

Investigating Location-aware Advertisements in Anycast IP Networks

Savvas Kastanakis

Lancaster University
s.kastanakis@lancaster.ac.uk

Ioana Livadariu

Simula Metropolitan Center
ioana@simula.no

Vasileios Giotsas

Cloudflare
vasilis@cloudflare.com

Neeraj Suri

Lancaster University
neeraj.suri@lancaster.ac.uk

ABSTRACT

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.

In this work, we focus on a new vector of the performance inefficiency of anycast and we aim to shed light on the prevalence and rationale behind selective announcements in anycast networks. Through empirical evidence, we identify their significant adoption, primarily driven by the geolocation-agnostic BGP best path selection process. In particular, we observe that 84.06% of anycast ASes announce at least one of their anycast prefixes to a specific subset of their neighbors, whereas 80% of these selective announcers announce 100% of their prefixes selectively. This research represents an initial step towards comprehending the root causes of selective announcements in anycast IP networks.

1 INTRODUCTION

IP anycast [32, 40, 46] is widely used to allow services to be transparently replicated across the Internet. By advertising, via BGP [47], the same IP prefix from each Point of Presence (PoP), traffic from any source on the Internet is routed to one of the PoPs according to properties of the path. Optimally users are associated with the nearest site (PoP), minimizing latency. Two of the most studied applications of IP anycast to date are root DNS servers [10, 12, 19, 20, 23, 34, 36, 41, 45, 48] and content delivery networks (CDNs) [7, 14, 15, 21, 22, 26]. This paper focuses explicitly on IP anycast for CDNs.

What makes IP anycast attractive when deploying a globally replicated service is the mental model that it seems to permit. In particular, as one adds more anycast replicas in locations with many clients, it is generally believed [9, 15, 20]

that: (1) overall client latency will decrease and (2) load from nearby clients will be more evenly distributed. Of course, inter-domain routing is not guaranteed to be optimal in terms of bandwidth, latency, or geographic proximity: at best, BGP can be relied upon for connectivity and policy-compliance.

The challenge in managing anycast is that BGP routing is not always what one would expect. Although user-to-site mapping generally follows geography [8, 31, 37], studies of routing have shown that actual network topology can vary [49], and recent observations have shown that the mapping depends heavily on the policies of many ASes, thus, can be unexpectedly chaotic [11, 30, 33, 37]. Lacking any useful information to distinguish between two or more anycast replicas, routers often select a distant, high-latency anycast site over the closer, low-latency one. Again, it is not surprising that inter-domain routing would not choose the best alternative, but it is surprising that the best alternative is often an unselected option [31, 35].

Additionally, an AS (anycast or unicast) may select to restrict the propagation of certain routes for traffic engineering purposes [16, 30, 50]. An anycast network might strategically employ selective announcements to fine-tune its routing dynamics for enhanced operational efficiency and user experience. By selectively announcing prefixes to specific neighboring networks, the anycast network can exert control over incoming traffic flows, directing them away from congested or underperforming nodes towards more optimal PoPs. This targeted approach allows the network to dynamically adapt to changing traffic patterns and network conditions, mitigating potential congestion issues and ensuring a smoother delivery of services to end-users.

In this study, our objective is to shed light on the prevalence of selectively announced prefixes within anycast networks and delve into the underlying reasons driving such practices. We present compelling evidence indicating that a significant portion of anycast ASes adopt selective prefix announcements as a traffic engineering strategy.

Specifically, our contributions are as follows:

We measure the adoption of selective announcements across all anycast ASes.

We derive a methodology to capture the AS-level catchment of anycast ASes, i.e., the direct neighbors that receive announcements for an anycast prefix.

We measure the extent at which selective anycast ASes decide to propagate their announcements across regional and global neighbors.

We discuss the potential root causes of the adoption of selective announcements in anycast networks.

We publish the artifacts (source code and data) of our study to facilitate future research on location-aware BGP path selection and anycast catchment prediction.

Our investigation indicates that the primary motivation behind this approach stems from a deficiency in backbone infrastructure for coordinating and managing routing announcements across PoPs. Consequently, anycast networks often resort to leveraging regional providers to establish connections between their PoPs. To the best of our knowledge, this is the first study of the interplay between selective announcements and IP anycasting on CDN operations.

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 [24, 28] 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.

2.1 BGP Routing Policies

The Gao-Rexford Model In the BGP selection process [18, 29], a router evaluates multiple routes for the same destination IP prefix from different AS neighbors, choosing the most preferable route. *Locpref*, the highest-priority metric, determines route selection, typically with customer routes assigned the highest *locpref* values due to revenue generation and provider routes the lowest due to cost. Gao and Rexford [24] demonstrated that this *locpref* ordering is crucial for global routing system convergence, termed the Gao-Rexford model.

The Valley-free Rule After selecting the best route, a router may propagate it to neighboring ASes following the *valley-free rule*. This rule dictates that a customer route can be shared with any neighbor AS, while routes from peers or providers can only be shared with customers. This policy prevents an AS from providing free transit to peers or providers, conserving resources and avoiding unnecessary traffic costs.



(a) Selective announcement per AS. (b) Selective announcement per location.

Figure 1: Selective Announcement Types.

2.2 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 [24] can model [8, 25, 27, 30, 39, 50]. 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. 1a, AS C announces prefix p only to its neighbor AS B and not to AS A.

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 [38]. As a result, ASes may prefer *peer* over *customer routes* for performance reasons, since for more than 90% peering paths outperform customer paths [6].

2.3 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 [50]. This lack of a backbone network means that each PoP operates relatively independently, connecting directly to upstream providers for internet connectivity in a location-dependent manner rather than an AS-dependent one. For example, in Fig. 1b, 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.

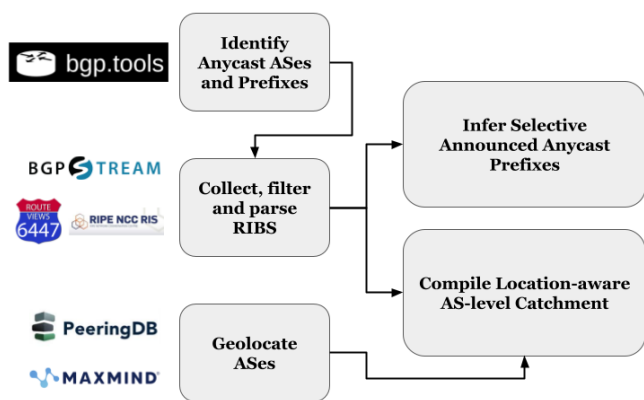


Figure 2: Tools used and overview of the methodology.

This decentralized architecture can lead to the occurrence of selective announced prefixes [30, 50]. 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.

3 DATA SOURCES AND METHODOLOGY

The goal of this paper is to inform for the selective routing policies of anycast ASes against their direct neighbors per geographical region. An overview of our methodology can be found in Fig. 2.

Our methodology 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 map all anycast prefixes to their respective origin ASes we utilize the bgp.tools platform [17]. This is an important step, which aims to narrow down our analysis only to anycast IP prefixes announced by anycast ASes.

Due to space constraints, we provide only the characteristics of the top anycast ASes based on their RUM uptime [1] in Table 1, nonetheless, our study investigates all anycast ASes on the Internet¹.

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

¹Specific information per anycast AS can be found in our online repository.

ASN	AS Name	RUM Uptime	# of PoPs
54113	Fastly	99.56	102
16509	Amazon	99.55	166
15133	Edgecast	99.53	59
20940	Akamai	99.49	183
13335	Cloudflare	99.33	197
60068	CDN77	99.20	63
16276	OVH	99.20	43
21859	Zenlayer	99.20	78
199524	G-Core	99.11	91
15169	Google	98.96	135
30081	Cachefly	98.62	66
22822	Edgio	97.75	72

Table 1: Top Anycast Networks based on RUM Uptime.

BGPStream API [44]. BGPStream utilizes the route collectors of RIPE RIS [42] and RouteViews [43] and in this project we make use of all available route collectors up to this date (63).

We follow the standard procedure to remove RIB entries that include paths with cycles, or entries which have been flagged as bogus, as well as remove prepended ASes in the AS-paths.

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 [50] to infer the selective announced prefixes in the interdomain routing system. Their methodology 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 [13], as well as, the routing tables for all anycast ASes collected in the previous step. We follow the logic and conventions of [50]: if a customer prefix is received through a peer/provider route, then this is a selective announced prefix. Respectively, a peer prefix is selectively announced if it is received through a provider route.

From a total of 691 anycast ASes, we found 110 ASes (15.94%) that do not advertise a single selective announced prefix. Hence, 84.06% of anycast ASes announce at least one anycast prefix to only a subset of their neighbors. In Fig. 3, we plot the CDF of selective announced prefixes per anycast origin AS. We find that 80% of the selective anycast ASes announce 100% of their prefixes selectively. I.e., *if an anycast AS employs selective routing policies for one pre x, it likely does the same for all of its pre xes*. Specifically, all the top CDNs mentioned in Table 1, announce 100% of their anycast

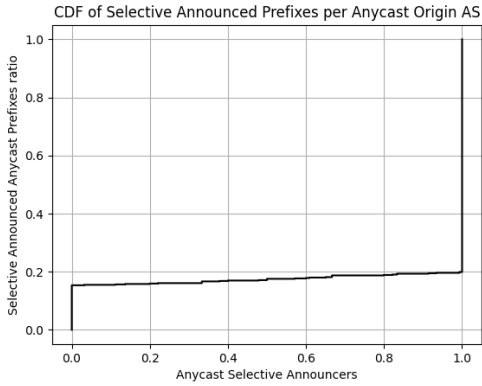


Figure 3: Selective announced prefixes per anycast AS.

prefixes selectively, while, the average selective announced prefix ratio across all anycast ASes is 82.5%.

Augment ASes with location specific characteristics To contextualize our findings within a geographical framework, we geolocate all ASes of the interdomain routing system (up to this date) into their respective countries and regions. By incorporating geolocation data, we aim to uncover regional trends and disparities in anycast deployment and routing behavior.

For the geolocation part, 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 all ASes using the MaxMind geolocation tool [2] through the RIPEstat API [4], as well as, the countries of the public peering locations of all available ASes by PeeringDB API [3]. To map countries to their respective regions, we use the United Nations dataset [5]. We leverage these data in the following step of our methodology (AS-level Catchment compilation), as well as, in our analysis in Section 4, where we classify the neighbors of an anycast AS into regionals or globals based on their geographic footprint.

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. 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 (e.g., in the same geographic region). 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 [5].

We group all the direct neighboring ASes (penultimate ASes in the AS paths) responsible for carrying traffic for a specific prefix according to the regions of the VPs.

This dataset looks as follows:

```
"Anycast ASx": {
  "anycast prefix p1 announced by ASx": {
    "region r1": {
      "receivers": ["ASy", "ASz"]
    }, "region r2": {
      "receivers": ["ASq"]
    }
  }
  "anycast prefix p2 announced by ASx": {
    "region r1": {
      "receivers": ["ASa", "ASz"]
    }
  }
}
```

In the above example, *ASx* is an anycast CDN which selectively announces its anycast prefixes to its neighbors. Prefix *p1* is announced to *ASy* and *ASz* when the VP locates in region *r1*, but the same prefix is advertised only to *ASq* when the VP locates in region *r2*.


We delve into the intricacies of the geolocation specific routing policies of anycast ASes in Section 4.

4 REGIONALITY OF DIRECT NEIGHBORS OF ANYCAST ASSES


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.

4.1 Definition of Regionality

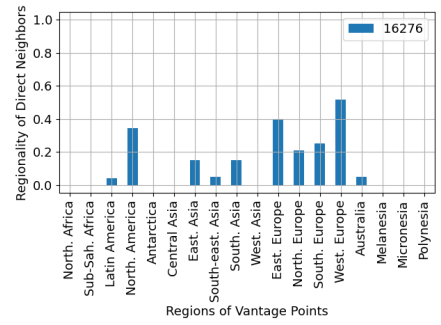
We label an AS as *regional* if: a) more than 90% of its prefixes are announced in a single region, and/or b) more than 90% of its peering links exist in the same region. For instance, *ASa* announces half of its prefixes in Italy and the other half in Spain, both countries falling under Southern Europe. Furthermore, *ASa* peers with three neighbors in: Portugal, Malta, and Greece, all located in Southern Europe. Thus, *ASa* qualifies as a regional AS.




(a) AS13335.




(b) AS15133.




(c) AS16276.




(d) AS16509.




(e) AS20940.




(f) AS21859.



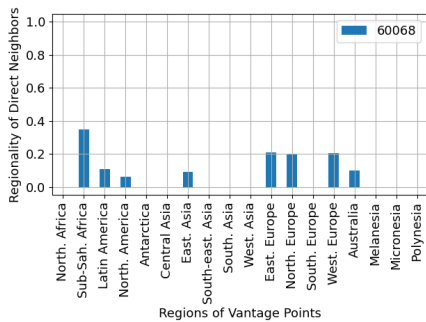
(g) AS22822.



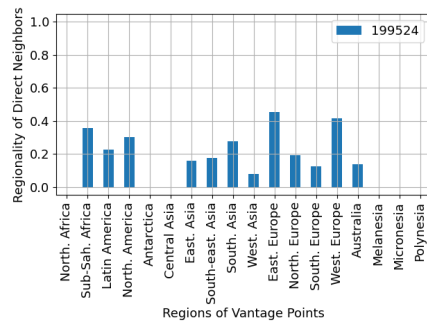
(h) AS30081.




(i) AS54113.



(j) AS60068.



(k) AS199524.



(l) AS15169.

Figure 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.

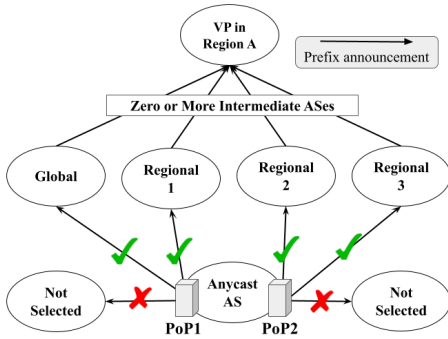


Figure 5: The Anycast AS connects to different sets of ASes in different PoPs. Large PoPs (e.g., PoP1) tend to be connected to both regional and global ASes, while regional PoPs (e.g., PoP2) are typically connected only to regional ASes. As a result, a VP in the region of a regional PoP is more inclined to reach the CDN through regional ASes.

By leveraging the AS-level catchment dataset (mentioned in Section 4) 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:

$$R = \frac{\# \text{ of regional neighbors that prefer the announcement}}{\# \text{ of all neighbors that prefer the announcement}}$$

As demonstrated in Fig. 5, the Anycast AS connects to different sets of ASes in different PoPs and selectively announces a prefix p 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 reach the CDN 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 CDN PoP 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.

4.2 Regionality Analysis

In Fig. 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 CDNs. Cloudflare

(AS13335) prefers regional ASes when the source locates at the Sub-Saharan Africa, the Southern Asia and the Western Europe in more than 50% of the times. Amazon (AS16509), Google (AS15169) and Cachefly (AS30081) 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 big 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 (AS60068), Edgecast (AS15133) and Zenlayer (AS21859) 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.

Our results highlight the need to further study the correlation between selective announcements and location-based routing policies and can lay the ground towards understanding the confounding factors of: a) the anycast inefficiencies and b) the location-agnostic BGP best-path selection process.

5 CONCLUSION

The opacity surrounding anycast operations has profound implications in our ability to predict, understand, and debug such networks effectively. Through the analysis of BGP data, we identify and characterize selective and region-specific announcements, introducing a novel metric, "*regionality*", to delineate varying CDN strategies.

Our findings indicate a substantial proportion of anycast ASes employing selective announcements in a *per location* or *per AS* basis, with the majority announcing all their prefixes selectively. In particular, we observe that 84.06% of anycast ASes announce at least one of their anycast prefixes to a specific subset of their neighbors, whereas 80% of these selective announcers announce 100% of their prefixes selectively.

Looking ahead, we aim to quantify the extent at which geography influences BGP routing policies in anycast ASes. By enhancing our predictive capabilities with location-specific information, we aspire to contribute towards a deeper understanding and improved management of anycast networks.

REFERENCES

- [1] 2024. cdnperf.com. <https://www.cdnperf.com/>. (2024). [Online; accessed 20-March-2024].
- [2] 2024. MaxMind API. <https://www.maxmind.com/>. (2024). [Online; accessed 20-March-2024].
- [3] 2024. PeeringDB. <https://www.peeringdb.com>. (2024).
- [4] 2024. RIPEStat API. <https://stat.ripe.net/app/launchpad>. (2024). [Online; accessed 20-March-2024].
- [5] 2024. United Nations Statistics Division. <https://unstats.un.org/unsd/methodology/m49/>. (2024).
- [6] 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.
- [7] Hussein A Alzoubi, Seungjoon Lee, Michael Rabinovich, Oliver Spatscheck, and Jacobus Van Der Merwe. 2011. A practical architecture for an anycast CDN. *ACM Transactions on the Web (TWEB)* 5, 4 (2011), 1–29.
- [8] 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.
- [9] Hitesh Ballani and Paul Francis. 2005. Towards a global IP anycast service. *ACM SIGCOMM Computer Communication Review* 35, 4 (2005), 301–312.
- [10] Hitesh Ballani, Paul Francis, and Sylvia Ratnasamy. 2006. A measurement-based deployment proposal for IP anycast. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. 231–244.
- [11] Ray Bellis. 2015. Researching F-root Anycast Placement Using RIPE Atlas. ripe blog https://labs.ripe.net/Members/ray_bellis/researching-f-root-anycast-placement-using-ripe-atlas. (2015).
- [12] Nevil Brownlee, KC Claffy, and Evi Nemeth. 2001. DNS Root/gTLD performance measurements. *USENIX LISA, San Diego, CA* (2001).
- [13] CAIDA. 2024. AS-Relationships Dataset. <https://publicdata.caida.org/datasets/as-relationships/>. (2024). [Online; accessed 20-March-2024].
- [14] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. 2013. Mapping the expansion of Google’s serving infrastructure. In *Proceedings of the 2013 conference on Internet measurement conference*. 313–326.
- [15] Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye. 2015. Analyzing the Performance of an Anycast CDN. In *Proceedings of the 2015 Internet Measurement Conference*. 531–537.
- [16] C Cardona, Pierre Francois, and Paolo Lucente. 2016. *Impact of BGP Itering on Inter-Domain Routing Policies*. Technical Report.
- [17] Ben Cartwright-Cox. 2024. BGP.Tools <https://bgp.tools/>. (2024).
- [18] Cisco. 2024. BGP best path selection algorithm. <https://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html>. (2024). [Online; accessed 28-March-2024].
- [19] Lorenzo Colitti, Erik Romijn, Henk Uijterwaal, and Andrei Robachevsky. 2006. Evaluating the effects of anycast on DNS root name servers. *RIPE document RIPE-393* 6 (2006).
- [20] Ricardo de Oliveira Schmidt, John Heidemann, and Jan Harm Kuipers. 2017. Anycast latency: How many sites are enough?. In *Passive and Active Measurement: 18th International Conference, PAM 2017, Sydney, NSW, Australia, March 30-31, 2017, Proceedings 18*. Springer, 188–200.
- [21] Xun Fan, Ethan Katz-Bassett, and John Heidemann. 2015. Assessing affinity between users and CDN sites. In *Traffic Monitoring and Analysis: 7th International Workshop, TMA 2015, Barcelona, Spain, April 21-24, 2015. Proceedings 7*. Springer, 95–110.
- [22] Ashley Flavel, Pradeepkumar Mani, David Maltz, Nick Holt, Jie Liu, Yingying Chen, and Oleg Surmachev. 2015. fFastRouteg: A Scalable fLoad-Awareg Anycast Routing Architecture for Modern fCDNs. In *12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)*. 381–394.
- [23] Marina Fomenkov and Bradley Huffaker. 2001. Macroscopic Internet Topology and Performance Measurements from the fDNSg Root Name Servers. In *15th Systems Administration Conference (LISA 2001)*.
- [24] Lixin Gao and Jennifer Rexford. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on networking* 9, 6 (2001), 681–692.
- [25] Phillipa Gill, Michael Schapira, and Sharon Goldberg. 2013. A survey of interdomain routing policies. *ACM SIGCOMM Computer Communication Review* 44, 1 (2013), 28–34.
- [26] Danilo Giordano, Danilo Cicalese, Alessandro Finamore, Marco Mellia, Maurizio M Munafò, Diana Zeaiter Joulblatt, Dario Rossi, et al. 2016. A First Characterization of Anycast Traffic from Passive Traces.. In *TMA*.
- [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] Juniper. 2024. BGP best path selection algorithm. <https://www.juniper.net/documentation/us/en/software/junos/vpn-l2/bgp/topics/concept/routing-protocols-address-representation.html>. (2024). [Online; accessed 28-March-2024].
- [30] Savvas Kastanakis, Vasileios Giotsas, Ioana Livadariu, and Neeraj Suri. 2023. Replication: 20 Years of Inferring Interdomain Routing Policies. In *Proceedings of the 2023 ACM on Internet Measurement Conference*. 16–29.
- [31] 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.
- [32] Dina Katabi and John Wroclawski. 2000. A framework for scalable global IP-anycast (GIA). *ACM SIGCOMM Computer Communication Review* 30, 4 (2000), 3–15.
- [33] Jan Harm Kuipers. 2015. Analyzing the K-root DNS anycast infrastructure. In *Twente Student Conference on IT*. 1–6.
- [34] Matthew Lentz, Dave Levin, Jason Castonguay, Neil Spring, and Bobby Bhattacharjee. 2013. D-mystifying the D-root Address Change. In *Proceedings of the 2013 conference on Internet measurement conference*. 57–62.
- [35] Zhihao Li, Dave Levin, Neil Spring, and Bobby Bhattacharjee. 2018. Internet anycast: performance, problems, & potential. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*. 59–73.
- [36] Jinjin Liang, Jian Jiang, Haixin Duan, Kang Li, and Jianping Wu. 2013. Measuring query latency of top level DNS servers. In *Passive and Active Measurement: 14th International Conference, PAM 2013, Hong Kong, China, March 18-19, 2013. Proceedings 14*. Springer, 145–154.
- [37] Ziqian Liu, Bradley Huffaker, Marina Fomenkov, Nevil Brownlee, and KC Claffy. 2007. Two days in the life of the DNS anycast root servers. In *Passive and Active Network Measurement: 8th International Conference, PAM 2007, Louvain-la-neuve, Belgium, April 5-6, 2007. Proceedings 8*. Springer, 125–134.
- [38] Aemen Lodhi, Amogh Dhamdhare, 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.

- [39] 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.
- [40] Christopher Metz. 2002. IP anycast point-to-(any) point communication. *IEEE Internet computing* 6, 2 (2002), 94–98.
- [41] Giovane CM Moura, Ricardo de O Schmidt, John Heidemann, Wouter B de Vries, Moritz Muller, Lan Wei, and Cristian Hesselman. 2016. Anycast vs. DDoS: Evaluating the November 2015 root DNS event. In *Proceedings of the 2016 Internet Measurement Conference*. 255–270.
- [42] RIPE NCC. 2024. Routing Information Service (RIS). <https://www.ripe.net/>. (2024). [Online; accessed 20-March-2024].
- [43] University of Oregon. 2024. Route Views Project. <http://www.routeviews.org/>. (2024). [Online; accessed 20-March-2024].
- [44] 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.
- [45] Jeffrey Pang, James Hendricks, Aditya Akella, Roberto De Prisco, Bruce Maggs, and Srinivasan Seshan. 2004. Availability, usage, and deployment characteristics of the domain name system. In *Proceedings of the 4th ACM SIGCOMM conference on Internet measurement*. 1–14.
- [46] Craig Partridge, Trevor Mendez, and Walter Milliken. 1993. *Host anycasting service*. Technical Report.
- [47] Yakov Rekhter, Tony Li, and Susan Hares. 2006. RFC 4271: A border gateway protocol 4 (BGP-4). (2006).
- [48] Sandeep Sarat, Vasileios Pappas, and Andreas Terzis. 2005. On the use of anycast in DNS. *ACM sigmetrics performance evaluation review* 33, 1 (2005), 394–395.
- [49] Neil Spring, Ratul Mahajan, and Thomas Anderson. 2003. The causes of path inflation. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. 113–124.
- [50] 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.