to a more distant AS than the actual neighbor in the observed path.

From Table 4.5, we can see that **(4) the most important factor that affects the inference process accuracy is the *geo-agnostic path selection* and the second most important factor is the *local preference violations***. We further demonstrate that geography plays a crucial role in routing modeling in Fig.

3.6, where the respective accuracy per path length when we have addressed the first-hop problem *and* taken geography into consideration is significantly higher than the vanilla model (2.75x increase in exact path match).

	CAIDA	ProbLink
Exact Path Match	84.1 %	82.1 %
Missing AS relationships	0.95 %	0.53 %
Valley-free violations	4.36 %	1.32 %
Local-pref violations	7.86 %	11.42 %
Shortest path violations	2.93 %	3.62 %
Geo-agnostic selection	8.54 %	10.31 %

Table 3.1: Key-factors that affect AS-path inference.

3.3 Conclusion

In Chapter 3 we explored how path diversity and first-hop prediction errors significantly reduce the precision of AS-path inferences. Additionally, we highlighted the inherent weaknesses of geo-agnostic models, which overlook the influence of geographical location on routing decisions. These insights underline the need for more sophisticated models that can adapt to the dynamic and complex nature of interdomain routing. The confounding factors explored in this chapter serve as a crucial step toward understanding the limitations of current models and identifying

areas for improvement.

As we move forward, Chapter 4 will investigate in more detail the local preference violations introduced in this chapter. In particular, we will re-examine the role of selective announcements and dynamic AS relationships in influencing routing decisions. By conducting a comprehensive replication study and analyzing the variability of local preferences across ASes, we will uncover how modern routing policies have evolved and the ways in which they deviate from traditional models. In Chapter 5, we will explore how integrating geo-location data can significantly enhance the fidelity of BGP simulations, providing a more accurate representation of routing dynamics across different regions of the world.

Chapter 4

Interdomain Routing Policies

Having explored the confounding factors that complicate BGP routing models, Chapter 4 now shifts the focus to rethinking how interdomain routing policies are modeled in the context of these challenges. Traditional routing models, which often assume static AS relationships and straightforward path selection, fall short of capturing the dynamic and selective nature of real-world routing policies. In this chapter, we will delve into the practice of selective route announcements, where ASes strategically announce or withhold routes to manage traffic flow. This practice has significant implications for both path prediction and network performance, particularly when combined with the dynamic variability in local preference settings across different ASes.

By revisiting a seminal study on routing policies (Wang and Gao, 2003) and conducting a replication analysis (Kastanakis, Giotsas, Livadariu, et al., 2023c), we aim to provide new insights into how these policies have evolved over the past two decades. Our findings will highlight the growing importance of frequent and adaptive policy inference, as well as the need for models that can accommodate the fluid and often unpredictable nature of interdomain routing. This chapter will propose a revised framework for modeling routing policies that accounts for both static and dynamic factors, offering a more accurate reflection of the current state of the Internet’s routing infrastructure.

4.1 Replication Overview

In 2003, Wang and Gao (Wang and Gao, 2003) presented an algorithm to infer and characterize routing policies as this knowledge could be valuable in predicting and debugging routing paths. They used their algorithm to measure the phenomenon of selectively announced prefixes, in which, ASes would announce their prefixes to specific providers to manipulate incoming traffic. Since 2003, the Internet has evolved from a hierarchical graph, to a flat and dense structure. Despite 20 years of extensive research since that seminal work, the impact of these topological changes on routing policies is still blurred. In this chapter we conduct a replicability study of the Wang and Gao paper, to shed light on the evolution and the current state of selectively announced prefixes. We show that selective announcements are persistent, not only across time, but also across networks. Moreover, we observe that neighbors of different AS relationships may be assigned with the same local preference values, and path selection is not as heavily dependent on AS relationships as it used to be.

4.1.1 Take-aways from the Original Paper

In 2003, Wang and Gao tackled the problem of inferring and characterizing the Internet routing policies. For the import routing policies, they observed that local preference values follow the Gao-Rexford model, namely customers are assigned with the highest *locpref* values while providers with the lowest *locpref* values. Additionally, they observed that ASes tend to assign *locpref* values based on next-hop instead of prefix. Nonetheless, 7 out of the 62 ASes had 10% or more neighbors that deviate from the Gao-Rexford model. It is unclear if those disparities were due to errors in the inference of AS relationships or unconventional *locpref* assignments.

Moreover, they described an algorithm to infer and characterize the export routing policies. In a nutshell, they studied whether customer or peer prefixes are reached through *customer* or *peer routes* respectively. If not, they were classified

as selectively announced routes. To determine export behavior, they searched for prefixes originated by customers or peers. If those prefixes had corresponding customer or peer routes (as defined above), it indicated that the customer or peer exported those prefixes to the provider or peer. Conversely, if such prefixes were absent or lacked customer or peer routes, it implied that the customer or peer did not export them directly to the provider or peer.

Specifically, they collected the public routing tables from RouteViews for a list of 16 ASes, and for each route they compared two AS relationships in the AS path: a) between the first AS and the origin AS (to characterize the prefix as customer/peer/provider), and b) between the first AS and the next-hop AS (to characterize the route as customer/peer/provider). If the customer prefix was announced through a *peer/provider route*, then they characterized the respective prefix as selectively announced prefix. They showed, that the percentage of selective announcements differs significantly between different ASes, with a range between 0% to 49%. The selectively advertised routes tend to be persistent, with only 17% of selectively advertised prefixes switching to non-selective within the period of a month. Last but not least, they found that the main cause for selective announcements is selective export policies, instead of other factors such as prefix splitting or prefix aggregation.

4.1.2 Replication Strategy

In Fig.4.1 we provide an overview of the datasets used in our study, for the import (left part) and export (right part) routing policies respectively.

As in the original Wang and Gao paper (Wang and Gao, 2003), our study of routing policies relies on AS relationship inferences. We use the current state-of-the-art AS relationships, made available by CAIDA (CAIDA, 2023). Moreover, we leverage data from Looking Glass (LG) servers (routeservers.org, 2023) to study import policies in Section 4.2. LG servers are interfaces to network devices that can be queried through web-based, telnet or ssh interfaces and allow users to query

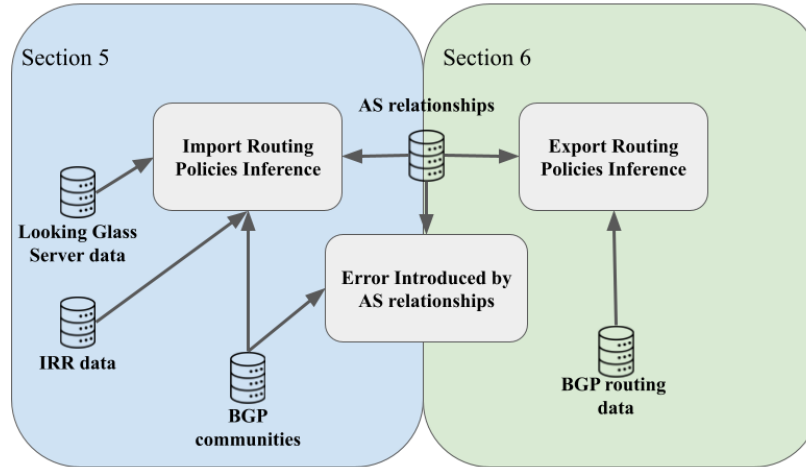


Figure 4.1: Overview of the data used for the import and export routing policies analysis.

BGP routing tables or measure traceroute paths from the perspective of the server’s location.

In contrast, with the import policies that can be directly observed by the AS that sets the *locpref* values, export policies have to be observed from the point of view of the neighbors that receive the announcement. The proposed way of Wang and Gao (Wang and Gao, 2003) to infer a customer’s export policies was to use the BGP table from its direct/indirect provider. We follow the same approach and use BGP data from the RouteViews (Oregon, 2023) and RIPE RIS projects (NCC, 2023) to infer the configuration of export policies and analyze the prevalence, the persistence and the causes of selective advertisements in Section 4.3.

In order to distinguish routes received from different neighbors, we use the conventions as in the original paper: 1) a *customer route* is a route received from a customer neighbor, 2) a *peer route* is a route received from a peer neighbor, 3) a *provider route* is a route received from a provider neighbor. Regarding the announced prefixes: 1) a *customer prefix* is a prefix originated by a direct/indirect customer neighbor, 2) a *peer prefix* is a prefix originated by a peer neighbor, 3) a *provider prefix* is a prefix originated by a provider neighbor.

Table 4.1: Characteristics of the ASes used in the import/export policies inference.

AS Number	AS Name	Degree	Location	AS Type
2495	Kansas Research and Education Network (KanREN)	19	USA (regional)	Educational/Research
6730	Sunrise	121	Europe	Cable/DSL/ISP
7922	Comcast	203	North America	Cable/DSL/ISP
53062	ACCESSOLINE TELECOM BACKBONE (GGT)	355	Brazil (regional)	Network Service Provider
62887	Whitesky Communications	82	USA (national)	Cable/DSL/ISP
3303	Swisscom	1194	Europe, USA	Cable/DSL/ISP
3257	GTT Communications	2831	Global	Network Service Provider
6939	Hurricane Electric	9780	Global	Network Service Provider
3549	Lumen AS	968	South America	Network Service Provider
37100	SEACOM	1133	Global	Network Service Provider
7018	AT&T	2438	North America	Network Service Provider
37271	Workonline Communications	344	Global	Network Service Provider
3292	TDC A/S (Tele Danmark)	360	Europe, USA	Cable/DSL/ISP
3741	Internet Solutions	806	Global	Network Service Provider
31027	GlobalConnect Group	344	Europe	Network Service Provider
852	TELUS Communications	474	North America	Network Service Provider
553	BelWü	1004	Germany (national)	Educational/Research
22548	NIC.BR	48	Brazil (national)	Non-Profit
5511	Orange	316	Global	Network Service Provider
6667	Elisa Corporation	501	Europe	Network Service Provider
1280	Internet Systems Consortium (ISC)	81	Global	Non-Profit
19653	CTS Communications Corp.	541	USA (national)	Cable/DSL/ISP
20751	AZISTA GmbH	28	Europe	Cable/DSL/ISP
2500	WIDE Project	22	Asia Pacific, USA	Educational/Research
5413	Daisy Communications	226	Europe	Cable/DSL/ISP
9009	M247	423	Global	Network Service Provider

4.2 Import Policies

4.2.1 Route Preference Among Provider, Customer and Peer Routes

The highest-priority metric when selecting the best path among all the available paths toward an IP prefix is the *locpref* attribute, that reflects how preferable is a route. Since *locpref* is not a transitive attribute, it is not possible to obtain *locpref* values through RouteViews and RIPE RIS route collectors. Instead, in (Wang and Gao, 2003) Wang and Gao queried the then-available LGs that provide a direct telnet interface to BGP routers of the ASes that deployed those servers. Such interfaces allow the querying of the full BGP Routing Information Base (RIB) along with the corresponding BGP attributes (both transitive and non-transitive). We replicate their methodology by querying the full routing table of the ASes that offer route server LGs at the moment of writing this paper. The selection of ASes in our study is in line with the original work of (Wang and Gao, 2003). Unfortunately, the original route server LGs used in (Wang and Gao, 2003) are not available online anymore, so we replicate the experiment with the currently available route server LGs. To this end, we compile a list of telnet and SSH LGs by parsing two resources: (a) PeeringDB (*PeeringDB* 2023), which is a voluntarily maintained database that aims to facilitate AS interconnection, and, (b) the `routeservers.org` website that provides a list of public route servers along with their access details (*Public Route Server* 2023).

In total we discovered route server LGs for 76 different ASes, of which 52 were offline and thus not accessible. For the remaining 24 route server LGs, 14 did not provide *locpref* values because the LG interface was running on an internal BGP router and the next-hop was to another router of the same AS. Therefore, we are able to collect *locpref* values only from 10 of the discovered LGs. The full list of all parsed LGs are available in (`routeservers.org`, 2023) to enable the repeatability of our experiments.

While the sample size of 10 LGs and their associated ASes may appear small, it is representative for several reasons. First, the ASes with accessible LGs span multiple geographic regions and include diverse types of networks, such as Tier-1 providers, Tier-2 ISPs, and regional networks. This diversity ensures that our dataset captures a wide range of routing behaviors and policy implementations. Second, the limited number of accessible LGs is a reflection of the inherent scarcity of publicly available LGs that provide detailed *locpref* values. Despite this limitation, the collected data enables us to observe key trends and validate our methodology. Lastly, by making our dataset publicly available, we facilitate future studies to extend our work and verify the representativeness of our results. Thus, while constrained by availability, the sample size is adequate for uncovering meaningful insights into *locpref*-based routing behaviors.

Table 4.2 shows the degree of consistency between *locpref* allocations and AS relationship types. We consider *locpref* allocations consistent with AS relationships if they reflect the Gao-Rexford (GR) ordering: $locpref_{p2c} > locpref_{p2p} > locpref_{c2p}$. For all routes, *locpref* allocations are consistent with AS relationships only in 83% of the cases, varying between 39-99% across the tested ASes. In contrast, in the original 2003 study the average consistency was above 99%, with only 2 ASes having consistency below 94% and 96%. Therefore, we observe that today *locpref* allocations have become significantly less conventional.

To better understand the deviations between the observed *locpref* values and the expected values according to the GR model, we examine the *locpref* consistency per relationship type. Peering relationships (p2p) appear to be deviating from the expected model at higher frequency. For instance, we observe that AS7922 (Comcast) uses the same *locpref* value between customers and peers, while AS3303 (Swisscom) uses the same *locpref* value between peers and providers. Note that when two different relationship types are assigned with the same *locpref* value, we assume that the relationship type that is consistent with the value is the one with the highest number of neighbors.

Table 4.2: List of ASes for which we extracted locpref values along with the percentage of routes that conform to the Gao-Rexford (GR) local preference model.

ASN	Customer	Provider	Peer	All routes	% Neighbors with one locpref value
2495	98%	98%	57%	99%	53%
3303	100%	99%	0%	44%	98%
5511	96%	N/A	98%	99%	63%
6730	98%	100%	100%	99%	88%
6939	86%	51%	100%	86%	86%
7922	98%	99%	0%	82%	71%
9009	42%	85%	100%	93%	66%
12779	92%	99%	30%	39%	85%
53062	90%	98%	96%	98%	76%
62887	99%	100%	10%	89%	97%
Average	90%	92%	59%	83%	70%

On average, only 59% of the p2p routes have a *locpref* value between the c2p and p2c values, nonetheless, this 59% is a mix of ASes that have either very high or very low compliance to the *locpref* model. This indicates a different peering strategy adopted by a decent number of ASes, nonetheless, categorizing ASes by their type and observing how their role affects their peering strategies is out of the scope of this work.

We hypothesize that the observed differences in consistency between our study and the original study can be explained by the advent of a much denser peering interconnection ecosystem, with different peering strategies that either did not exist 20 years ago, or if they existed they were much less popular (Böttger et al., 2018). Another reason is that given the size of such networks (e.g., Comcast or Swisscom), the business relationships they establish with their neighbors might be more complex than what the GR model can cope and describe (Giotsas, Luckie, et al., 2014).

While LGs provide a unique view of ground-truth *locpref* assignments in the control-plane, the small number of available route server LGs makes it hard to generalize the observations. Similarly to the original study, we try to complement the LG *locpref* data with data extracted from Internet Routing Registries (IRR), where operators often document their intended *locpref* values. We parse the IRR data available in RADB (Network, 2021), and we extract *locpref* documentation for 32 ASes that are also visible in the RouteViews BGP AS paths and have at least 50 neighbors. We extract *locpref* configurations either described in the `remarks` section of the IRR records, or expressed through the `pref` attribute of the Routing Policy Specification Language (RPSL). Table 4.3 summarizes our results. IRR *locpref* policies are generally more consistent with the GR model compared to the *locpref* allocations extracted from LGs. This is most likely due to the difference between actual control-plane configurations that actively affect routing decisions, and abstract policies described for documentation purposes.

4.2.2 Consistency of *locpref* with next-hop

```
route-map prefix-import permit 10
    match ip address prefix-filter
    set local-preference 200
```

Network operators might set their *locpref* based on next-hop or on prefix. For example, in the above configuration, the `match ip address prefix-filter` statement, specifies that the above rule should match routes that pass a specific prefix-filter. If, instead, we identify a next-hop specific rule (as depicted below), e.g., `match ip next-hop 203.0.113.1`, then the *locpref* for this route is set based on the next-hop AS.

```
route-map nexthop-import permit 10
    match ip next-hop 203.0.113.1
    set local-preference 200
```


In the last column of Table 4.2, we observe that only two ASes assign only a single *locpref* value to more than 90% of their neighbors. Instead, on average ASes assign more than one *locpref* values for 30% of their neighbors. Therefore, while ASes tend to assign *locpref* based on next-hop instead of prefix, we still see a non-trivial number of per-prefix *locpref* allocations. Here, it is worth to mention, that even though the GR model requires an AS to base its *locpref* value based on the business relationships with the next-hop on the AS path, nothing prevents an AS from basing its routing decisions on distant ASes along the AS path as well, (e.g., by prioritizing customer paths that do not traverse a distant, undesirable AS over customer paths that do traverse that AS) (Gill, Schapira, and Goldberg, 2013).

4.2.3 Error Introduced by AS relationships

Since our work relies on the inferred AS relationships, we verify them as in the original paper by comparing the inferred AS relationships against BGP communities. The BGP communities is an optional numerical attribute that is used to attach metadata on a route announcement. Among other types of metadata, many operators use BGP communities to annotate the relationship type of the neighbor from which a prefix was received (Donnet and Bonaventure, 2008).

The values of BGP communities and their corresponding meanings are arbitrary, but many AS operators document the use of their BGP communities either in IRR or in their websites. These values are 32-bit integers divided into two parts. The top 16 bits typically correspond to the 16-bit AS number of the AS that sets the community. The bottom 16-bits correspond to the actual meaning of the community. For example, the BGP community 3303:1000 is used by AS3303 to denote customer routes, while the community 3303:1004 is used by AS3303 to denote peering routes.

Step 1: Compile a list of relationship-tagging BGP communities. We manually compiled the BGP communities values and their corresponding meanings for 11 of the ASes listed in Table 4.1, and we keep the communities that are used to annotate relationship types. The documentation of the corresponding BGP

Table 4.3: Typical locpref assignments for 32 ASes which are selected from IRR.

ASN	% of typical locpref	ASN	% of typical locpref
1887	100%	20845	100%
2118	100%	20850	100%
5408	89%	21483	83%
6730	100%	24739	100%
6799	100%	35566	93%
8280	100%	39775	100%
8342	100%	43893	100%
8343	100%	44946	100%
8369	92%	47764	92%
8371	100%	49673	100%
9032	96%	50639	100%
12695	100%	52075	100%
12713	100%	60476	100%
15290	100%	199081	100%
15544	100%	199860	100%
16559	100%	396298	100%

communities have been extracted from IRR and the websites of AS operators.

Step 2: Map community to AS relationship. After collecting a list of relationship-tagging BGP communities, we parse BGP updates from RouteViews and RIPE RIS and we search for routes annotated with one or more of the collected BGP communities. We then map the attached communities to a link in the corresponding AS path by matching the first 16-bits of a relationship-tagging BGP communities value with an AS number in the path. More details on this methodology are described in the Appendix of (Wang and Gao, 2003).

Table 4.4 summarizes our validation statistics. For most of the tested ASes

Table 4.4: Validation results of AS relationships based on BGP Communities.

AS Number	% Validated Customers	% Validated Peers
1239	99% (271/272)	100% (16/16)
3292	99% (108/109)	99% (236/237)
3303	98% (45/46)	99% (823/829)
3257	99% (1497/1501)	99% (21/22)
3549	99% (945/959)	92% (12/13)
5511	98% (143/146)	95% (40/42)
6667	100% (16/16)	100% (416/416)
6730	100% (4/4)	100% (80/80)
7018	99% (2327/2335)	100% (34/34)
9009	100% (95/95)	100% (26/26)
12779	100% (28/28)	100% (724/724)

the inferred relationships agree with the BGP community tags for 99% of their AS links, which means that the error rate of the inferred relationships is negligible and we can interpret our observations as an outcome of routing policies and not as an artifact of erroneous AS relationship inference. We are restricted to the list of ASes that provide BGP feeds and routing information (*Public Route Server 2023*), among which we select the largest networks in terms of customer cone size and number of interconnections. This approach is in line with the methodology of the original work (Wang and Gao, 2003).

4.3 Export Policies

Export policies implemented by an AS play a major role in how prefixes are announced to its neighboring ASes. Usually, an upstream provider announces all of its prefixes to its customers. A customer on the contrary, may advertise its prefixes

either to all of its providers, or a subset of them for traffic engineering purposes. Figure 4.2 shows how AS13335 announces its prefixes. AS13335 is customer of both AS3549 and AS9498. However, AS13335 announces prefix p only to its direct provider AS3549, hence, AS9498 learns about prefix p via his peer (AS3257). Peers also have control over their prefix announcements to neighbors.

4.3.1 Export to Provider

Here, we first describe the algorithm used in (Wang and Gao, 2003) to infer the export policies that customers use to advertise their prefixes to direct/indirect providers. Then we study the prevalence, the persistence and the causes of these prefixes.

To study export policies we use the following datasets: a) the inferred CAIDA AS relationships (CAIDA, 2023) and b) the routing tables of 21 ASes (listed in Table 4.1) via the BGPStream API (Orsini et al., 2016) for different time periods.

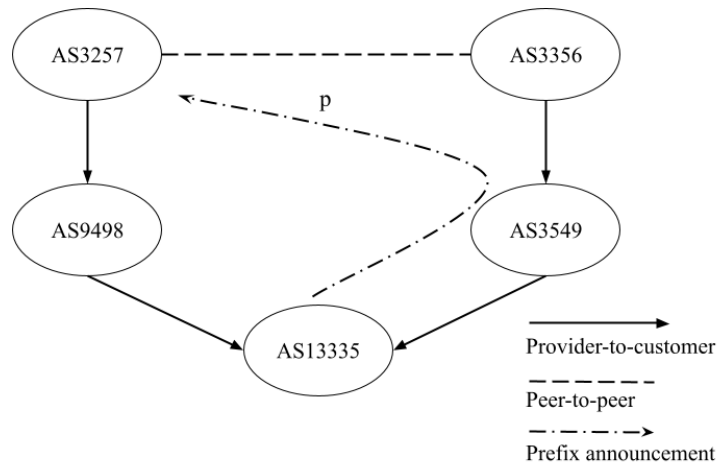


Figure 4.2: The export policies of AS13335, can be observed by its indirect provider AS3257. AS13335 announces prefix p to provider AS3549, but not to AS9498.

Algorithm 1: Algorithm for inferring export policies**Input:**

AS-relationships graph G
 AS o which originates prefixes P
 Routing table from the viewpoint of AS u

Output:

Whether P contains SA prefixes

Phase 1: Compute the Customer Cone of AS u

$CC = \{ \}$
 $S = \{u\}$
 helper = $\{ \}$
 while S is not empty:
 $s = S.pop()$
 if s not in helper:
 helper.add(s)
 $CC.add(s)$
 for each customer c of s :
 $S.add(c)$
 go to Phase 2

Phase 2: Determine if AS o is a customer of AS u

if o is in CC :
 go to Phase 3
 else:
 P does not contain SA prefixes

Phase 3: Determine if P contains SA prefixes

for each prefix p originated by AS o :
 if next hop AS w is not in CC :
 p is a SA prefix, P contains SA prefixes
 else:
 p is not a SA prefix
 if there is no SA prefix in P :
 P does not contain SA prefixes

4.3.1.1 Inference Algorithm

The direct way to observe the export policies of a customer is to use the BGP table from its providers (direct/indirect), since there is no discrete value (such as *locpref* in import policies) that describes the export preferences of an AS.

A customer can export its prefixes to all or a subset of its providers. If a direct/indirect provider receives a prefix originated by a customer AS (*customer prefix*) through a *peer/provider route*, this is a selectively announced prefix (*SA prefix*) and the origin AS is a selectively announced origin (*SA origin*). From the provider’s point of view, the best routes to *customer prefixes* are sufficient to capture the *SA prefixes* and *SA origins*. In a provider’s routing table, if a *customer route* to a prefix exists, the route should have the highest *locpref* according to the G-R model. Otherwise, the best routes are either *peer* or *provider routes*.

The process of inferring export policies (Algorithm 1), starts by computing the customer cone (*CC*) of an AS_{vp} (Phase 1). To that end, we collect all direct/indirect customers of AS_{vp} by using the Depth-First Search (DFS) algorithm on a directed AS topology graph composed of only p2c AS links. In the next phase (Phase 2), we parse the BGP table of AS_{vp} and extract all routes originated by ASes that belong in the *CC* of AS_{vp} (i.e., all *customer prefixes*). Finally in Phase 3, if the *customer prefixes* are learned from a *peer/provider route*, then, we characterize the prefix as *SA prefix* and the origin AS as *SA origin*.

4.3.1.2 Prevalence of *SA Prefixes*

We explore the existence of *SA prefixes* and *SA origins* on the 1st of April, 2023¹. We collect the routing tables of 21 ASes described in Table 4.1, using all available route collectors from the projects Routeviews and RIPE RIS (see Section 4.1).

Note here, that *SA prefixes* for a provider may be due to the selective

¹Similar duration as in the original Wang and Gao study (Wang and Gao, 2003).

Table 4.5: % of SA prefixes and SA origins observed by 21 ASes.

AS number	% of SA prefixes	% of SA origins
3303	69	45
3257	54	55
6939	44	44
3549	32	26
7018	28	17
37100	27	18
37271	21	12
3741	19	12
31027	13	09
852	13	09
3292	12	10
553	05	02
22548	04	03
5511	02	03
6667	0.01	0.01
1280	0.001	0.001
19653	0.0001	0.0001
20751	0	0
2500	0	0
5413	0	0
9009	0	0




Figure 4.3: CDF of all SA origins in 2023.

announcement policies of the origin or intermediate ASes. For instance, in Figure 4.2, the *SA prefix* for AS3257, may be due to the selective announcement policies employed by AS9498 as well. We study this possibility in Subsection 4.3.1.5.

Table 4.5 shows the percentage of the total *SA prefixes* and *SA origins* observed in the routing tables of each AS we considered in our study². We find that ASes like Swisscom (3303), GTT Communications (3257), Hurricane Electric (6939), and AT&T (7018), observe a significant portion of *SA prefixes*. For Swisscom 69% of the observed prefixes are *SA prefixes*, which means that a high portion of prefixes are reached through a *peer/provider route*, rather than a *customer route*, as expected based on the AS relationships.

Next, we examine *SA prefixes* from the point of view of the customers that announce them, namely the *SA origins*. In Figure 4.3 we plot the CDF of *SA prefixes* per *SA origin* for all the ASes in our study. We find that more than 75% of *SA origins* announce all of their prefixes selectively. Note that in the 2003 study, the authors analyzed only eight *SA origins* which were common among three of the

²In the original paper 16 ASes were studied based on the availability of LG servers and which ASes were peering with Routeviews collectors. We follow a similar approach in our paper.

AS_{vp} providers. Among those eight *SA origins*, none advertised all of their prefixes selectively, and only two *SA origins* advertised more than 90% of their prefixes selectively. However, as we show in Section 4.3.1.4 even in 2003 the majority of *SA origins* advertised 100% of their prefixes selectively, and the result of the original paper is probably due to under-sampling bias.

4.3.1.3 Verification of *SA Prefixes*

In this section we verify the AS-relationships used on the export inference Algorithm 1. To verify *SA prefixes*, we first verify the AS relationships for direct and then for indirect customers respectively.

Step 1: Verify AS-relationships between an AS and its direct customers.

In Subsection 4.2.3, we verify the AS relationships between 11 ASes in Table 4.4 and their neighboring ASes using relationship-tagging community values. The inference error is very small, therefore we can be confident in the AS inferences.

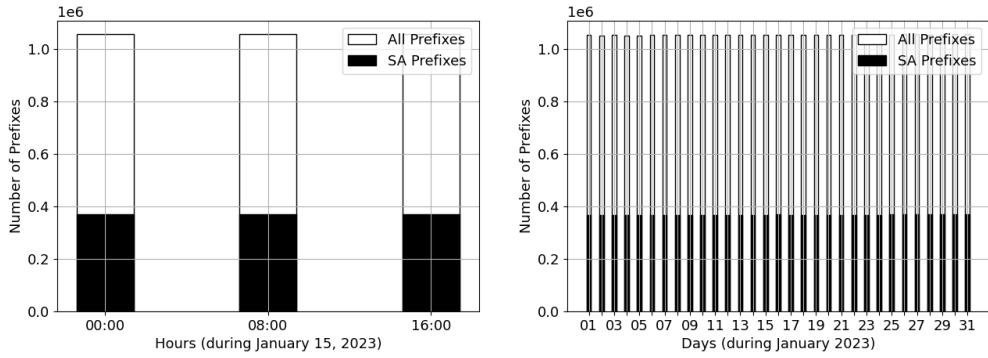
Step 2: Verify AS-relationships between an AS and its indirect customers.

We follow the approach of the original paper to verify all AS relationships between a provider and its indirect customers. For each *SA prefix* observed by an AS_{vp} , we search all the BGP routing tables to find if there is an AS path between AS_{vp} and the origin AS that traverses only p2c links. In that case we consider that we have at least one *active customer route* between the two ASes, and the *SA prefix* is verified. We only consider *SA origins* with a high number of observed prefixes.

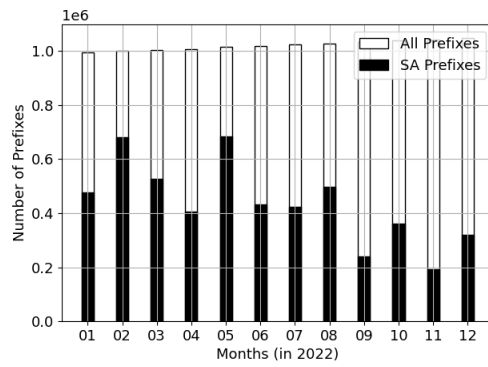
Table 4.6 shows that for most ASes, more than 80% of *SA prefixes* are verified. In contrast with the original paper, we consider both routes with typical and atypical *locpref*. The average conformance in *locpref* settings is 83% as shown in Table 4.2, which explains the average 19.2% of unverified *SA prefixes* per *SA origin* in our study.

Table 4.6: Percentage of SA prefixes verified per AS.

AS number	% of verified SA prefixes (# of total SA prefixes)
6667	97 (89)
3741	89 (449)
37100	83 (1432)
3292	83 (7076)
31027	82 (7887)
3303	82 (1955)
852	81 (7775)
7018	80 (14228)
553	80 (2339)
37271	80 (9748)
22548	80 (2329)
5511	76 (584)
3549	74 (19842)
6939	74 (24231)
3257	70 (44803)




(a) SA prefixes for AS7018 on Jan. 15, (b) SA prefixes for AS7018 on Jan. 2023.



(c) SA prefixes for AS7018 in 2022.

Figure 4.4: Daily and monthly persistence of SA prefixes for AS7018.



(a) CDF of SA prefixes uptime in Jan. (b) CDF of SA prefixes uptime in 2023.

Figure 4.5: Daily and monthly uptimes of SA prefixes for AS7018.

4.3.1.4 Persistence of SA prefixes

Network operators may configure their export policies using different patterns over longer time periods. This could in turn affect the behavior of *SA prefixes*. Having identified the prevalence of *SA prefixes*, we focus further on characterizing the persistence of these prefixes. To this end, we collect two families of datasets using BGPStream (Orsini et al., 2016) from all the routing collectors. The first family covers short-time periods for AS7018 and the collection method resembles the one employed by Wang and Gao (Wang and Gao, 2003). We go one step beyond and characterize the *SA prefixes* persistence over the span of 20 years for AS3259, AS3292, AS3549, AS5511 and AS7018.

For our short-term period analysis, we focus only on AT&T (AS7018) since it has a large number of *SA prefixes*. In the original paper the analysis focused only on AS1 as it had one of the highest number of SA prefixes. In our case the equivalent AS is AS7018 since AS1 is not as well-connected anymore. We thus fetch routing tables of AS7018 for: a) the 15th of January 2023, b) all 31 days of January 2023, and c) the 1st day of each month during 2022. We show in Figure 4.4a the number of *SA prefixes* during the 15th of January 2023, for AS7018, while Figure 4.4b illustrates the number of such prefixes for every day of January 2023. Same as in the original study, we find that the contribution of *SA prefixes* is consistent during the period of a day and the period of a month. When expanding the measurement period to one year, we find that *SA prefixes* exhibit an unstable behavior (see Figure 4.4c). This could either be explained due to customers switching their export policies for an *SA prefix*, or due to providers switching their import policies. In the 20-year longitudinal analysis, apart from AS7018 we also include AS3257, AS3292, AS3549 and AS5511. Our data comprises of routing tables from the 1st of April of each year between 2003 and 2023.³

To find out how export policies affect the *SA prefixes*, we follow the approach of

³*SA origins* findings are omitted due to similar insights with the *SA prefix* findings.

the original paper. We define *SA prefix uptime* as the times an *SA prefix* appears during the measurement time window. For example, an *SA prefix* can have a minimum of 1 day uptime and a maximum of 31 days uptime in a month, or a minimum of 1 month and a maximum of 12 months uptime in a year, depending on which view of the data we examine. Figure 4.5a shows the distribution of the SA prefix uptime for January 2023. More than 90% of the *SA prefixes* are stable for the entire month, since most of the SA prefixes have an uptime of 31 days. On the contrary, when we study the monthly uptime for 2022, the results range from 1 to 12 months as shown in Figure 4.5b. This instability in *SA prefix* ratio is a strong indicator that the modeling of BGP routing policies should be conducted with high periodicity, due to the dynamic nature of AS connectivity and the derived routing policies.

Figure 4.6 shows the boxplot distribution of the *SA prefix* ratio for the 5 major AS providers over the last two decades. The ratio of *SA prefixes* is highly dynamic from year to year. However, there is a pronounced jump in the median ratio of *SA prefixes* after 2006, which stayed consistently above 0.25 since then. These results are evidence that the ratio of *SA prefixes* can be sensitive to topological and policy changes, and the assessment of *SA prefixes* should be updated regularly. The maximum SA prefix ratio was observed in 2021 by AS3549 (a high-centrality network), with a value of 80.9%.

In Figure 4.7, we plot the CDF of *SA prefix* ratio from the customers point of view, over the last two decades. We observe that the distribution of *SA prefix* ratios is consistently skewed toward 100%, meaning that the majority of *SA origins* announce only *SA prefixes*. When comparing the providers point of view in Figure 4.5 with the customers point of view in Figure 4.7, we can see that the left skew in the fraction of ASes that advertise 100% of their prefixes selectively is correlated with shifts in the median *SA prefix* ratio observed by providers. Therefore, **we present evidence that *SA prefixes* are mainly an outcome of selective export policies, and not selective import policies from the provider.**

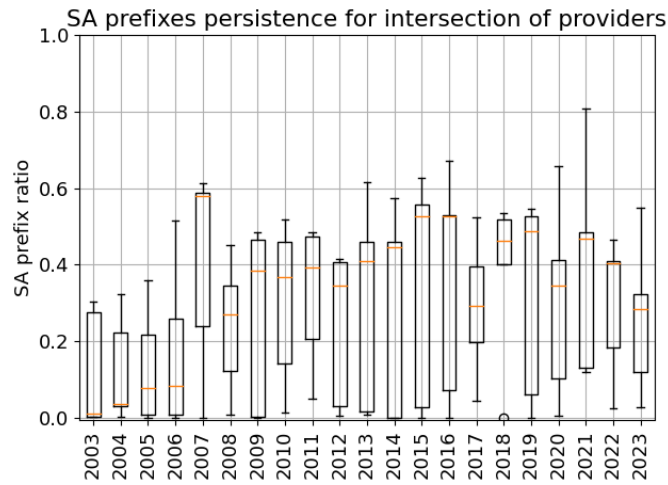


Figure 4.6: SA prefix ratio over the last 20 years for AS3257, AS3292, AS3549, AS5511 and AS7018.

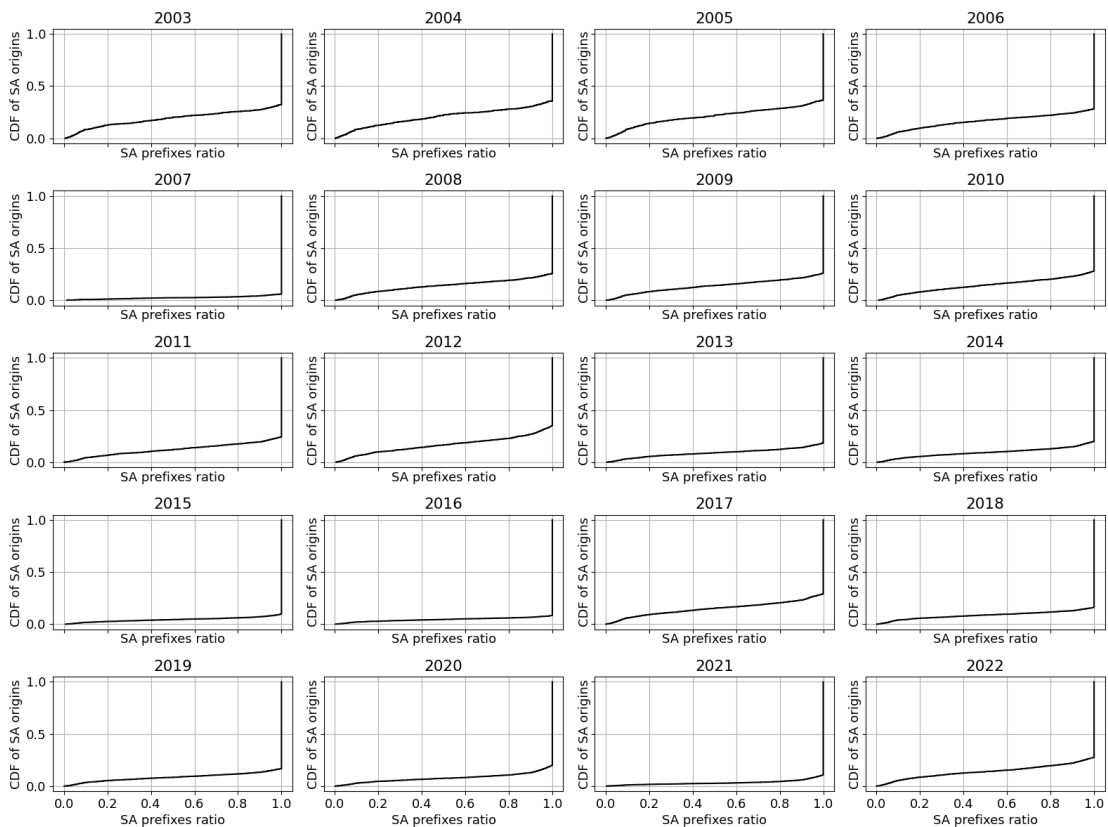


Figure 4.7: CDF of union of SA origins over the last 20 years for AS3257, AS3292, AS3549, AS5511 and AS7018.

4.3.1.5 Causes of SA prefixes

As mentioned in Subsection 2.1.1, customers may connect to multiple providers for traffic engineering purposes and/or to make the reachability to their prefixes resilient to link or node failures. Intuitively, it is unlikely for *single-homed* ASes to apply selective announcements, otherwise not all of their routes will be globally reachable, however, their upstream providers in the AS path, which are *multi-homed*, may apply selective policies.

For AS3257, AS3292, AS3549, AS5511 and AS7018, we examine how many SA origins are *multi-homed*. From Table 4.7, we observe that, at most, 1 out of 3 customers is *single-homed*. For these ASes, an intermediate AS in the path applies selective export policies rather than the *SA origin*. Compared to the original paper, we observe a $\sim 10\%$ increase in the *single-homed* customers for AS3549 and AS7018, since 2003. Apart from selective announcements though, there are other factors that may give rise to *SA prefixes*.

Case 1: Prefix Splitting. Network operators may split a prefix into more specific prefixes for resilient traffic engineering. Assume an AS_0 originates a /23 prefix p_0 , to which it wants to load-balance traffic between two of its providers, AS_1 and AS_2 . At the same time AS_0 wants to ensure that if a link to one of its providers

Table 4.7: % (#) of multi-homed and single-homed SA origins for AS3257, AS3292, AS3549, AS5511 and AS7018

AS number	% (#) of single-homed SA origins	% (#) of multi-homed SA origins
3257	37.5 (16810)	62.5 (27993)
3292	32.2 (2280)	67.8 (4796)
3549	36.3 (7207)	63.7 (12635)
5511	14.0 (82)	86.0 (502)
7018	33.2 (4718)	66.8 (9510)

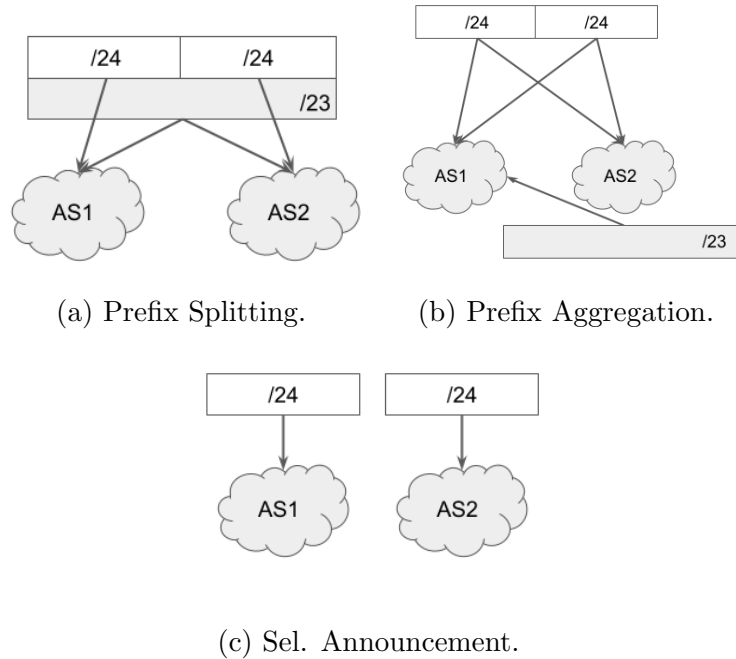


Figure 4.8: Causes of a SA prefix

fails, traffic flows to p_0 will not be disrupted. To that end, AS_0 splits the $/23$ prefix to two more specific $/24$ prefixes, and advertises the $/23$ prefix to both AS_1 and AS_2 providers, and each $/24$ to a different provider. In that case AS_1 and AS_2 will not see the $/23$ as an *SA Prefix*, but they will see one of the covered $/24$ prefixes as an *SA Prefix*. Figure 4.8a illustrates this case.

Case 2: Prefix Aggregation. An *SA prefix* may arise due to prefix aggregation along the path. Lets assume that an origin AS_0 originates two consecutive $/24$ prefixes p_1 and p_2 , and advertises both to its two providers AS_1 and AS_2 . AS_2 may opt to aggregate the two consecutive prefixes to their covering $/23$ prefix p_0 in order to conserve memory in the routing table, while AS_1 does not aggregate p_1 and p_2 . AS_1 may then receive p_0 through AS_2 (directly or indirectly). Since it has not received the $/23$ prefix directly from AS_0 it will appear as an *SA prefix*. This is illustrated in Figure 4.8b.

Case 3: Selective Announcement. An origin AS_0 may load balance non-consecutive prefixes. In that case no prefix splitting or aggregation is possible, and

Table 4.8: Causes of SA prefixes

AS number	% of prefix splitting	% of prefix aggregation
3257	1.3	0.03
3292	0.02	0
3549	0.4	0.1
5511	1.6	0
7018	0.3	0.3

the load-balanced prefixes are advertised selectively to different providers. In that case, each provider will learn the prefixes it does not receive directly as *SA prefixes*. This example is shown in 4.8c.

We study whether prefix splitting and aggregation are the main reasons of *SA prefixes*. For prefix splitting, we study how many *SA prefixes* can be aggregated by a non-SA prefix of the same origin AS. For prefix aggregation, we observe how many *SA prefixes* in the routing table of AS_{vp} can be aggregated in the BGP tables of the remaining ASes. Table 4.8 shows that both ratios are negligible, therefore, the main cause of *SA prefixes* cannot be prefix splitting or prefix aggregation.

4.3.2 Export to Peer

In this section, we use the algorithm described in Section 4.3.1.1 with minor tweaks, to infer export policies that peers use to advertise their prefixes to other peers. Specifically, we study whether the ASes of Table 4.1 reach their peers' prefixes through *peer* or *provider routes*. If an AS reaches a peer's prefix through a *provider route*, the prefix is *SA prefix*. To account for peer *SA prefixes*, we make the following changes in Algorithm 1. In Phase 2, instead of checking the customer relationship between the origin AS and AS_{vp} , we test whether the two ASes are peers. In Phase 3, instead of checking if the next-hop belongs in the *CC* of the AS_{vp} , we test whether AS_{vp} belongs in the *CC* of the next-hop. Specifically, we check whether the route

towards a *peer prefix* is a *provider route*, by studying whether the next-hop is an upstream provider of AS_{vp} .

In Table 4.9, we observe that selective announcements is not a phenomenon prevalent among peer AS networks, since more than 90% of the *peer prefixes* are reached through *peer routes* rather than upstream providers' links. However, it should be highlighted that this result may not be representative of the actual selective advertisement practices between peers, because routes exchanged over peering links have limited visibility due to the valley-free rule (R. Oliveira et al., 2009).

This limited visibility due to the fact that BGP feeds are mostly provided by high-tier ASes and some geographic areas are poorly covered. Furthermore, two-thirds of all contributing ASes configure their connection with the BGP collector as a p2p link, which means they advertise only routes learned from customers. Theoretically optimal placement of BGP monitors might mitigate this incompleteness (R. Oliveira et al., 2009), but in practice ASes participate voluntarily in such data collection projects so optimal placement is not possible. Some researchers suggest highly distributed traceroute monitoring infrastructures (Shavitt and Shir, 2005; Chen et al., 2009) are a promising approach to discover invisible AS links, yet the visibility improvement so far is limited compared with the links discovered at just a single IXP by Ager et al. (Ager et al., 2012). Therefore, a selectively advertised peering prefix may be invisible to the BGP collectors, especially when the peering link is not adjacent to the peer of the BGP collector.

To study how much the incompleteness of the AS graph affects our results, we plot in Fig. 4.9 the unique SA prefixes that we can identify by incrementally adding vantage points. We observe that the unique SA prefixes size increases, especially when we include high centrality ASes with wide geographical coverage (e.g., AS3257). It is worth noting that the rate of increase does appear to plateau, indicating a reasonable lower bound in our results. This is in line with the observations of related works (Dhamdhere and Dovrolis, 2008; R. V. Oliveira et al.,

Table 4.9: % of peer SA prefixes and % peer SA origins.

AS number	% of SA prefixes	% of SA origins
AS5413	9.1	0.4
AS3741	8.2	1.1
AS3303	7.6	1
AS19653	6.6	0.8
AS37100	4.6	0.8
AS553	3.7	0.5
AS6667	2.8	0.2
AS852	2.7	0.2
AS6939	2.1	1.5
AS3292	1.3	0.1
AS37271	1.1	0.1
AS1280	1.1	0.06
AS31027	0.5	0.07

2008; Milolidakis, 2022), which suggest that the fraction of visible links increases linearly with the fraction of used monitors, so, the estimated population size of these links should be viewed as a lower bound on the actual population size.




Figure 4.9: Progressive enumeration of unique SA prefixes by incrementally adding vantage points. The rate of increase does appear to plateau, indicating that our results are a lower-bound estimation of the SA prefix ratio, due to the incompleteness problem of the AS graph.

In this chapter, we replicate the study of Wang and Gao (Wang and Gao, 2003) and conduct a longitudinal analysis on the evolution of selective announcements over the past twenty years. Our results provide experimental evidence that a large part of their findings still holds true today, such as a) the persistence of *SA prefixes* over time, and, b) the prevalence of *SA prefixes* across different ASes. On the contrary, the assignment of *locpref* settings among ASes, is significantly less conforming to AS relationships, especially for peering links.

4.4 Conclusion

In Chapter 4, we critically re-examined interdomain routing policies through the lens of evolving practices such as selective route announcements and dynamic local preferences. Our findings from the replication study revealed significant shifts in how ASes manage their routing policies, indicating that the rigid models of the past are no longer sufficient to capture the complexity of modern interdomain routing. The use of selective announcements has increased, and ASes are now adopting more nuanced and variable local preference strategies. This chapter reinforced the need for adaptive models that can keep pace with the fluid and evolving nature of AS behaviors.

In Chapter 5, we will introduce a critical new dimension to our routing models—geolocation. While we have so far focused on the policy aspects of routing, geographical location plays a crucial role in determining how ASes interact and route traffic (as introduced in Chapter 3). Traditional models often overlook this aspect, leading to inaccuracies in path predictions (Anwar et al., 2015; Kastanakis, Giotsas, and Suri, 2022; Kastanakis, Giotsas, Livadariu, et al., 2023c; Kastanakis, Giotsas, Livadariu, et al., 2024). In the next chapter, we will explore how integrating geolocation data can significantly enhance the fidelity of BGP simulations, providing a more accurate representation of routing dynamics across different regions of the world.

Chapter 5

Location Specific Announcements

Chapter 5 marks a pivotal shift in our approach to BGP modeling by introducing geolocation as a critical factor in AS-path prediction. Up until this point, our exploration of routing models has primarily focused on the policy-driven aspects of interdomain routing, such as selective announcements and local preferences. However, geographical location is another essential component that shapes routing decisions, as the physical distance between ASes, regional traffic patterns, and geopolitical factors can all influence the paths that data takes through the network.

Towards that goal, we will study the routing paradigm of anycasting. Anycast routing offers transparent service replication by distributing traffic across multiple Points of Presence (PoP). By advertising the same IP prefix from each PoP via BGP, traffic is routed to the nearest server, minimizing user latency. Despite its perceived benefits, prior research suggests IP anycast often falls short, with clients routed to distant replicas, increasing latency. Selective announcements made by anycast ASes contribute to this inefficiency, serving as a traffic engineering strategy to control incoming traffic flows.

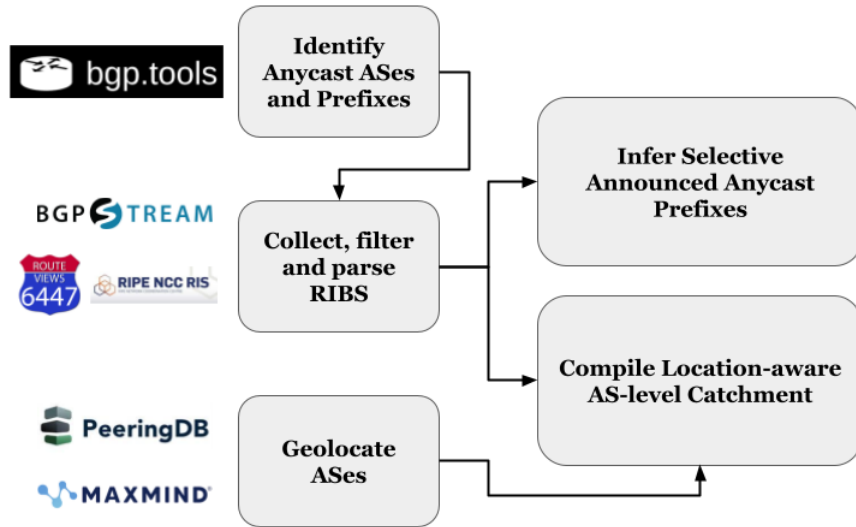


Figure 5.1: Tools used and overview of the methodology.

5.1 Methodology Overview

The goal in this chapter is to study the selective routing policies of all anycast ASes in the interdomain routing system per geographical region. An overview of our methodology can be found in Fig. 5.1 and consists of the following steps:

Identify Anycast ASes and Prefixes To initiate our study, we compile a map of anycast ASes and their associated anycast prefixes. To that end, we leverage the `bgp.tools` anycast prefixes and ASes dataset (Cartwright-Cox, 2024b) (the methodology to detect such prefixes is described in (Cartwright-Cox, 2024a)). This is an important step, in which we narrow down our analysis only to anycast IP prefixes announced by anycast ASes.

In Table 5.1, we provide only the characteristics of the top anycast ASes based on their RUM uptime (*cdnperf.com* 2024), but our methodology applies to every anycast AS. RUM (Real User Monitoring) uptime measures service availability and performance based on actual user experiences, providing a realistic assessment of network reliability. This makes it suitable for ranking anycast ASes, as it reflects user-centric metrics across diverse geographical locations.

Although our dataset spans all anycast ASes, presenting the full details for

ASN	AS Name	RUM Uptime	# of PoPs
13335	Cloudflare	99.43	197
16509	Amazon	99.37	166
15133	BytePlus	99.30	59
54113	Fastly	99.20	102
20940	Akamai	99.19	183
60068	CDN77	98.65	63
16276	OVH	99.20	43
21859	Zenlayer	99.20	78
199524	G-Core	99.14	91
15169	Google	98.96	135
30081	Cachefly	98.62	66
22822	Edgio	97.75	72

Table 5.1: Top Anycast Networks based on RUM Uptime.

every AS would complicate our discussion without providing substantial new insights beyond those highlighted by the top-ranked ASes. For simplicity and clarity, we focus on the most illustrative examples, as the selected ASes effectively demonstrate the critical trends in RUM uptime and performance. However, to ensure transparency and reproducibility, the complete dataset and methodology covering all anycast ASes are available in our online repository (Kastanakis, 2024).

Collect, Filter and Parse RIBS In this step, we collect the BGP routing tables of all (691) anycast ASes (identified in the previous step) on the 1st of November, 2023, through the BGPStream API (Orsini et al., 2016) which includes the route collectors of RIPE RIS (NCC, 2023) and RouteViews (Oregon, 2023). To gain the best view in terms of geographical distribution and coverage, in this project we make use of all available public route collectors up to this date (63). Note, that some of the BGP collector peers may contribute only partial routing tables (for example they may send only prefixes received by their customers to the BGP collector) (R. Oliveira et al.,

2009). For those ASes, we may overestimate the number of selectively advertised anycast prefixes. Therefore, our results should be considered as an upper bound of selective advertisements of anycast prefixes.

Infer Selective Announced Prefixes Anycast ASes often employ selective announcement strategies, where they announce subsets of their anycast prefixes to specific neighboring ASes.

In this step, we follow the approach of (Wang and Gao, 2003) to infer the selective announced prefixes in the interdomain routing system. The methodology of inferring selective announcements (also described in 5.2.1) relies on the assumption that if a prefix is received through a more expensive route than what is expected, either the origin AS or an intermediate AS in the AS-path applied selective export routing policies.

To label a selective announced prefix, we use the state-of-the-art AS-relationships (CAIDA, 2023) as well as the routing tables for all anycast ASes collected in the previous step. We follow the logic and conventions of (Wang and Gao, 2003): *if a customer prefix is received through a peer/provider route, then this is a selective announced prefix. Similarly, a peer prefix is selectively announced if it is received through a provider route.*

From a total of 691 anycast ASes, we found that 581 ASes (84.06%) announce at least one anycast prefix to only a subset of their neighbors. In Fig. 5.2, we plot the CDF of selective announced prefixes per anycast origin AS. We find that 80% of the selective anycast ASes announce *all* of their prefixes selectively. Specifically, all the top anycast ASes mentioned in Table 5.1, announce 100% of their anycast prefixes selectively, while, the average selective announced prefix ratio across all anycast ASes is 82.5%¹.

Augment ASes with location specific characteristics To contextualize our findings within a geographical framework, we geolocate *all* ASes, up to this date,

¹Due to space constraints we refer the reader to (Wang and Gao, 2003; Kastanakis, Giotsas, Livadariu, et al., 2023c) for the validity of the selective announcements inference algorithm.




Figure 5.2: Selective announced prefixes per anycast AS.

into their respective countries and regions. By incorporating geolocation data, we aim to uncover regional trends and disparities in anycast deployment and routing behavior.

When geolocating an AS, we consider: a) *the prefixes that an AS announces* as well as b) *its public peering locations*. Towards that goal, we extract country-level information for all prefixes announced by an AS from MaxMind (*MaxMind API* 2024) through the RIPEstat API (*RIPEstat API* 2024) as well as the countries of the public peering locations from PeeringDB API (*PeeringDB* 2023). We further use the United Nations dataset (*United Nations Statistics Division* 2024) to map countries to their respective regions. We leverage these data in the following step of our methodology as well as in our analysis in Section 5.3, where we classify the neighbors of an anycast AS into *regionals* or *globals* based on their geographic footprint.

The geolocation of infrastructure components in our study relies on existing geolocation datasets (*MaxMind API* 2024; *United Nations Statistics Division* 2024). While the state-of-the-art geolocation techniques are not 100% reliable,

these datasets are generally accurate enough on the country-level granularity and appropriate for understanding broad geographic trends (Winter et al., 2019). In our analysis, the primary goal is to capture regional trends and spatial dynamics rather than pinpoint precise country-level locations. Thus, the inherent limitations of geolocation data have minimal impact on the broader conclusions of our study. By using widely adopted and validated datasets, we ensure that the data is robust enough to support the insights derived from our work, even if some granular inaccuracies remain.

Compile the Location-aware AS-level Catchment A catchment refers to the geographic region served by a specific PoP and represents the set of prefixes that are routed to that particular anycast site. When a user sends a request to an anycast IP address, the routing infrastructure directs that request to the PoP that is topologically closest to the user in terms of network hops.

In this work, instead of mapping IP prefixes to PoPs we map IP prefixes to ASes, which are the direct neighbors responsible for routing traffic to and from the anycast site at an AS-level granularity, namely, the *AS-level catchment*. Moreover, we rely on the intuition that when anycast routing is deployed, the nearest PoP site to the end-user is going to attract the traffic. Therefore, when the vantage point AS (VP) routes to an anycast prefix, usually this route is going to be approximate to the origin AS in terms of geographic distance. To achieve this:

- We geolocate all vantage point ASes (VPs) using MaxMind, based on the IP addresses of the routers providing the Routing Information Base (RIB) entries.
- We map the country of each VP to its respective region through the UN dataset (*United Nations Statistics Division* 2024).
- We group all the direct neighboring ASes responsible for carrying traffic for a specific prefix according to the regions of the VPs.

We delve into the intricacies of the geolocation specific routing policies of anycast ASes in Section 5.3.

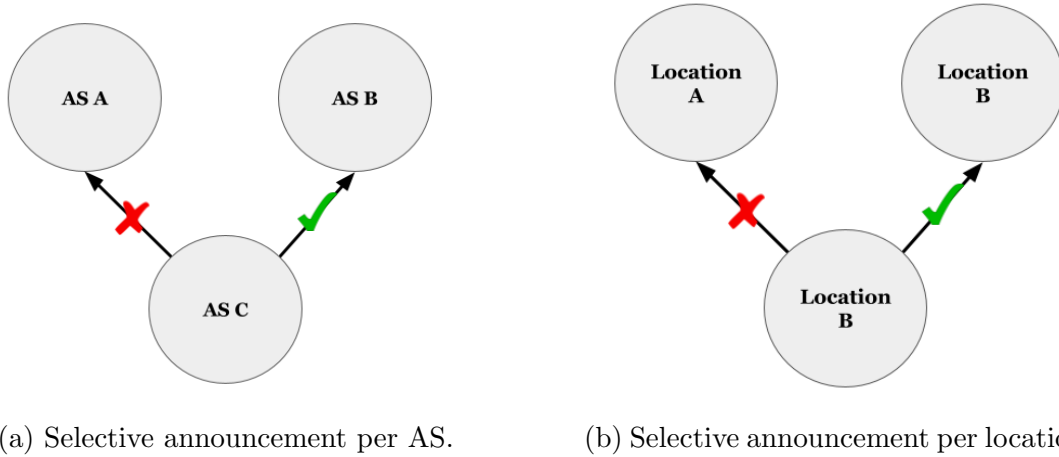


Figure 5.3: Selective Announcement Types.

5.2 Selective Announcements in Anycast IP Networks

Inter-domain routing does not follow the shortest path principle, but its based on the economic, performance or security needs of the organization. ASes independently define their routing policies Gao and Rexford, 2001; Huston, 1999 in order to select routes to a certain destination when multiple routes are available, and to decide to which neighbors to propagate the routes they know.

5.2.1 Selective Announcements per AS

In theory, routing policies need to follow the Gao-Rexford model and the valley-free rule in order to be *safe* to converge to a stable state under any link or node failure.

Nonetheless, network operators can arbitrarily configure their policies, without any coordination with their neighbors, therefore, a number of ASes might not follow the Gao-Rexford rules and that routing policies are more complex than what the state-of-the-art Gao and Rexford, 2001 can model Wang and Gao, 2003; Anwar et al., 2015; Mazloum et al., 2014; Giotsas and Zhou, 2012; Kastanakis, Giotsas, Livadariu, et al., 2023c; Gill, Schapira, and Goldberg, 2013. Moreover, an AS may

select to further restrict the propagation of certain routes to specific neighbors for traffic engineering purposes. By selectively advertising routes to different neighbors an AS may be able to control the links which will carry traffic for a specific route. For example, in Fig. 5.3a, *AS C* announces their prefixes to its neighbor *AS B* but not to *AS A*.

5.2.2 Selective Announcements per Location

Anycast ASes offer enhanced reliability and performance by directing users to the nearest PoP, reducing latency and improving overall user experience. However, a notable characteristic of some anycast ASes is the absence of a centralized backbone infrastructure to interconnect these dispersed PoPs Wang and Gao, 2003. For example, in Fig. 5.3b, the origin AS has a PoP in *Location B* and decides to announce prefixes from this PoP only to ASes that operate in *Location B* and not to distant ASes (e.g., ASes that operate in *Location A*), for latency optimization reasons.

This decentralized architecture can lead to the occurrence of selective announced prefixes Wang and Gao, 2003; Kastanakis, Giotsas, Livadariu, et al., 2023c. In the case of anycast ASes, this selective advertising can arise from the fact that each PoP connects separately to regional upstream providers, often resulting in varying routing policies and capabilities across different PoPs. Since there is no centralized backbone to manage and coordinate routing announcements, each PoP may independently determine which routes to advertise based on factors such as network capacity, peering agreements, and traffic optimization strategies.

In this work, we focus on the scenario of location dependent selective announcements. To the best of our knowledge, this is the first effort to study the phenomenon of selective announced anycast prefixes due to geolocation factors.

5.3 Regionality of Receivers of Selective Announcements

In this section, we aim to answer whether anycast ASes deploy selective announcements based on the geographic scope of their direct neighbors. To do so, we quantify the *regionality* levels of the direct neighbors of all anycast ASes per the regions of the vantage points.

5.3.1 Definition of Regionality

We label an AS as *regional* if: a) more than 90% of its prefixes are announced in a single region (for details on regions see *United Nations Statistics Division 2024*), and/or b) more than 90% of its peering links exist in the same region. For instance, *ASa* announces half of its prefixes in Italy and the other half in Spain. Both countries are in the Southern Europe region. Furthermore, *ASa* peers with three neighbors in: Portugal, Malta, and Greece, all located in Southern Europe. Thus, *ASa* qualifies as a regional AS.

By leveraging the *AS-level catchment* dataset (mentioned in Section 5.1) and the methodology of classifying ASes based on their geographic scope, we quantify how many direct neighbors of anycast ASes (i.e., ASes which receive announcements for anycast prefixes of the anycast AS) are *regional* ASes and how many are *global* ASes. We define as *regionality R* the portion of regional neighbors that prefer and use an anycast announcement against the total number of neighbors that prefer and use the announcement.

As demonstrated in Fig. 5.5, the Anycast AS connects to different sets of ASes in different PoPs and selectively announces its prefixes only to 4 out of its 6 direct neighbors, when the source AS locates in *Region A*. Additionally, large PoPs (e.g., PoP1) tend to be connected to both regional and global ASes, while regional PoPs (e.g., PoP2) are typically connected only to regional ASes. As a result, a VP in the region of a PoP with solely regional transit providers (i.e., PoP2) is more inclined to

5.3. Regionality of Receivers of Selective Announcements

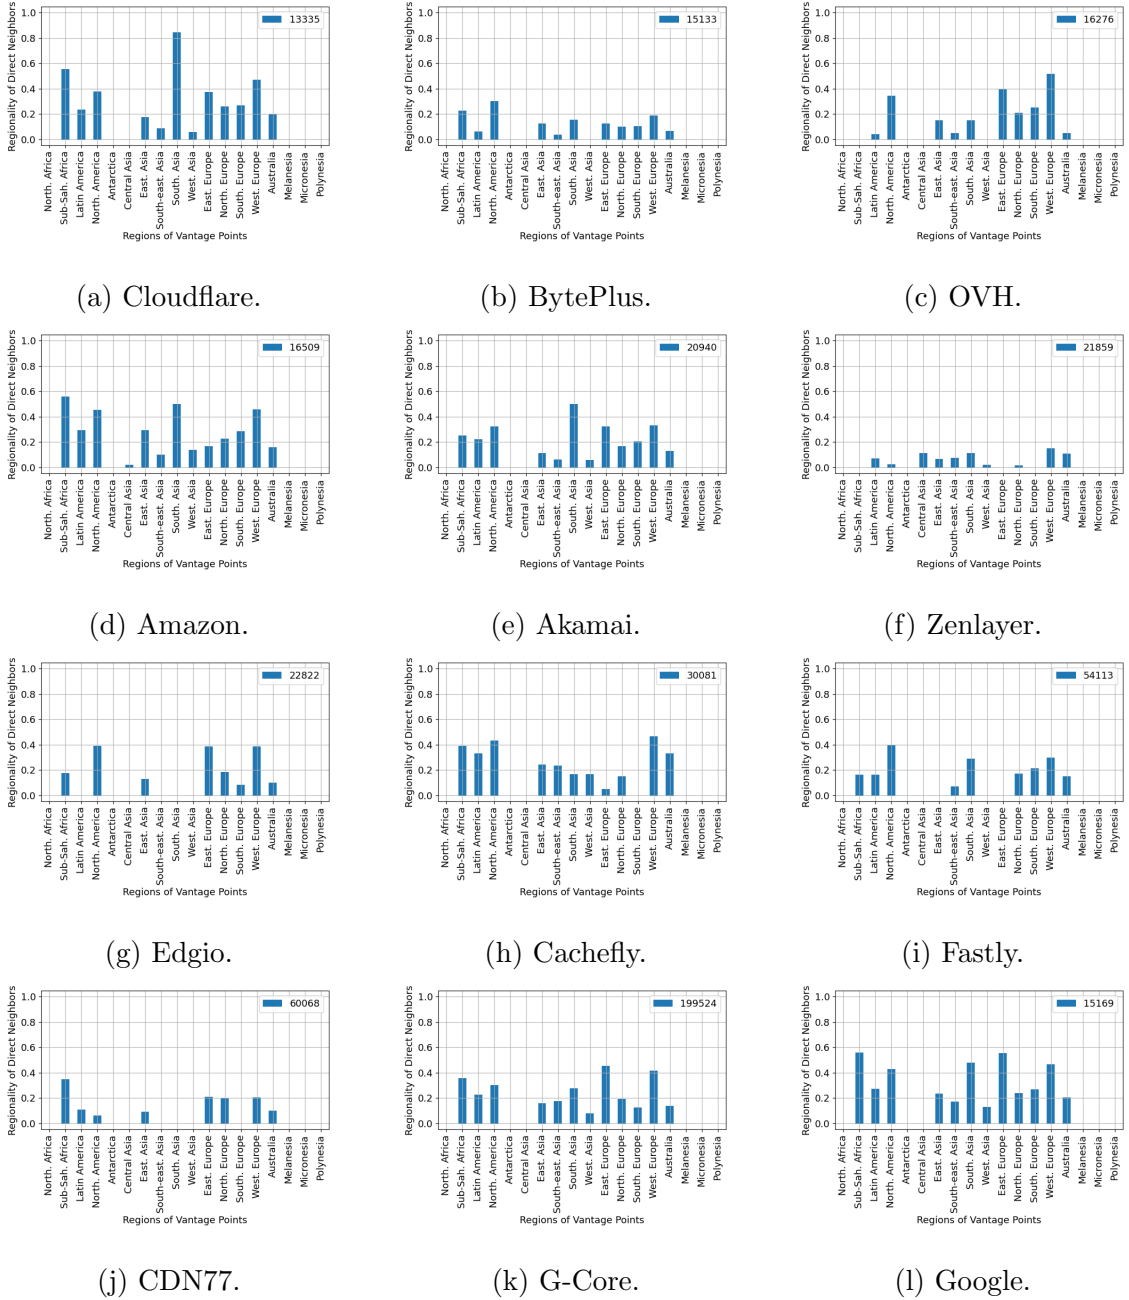


Figure 5.4: Regionality levels of the direct neighbors of the top Anycast ASes. In specific regions, big CDNs (e.g., Google, Cloudflare, G-Core, Amazon) rely on regional neighbors to carry their traffic to/from the rest of the Internet. This could be due to regulatory considerations, missing backbone or strategic business interests.

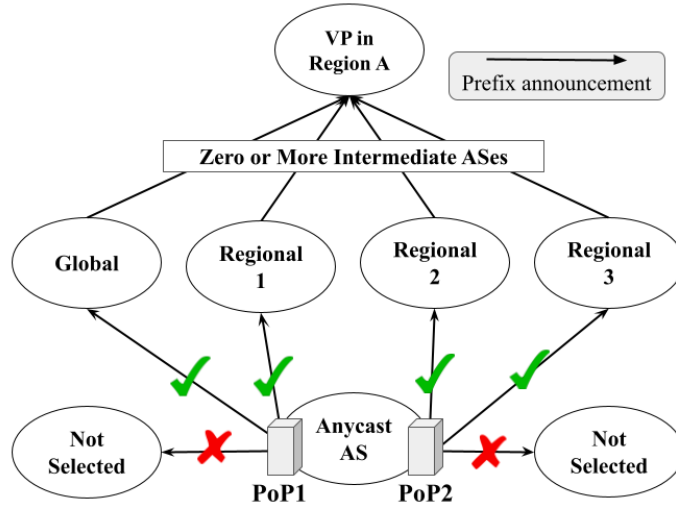


Figure 5.5: The Anycast AS connects to different ASes in different PoPs. Large PoPs (e.g., PoP1) tend to be connected to both regional and global ASes, while regional PoPs (e.g., PoP2) are typically connected only to regional ASes.

reach the anycast AS through regional ASes. This allows us to infer more accurately the providers that would transit the traffic from a VP in Region A to the PoP of the anycast AS in that region.

In the following step, we measure the extent at which anycast ASes rely on regional ASes to carry their traffic to/from the rest of the Internet.

5.3.2 Regionality Analysis

In Fig. 5.4, we observe that selecting a regional upstream provider or peer to carry the traffic to and from an anycast prefix is a common practice among anycast ASes. Cloudflare prefers regional ASes when the source locates at the Sub-Saharan Africa, the Southern Asia and the Western Europe in more than 45% of the times. Amazon, Google and Cachefly exhibit more than 40% of regionality levels when the source AS locates in the aforementioned regions.

In regions like Sub-Saharan Africa, Western Europe, South Asia and Northern America, we observe high levels of regionality by large CDN players like Google, Amazon, Akamai and Cloudflare. Factors such as: a) specific infrastructure

footprint, b) lower transit fees and c) strict regulatory conditions, drive the need for efficient traffic routing to ensure optimal user experience. Companies invest in local infrastructure and peering agreements to gain market penetration, comply with regulations, and enhance network resilience. By optimizing performance and ensuring redundancy, they can deliver faster content delivery and mitigate the impact of network disruptions, ultimately improving the quality of service for users in these regions.

On the other hand, ASes like CDN77, Edgecast and Zenlayer rely mainly on global ASes to carry their traffic to and from the rest of the Internet. A possible reasoning of this behavior is that these ASes prioritize global scalability and reach over regional optimization, especially if their services cater to a broad and geographically diverse user base.

As shown above, a part of the studied anycast ASes have high *regionality* levels while others prefer global providers. This highlights the need to further investigate the correlation between selective announcements and location-based routing policies. Our results can lay the grounds for understanding the confounding factors of: a) the anycast inefficiencies and b) the location-agnostic BGP best-path selection process.

5.4 Conclusion

Our investigation into anycast IP networks has shed light on the growing role of selective announcements and the significance of geolocation in modern routing practices. By analyzing how different. The findings illustrate that Anycast, with its inherent capability to route traffic based on proximity and strategic announcements, is not only a tool for enhancing performance but also a critical component for managing traffic across diverse and distributed infrastructures. Furthermore, the regional variations observed in the propagation and reception of selective announcements highlight the nuanced interplay between network topology and geographic factors, underscoring the necessity for location-aware models in

accurately predicting routing behaviors.

In the final chapter, we will reflect on the broader implications of our findings and discuss future research directions. While this thesis has taken important steps towards improving the accuracy of BGP simulations and understanding the complexities of interdomain routing, there are still many challenges that lie ahead. By summarizing the key contributions of this work, we aim to lay the groundwork for ongoing advancements in the field of Internet routing.

Chapter 6

Conclusion

This dissertation has systematically addressed the hypothesis that modern AS-path inference can be significantly improved by tackling topology incompleteness, integrating geolocation-aware paradigms, and refining interdomain routing policy models. Through the analysis of confounding factors, we identified key limitations such as missing links, selective announcements, and evolving AS behaviors, quantifying their impact on AS-path inference.

The replication study of interdomain routing policies over two decades demonstrated a clear shift towards more dynamic and complex behaviors, underscoring the necessity of adaptive models that align with current practices. The introduction of geolocation-aware models provided compelling evidence that incorporating spatial information into BGP simulations enhances accuracy, particularly in anycast routing scenarios, where regional trends and selective advertisements play a critical role.

The above findings validate the overarching hypothesis presented in Section 1 and highlight the practical value of the proposed methodologies for real-world applications.

6.1 Summary of Research Work

The opacity surrounding routing operations has profound implications in our ability to predict, understand, and debug the interdomain routing system effectively. This thesis investigates the longstanding challenges in interdomain routing inference models, with a specific focus on: 1) the confounding factors that affect the accuracy of AS path predictions, such as the complex routing policies and the geolocation-agnostic nature of contemporary routing models, 2) the phenomenon of selective route announcements, and 3) the impact of geographical factors on routing decisions.

In Chapter 3, apart from investigating the well known and studied problems on interdomain routing modeling, such as the path diversity and the first-hop inference problem, we identify and quantify an extra set of confounding factors. These factors, such as the routing policy violations between ASes or the geo-agnostic nature of contemporary models significantly impact the accuracy of AS-path predictions.

In Chapter 4, we conduct a longitudinal replicability study of the seminal Wang and Gao paper (Wang and Gao, 2003), to shed light on the evolution and the current state of selectively announced prefixes. We show that selective announcements are persistent, not only across time, but also across networks. Moreover, we observe that neighbors of different AS relationships may be assigned with the same local preference values, and path selection is not as heavily dependent on AS relationships as it used to be.

In Chapter 5, we identify and characterize selective and region-specific announcements, introducing a novel metric, "regionality", to delineate varying anycast strategies. Our findings indicate a substantial proportion of anycast ASes employing selective announcements in a per location basis, with the majority announcing all their prefixes selectively.

The above results highlight the need for BGP policy inference to be conducted as a high-periodicity process to account for the dynamic nature of AS connectivity and the derived policies. Moreover, in this study we demonstrate that geographic factors, such as the regional traffic patterns of anycast ASes or the location of

certain PoPs, significantly influence the AS path selection. Traditional BGP models, which ignore these factors, are unable to capture these intricacies and can lead to suboptimal routing inference accuracy. This study can serve as a guide for network operators and researchers, who aim to optimize their routing strategies, enhance their security mechanisms in the periphery of their network, or improve the accuracy of interdomain path predictions.

6.2 Discussion

Routing models are mainly used to address operational, performance or security research questions. Apart from the nature of these questions though, there are other factors which influence the overall setup and fidelity of the experiments.

The nature of the question One key reason is the need to predict the impact of routing changes on the network. For example, when an AS modifies its routing policies or adds new interconnections, the goal of the network operator, in this case, is to simulate and predict how these changes will affect the overall routing decisions. Another key factor is the optimization of the network performance. Operators in large CDNs can use simulations to ensure that traffic is efficiently distributed across their network, paying specific attention in minimizing latency, maximizing throughput, and improving the overall quality of service (QoS). Furthermore, routing models can be used in addressing specific 'what-if' questions in BGP routing security. Network operators and researchers can simulate routing scenarios to assess vulnerabilities, such as BGP hijacking or route leaks, and evaluate the resilience of (their) networks to such threats, or deploy and test new defense mechanisms (proactive or reactive). Lastly, from an economic perspective, BGP simulations can help in evaluating strategies such as whether to peer directly with another AS or purchase transit services from an upstream provider. These controlled experiments allow operators to weigh the costs and benefits of different interconnection strategies before making financial commitments.

Access to the source The privileges and the computational power that a researcher or network operator has access to also shape the simulation’s design and results. For example, a network operator may have direct access to internal routing data and policies within their own AS, providing valuable insights that improve the accuracy of path inferences. This is especially useful for simulating scenarios that are deeply influenced by internal factors, such as the AS’s traffic engineering policies or the specific configurations of routers and PoPs.

In addition to having access to internal policies, network operators in specific ASes also have the computational power and the necessary access rights to perform detailed operations, such as running traceroutes to specific PoPs and gain insights into latency, performance, and bottlenecks that an outsider would not easily observe. In contrast, a researcher without such direct access to the source might be restricted to leverage public data or inferred topologies, and rely solely on AS-level information. This can limit the scope of their experiments to more abstract and less detailed models, as they cannot directly probe the internal structures of specific networks.

The scope of the question The scope of a simulation strategy also varies depending on the question being asked. Some simulations might focus specifically on ingress traffic to a single AS, aiming to understand how traffic enters and traverses the network. Others may have a broader geographic or functional focus, such as analyzing the catchment area of a PoP or modeling the flow of traffic through specific Internet Exchange Points (IXPs). This scope can influence how detailed the simulation needs to be, and whether PoP-level or AS-level granularity is required.

The insights from this thesis Addressing these complexities requires innovative approaches. One effective method would be to foster closer collaboration between the academia and the industry, where the latter shares more insights into their operational aspects. This kind of data-sharing could significantly enhance the accuracy of BGP simulations by providing real-world grounding to theoretical models. Today, much of the research in this area relies on limited or inferred data sets and suffers from incomplete ground-truth data, which can lead to inaccuracies

when modeling the behavior of large-scale networks, especially those involving CDNs or Tier-1 ISPs. Another promising approach is to use hybrid modeling techniques that incorporate both AS-level and PoP-level granularity. In specific use-cases where PoP-level granularity becomes crucial, such as, at certain network hubs or PoPs that serve as key gateways for regional or international traffic, more precise modeling can capture how traffic is handled and routed, leading to more accurate predictions about performance and congestion.

6.3 Future Directions

Future work on routing modeling can put more emphasis in the inference of locpref allocations independently of AS relationships, given the scarcity of large-scale locpref data which are only available from a limited number of LGs. Our results support the need for more flexible routing models, that would allow routes from more “expensive” neighbors (peers/providers) to be selected as the best, and that will infer specific paths in a finer-granularity than the AS-level (Kastanakis, Giotsas, and Suri, 2022; Kastanakis, Giotsas, Livadariu, et al., 2023c).

Furthermore, as we pinpoint in this work, integrating geolocation data can significantly improve the accuracy of routing path predictions. However, more advanced models that can dynamically adjust based on real-time geographical and topological data could further enhance the prediction fidelity, particularly as the Internet structure evolves (Kastanakis, Giotsas, Livadariu, et al., 2024).

Finally, Artificial Intelligence (AI) tools and models can be leveraged to enhance the predictive power of interdomain routing inference. As we demonstrate in this work, routing behaviors become increasingly complex and the contemporary models struggle to capture the dynamic nature of AS-level path selections, particularly when influenced by evolving policies, traffic engineering, and performance optimization goals. Machine learning models, trained on historical BGP data and real-time network conditions, could provide more accurate predictions by recognizing patterns

and anticipating changes in routing decisions. Such models could adapt more quickly to policy shifts, outages, or congestion, offering network operators a valuable tool for preemptively addressing routing inefficiencies and potential security threats.

6.4 The Ph. in my Ph.D.

*“In every systematic inquiry (methodos) where there are **first principles**, or causes, or elements, knowledge and science result from acquiring knowledge of these; for we think we know something just in case we acquire knowledge of the primary causes, the primary first principles, all the way to the elements.*

–Aristotle”

This Ph.D. journey started on October 1st, 2020, in the School of Computing and Communications (SCC) at Lancaster University (LU). The progress of my Ph.D. at LU was monitored and evaluated through a series of review panels throughout the program’s duration. The first annual review panel involved presenting the literature and identifying the open research questions related to my Ph.D. topic. The title of this presentation was ‘*Impact Estimation of Traffic Misdirection Attacks*’. In the accompanying document, I outlined how BGP routing models could be used to estimate the impact of traffic misdirection attacks (i.e., prefix hijacks and route leaks) proactively, allowing for the optimal selection and deployment of defense mechanisms reactively when such attacks occur. Little did I know that after weeks of rigorous simulations I would realise that the contemporary BGP routing models were prone to errors (as discussed in this document) and that the following three years of my studies would have nothing to do with prefix hijacks or route leaks, rather than continuously trying to debug the best path selection process in interdomain routing models.

Throughout these four years, I experienced many stages of understanding how research unfolds. There were the “*let’s solve everything*” moments, typical of my first year, and the “*why did I get myself into this*” moments, which characterized

my second year. On the one hand there were moments of success — the travel opportunities in Nice and Montreal, my 2-months internship in SimulaMet in Oslo (Ioana Livadariu thank you for having me in your group), the moments where the results were publishable, the applaud in the presentations I delivered. On the other hand there were the frustrations, like spending three months on a controlled experiment only to achieve a zero-percent increase in accuracy or all the '*We regret to inform you that your paper has not been accepted..*' emails that I received.

Ultimately, this experience has shown me that the value of research lies not only in the tangible results but in the first principles of research: the perseverance, the adaptability, and the deepening of critical thinking that come from navigating the unknown. In this way, the last four years have shaped not just the outcome of my Ph.D. but my entire understanding of what it means to be a researcher.

References

- Ager, Bernhard et al. (2012). “Anatomy of a large European IXP”. In: *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pp. 163–174.
- Ahmed, Adnan et al. (2017). “Peering vs. transit: Performance comparison of peering and transit interconnections”. In: *2017 IEEE 25th International Conference on Network Protocols (ICNP)*. IEEE, pp. 1–10.
- Anwar, Ruwaifa et al. (2015). “Investigating interdomain routing policies in the wild”. In: *Proceedings of the 2015 Internet Measurement Conference*, pp. 71–77.
- Apostolaki, Maria, Aviv Zohar, and Laurent Vanbever (2017). “Hijacking bitcoin: Routing attacks on cryptocurrencies”. In: *2017 IEEE symposium on security and privacy (SP)*. IEEE, pp. 375–392.
- Augustin, Brice, Balachander Krishnamurthy, and Walter Willinger (2009). “IXPs: mapped?” In: *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*, pp. 336–349.
- Bailey, Joseph P (1997). “The economics of Internet interconnection agreements”. In.
- Ballani, Hitesh and Paul Francis (2005). “Towards a global IP anycast service”. In: *ACM SIGCOMM Computer Communication Review* 35.4, pp. 301–312.
- Böttger, Timm et al. (2018). “The Elusive Internet Flattening: 10 Years of IXP Growth”. In: *CoRR*.
- Bush, Randy and Rob Austein (2013). *The resource public key infrastructure (RPKI) to router protocol*. Tech. rep.

-
- Butler, Kevin et al. (2009). “A survey of BGP security issues and solutions”. In: *Proceedings of the IEEE* 98.1, pp. 100–122.
- CAIDA (2022). *The CAIDA UCSD AS to Organization Mapping Dataset*. https://www.caida.org/data/as_organizations. [Online; accessed 16-May-2022].
- (2023). *AS-Relationships Dataset*. <https://publicdata.caida.org/datasets/as-relationships/>. [Online; accessed 11-May-2023].
- Cardona, C, Pierre Francois, and Paolo Lucente (2016). *Impact of BGP filtering on Inter-Domain Routing Policies*. Tech. rep.
- Cartwright-Cox, Ben (2024a). *BGP.Tools* <https://bgp.tools/kb/anycatch>.
- (2024b). <https://github.com/bgptools/anycast-prefixes>.
- cdnperf.com* (2024). <https://www.cdnperf.com/>. [Online; accessed 20-March-2024].
- Chen, Kai et al. (2009). “Where the sidewalk ends: Extending the Internet AS graph using traceroutes from P2P users”. In: *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, pp. 217–228.
- Chiu, Yi-Ching et al. (2015). “Are We One Hop Away from a Better Internet?” In: *Proceedings of the 2015 Internet Measurement Conference*. IMC ’15. Tokyo, Japan: Association for Computing Machinery, pp. 523–529. ISBN: 9781450338486. DOI: 10.1145/2815675.2815719. URL: <https://doi.org/10.1145/2815675.2815719>.
- CISCO (2022). *BGP Best Path Selection Algorithm*. <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>. [Online; accessed 16-May-2022].
- Cisco (2021). *2021 Global Networking Trends Report*. [Online; accessed 10-June-2021]. URL: <https://www.cisco.com/c/en/us/solutions/enterprise-networks/2021-networking-report-preview.html>.
- (2023). *BGP best path selection algorithm*. <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>. [Online; accessed 11-May-2023].

- Cisco (2024). *2024 Global Networking Trends Report*. [Online; accessed 06-Sep-2024].
URL: https://www.cisco.com/c/dam/global/en_uk/solutions/enterprise-networks/2024-global-networking-trends.pdf.
- Cunha, Ítalo et al. (2016). “Sibyl: a practical Internet route oracle”. In: *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*, pp. 325–344.
- Dhamdhere, Amogh and Constantine Dovrolis (2008). “Ten years in the evolution of the Internet ecosystem”. In: *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pp. 183–196.
- (2010). “The internet is flat: modeling the transition from a transit hierarchy to a peering mesh”. In: *Proceedings of the 6th International Conference*, pp. 1–12.
- Dimitropoulos, Xenofontas et al. (2005). “Inferring as relationships: Dead end or lively beginning?” In: *International Workshop on Experimental and Efficient Algorithms*. Springer, pp. 113–125.
- Donnet, Benoit and Olivier Bonaventure (2008). “On BGP communities”. In: *ACM SIGCOMM Computer Communication Review* 38.2, pp. 55–59.
- Gao, Lixin (2001). “On inferring autonomous system relationships in the Internet”. In: *IEEE/ACM Transactions on networking* 9.6, pp. 733–745.
- Gao, Lixin and Jennifer Rexford (2001). “Stable Internet routing without global coordination”. In: *IEEE/ACM Transactions on networking* 9.6, pp. 681–692.
- Gigis, Petros et al. (2021). “Seven years in the life of Hypergiants’ off-nets”. In: *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pp. 516–533.
- Gill, Phillipa, Martin Arlitt, et al. (2008). “The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse?” In: *Passive and Active Network Measurement: 9th International Conference, PAM 2008, Cleveland, OH, USA, April 29-30, 2008. Proceedings 9*. Springer, pp. 1–10.
- Gill, Phillipa, Michael Schapira, and Sharon Goldberg (2012). “Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data”. In: *ACM SIGCOMM Computer Communication Review* 42.1, pp. 40–46.

- (2013). “A survey of interdomain routing policies”. In: *ACM SIGCOMM Computer Communication Review* 44.1, pp. 28–34.
- Giotsas, Vasileios, Matthew Luckie, et al. (2014). “Inferring complex AS relationships”. In: *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 23–30.
- Giotsas, Vasileios, George Nomikos, et al. (2020). “O peer, where art thou? Uncovering remote peering interconnections at IXPs”. In: *IEEE/ACM Transactions on Networking* 29.1, pp. 1–16.
- Giotsas, Vasileios and Shi Zhou (2012). “Valley-free violation in internet routing—analysis based on bgp community data”. In: *2012 IEEE International Conference on Communications (ICC)*. IEEE, pp. 1193–1197.
- Giotsas, Vasileios, Shi Zhou, et al. (2013). “Inferring multilateral peering”. In: *Proceedings of the ninth ACM conference on Emerging networking experiments and technologies*, pp. 247–258.
- Goldberg, Sharon (2014). “Why is it taking so long to secure internet routing?” In: *Communications of the ACM* 57.10, pp. 56–63.
- Huston, Geoff (1999). “Interconnection, peering, and settlements”. In: *proc. INET*. Vol. 9, p. 1.
- (2023). *IPv4 Address Report*. <https://www.potaroo.net/tools/ipv4/>.
- Huston, Geoff, Mattia Rossi, and Grenville Armitage (2010). “Securing BGP—A literature survey”. In: *IEEE Communications Surveys & Tutorials* 13.2, pp. 199–222.
- Jin, Yuchen et al. (2019). “Stable and Practical AS Relationship Inference with ProbLink.” In: *NSDI*. Vol. 19, pp. 581–598.
- Jin, Zitong et al. (2020). “Toposcope: Recover as relationships from fragmentary observations”. In: *Proceedings of the ACM Internet Measurement Conference*, pp. 266–280.

- Juniper (2023). *BGP best path selection algorithm*. <https://www.juniper.net/documentation/us/en/software/junos/vpn-12/bgp/topics/concept/routing-protocols-address-representation.html>. [Online; accessed 11-May-2023].
- Kastanakis, Savvas (2024). <https://github.com/kastanakis/inferring-location-based-routing-policies>.
- Kastanakis, Savvas, Vasileios Giotsas, Ioana Livadariu, et al. (2023a). *Online GitHub Repository*. <https://github.com/kastanakis/replicating-selective-announcements-inference>. [Online; accessed 24-September-2023].
- (2023b). *Online GitHub Repository*. <https://github.com/kastanakis/inferring-location-based-routing-policies>. [Online; accessed 24-September-2023].
- (2023c). “Replication: 20 Years of Inferring Interdomain Routing Policies”. In: *Proceedings of the 2023 ACM on Internet Measurement Conference*, pp. 16–29.
- (2024). “Investigating Location-aware Advertisements in Anycast IP Networks”. In: *Proceedings of the 2024 Applied Networking Research Workshop*, pp. 15–22.
- Kastanakis, Savvas, Vasileios Giotsas, and Neeraj Suri (2022). “Understanding the confounding factors of inter-domain routing modeling”. In: *Proceedings of the 22nd ACM Internet Measurement Conference*, pp. 758–759.
- Kastanakis, Savvas, Pavlos Sermpezis, Vasileios Kotronis, and Xenofontas Dimitropoulos (2018). “CABaRet: Leveraging recommendation systems for mobile edge caching”. In: *Proceedings of the 2018 Workshop on Mobile Edge Communications*, pp. 19–24.
- Kastanakis, Savvas, Pavlos Sermpezis, Vasileios Kotronis, Daniel Menasché, et al. (2020). “Network-aware recommendations in the wild: Methodology, realistic evaluations, experiments”. In: *IEEE Transactions on Mobile Computing* 21.7, pp. 2466–2479.
- Khan, Akmal et al. (2013). “As-level topology collection through looking glass servers”. In: *Proceedings of the 2013 conference on Internet measurement conference*, pp. 235–242.

- Klöti, Rowan et al. (2016). “A comparative look into public IXP datasets”. In: *ACM SIGCOMM Computer Communication Review* 46.1, pp. 21–29.
- Kotronis, Vasileios et al. (2016). “Stitching inter-domain paths over IXPs”. In: *Proceedings of the Symposium on SDN Research*, pp. 1–12.
- Labovitz, Craig et al. (2010). “Internet inter-domain traffic”. In: *ACM SIGCOMM Computer Communication Review* 40.4, pp. 75–86.
- Lepinski, Matthew and Kotikalapudi Sriram (2017). *Rfc 8205: Bgpsec protocol specification*.
- Li, Zhihao et al. (2018). “Internet anycast: performance, problems, & potential”. In: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pp. 59–73.
- Lodhi, Aemen, Amogh Dhamdhere, and Constantine Dovrolis (2014). “Open peering by Internet transit providers: Peer preference or peer pressure?” In: *IEEE INFOCOM 2014-IEEE Conference on Computer Communications*. IEEE, pp. 2562–2570.
- Luckie, Matthew et al. (2013). “AS relationships, customer cones, and validation”. In: *Proceedings of the 2013 conference on Internet measurement conference*, pp. 243–256.
- Lychev, Robert, Sharon Goldberg, and Michael Schapira (2013). “BGP security in partial deployment: Is the juice worth the squeeze?” In: *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, pp. 171–182.
- Madhyastha, Harsha V et al. (2009). “iPlane Nano: Path Prediction for Peer-to-Peer Applications.” In: *NSDI*. Vol. 9, pp. 137–152.
- Mahajan, Ratul, David Wetherall, and Tom Anderson (2002). “Understanding BGP misconfiguration”. In: *ACM SIGCOMM Computer Communication Review* 32.4, pp. 3–16.
- Mao, Zhuoqing Morley, Lili Qiu, et al. (2005). “On AS-level path inference”. In: *Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pp. 339–349.

- Mao, Zhuoqing Morley, Jennifer Rexford, et al. (2003). “Towards an accurate AS-level traceroute tool”. In: *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 365–378.
- Marcos, Pedro et al. (2018). “Dynam-IX: A dynamic interconnection exchange”. In: *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*, pp. 228–240.
- MaxMind API (2024). <https://www.maxmind.com/>. [Online; accessed 20-March-2024].
- Mazloun, Riad et al. (2014). “Violation of interdomain routing assumptions”. In: *Passive and Active Measurement: 15th International Conference, PAM 2014, Los Angeles, CA, USA, March 10-11, 2014, Proceedings 15*. Springer, pp. 173–182.
- Milolidakis, Alexandros (2022). “Understanding the Capabilities of Route Collectors to Observe Stealthy Hijacks: Does adding more monitors or reporting more paths help?” PhD thesis. KTH Royal Institute of Technology.
- Moura, Giovane CM et al. (2016). “Anycast vs. DDoS: Evaluating the November 2015 root DNS event”. In: *Proceedings of the 2016 Internet Measurement Conference*, pp. 255–270.
- Mühlbauer, Wolfgang, Anja Feldmann, et al. (2006). “Building an AS-topology model that captures route diversity”. In: *ACM SIGCOMM Computer Communication Review* 36.4, pp. 195–206.
- Mühlbauer, Wolfgang, Steve Uhlig, et al. (2007). “In search for an appropriate granularity to model routing policies”. In: *ACM SIGCOMM Computer Communication Review* 37.4, pp. 145–156.
- Murphy, Sandra (2006). *BGP security vulnerabilities analysis*. Tech. rep.
- NCC, RIPE (2022). *Routing Information Service (RIS)*. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>. [Online; accessed 16-May-2022].

- (2023). *Routing Information Service (RIS)*. <https://www.ripe.net/>. [Online; accessed 12-May-2023].
- Network, Merit (2021). *Merit RADb*. <https://www.radb.net/>.
- Norton, William B (2001). “Internet service providers and peering”. In: *Proceedings of NANOG*. Vol. 19, pp. 1–17.
- Oliveira, Ricardo et al. (2009). “The (in) completeness of the observed Internet AS-level structure”. In: *IEEE/ACM Transactions on Networking* 18.1, pp. 109–122.
- Oliveira, Ricardo V et al. (2008). “In search of the elusive ground truth: the Internet’s AS-level connectivity structure”. In: *ACM SIGMETRICS Performance Evaluation Review* 36.1, pp. 217–228.
- Oregon, University of (2022). *Route Views Project*. <http://www.routeviews.org/routeviews/>. [Online; accessed 16-May-2022].
- (2023). *Route Views Project*. <http://www.routeviews.org/>. [Online; accessed 12-May-2023].
- Orsini, Chiara et al. (2016). “BGPStream: a software framework for live and historical BGP data analysis”. In: *Proceedings of the 2016 Internet Measurement Conference*, pp. 429–444.
- PeeringDB* (2023). <https://www.peeringdb.com>.
- Public Route Server* (2023). <https://www.routeservers.org/>.
- Qiu, Jian and Lixin Gao (2006). “Cam04-4: As path inference by exploiting known as paths”. In: *IEEE Globecom 2006*. IEEE, pp. 1–5.
- Quoitin, Bruno and Steve Uhlig (2005). “Modeling the routing of an autonomous system with C-BGP”. In: *IEEE network* 19.6, pp. 12–19.
- Rekhter, Yakov, Tony Li, and Susan Hares (Jan. 2006). *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. URL: <https://www.rfc-editor.org/rfc/rfc4271.txt>.
- RIPESat API* (2024). <https://stat.ripe.net/app/launchpad>. [Online; accessed 20-March-2024].
- Roughan, Matthew, Simon Jonathan Tuke, and Olaf Maennel (2008). “Bigfoot, sasquatch, the yeti and other missing links: what we don’t know about the

- as graph”. In: *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pp. 325–330.
- Roughan, Matthew, Walter Willinger, et al. (2011). “10 lessons from 10 years of measuring and modeling the internet’s autonomous systems”. In: *IEEE Journal on Selected Areas in Communications* 29.9, pp. 1810–1821.
- routesservers.org (2023). *Public Route Servers*. <https://www.routesservers.org/>. [Online; accessed 11-May-2023].
- Schlinker, Brandon et al. (2017). “Engineering egress with edge fabric: Steering oceans of content to the world”. In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pp. 418–431.
- Sermpezis, Pavlos, Savvas Kastanakis, et al. (2019). “Towards qos-aware recommendations”. In: *arXiv preprint arXiv:1907.06392*.
- Sermpezis, Pavlos and Vasileios Kotronis (2019). “Inferring catchment in internet routing”. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 3.2, pp. 1–31.
- Sermpezis, Pavlos, Vasileios Kotronis, Konstantinos Arakadakis, et al. (2021). “Estimating the Impact of BGP Prefix Hijacking”. In: *2021 IFIP Networking Conference (IFIP Networking)*. IEEE, pp. 1–10.
- Sermpezis, Pavlos, Vasileios Kotronis, Petros Gigis, et al. (2018). “ARTEMIS: Neutralizing BGP hijacking within a minute”. In: *IEEE/ACM Transactions on Networking* 26.6, pp. 2471–2486.
- Shao, Xiaozhe and Lixin Gao (2021). “Policy-rich interdomain routing with local coordination”. In: *Computer Networks* 197, p. 108292.
- Shavitt, Yuval and Eran Shir (2005). “DIMES: Let the Internet measure itself”. In: *ACM SIGCOMM Computer Communication Review* 35.5, pp. 71–74.
- Singh, Rachee et al. (2021). “PredictRoute: a network path prediction toolkit”. In: *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5.2, pp. 1–24.

- Society, The Internet (2019). *Consolidation in the Internet Economy*. <https://www.internetsociety.org/wp-content/uploads/2022/12/2019-Internet-Society-Global-Internet-Report-Consolidation-in-the-Internet-Economy.pdf>.
- Subramanian, Lakshminarayanan et al. (2002). “Characterizing the Internet hierarchy from multiple vantage points”. In: *Proceedings. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*. Vol. 2. IEEE, pp. 618–627.
- Tian, Zhihong et al. (2019). “A data-driven method for future Internet route decision modeling”. In: *Future Generation Computer Systems* 95, pp. 212–220.
- United Nations Statistics Division (2024). <https://unstats.un.org/unsd/methodology/m49/>.
- Wang, Feng and Lixin Gao (2003). “On inferring and characterizing Internet routing policies”. In: *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, pp. 15–26.
- Wikipedia (2022). *Monte Carlo method*. https://en.wikipedia.org/wiki/Monte_Carlo_method. [Online; accessed 16-May-2022].
- Winter, Philipp et al. (2019). “Geo-locating BGP prefixes”. In: *2019 Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, pp. 9–16.
- Wohlfart, Florian et al. (2018). “Leveraging interconnections for performance: the serving infrastructure of a large CDN”. In: *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pp. 206–220.
- Wu, Tianhao et al. (2022). “RouteInfer: Inferring Interdomain Paths by Capturing ISP Routing Behavior Diversity and Generality”. In: *Passive and Active Measurement: 23rd International Conference, PAM 2022, Virtual Event, March 28–30, 2022, Proceedings*. Springer, pp. 216–244.
- Yap, Kok-Kiong et al. (2017). “Taking the edge off with espresso: Scale, reliability and programmability for global internet peering”. In: *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, pp. 432–445.

References

Zhang, Zheng et al. (2007). “Practical defenses against BGP prefix hijacking”. In:
Proceedings of the 2007 ACM CoNEXT conference, pp. 1–12.