

20 Years of Inferring Inter-domain Routing Policies

Savvas Kastanakis
Lancaster University
s.kastanakis@lancaster.ac.uk

Vasileios Giotsas
Cloudflare
vasilis@cloudflare.com

Ioana Livadariu
Simula Metropolitan Center
ioana@simula.no

Neeraj Suri
Lancaster University
neeraj.suri@lancaster.ac.uk

ABSTRACT

In 2003, Wang and Gao [63] 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 paper we conduct a replicability study of the Wang and Gao paper [63], 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. Our 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.

CCS CONCEPTS

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

KEYWORDS

Internet Routing Policies; Selective BGP Announcements

ACM Reference Format:

Savvas Kastanakis, Vasileios Giotsas, Ioana Livadariu, and Neeraj Suri. 2023. 20 Years of Inferring Inter-domain Routing Policies. In *Proceedings of the 2023 ACM Internet Measurement Conference (IMC '23)*, October 24–26, 2023, Montreal, QC, Canada. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3618257.3624799>

1 INTRODUCTION

The Internet is a collation of thousands of networks (Autonomous Systems (ASes)), each of which belongs to an organization, be it an Internet Service Provider (ISP), a university, or a company. To learn how to reach remote network addresses (IP prefixes), ASes exchange routing messages with each other through the Border Gateway Protocol (BGP), which is the de-facto protocol for routing in the AS graph (inter-domain routing). BGP messages (announcements) include information on which routes should be followed for an AS to reach an IP prefix. Such routes are sequences of AS hops, generally referred to as AS paths.

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 [20, 28] 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 [14].

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 routing policies of the Internet ASes could significantly improve network operations.

In 2003, Wang and Gao measured inter-domain routing policies in the wild to inform the modeling of such policies [63]. They showed that routing policies are more complex than what the state-of-the-art [20] could model. Since 2003, the Internet topology has experienced fundamental changes in interconnection practices, such as the flattening of the Internet hierarchy and the dominance of Content-Distribution Networks (see Section 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 [39, 41] and there has been noticed a significant performance difference between peering and transit interconnections [4], 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 [12, 14, 19, 20, 23, 33, 43, 51, 55, 59, 61, 65], 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 [5, 43]. To update our understanding of routing policies and inform the current efforts

to model inter-domain routing, we conduct an in-depth replication study of the Wang and Gao paper [63].

Specifically, our contributions are as follows:

- We answer to what extent the findings of [63] are still valid. The usage of selective announcements has increased up to 30%, but with significant variability across ASes. The assignment of Local Preference exhibits higher variability than 20 years ago.
- We conduct a longitudinal study on the evolution of selective announcements. We observe a median increase of more than 20% after 2007, but with pronounced yearly fluctuations.
- We discuss the potential root causes of the differences between the routing policies today compared to 2003, and the implications for inferring routing paths.
- We publish the artifacts (source code and data) of our study to facilitate future research on inter-domain routing modeling and further replicability.

As a final remark, our goal is not to propose a new interdomain routing model, or infer more accurately the routing policies in the Internet, but to pinpoint the intersections and disparities of the results our replication effort against the findings and insights of the original work [63]. Finally, we believe that our results can aid in the understanding of a variety of interdomain routing applications, such as the measurement of the RPKI adoption [17, 49], fine-grained interdomain policy learning [62, 67], interdomain routing verification [10], privacy-preserving routing [13], discovering caching policies in the wild [18, 34] and studying routing attacks [56].

1.1 Ethical considerations

This study does not raise any ethical issues. The datasets we use in this study are publicly available. The data collection from route server looking glasses uses a low rate of queries to keep measurement traffic low.

2 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 AS Business Relationships

BGP is a policy-based protocol (rather than a shortest-path 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 [20, 28].

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 [19]: (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% [30, 31, 40].

2.2 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 best (most preferable) route. The BGP route selection process is comprised of the following steps [11, 32]. 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. *locpref* values are arbitrary.
- (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 [24].

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 [20]. For this reason this *locpref* allocation pattern is also referred to as the Gao-Rexford model.

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%

Table 1: Peer Link Statistics of the Internet, 1998–2023, as observed in the CAIDA AS Relationships Graph.

2.3 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: (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.

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 [22]. Table 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 of Internet Service Providers (ISPs). A small clique of international ISPs (Tier-1) were

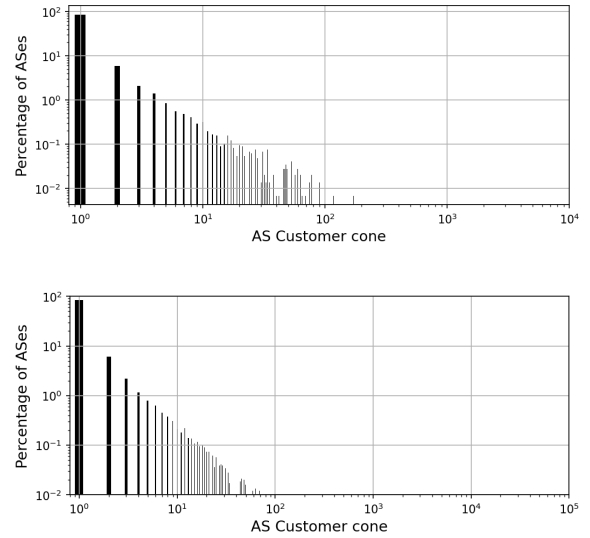


Figure 1: Customer Cone (CC) sizes in 2003 (top), and 2023 (bottom) exhibit similar power-law distributions.

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 1) due to the rise of Content Providers (CPs) and Content Delivery Networks (CDNs) [5, 15, 22, 36, 38]. 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 [21, 26, 54, 64, 66]. In this flattening topology era, Internet Exchange Points (IXPs) emerged and played a key role in enabling large CDNs to bypass Tier-1 ISPs [6, 36, 37]. 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 [7].

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 [8] to plot the CC distribution in 2003 and 2023 in Figure 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 [29]. 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 [60].

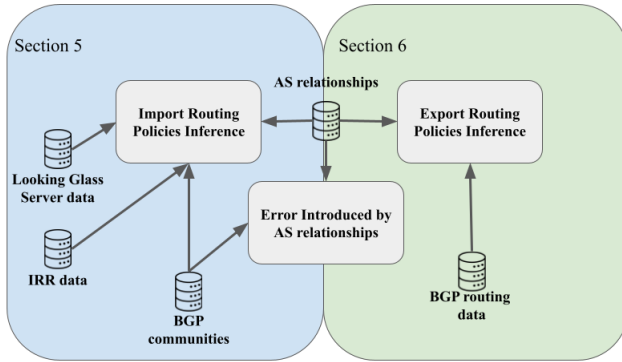


Figure 2: Overview of the data used for the import and export routing policies analysis.

4 REPLICATION OVERVIEW

In this section, we first describe the research methodology of Wang and Gao [63] and our strategy to replicate their work in this paper, and conclude with an inherent limitation of such studies, the AS graph incompleteness problem.

4.1 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. 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.

4.2 Replication Strategy

In Fig.2 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 [63], our study of routing policies relies on AS relationship inferences. We use the current state-of-the-art AS relationships, made available by CAIDA [8]. Moreover, we leverage data from Looking Glass (LG) servers [53] to study import policies in Section 5. LG servers are interfaces to network devices that can be queried through web-based, telnet or ssh interfaces and allow users to query 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 [63] 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 [46] and RIPE RIS projects [44] to infer the configuration of export policies and analyze the prevalence, the persistence and the causes of selective advertisements in Section 6.

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.

4.3 The incompleteness of the AS graph

Our study relies on BGP data, LG 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* [47]. 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 [35, 47, 52].

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. 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 6.2.

5 IMPORT POLICIES

5.1 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 [63] Wang and Gao queried the then-available LGs that provide a direct telnet interface

Table 2: 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

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 [63]. Unfortunately, the original route server LGs used in [63] 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 [1], 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 [2].

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 [53] to enable the repeatability of our experiments.

Table 3 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.

On average, only 59% of the p2p routes have a *locpref* value between the c2p and p2c values, nonetheless, this 59% is a mix

Table 3: 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%

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 [7]. 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 [25].

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 [45], 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 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.

5.2 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 rule should match routes that pass a specific prefix-filter. If, instead, we identify a next-hop specific rule, e.g., *match ip next-hop 203.0.113.1*, then the *locpref* for this route is set based on the next-hop AS.

Table 4: 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%

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

In the last column of Table 3, 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) [24].

5.3 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 [16].

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

Table 5: 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)

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 2, and we keep the communities that are used to annotate relationship types. The documentation of the corresponding BGP 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 [63].

Table 5 summarizes our validation statistics. For most of the tested ASes 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 [2], 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 [63].

6 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 3 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.

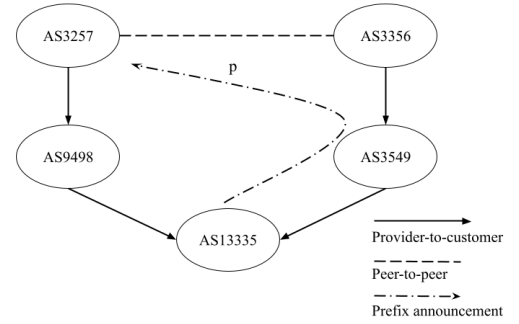


Figure 3: The export policies of AS13335, can be observed by its indirect provider AS3257. AS13335 announces prefix p to provider AS3549, but not to AS9498.

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

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

6.1 Export to Provider

Here, we first describe the algorithm used in [63] 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 [8] and b) the routing tables of 21 ASes (listed in Table 2) via the BGPStream API [50] for different time periods.

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

Table 6: % of SA prefixes and SA origins observed by 21 ASes, listed in decreasing order of SA prefixes.

6.1.1 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_{op} (Phase 1). To that end, we collect all direct/indirect customers of AS_{op} 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_{op} and extract all routes originated by ASes that belong in the *CC* of AS_{op} (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*.

6.1.2 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

¹Similar duration as in the original Wang and Gao study [63].

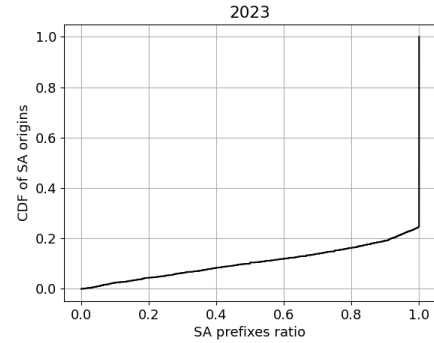


Figure 4: CDF of all SA origins in 2023.

in Table 2, using all available route collectors from the projects Routeviews and RIPE RIS (see Section 4).

Note here, that *SA prefixes* for a provider may be due to the selective announcement policies of the origin or intermediate ASes. For instance, in Figure 3, the *SA prefix* for AS3257, may be due to the selective announcement policies employed by AS9498 as well. We study this possibility in Subsection 6.1.5.

Table 6 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*. For each one of the 21 ASes in our study, we store all observed prefixes per origin AS and all *SA prefixes* per *SA origin* in a key-value format. Hence, we have 21 data structures that describe both the observed prefixes per origin AS, and the characterized *SA prefixes* per *SA origin*. In Figure 4 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 AS_{op} 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 6.1.4 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.

6.1.3 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 5.3, we verify the AS relationships

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 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)

Table 7: Percentage of SA prefixes verified per AS.

between 11 ASes in Table 5 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_{op} , we search all the BGP routing tables to find if there is an AS path between AS_{op} 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 7 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 local preference. The average conformance in *locpref* settings is 83% as shown in Table 3, which explains the average 19.2% of unverified SA prefixes per SA origin in our study.

6.1.4 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 [50] 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 [63]. 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 5a the number of SA prefixes during the 15th of January 2023, for AS7018, while Figure 5b 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 5c). 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 collected data comprises of routing tables from the 1st of April of each year between 2003 and 2023. Note that SA origins findings are omitted due to similar insights with the SA prefix findings.

To find out how export policies affect the SA prefixes, we follow the approach of 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 6a 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 6b. 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 7 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 8, 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 6 with the customers point of view in Figure 8, 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.**

6.1.5 Causes of SA prefixes.

As mentioned in Subsection 2.1, customers may connect to multiple providers (*multi-homed*) 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

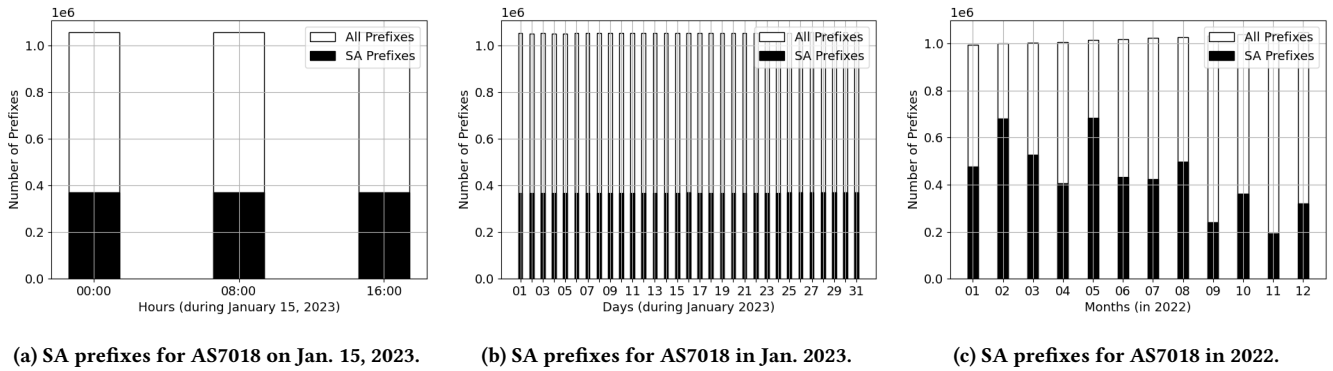


Figure 5: Daily and monthly persistence of SA prefixes for AS7018.

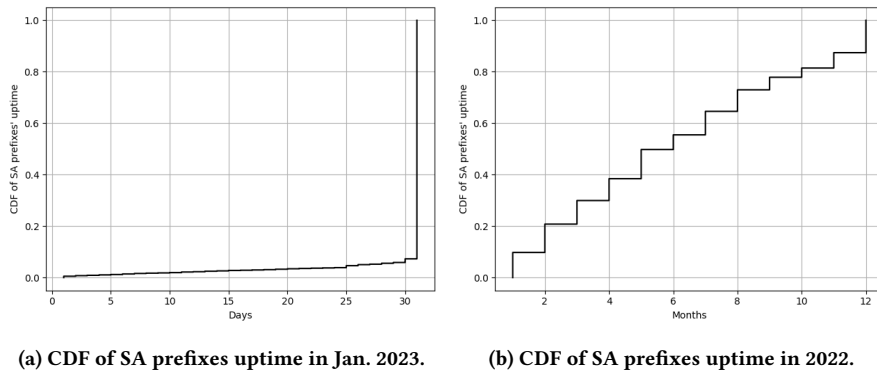


Figure 6: Daily and monthly uptimes of SA prefixes for AS7018.

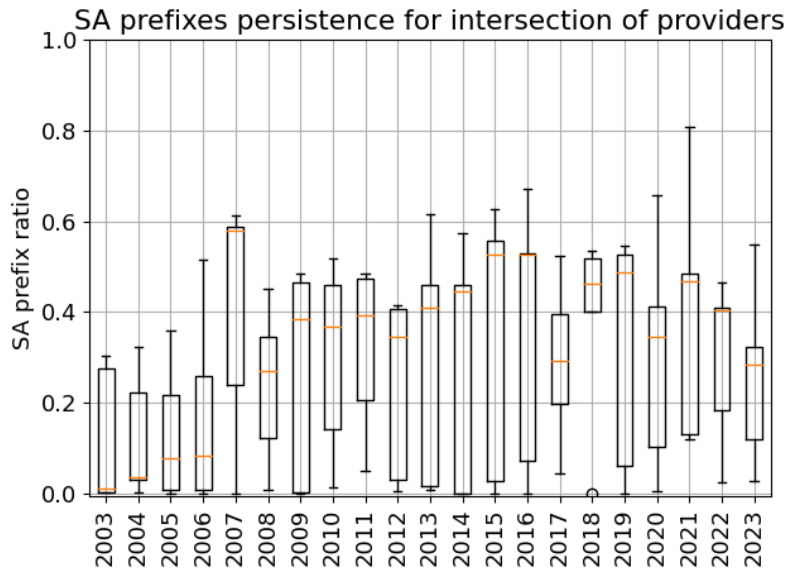


Figure 7: SA prefix ratio over the last 20 years for AS3257, AS3292, AS3549, AS5511 and AS7018.

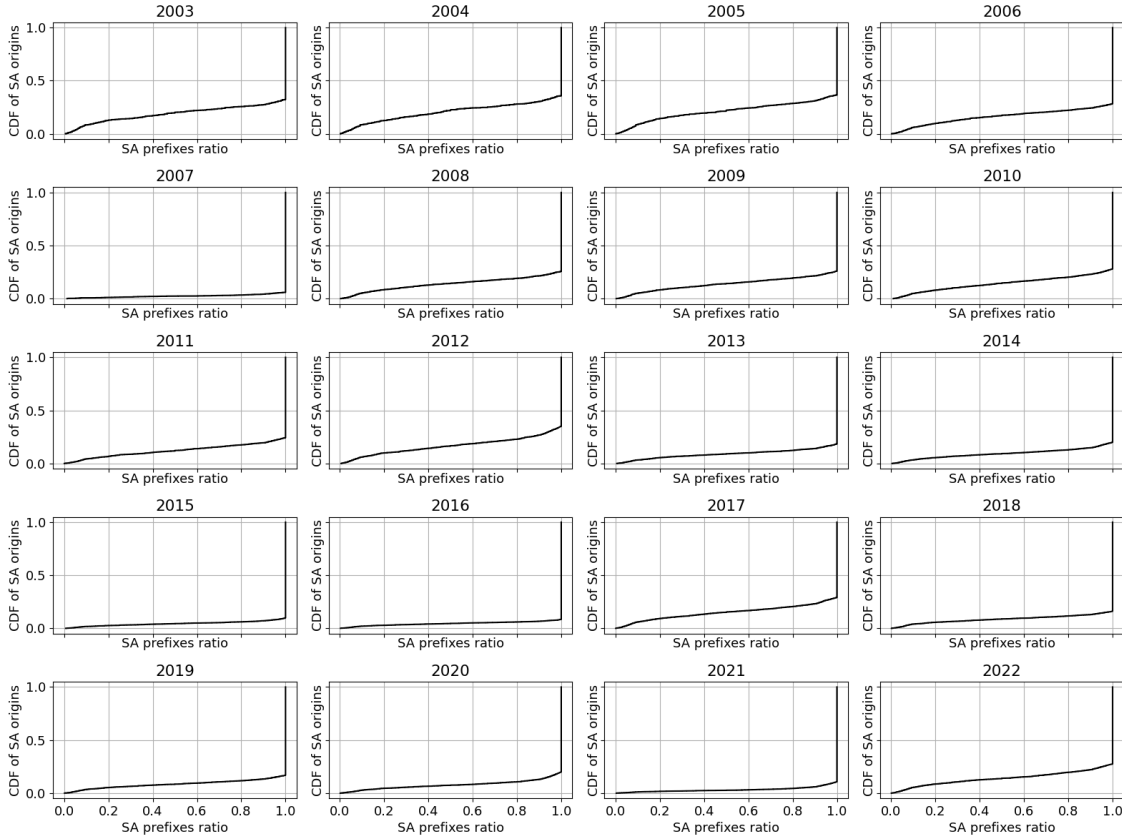


Figure 8: CDF of union of SA origins over the last 20 years for AS3257, AS3292, AS3549, AS5511 and AS7018.

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.

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)

Table 8: % (#) of multi-homed and single-homed SA origins for AS3257, AS3292, AS3549, AS5511 and AS7018

For AS3257, AS3292, AS3549, AS5511 and AS7018, we examine how many SA origins are *multi-homed*. From Table 8, 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 ~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

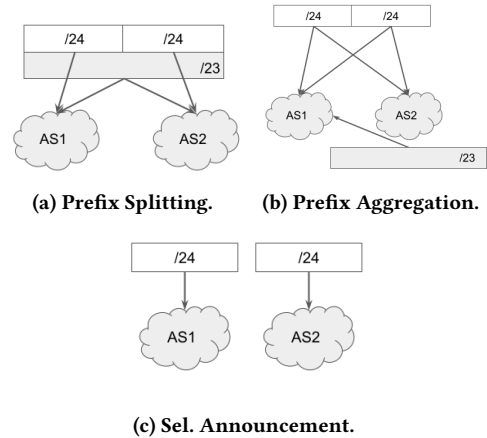


Figure 9: Causes of a SA prefix

an AS₀ originates a /23 prefix p_0 , to which it wants to load-balance traffic between two of its providers, AS₁ and AS₂. At the same time AS₀ wants to ensure that if a link to one of its providers fails, traffic flows to p_0 will not be disrupted. To that end, AS₀ splits the /23 prefix to two more specific /24 prefixes, and advertise the /23 prefix to both AS₁ and AS₂ providers, and each /24 to a different provider.

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

Table 9: Causes of SA prefixes

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 9a 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 9b.

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 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 9c.

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_{op} can be aggregated in the BGP tables of the remaining ASes. Table 9 shows that both ratios are negligible, therefore, the main cause of SA prefixes cannot be prefix splitting or prefix aggregation.

6.2 Export to Peer

In this section, we use the algorithm described in Section 6.1.1 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 2 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_{op} , 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_{op} , we test whether AS_{op} 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_{op} .

In Table 10, 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

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

Table 10: % of peer SA prefixes and % peer SA origins.

advertisement practices between peers, because routes exchanged over peering links have limited visibility due to the valley-free rule [47].

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 [47], 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 [9, 58] 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. [3]. 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. 10 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 [14, 48], 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.

7 RELATED WORK

Since the original study was published in 2003, we are not aware of other papers that reproduce or replicate the full methodology. However, inter-domain routing has been extensively studied over the past two decades, and while our knowledge and understanding of the routing ecosystem has been enriched significantly, there are still gaps remaining to be filled.

In theory, routing policies need to follow the GR model [20] in order to be *safe* to converge to a stable state under any link or

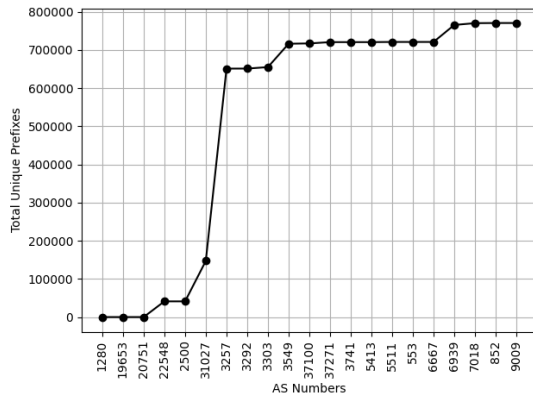


Figure 10: Progressive enumeration of unique SA prefixes by incrementally adding vantage points.

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 model. This has been indeed observed both by Internet measurement studies [5, 27, 42] and reported by network operators in a 2013 survey [24]. In this survey, 32% of the participating networks do not follow the GR model completely. In 16 out of 97 networks, *peer* or *provider* routes are preferred over *customer* routes and in 21 out of 97 networks, *peer* and *provider* routes are announced to peers and providers. In the study published by Anwar et al. [5], 34% of routing decisions in the Internet routing system cannot be explained by the Gao-Rexford model.

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 [39]. As a result, ASes may prefer *peer* over *customer* routes for performance reasons [4]. 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 [57] 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.

8 CONCLUSION

In this paper, we replicate the study of Wang and Gao [63] 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. Future work 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 route server LGs.

Our results support the need for a more flexible routing model, that would allow routes from more “*expensive*” neighbors (i.e., peers/providers) to be selected as the best, rather than follow a strict *customer-over-all* rule. Additionally, our work 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. Concluding, our findings highlight the need for BGP policy inference to be conducted with high-periodicity in order to account for the dynamic nature of AS connectivity and the derived routing policies.

ACKNOWLEDGMENTS

This research was supported, in part, by the UKRI Trustworthy Autonomous Systems Node in Security [EPSRC grant EP/V026763/1] and the Norwegian Research Council [grant #288744].

REFERENCES

- [1] 2023. PeeringDB. <https://www.peeringdb.com>. (2023).
- [2] 2023. Public Route Server. <https://www.routeservers.org/>. (2023).
- [3] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. 2012. Anatomy of a large European IXP. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. 163–174.
- [4] Adnan Ahmed, Zubair Shafiq, Harkeerat Bedi, and Amir Khakpour. 2017. Peering vs. transit: Performance comparison of peering and transit interconnections. In *2017 IEEE 25th International Conference on Network Protocols (ICNP)*. IEEE, 1–10.
- [5] Ruwaifa Anwar, Haseeb Niaz, David Choffnes, Ítalo Cunha, Phillipa Gill, and Ethan Katz-Bassett. 2015. Investigating interdomain routing policies in the wild. In *Proceedings of the 2015 Internet Measurement Conference*. 71–77.
- [6] Brice Augustin, Balachander Krishnamurthy, and Walter Willinger. 2009. IXPs: mapped?. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*. 336–349.
- [7] Timm Böttger, Gianni Antichi, Eder Leão Fernandes, Roberto di Lallo, Marc Bruyere, Steve Uhlig, and Ignacio Castro. 2018. The Elusive Internet Flattening: 10 Years of IXP Growth. *CoRR* (2018).
- [8] CAIDA. 2023. AS-Relationships Dataset. <https://publicdata.caida.org/datasets/as-relationships/>. (2023). [Online; accessed 11-May-2023].
- [9] Kai Chen, David R Choffnes, Rahul Potharaju, Yan Chen, Fabian E Bustamante, Dan Pei, and Yao Zhao. 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*. 217–228.
- [10] Marco Chiesa, Daniel Demmler, Marco Canini, Michael Schapira, and Thomas Schneider. 2017. SIXPACK: Securing internet exchange points against curious onlookers. In *Proceedings of the 13th International Conference on emerging Networking EXperiments and Technologies*. 120–133.
- [11] Cisco. 2023. BGP best path selection algorithm. <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>. (2023). [Online; accessed 11-May-2023].
- [12] Ítalo Cunha, Pietro Marchetta, Matt Calder, Yi-Ching Chiu, Bruno VA Machado, Antonio Pescapè, Vasileios Giotsas, Harsha V Madhyastha, and Ethan Katz-Bassett. 2016. Sibyl: a practical Internet route oracle. In *13th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 16)*. 325–344.
- [13] Daniel Demmler. 2022. Towards practical privacy-preserving protocols. *it-Information Technology* 64, 1-2 (2022), 49–53.
- [14] Amogh Dhamdhere and Constantine Dovrolis. 2008. Ten years in the evolution of the Internet ecosystem. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*. 183–196.
- [15] Amogh Dhamdhere and Constantine Dovrolis. 2010. The internet is flat: modeling the transition from a transit hierarchy to a peering mesh. In *Proceedings of the 6th International Conference*. 1–12.
- [16] Benoit Donnet and Olivier Bonaventure. 2008. On BGP communities. *ACM SIGCOMM Computer Communication Review* 38, 2 (2008), 55–59.
- [17] Ben Du, Gautam Akiwate, Thomas Krenc, Cecilia Testart, Alexander Marder, Bradley Huffaker, Alex C Snoeren, and Kimberly C Claffy. 2022. IRR Hygiene in the RPKI Era. In *International Conference on Passive and Active Network Measurement*. Springer, 321–337.
- [18] Chengyu Fan, Susmit Shannigrahi, Christos Papadopoulos, and Craig Partridge. 2020. Discovering in-network caching policies in ndn networks from a measurement perspective. In *Proceedings of the 7th ACM Conference on Information-Centric Networking*. 106–116.

- [19] Lixin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on networking* 9, 6 (2001), 733–745.
- [20] Lixin Gao and Jennifer Rexford. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on networking* 9, 6 (2001), 681–692.
- [21] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. 2021. Seven years in the life of Hypergiants' off-nets. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 516–533.
- [22] Phillipa Gill, Martin Arlitt, Zongpeng Li, and Anirban Mahanti. 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, 1–10.
- [23] Phillipa Gill, Michael Schapira, and Sharon Goldberg. 2012. Modeling on quicksand: Dealing with the scarcity of ground truth in interdomain routing data. *ACM SIGCOMM Computer Communication Review* 42, 1 (2012), 40–46.
- [24] Phillipa Gill, Michael Schapira, and Sharon Goldberg. 2013. A survey of interdomain routing policies. *ACM SIGCOMM Computer Communication Review* 44, 1 (2013), 28–34.
- [25] Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, and KC Claffy. 2014. Inferring complex AS relationships. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. 23–30.
- [26] Vasileios Giotsas, George Nomikos, Vasileios Kotronis, Pavlos Sermpezis, Petros Gigis, Lefteris Manassakis, Christoph Dietzel, Stavros Konstantaras, and Xenofontas Dimitropoulos. 2020. O peer, where art thou? Uncovering remote peering interconnections at IXPs. *IEEE/ACM Transactions on Networking* 29, 1 (2020), 1–16.
- [27] Vasileios Giotsas 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, 1193–1197.
- [28] Geoff Huston. 1999. Interconnection, peering, and settlements. In *proc. INET*, Vol. 9, 1.
- [29] Geoff Huston. 2023. IPv4 Address Report. <https://www.potaroo.net/tools/ipv4/>. (2023).
- [30] Yuchen Jin, Colin Scott, Amogh Dhamdhere, Vasileios Giotsas, Arvind Krishnamurthy, and Scott Shenker. 2019. Stable and Practical AS Relationship Inference with ProbLink. In *NSDI*, Vol. 19. 581–598.
- [31] Zitong Jin, Xingang Shi, Yan Yang, Xia Yin, Zhiliang Wang, and Jianping Wu. 2020. Toposcope: Recover as relationships from fragmentary observations. In *Proceedings of the ACM Internet Measurement Conference*. 266–280.
- [32] Juniper. 2023. BGP best path selection algorithm. <https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/bgp/topics/concept/routing-protocols-address-representation.html>. (2023). [Online; accessed 11-May-2023].
- [33] 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*. 758–759.
- [34] Savvas Kastanakis, Pavlos Sermpezis, Vasileios Kotronis, Daniel Menasché, and Thrasylvoulos Spyropoulos. 2020. Network-aware recommendations in the wild: Methodology, realistic evaluations, experiments. *IEEE Transactions on Mobile Computing* 21, 7 (2020), 2466–2479.
- [35] Akmal Khan, Taekyoung Kwon, Hyun-chul Kim, and Yanghee Choi. 2013. AS-level topology collection through looking glass servers. In *Proceedings of the 2013 conference on Internet measurement conference*. 235–242.
- [36] Rowan Klöti, Bernhard Ager, Vasileios Kotronis, George Nomikos, and Xenofontas Dimitropoulos. 2016. A comparative look into public IXP datasets. *ACM SIGCOMM Computer Communication Review* 46, 1 (2016), 21–29.
- [37] Vasileios Kotronis, Rowan Klöti, Matthias Rost, Panagiotis Georgopoulos, Bernhard Ager, Stefan Schmid, and Xenofontas Dimitropoulos. 2016. Stitching interdomain paths over IXPs. In *Proceedings of the Symposium on SDN Research*. 1–12.
- [38] Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. 2010. Internet inter-domain traffic. *ACM SIGCOMM Computer Communication Review* 40, 4 (2010), 75–86.
- [39] Aemen Lodhi, 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, 2562–2570.
- [40] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and KC Claffy. 2013. AS relationships, customer cones, and validation. In *Proceedings of the 2013 conference on Internet measurement conference*. 243–256.
- [41] Pedro Marcos, Marco Chiesa, Lucas Müller, Pradeeban Kathiravelu, Christoph Dietzel, Marco Canini, and Marinho Barcellos. 2018. Dynam-IX: A dynamic interconnection exchange. In *Proceedings of the 14th International Conference on emerging Networking Experiments and Technologies*. 228–240.
- [42] Riad Mazloum, Marc-Olivier Buob, Jordan Auge, Bruno Baynat, Dario Rossi, and Timur Friedman. 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, 173–182.
- [43] Wolfgang Mühlbauer, Anja Feldmann, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2006. Building an AS-topology model that captures route diversity. *ACM SIGCOMM Computer Communication Review* 36, 4 (2006), 195–206.
- [44] RIPE NCC. 2023. Routing Information Service (RIS). <https://www.ripe.net/>. (2023). [Online; accessed 12-May-2023].
- [45] Merit Network. 2021. Merit RADb. (2021). <https://www.radb.net/>.
- [46] University of Oregon. 2023. Route Views Project. <http://www.routeviews.org/>. (2023). [Online; accessed 12-May-2023].
- [47] Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. 2009. The (in) completeness of the observed Internet AS-level structure. *IEEE/ACM Transactions on Networking* 18, 1 (2009), 109–122.
- [48] Ricardo V Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang. 2008. In search of the elusive ground truth: the Internet's AS-level connectivity structure. *ACM SIGMETRICS Performance Evaluation Review* 36, 1 (2008), 217–228.
- [49] Leo Oliver, Gautam Akiwate, Matthew Luckie, Ben Du, and kc claffy. 2022. Stop, DROP, and ROA: Effectiveness of Defenses through the lens of DROP. In *Proceedings of the 22nd ACM Internet Measurement Conference*. 730–737.
- [50] Chiara Orsini, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. 2016. BGPStream: a software framework for live and historical BGP data analysis. In *Proceedings of the 2016 Internet Measurement Conference*. 429–444.
- [51] Bruno Quoitin and Steve Uhlig. 2005. Modeling the routing of an autonomous system with C-BGP. *IEEE network* 19, 6 (2005), 12–19.
- [52] Matthew Roughan, 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*. 325–330.
- [53] routeservers.org. 2023. Public Route Servers. <https://www.routeservers.org/>. (2023). [Online; accessed 11-May-2023].
- [54] Brandon Schlinker, Hyojeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V Madhyastha, Italo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. 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*. 418–431.
- [55] Pavlos Sermpezis and Vasileios Kotronis. 2019. Inferring catchment in internet routing. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 3, 2 (2019), 1–31.
- [56] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. 2018. ARTEMIS: Neutralizing BGP hijacking within a minute. *IEEE/ACM Transactions on Networking* 26, 6 (2018), 2471–2486.
- [57] Xiaozhe Shao and Lixin Gao. 2021. Policy-rich interdomain routing with local coordination. *Computer Networks* 197 (2021), 108292.
- [58] Yuval Shavitt and Eran Shir. 2005. DIMES: Let the Internet measure itself. *ACM SIGCOMM Computer Communication Review* 35, 5 (2005), 71–74.
- [59] Rachee Singh, David Tench, Phillipa Gill, and Andrew McGregor. 2021. PredictRoute: a network path prediction toolkit. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5, 2 (2021), 1–24.
- [60] The Internet Society. 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>. (2019).
- [61] Zhihong Tian, Shen Su, Wei Shi, Xiaojiang Du, Mohsen Guizani, and Xiang Yu. 2019. A data-driven method for future Internet route decision modeling. *Future Generation Computer Systems* 95 (2019), 212–220.
- [62] Peter Vranx, Pasquale Gurzi, Abdel Rodriguez, Kris Steenhaut, and Ann Nowé. 2015. A reinforcement learning approach for interdomain routing with link prices. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 10, 1 (2015), 1–26.
- [63] Feng Wang and Lixin Gao. 2003. On inferring and characterizing Internet routing policies. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*. 15–26.
- [64] Florian Wohlfart, Nikolaos Chatzis, Caglar Dabanoglu, Georg Carle, and Walter Willinger. 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*. 206–220.
- [65] Tianhao Wu, Jessie Hui Wang, Jilong Wang, and Shuying Zhuang. 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, 216–244.
- [66] Kok-Kiong Yap, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holliman, Gary Baldus, Marcus Hines, Taeewon Kim, Ashok Narayanan, Ankur Jain, 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*. 432–445.
- [67] Xiaoyang Zhao, Chuan Wu, and Franck Le. 2020. Improving inter-domain routing through multi-agent reinforcement learning. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 1129–1134.