Too Remote to Be Local: Latency Inflation in Anycast due to Remote Peering

Remi Hendriks*, Stefano Servillo†, Francesca Cuomo†, Cristian Hesselman*‡,
Roland van Rijswijk-Deij*, Savvas Kastanakis*
*University of Twente, Netherlands ‡SIDN Labs, Netherlands
Emails: {remi.hendriks, r.m.vanrijswijk, c.e.w.hesselman, s.kastanakis}@utwente.nl
†University of Rome "Sapienza", Italy Emails: {stefano.servillo, francesca.cuomo}@uniroma1.it

***Abstract*—Remote peering (RP) enables networks to reach Internet Exchange Points (IXPs) without physical presence in the same datacenter/location, offering cost-effective interconnection. However, in the context of anycast, where routing efficiency relies on the assumption that traffic is routed to the geographically nearest instance, RP can undermine this principle by introducing hidden performance costs and inflating paths. This paper presents a hybrid methodology, combining data-plane measurements with control-plane insights, to detect and geolocate RP links at scale using IP-level geolocation, AS-level mapping, and IXP facility data. Using 3 million traceroutes targeting 13,735 anycast /24-prefixes, we quantify the impact of RP on anycast routing.**

Our analysis shows that whilst RP is widespread, it often connects ASes geographically close to the IXP facility. However, for cases where it is "remote" we find it frequently results in long detours, elevated latency, and sub-optimal anycast site selection. Finally, we identify key ASes and IXP regions that play a central role in driving RP-related detours and locality violations in anycast routing. Our findings demonstrate at Internet scale that RP undermines anycast locality by increasing the likelihood of detours and inflated paths, underscoring the need for improved visibility into interconnection practices that impact latency-sensitive services.

I. Introduction

Border Gateway Protocol (BGP) is the de facto inter-domain routing protocol of the Internet which enables ASes to exchange reachability information [1]. A key venue for this interconnection are Internet Exchange Points (IXPs), critical components of today's Internet infrastructure [2]–[4]. Traditionally, IXPs were designed as physical layer-2 (L2) interconnection hubs, where ASes would colocate at a common facility to exchange traffic locally [5]. However, as the demands for interconnection increased, a shift occurred, termed as the *Flattening of the Internet* [6]–[8]. Due to the rise of Content Providers (CPs) and Content Delivery Networks (CDNs), the Internet evolved into a mesh interconnection network with a dense topology. Big Internet players (*e.g.*, 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. In this context, IXPs played a central role in enabling large CDNs and CPs to bypass Tier-1 ISPs, providing an efficient interconnection infrastructure that supported the flattening trend. As a result, the IXP ecosystem has undergone a fundamental shift in peering practices to respond to the increased performance and connectivity requirements [7].

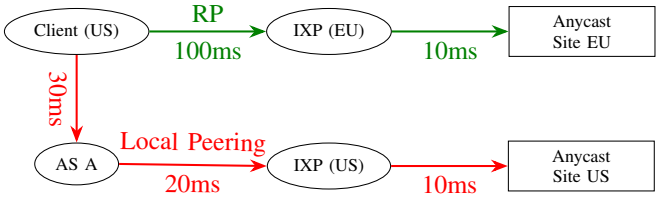

Fig. 1. A client in the US sub-optimally connects to an anycast service through a topologically short RP link reaching Site EU (110 ms) whilst a lower latency link exists reaching Site US (60 ms) through local peering.

To accommodate this changing landscape and meet increasing demands for performance and reachability, the IXP ecosystem itself has evolved. One notable development is the rise of *Remote Peering* (RP) [9]. RP is a form of peering in which parties establish a peering relationship over a distance. Through RP, an AS connects to an IXP without requiring physical presence, often leveraging L2 transport from a reseller or partners. Whilst RP reduces operational costs, it introduces ambiguity in topology visibility, peering policies, and Traffic Engineering (TE) [10]–[12].

This hidden complexity is particularly impactful in services that rely on accurate network proximity, such as anycast [13]. Anycast is a widely adopted technique in CDNs, DNS resolution, and DDoS mitigation, where the same IP prefix is announced from multiple geographically distributed locations. Anycast relies on BGP to route users to the "nearest" instance, typically assumed to be geographically close. However, RP challenges this assumption: it introduces misleading topology signals, allowing BGP speaking routers to select remote peers that may appear close in the AS-path but are physically distant in geographical terms [10], [11]. As illustrated in Fig. 1, a client trying to reach an anycast service may be routed to Site EU via remote peering, even though Site US is geographically closer. As a result, RP not only affects interconnection visibility, but also undermines anycast-based service performance and predictability.

In this work, we investigate the phenomenon of RP from both a control-plane and a data-plane perspective. While prior work has established accurate RP inference methods [11], [14], there remains a striking lack of visibility into its geographic and performance consequences. Our goal is to bridge this gap.

Specifically, our contributions are the following:

- We develop a traceroute-based methodology to detect and geolocate RP connections.
- We apply our method to millions of measured paths toward anycast prefixes, analyzing how RP affects anycast site selection and revealing that RP is widespread and can cause significant latency and geographic detours.
- We identify key ASes and regions most affected by RP.
- We release all code and data as open-source artifacts to support reproducibility and future research [15].

By improving visibility as to where RP occurs and its impact on anycast routing, we provide insights for anycast operators and researchers to troubleshoot undesirable anycast routing.

II. Background and Related Work

Although the rise of RP has been acknowledged in past measurement studies [10]–[12], [14], [16], it has largely been treated as a side effect of broader trends like Internet flattening or IXP expansion. This gap in the literature leaves network operators with limited tools to expose where and how RP manifests, and what its effects and implications are for latency-critical systems. To contextualize our contributions, we review prior work on remote peering, its detection, and its implications across interconnection and routing systems.

Prior studies have highlighted the growing prevalence of RP and its unintended consequences across Internet routing and interconnection. Castro et al. [16] introduced the first systematic method to detect RP using traceroute data and IXP/facility information. They showed that RP is already widespread, enabling ASes to interconnect without physical colocation and decoupling geographic proximity from peering. Their findings suggest RP contributes to increased interconnection without Internet flattening. However, their approach is limited by coarse-grained location data and reliance on public databases, leaving the geographic and performance impact of RP largely unexamined.

Mazzola et al. [10], [11] examine the latency implications of RP using active measurements from major IXPs, showing that BGP often prefers RP routes due to shorter AS paths even when they yield higher latency than their local peers. Their work underscores the complexity RP introduces into TE. Similarly, Bian et al. [12] analyze how RP distorts anycast-based service by breaking BGP's assumption of topological proximity, revealing that at least 19.2% of anycast prefixes are misrouted due to invisible RP paths.

The current state-of-the-art in RP detection is the work by Nomikos et al. [14], who proposed a multi-layered inference methodology combining latency measurements, IP-to-AS mapping, colocation data, and port capacity information. Their approach achieves high accuracy and coverage across 30 major IXPs, offering a robust framework for identifying RP relationships at scale. However, their analysis remains primarily topological, focusing on the classification of AS-level remote peers without localizing the physical endpoints of RP connections or evaluating their operational impact.

Algorithm 1 Identify candidate Remote Peering (cRP) Links

```
Initialize cRPs ← ∅                                   ▷ Candidate RP links
for all hop h_i, hop h_{i+1} in T do
    if TTL(h_i) + 1 ≠ TTL(h_{i+1}) then
        continue                                      ▷ Exclude non-consecutive hops
    end if
    if Address(h_i) or Address(h_{i+1}) ∈ Bogons then
        continue                                      ▷ Exclude bogus IP addresses
    end if
    if City(h_i) = City(h_{i+1}) then
        continue                                      ▷ Exclude same city links
    end if
    if AS(h_i) = AS(h_{i+1}) then
        continue                                      ▷ Exclude same AS links
    end if
    if Org(h_i) and Org(h_{i+1}) = Org(Origin AS) then
        continue                                      ▷ Exclude same ORG links
    end if
    if AS(h_i) and AS(h_{i+1}) ∉ same IXP then
        continue                                      ▷ Ensure shared IXP membership
    end if
    Add link (h_i, h_{i+1}) to cRPs
end for
return cRPs
```

While these efforts reveal important structural and performance effects of RP, they generally treat detection and impact separately. Additionally, while latency measurements can offer hints about remote connectivity, they are often insufficient on their own: Factors such as network congestion and asymmetric routing can distort latency-based inferences, making them unreliable as a sole indicator of RP [14]. Our work unifies these threads, offering a measurement framework that jointly identifies RP instances, localizes their geographic footprint, and quantifies their prevalence and effects in the context of anycast-based services.

III. Measurement Architecture

We leverage extensive traceroute data toward 13,735 anycast /24-prefixes to systematically infer RP links at IXPs, focusing on their geographical and topological characteristics. The measurements were conducted on April 19, 2025, using CAIDA's active measurement infrastructure, Archipelago [17] (Ark), providing 274 Vantage Points (VPs) in 178 distinct ASes at the time of measurement. To obtain target anycast addresses, we use data from LACeS [18], a daily anycast census, dated April 19, 2025. We use prefixes confirmed using iGreedy [19], where the census provides high detection accuracy and city-level geolocation data of anycast sites. Specifically, the inference steps, as outlined in Algorithm 1, are as follows:

Step 1: Initialization. We initialize an empty set of candidate RP (cRP) links to be populated through subsequent filtering steps. These cRP links represent hops that potentially exhibit RP characteristics and are refined through the pipeline below.

Step 2: Link Identification. We iterate over pairs of consecutive hops where the time-to-live (TTL) difference is exactly one, indicating a direct router-level connection. This condition ensures that the link corresponds to a true forwarding step rather than an unrelated or spurious IP reply.

Step 3: Bogus Filtering. Links containing bogon addresses, *i.e.*, private or unallocated IP ranges, are excluded as they have no geographic or AS information. To identify these addresses, we use Team Cymru's bogon prefix list [20].

Step 4: Geographical Filtering. To reflect the inherent geographic separation of RP, we retain only those links where consecutive hops are located in different cities. This consideration ensures that we capture links where the underlying ASes are not co-located, which is a core characteristic of RP , as it typically involves connectivity between a local and a remotely connected participant at an IXP.

We locate hops using Hoiho [21] that extracts location hints from `PTR` records and IPInfo's [22] geolocation database as `PTR` record location hints are often unavailable. These locations are validated using captured round-trip times (RTT) from VPs using the speed-of-light in fiber optics, as both methods may yield incorrect geolocations (*e.g.*, due to outdated `PTR` records or commercial dataset inaccuracies). For hops near the anycast destination we leverage captured RTTs from multiple VPs. If any VP detects a speed-of-light violation, we conservatively mark the hop location as invalid for all traceroutes. We prefer valid PTR locations over IPInfo locations.

Step 5: AS Filtering. We discard links between same-AS hops as RP links are inter-domain connections. AS numbers for each hop are inferred using CAIDA's prefix2as dataset [23], which maps IP prefixes to ASes from global BGP tables. However, since IXPs operate internal IP ranges (*e.g.*, for peering LANs), not visible in global BGP tables, we use PeeringDB [24] that documents IP assignments within IXPs.

Step 6: ORG Filtering. To avoid false positives, we require that at least one hop belong to an organization different from the anycast AS, preventing internal backbone or peering links from being misclassified as RP. We use CAIDA's AS-to-organization dataset [25] to map ASes to organizations, which is especially important since large operators often manage multiple ASes or shared infrastructure.

Step 7: IXP Filtering. Finally, as RP connects ASes at an IXP, we only consider links connecting ASes that share membership of the same IXP using data from PeeringDB [24].

All other links that do not meet these criteria are treated as non-cRP; they constitute the control group in our analysis, against which we compare the characteristics of cRP links.

IV. Analysis

In this section, we examine the impact of RP on anycast routing using data from a large traceroute campaign to 13,735 anycast prefixes (ranging from regional DNS services to global CDNs), averaging 244 VPs per prefix. We structure our analysis in two parts: a) data-plane analysis, where we extract and characterize RP links based on traceroute path features; and b) control-plane analysis, where we identify RP links that consistently cause BGP paths to deviate from locality.

Our findings highlight that RP is both prevalent and operationally significant as it can lead to routing paths that violate locality expectations in anycast. We quantify this effect across several dimensions, *i.e.*, geographic distance, latency inflation, and detour ratios, by comparing observed paths to inferred optimal ones to the nearest anycast site (as located by LACeS [18]). This allows us to surface weak points in the interconnection fabric. By highlighting these problematic cases, our analysis provides actionable insight for network operators to detect and address inefficient RP along AS paths that may degrade their anycast performance.

A. Data plane Analysis

Identifying Candidate RP Links from Traceroutes. We begin by analyzing traceroute data to detect potential RP links based on structural, geographic, and topological features. From a total of 3,340,490 traceroutes, we extract 21,061,981 hop pairs (*i.e.*, links) across all paths. First, we eliminate 3.0 M links that contain bogon IPs, leaving 18.1 M links (86.1%)

Next, we discard 4.6 M links with missing or invalid location data and 9.2 M links where both hops are located in the same city. After applying these filters, we retain 4.4 M links (20.9%) that are confidently geolocated to distinct cities.

Following this step we remove inter-domain links (*i.e.*, both hops are inside the same AS), further narrowing the set to 1.0 M links (4.7%). As mentioned, we do not consider links between distinct ASes that belong to the same organization. This step removes only 711 links, indicating that the majority of cRP links are inter-organizational. Finally, we filter out 0.5 M links connecting ASes that do not share membership for any IXP. This leaves us with our final set of cRP, consisting of 438,391 links (2.0%). Here, we note that ideally, operator-confirmed RP cases would provide ground truth for validating our inference pipeline. However, such data is rarely available at scale due to the opaque nature of reseller arrangements and limited transparency of IXP memberships.

Table I provides a step-by-step breakdown of this filtering pipeline, showing the number and percentage of links retained at each stage. The resulting set of cRP links forms the basis of our subsequent analysis and includes 3,010 unique AS links, 10,036 unique IP links from a total of 407,238 traceroutes.

RTT and Distance Characteristics of cRP Links To better understand the performance implications of cRP links, we examine their RTT and geographic distance characteristics. Figure 2 presents the cumulative distribution functions (CDFs) of RTT deltas (left) and inter-hop distances (right) for cRP (red) and non-CRP (blue) links.

First, looking at the RTT deltas we find roughly 40% of non-cRP links and 20% of RP links have negative RTT deltas (as the CDF starts at 0.40 and 0.20 respectively), motivating our decision to not consider RTT for inferring RP. Next, we observe a significant increase in mean latency for cRP links (13.22 ms compared to 3.64 ms for non-CRP links) indicating

TABLE I
BREAKDOWN OF TRACEROUTES AND LINKS AFTER THE VARIOUS FILTERING STAGE FROM ALGORITHM 1

Step	Links		Unique AS-links		Unique IP-links		Traceroutes	
Start (line 3)	21,061,981	(100.0%)	18,101	(100.0%)	2,136,351	(100.0%)	3,340,490	(100.0%)
Bogus Filtering (line 6)	18,141,520	(86.1%)	18,101	(100.0%)	2,098,527	(98.2%)	3,303,531	(98.9%)
Geographical Filtering (line 9)	4,400,608	(20.9%)	4,729	(26.1%)	72,645	(3.4%)	1,989,098	(59.5%)
AS Filtering (line 12)	986,560	(4.7%)	4,254	(23.6%)	19,939	(0.9%)	831,878	(24.9%)
ORG Filtering (line 15)	985,849	(4.7%)	4,248	(23.5%)	19,679	(0.9%)	831,479	(24.9%)
IXP Filtering (line 18)	438,391	(2.0%)	3,010	(16.6%)	10,036	(0.5%)	407,238	(12.2%)

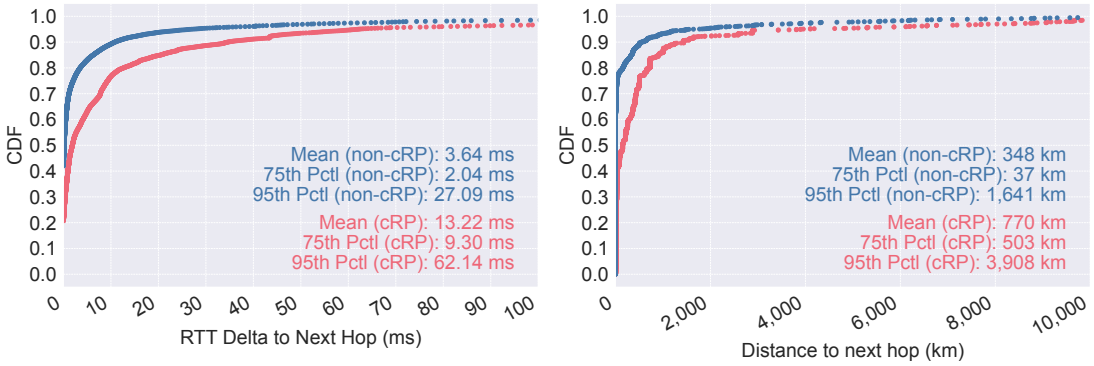

Fig. 2. CDF of RTT deltas (left) and geographical inter-hop distances (right) for cRP (red) and non-cRP (blue) links.

that RP on average incur higher latencies. This increase can mostly be attributed to the long-tail distribution, where we find 62.14 ms at the 95th percentile for cRP links compared to 27.09 ms for the non-cRP links.

As for the inter-hop distances, derived from IP-level geolocation (see Step 5), we observe a similar trend as with the RTT deltas with a significant increase in mean distances for cRP links that is especially visible in the long-tail distribution.

Interestingly, we observe that the majority of cRP links travel short distances and incur small latency increases, indicating that RP most often connects ASes in close proximity to the IXP facility. However, the differences in the long-tail show that there are a significant number of RP links connecting geographically distant ASes that incur high latency overhead.

Suboptimal Site Reachability To quantify the performance impact of RP on anycast routing, we plot the distance difference between the penultimate hop (p-hop) and the nearest anycast site (as inferred by the LACeS census [18]). We use the p-hop as a proxy for the physical location of the anycast site reached as done by [26], [27]. Although latency and packet loss are important indicators, packet loss cannot be measured with Ark at scale, and latency alone is unreliable [11], [14]; we therefore define sub-optimality by distance to the nearest anycast site, reflecting the core principle of locality in anycast.

Figure 3 presents the cumulative distribution of optimal and sub-optimal traceroutes for all traceroutes, traceroutes encountering cRP links (cRP traceroutes), and those that do not (non-cRP traceroutes). First, we observe that for roughly 40% of cRP traceroutes and 50% of non-CRP traceroutes, the p-hop is closer than the reached anycast site. This is unsurprising, as it indicates the VP initiating the traceroute reached the nearest anycast site, and the p-hop is in between. Overall, for ∼70% of all traceroutes and ∼63% of cRP traceroutes, the VP reaches a sub-optimal anycast site.

It is important to emphasize though, that RP is not the sole cause of suboptimal site selection; factors such as routing policies, traffic engineering, or peering constraints also play a role. Nevertheless, our results demonstrate that RP consistently amplifies the risk of detours and locality violations, making it a distinct contributor to suboptimal routing.

Inflation in Path Geometry A client reaching an "optimal" site does not guarantee good routing, *e.g.*, a client may take a long path to reach a nearby anycast site. To further understand how RP affects routing efficiency, we analyze how far traffic actually travels compared to how far it needs to go. Specifically, we compare the total path length of traceroutes (*i.e.*, the sum of distances between each hop) to the straight-line (geodesic) distance from the VP to the p-hop. Whilst this analysis has limitations, *i.e.*, the path length is a lower-bound and it assumes a straight-line between the VP and the p-hop, it gives an indication as to the detour that clients travel to reach their destination. Figure 4 visualizes this comparison.

The x-axis shows the direct distance to the p-hop (proxy for the anycast site reached), while the y-axis shows the measured path length. The green, orange, and red lines indicate a detour ratio of 1 (*i.e.*, the distance traveled is equal to the distance to the p-hop), 2, and 4, respectively. Next, we visualize non-cRP traceroutes using orange markers and cRP traceroutes using blue markers. Overall, we find that most traceroutes are near the green line with a detour ratio close to 1, indicating a straight path to reach the destination. However, for the traceroutes that take a large detour (above the orange and

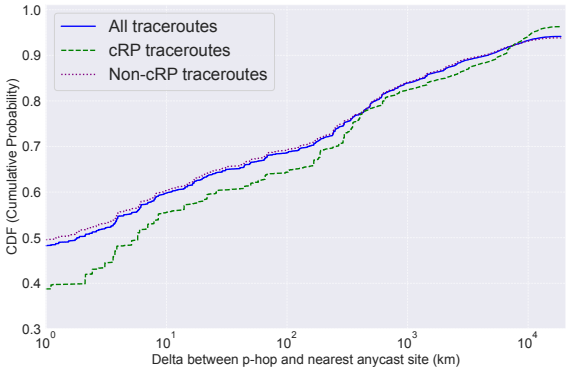

Fig. 3. CDF for the distance in km (logarithmic scale) between the nearest anycast site and the actual anycast site reached (based on the p-hop location).

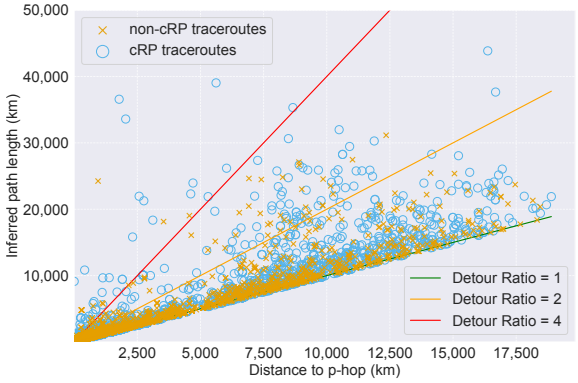

Fig. 4. Relationship between path length and distance to the p-hop for cRP (orange) and non-CRP (blue). The green line indicates direct routing (detour ratio = 1), with orange and red lines showing detour ratios of 2 and 4.

TABLE II
TOP 10 MOST FREQUENT ASNS IN CRP LINKS

ASN	AS name	Occurrence	Unique IP-links
11537	INTERNET2	91,554	198
15169	GOOGLE	83,197	194
13335	CLOUDFLARE	82,866	300
11164	INTERNET2-I2PX	73,349	90
6939	HURRICANE ELECTRIC	19,154	924
680	DFN	18,855	67
2603	NORDUNET	17,100	231
3491	PCCW	15,883	743
766	REDIRIS	14,270	21
395400	UNIVERSITY OF GUAM	12,854	36

TABLE III
TOP 10 MOST FREQUENT REGIONAL LINKS FOR UNIQUE CRP IP-LINKS

cRP start region	cRP end region	Unique IP-links
Northern America	Northern America	2,727
Western Europe	Western Europe	1,024
Northern Europe	Northern Europe	420
Northern Europe	Western Europe	407
Eastern Asia	Eastern Asia	397
Eastern Asia	South-eastern Asia	303
Western Europe	Eastern Europe	177
Southern Asia	South-eastern Asia	170
Western Europe	Northern America	159
Northern America	Eastern Asia	158

red lines), we observe a clear bias toward cRP traceroutes with some outliers near the 40,000km line (for reference the circumference of the earth is 40,075 km). Investigating these outliers, we find they mostly belong to Microsoft (AS 8075) where we observe routing loops sending traffic back-and-forth between continents reaching RTT values up to tens of seconds.

These findings demonstrate that RP impacts both anycast site selection and path geometry, with the detour ratio quantifying routing inefficiencies beyond mere destination choice.

B. Control plane Analysis

AS-level Participation in Remote Peering To better understand who participates in RP, we analyze the top individual ASes observed in cRP links. Table II lists the top ten ASNs by frequency of appearance and IP-level diversity. Notably, ASes such as AS 11537 (INTERNET2), AS 15169 (GOOGLE), and AS 13335 (CLOUDFLARE) appear in tens of thousands of links. These ASes may be connected to remote peers or themselves leveraging RP to optimize reachability. Their centrality in the control plane highlights the operational influence of a few large players in shaping the RP landscape.

Regional Characteristics of Remote Peering Links We classify cRP links based on regional locations using the United Nations geoscheme [1]. This regional grouping is supported by recent work showing that anycast ASes often restrict prefix announcements to nearby ASes to maintain locality [28].

Table III lists the ten most frequent regional pairings for cRP links. We color inter-regional as yellow, and inter-continental links as red. Strikingly, the majority of these links occur within the same sub-region, suggesting that RP is not limited to long-haul intercontinental peering but is also deeply embedded within regional ecosystems. This counters the naive assumption that RP always connects distant geographies. However, a closer look reveals a subtler inefficiency: many of these ”intra-region” RP links still result in geographically long paths (*e.g.*, a link within ‘Northern America‘ may span thousands of kilometers from West to East coast).

Latency-Weighted View of RP Intensity RP link frequency refers to the number of times a specific RP link is observed across our dataset. Figure 5 provides a latency-weighted heatmap of RP link frequency, where each cell represents a unique AS pair observed in the final cRP link set. The color intensity encodes the product of the mean RTT and the number of observed RP links between the two ASes. Darker cells indicate link pairs with low frequency or low latency, while brighter (yellow) cells highlight pairs that are both frequent and costly in latency, thus posing a higher risk for performance degradation. Importantly, blank cells marked with ‘×’ represent AS pairs with no candidate RP links in our dataset due to absence of peering relationships.

Among the brightest entries, certain ASes consistently stand

[1]https://unstats.un.org/unsd/methodology/m49/overview

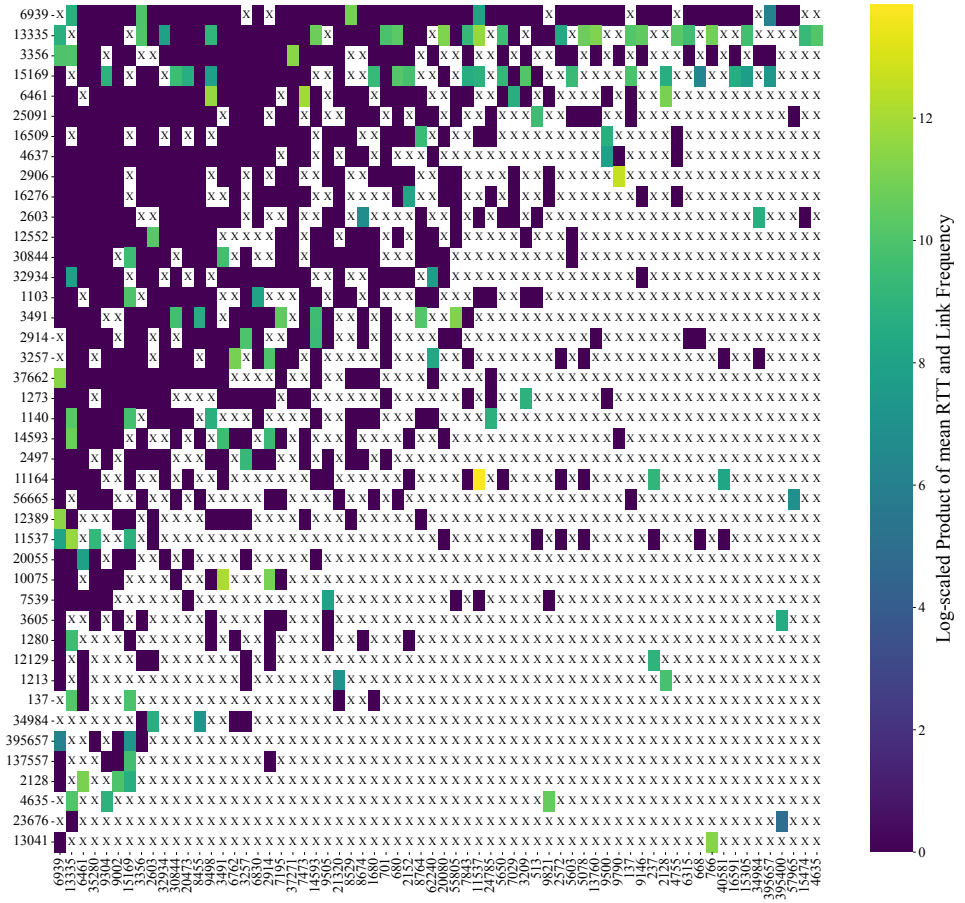

Fig. 5. Latency-weighted heatmap of RP intensity. Each cell represents an AS pair, with color intensity proportional to the product of mean RTT and link frequency. Brighter cells highlight AS relationships that are both frequent and high-latency, indicating potential contributors to routing inefficiencies.

out for their role in high-latency remote peering. Notably, AS 13335 (Cloudflare) and AS 15169 (Google) are associated with some of the most intense hotspots. The RP links that these ASes participate in, exhibit both high frequency and high latency, often resulting in traffic being steered to distant sites instead of local replicas. This pattern is especially problematic in anycast deployments, where proximity is critical. Another frequently highlighted participant is AS 11537 (INTERNET2), which appears in several high-intensity RP relationships, particularly with Cloudflare and its own sibling AS 11164. These links contribute to consistent detours, likely due to transcontinental paths that obscure proximity.

Our analysis offers network operators a clearer view of where attention is most needed. Rather than auditing every remote link, operators can use this insight to prioritize interconnection with high-impact ASes. In particular, those ASes that consistently appear in bright regions should be considered *candidates for policy tuning or routing adjustments*. Approaches such as BGP community tagging [11], local-pref tuning, or selective peer de-preferencing [8] can help steer traffic toward more geographically appropriate paths. These are not abstract concerns: the latency costs associated with these ASes are **real, measurable, and avoidable with informed engineering decisions**.

V. Limitations

This section outlines the main limitations and challenges in analyzing RP, highlighting complexities that must be addressed to enable accurate inference and meaningful classification. RP is a multifaceted concept that spans operational, topological, and policy dimensions within the IXP ecosystem [10], [14], [16]. For instance, Mazzola et al. [10] observe that some IXP and reseller representatives define RP by connection type (*e.g.*, via shared ports), while others emphasize geographic distance. Similarly, Nomikos et al. [14] note the lack of tools to classify peers as remote or local and introduce notions such as local-indirect peering and wide-area IXPs. These perspectives underscore the limitations of binary classifications and the need for a multidimensional taxonomy.

Additionally, RP inference often relies on auxiliary sources such as PeeringDB [24], Euro-IX [29], and the Packet Clearing House (PCH) [30]. Although these datasets offer valuable insights, they are self-reported, inconsistently maintained, and regionally inconsistent in coverage [11]. In some IXPs more than 70% of member interfaces cannot be confidently labeled due to missing or conflicting data [10]. This gap hampers not

only automated inference accuracy but also reproducibility and cross-study comparability.

Finally, resellers are essential actors in enabling RP but operate with little transparency. A single physical port at an IXP may be shared among multiple remote ASes via VLANs or other L2 multiplexing mechanisms, none of which are observable from the control or data plane. This makes it virtually impossible to infer the true physical or topological connection of a peer using BGP data or latency measurements [14].

VI. Conclusion

In this work, we revisited RP through the lens of anycast routing, introducing a geographic mapping approach that enables systematic assessment of its operational impact on latency-critical services. Our findings highlight the underappreciated costs of RP for latency-sensitive services and underscore the need for more transparent, locality-aware policies.

With prior analysis platforms no longer maintained, a key step forward is building sustainable, continuously updated systems that combine control- and data-plane measurements, IXP membership data, and crowd-sourced tagging of RP practices. We also plan to investigate standardization mechanisms, e.g., RP labeling via BGP communities or IXP route server metadata, to improve transparency and auditability. Finally, as hyperscalers expand their footprint, we aim to study how RP interacts with cloud-centric routing strategies, where cost and redundancy often come at the expense of proximity.

Acknowledgements

This research received funding from the Dutch Research Council (NWO) under the projects INTERSECT (NWA.1160.18.301), UPIN, CATRIN (NWA.1215.18.003), and was partially supported by SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union - NextGenerationEU.

References

[1] Y. Rekhter, S. Hares, and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, Jan. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4271

[2] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large european ixp," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, 2012, pp. 163–174.

[3] B. Augustin, B. Krishnamurthy, and W. Willinger, "Ixps: mapped?" in *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement*, 2009, pp. 336–349.

[4] N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann, "On the benefits of using a large ixp as an internet vantage point," in *Proceedings of the 2013 conference on Internet measurement conference*, 2013, pp. 333–346.

[5] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, "There is more to ixps than meets the eye," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 5, pp. 19–28, 2013.

[6] T. Böttger, G. Antichi, E. L. Fernandes, R. di Lallo, M. Bruyere, S. Uhlig, and I. Castro, "The elusive internet flattening: 10 years of ixp growth," *CoRR*, 2018.

[7] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse?" in *Passive and Active Network Measurement: 9th International Conference, PAM 2008, Cleveland, OH, USA, April 29-30, 2008. Proceedings 9*. Springer, pp. 1–10.

[8] S. Kastanakis, V. Giotsas, I. Livadariu, and N. Suri, "Replication: 20 years of inferring interdomain routing policies," in *PROCEEDINGS of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 16–29.

[9] W. B. Norton, "The great remote peering debate," https://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2012/9/18_The_Great_Remote_Peering_Debate.html, 2012, accessed: 2025-06-02.

[10] F. Mazzola, P. Marcos, I. Castro, M. Luckie, and M. Barcellos, "On the latency impact of remote peering," in *International Conference on Passive and Active Network Measurement*, 2022, pp. 367–392.

[11] F. Mazzola, A. Setti, P. Marcos, and M. Barcellos, "Analyzing remote peering deployment and its implications for internet routing," *IEEE/ACM Transactions on Networking*, 2024.

[12] R. Bian, S. Hao, H. Wang, A. Dhamdere, A. Dainotti, and C. Cotton, "Towards passive analysis of anycast in global routing: Unintended impact of remote peering," *ACM SIGCOMM Computer Communication Review*, vol. 49, no. 3, pp. 18–25, 2019.

[13] K. E. Lindqvist and J. Abley, "Operation of Anycast Services," RFC 4786, 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4786

[14] G. Nomikos, V. Kotronis, P. Sermpezis, P. Gigis, L. Manassakis, C. Dietzel, S. Konstantaras, X. Dimitropoulos, and V. Giotsas, "O peer, where art thou? uncovering remote peering interconnections at ixps," in *Proceedings of the Internet Measurement Conference 2018*, pp. 265–278.

[15] "Release of code and data." https://rhendriks.github.io/anycast-rp/, 2025.

[16] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois, "Remote peering: More peering without internet flattening," in *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*, 2014, pp. 185–198.

[17] CAIDA, "The caida archipelago (ark) measurement infrastructure," https://www.caida.org/projects/ark/, 2024, accessed: 2025-06-21.

[18] R. Hendriks, M. Luckie, M. Jonker, R. Sommese, and R. van Rijswijk-Deij, "Laces: an open, fast, responsible and efficient longitudinal anycast census system," 2025. [Online]. Available: https://arxiv.org/abs/2503.20554

[19] D. Cicalese, D. Joumblatt, D. Rossi, M.-O. Buob, J. Augé, and T. Friedman, "A fistful of pings: Accurate and lightweight anycast enumeration and geolocation," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 2776–2784.

[20] T. Cymru, "Team cymru bogon reference," https://www.team-cymru.com/bogon-reference-http, accessed: 2025-06-20.

[21] M. Luckie, B. Huffaker, A. Marder, Z. Bischof, M. Fletcher, and K. Claffy, "Learning to extract geographic information from internet router hostnames," in *Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies*, 2021, pp. 440–453.

[22] IPInfo, "Ipinfo: Ip address geolocation and network intelligence api," https://ipinfo.io, 2025, accessed: 2025-06-20.

[23] S. D. CAIDA, University of California, "Routeviews prefix to as mappings dataset (pfx2as)," https://catalog.caida.org/dataset/routeviews_prefix2as, accessed: 2025-06-20.

[24] "PeeringDB," https://www.peeringdb.com, 2025.

[25] CAIDA, "Caida as-to-organization dataset," https://publicdata.caida.org/datasets/as-organizations/, 2025, accessed: 2025-06-20.

[26] R. Hendriks, T. Betzer, B. Du, R. Sommese, M. Jonker, and R. van Rijswijk-Deij, "Locating and enumerating anycast: a comparison of two approaches," in *Proceedings of the 2025 Applied Networking Research Workshop*, ser. ANRW '25. ACM, 2025, p. 99–105.

[27] M. Zhou, X. Zhang, S. Hao, X. Yang, J. Zheng, G. Chen, and W. Dou, "Regional ip anycast: deployments, performance, and potentials," in *Proceedings of the ACM SIGCOMM 2023 Conference*, pp. 917–931.

[28] S. Kastanakis, V. Giotsas, I. Livadariu, and N. Suri, "Investigating location-aware advertisements in anycast ip networks," in *Proceedings of the 2024 Applied Networking Research Workshop*, ser. ANRW '24. ACM, 2024, p. 15–22.

[29] Euro-IX, "Ixp database (ixpdb)," https://ixpdb.euro-ix.net/en/, 2024, accessed: 2025-06-07.

[30] Packet Clearing House, "Packet clearing house (pch)," https://www.pch.net/, 2025, accessed: 2025-06-07.