Meta-Peering: Automating ISP Peering Decision Process

Md Ibrahim Ibne Alam[1], Anindo Mahmood[2], Prasun K. Dey[3], Murat Yuksel[2], and Koushik Kar[1]
[1]Rensselaer Polytechnic Institute, Troy, NY
[2]University of Central Florida, Orlando, FL
[3]The MathWorks, Inc., Natick, MA

***Abstract*—Peering between Internet Service Providers (ISPs) is playing an increasingly critical role in Internet traffic exchange. As content delivery networks continue to expand, major content ISPs are increasingly opting for peering arrangements over transit services to facilitate faster exchange of traffic. The satisfaction of the ISP pair and the longevity of the peering arrangement depend on the stability and performance of these peering relationships. We introduce *meta-peering*, a term which refers to the set of tools needed to help and automate the ISP peering process – starting with identifying a list of ISPs that are likely to peer, writing router rules to establish BGP sessions with them, and extending the service to monitor all these sessions for notifying any major outages or peering agreement violations.**

In this paper, we first make a thorough analysis of recent trends in ISP peering and describe how *meta-peering* can be implemented by integrating some of the existing tools. We mainly focus on instrumenting the automation of the peer selection process with an aim to identifying potential peering partners and peering locations to exchange traffic. Using these direct peering links greatly reduces energy consumption as traffic takes much shorter paths to their destinations, going through reduced number of intermediary devices (e.g., routers, switches) compared to elongated transit routes, consequently reducing the environmental impact. Utilizing PeeringDB and CAIDA datasets to identify possible peering points for ISP pairs, we consider ISPs' internal policies to generate a list of acceptable peering contracts (APCs). We design two methodologies to rank order each ISP in the APC list and offer guidance on which ones would be stable and beneficial for the potential peers. A study of more than 3,000 ISP pairs (mostly active in North America) shows that our peer selection methods can attain around 80% accuracy in predicting peering relations.

***Index Terms*—Peering; Internet Service Provider; Internet eXchange Point; Network Management; Traffic Engineering.**

I. Introduction

WITH almost 120K Autonomous Systems (ASes) [3] around the world, it is nearly impossible for an individual AS to connect with all others by establishing unique physical fiber cable to ensure its global reachability. Organizations capable of maintaining enormous network backbone are extremely rare [4], restrictive in nature, and require other ASes to meet stringent conditions before committing to connect [5]. In essence, Internet Service Providers (ISPs) have to collaborate and establish interconnections with each other for global connectivity. Interconnections between ISPs depend on multiple aspects such as the size of their ASes, geo-coverage, traffic-offload costs, and user count, namely the 'customer cone'. A smaller ISP (based on its customer cone and/or market capital), generally purchases *transit* service from a larger provider. An ISP can also choose to exchange its traffic directly with another ISP in a settlement-free manner. This type of agreement is known as *peering* and it is typically a "sender keeps it all" deal. This model allows the sender ISP to keep all the money it charges from its end-customers, and hand the traffic over to its peer; while the peer has to carry the traffic towards the destination without charging the origin ISP.

Peering is often preferred over transit for better control on routing, low latency, and most importantly, slashing cost. Since peering traffic does not have to traverse a transit ISP's network and can reach the destination AS directly, the propagation delays for peering paths are generally smaller. For instance, by leveraging peering paths, 68% of ASes connected to 920 access ISPs experienced 10 ms improvements in latency, and for 91% of those ASes, peering paths outperform transits [6]. Lack of peering causes extraneous traffic detours and often results in an increased path stretch, like what African or Latin American ISPs are dealing with [7], [8]. Such fragmented routing causes local traffic to unnecessarily traverse other continents (in the above cases, Europe and US respectively) and degrades the end-to-end performance. This is sometimes true for the North American ISPs as well. A study [9] found that mobile client traffic from AT&T Seattle enters Google's network in Bay Area due to absence of a peering point in close proximity. The impact is particularly significant in the case of content delivery networks (CDNs) whose traffic dominate the modern Internet. Peering allows CDNs to cache content closer to the eyeball network meaning the traffic has to travel significantly smaller routes to the end-users. Using transit routes would mean each time a content is requested, that traffic will need to flow through the entire path stretching from the content provider to the end-user significantly increasing latency as well as the load on the networking infrastructure, which will greatly increase the energy consumption [10], [11]. Despite its reduction of path lengths, settlement-free peering does not always mean traffic

This manuscript is based on preliminary work published in [1] and [2]. The key additions are: 1) The GEO-PP method is extended to include another metric termed similarity score; 2) The performance analysis now includes two much larger datasets of ASNs; 3) Performance analysis in terms of internal routing cost is added; 4) Impact of Machine Learning on peering prediction with the determined metrics as input is studied; 5) An extensive discussion of the evolution of ISP peering is added.

Dr. Prasun K. Dey was with University of Central Florida during most of this work.

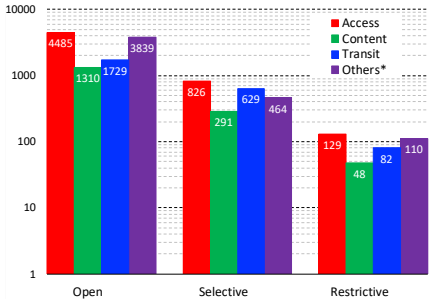

Fig. 1: ISPs' peering policies in PeeringDB (* = Educational, Non-profit, Enterprise, etc.)

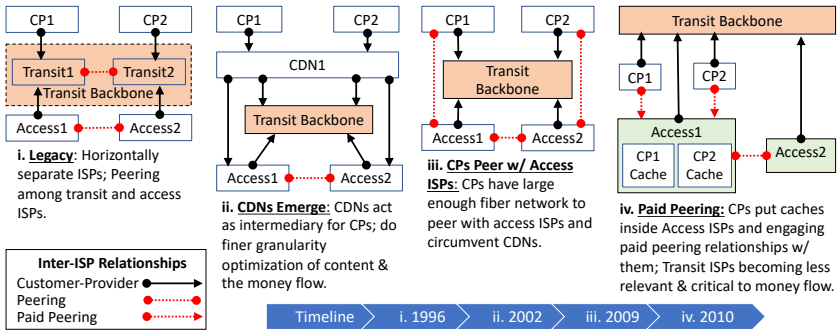

Fig. 2: Peering evolution [5], [14]–[17]

delegation at no cost [12]. However, depending on the transit fee and the total cost involved, narrowed down to per unit, an 1 Gbps peering connection can be cheaper than transit when exchanging at least 100 Mbps traffic [13].

Integrating peering with inter-AS policy routing decisions is a complex process. Whether peering or transit, an ISP has to establish new Border Gateway Protocol (BGP) sessions with each of its neighbor ISPs and monitor them. Blending BGP with Interior Gateway Protocol (IGP) without conflicting the external and internal policies is quite complex. Errors due to human involvement increase the chance of network outages, and thus maintaining these sessions is time consuming and cumbersome. To ease the process by reducing human coordination delay, ISPs have already implemented tools for automating the BGP peering process establishment [18], [19]. Existing tools (e.g., UnivMon [20] , sFlow, and IPFIX) can also perform network monitoring [21] and identify the routes causing outage so that stratagems like manual reconfiguration or temporary shutdown can be initiated.

Although BGP peering process has been notably automated, choosing the right ISP peer and Point-of-Presences (PoPs) is more challenging to automate due to the following reasons:

- 99% of peering is handshake [22]. No fixed written rules.
- Although modern switches (e.g., BIG-IP 2000S) offer better management and application performance by selecting the best route for both in- and out-bound traffic [23], peering with multiple ISPs is a hassle.
- To minimize "bit miles" [24], ISPs end up choosing suboptimal PoPs in peering deals.
- ASes follow different infrastructure-specific peering policies. Based on PeeringDB data of 15,078 ASes (Fig. 1), most ISPs are *open* to peer, while very few are *restrictive*. Finding ISPs willing to peer may be easy, but motivating ISPs from the Selective and Restrictive groups is difficult.

Considering these issues and the recent automation efforts, potential peers need to be identified based on the estimated traffic, customer cone size, peering policy, and cost of peering at a PoP. Since one of the key motives behind peering is to reduce cost, there has been extensive research on game-theoretic modeling of peering [25], and understanding the economics behind pricing where multiple ISPs are involved [26]. However, the topic of automated peer selection on a global scale using (publicly) available data has not received enough attention. The few works that address issues broadly related to automated peer selection only consider some specific types of peering (i.e., remote peering), certain ISP types (access-content), or are confined to a local area [27]–[29]. In this paper, we aim to address this gap and focus on answering a key question in the Internet peering: "To what extent can the selection of peers be automated?"

Partial automation of ISP peering is already being taken up by industry leaders. For example, Meta has deployed an online peering automation framework [30]. Though Meta's network is an outlier as it is quite large and, hence, sturdy against potential harms from mistakingly admitting peering requests, it clearly shows the appetite for more automation in peering. We envision the ISP peering process between a *requester* (an ISP that initiates the peering process) and a *candidate* (an ISP which satisfies *requester*'s requirements) to be eventually entirely automated, where the system suggests a list of potential *candidate* peers for a *requester* ISP, identifies feasible PoP locations, and if both parties agree, automatically generates BGP configuration. We consider the entirety of every tool, algorithm and other necessary components needed for peering automation as *Meta-Peering*. This requires several major innovations, such as tools to help ISP administrators make efficient peering decisions, negotiation protocols for accommodating peering strategies and policies, standardization and systematization of resolving intra- and inter-ISP routing policy conflicts with peering decisions, and defence against attacks to the automated peering process. In this work, we focus on the first of many steps along this long but important road: *establishing a quantifiable system for selecting potential peer ISPs*. Major contributions of our work include:

- detailing the *meta-peering* concept and breaking down the peering process into four phases;
- an approximate timeline of the evolution of peering;
- identifying the most frequently asked peering requirements, ISPs' PoP frequencies, and peering points;
- optimization problem formulation for selecting the best peering deals with another ISP;
- methodology to estimate the peer ISP's traffic amount;
- introduction of a new metric called *felicity score* for a pair of ISPs to quantify their peering possibility; and
- a publicly available web application[1] for access to recommended peering deals generated using our algorithms.

The rest of the paper is organized as follows: Section II discusses the peering evolution and discusses our motivation

[1]Meta-peering website: http://metapeering.net.

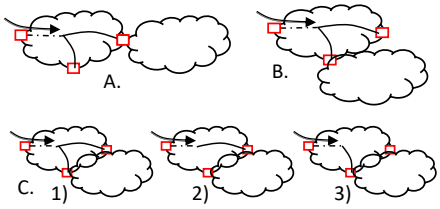

Fig. 3: Possible PoP locations

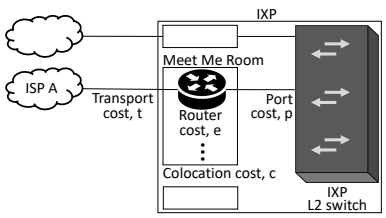

Fig. 4: Peering at IXP and related costs

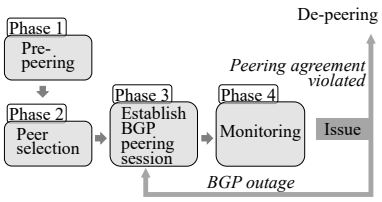

Fig. 5: Peering phases

for automated peering. Section III details *meta-peering* and explains the peering phases. Section IV expounds the framework with related terminology and develops the ISP network model. Section V introduces and describes two meta-peering methods, namely *GEOgraphic Peering Prediction* (GEO-PP) and *Routing Cost Aware Peering Prediction* (RCA-PP). Numerical evaluations are presented in Section VI, followed by summary in Section VII. Also, to aid in the flow of reading, Table I provides some of the common acronyms used throughout the paper.

TABLE I: Summary of Commonly Used Acronyms

Notation	*Description*
ISP	Internet Service Provider
ASN	Autonomous System Number
CP	Content Provider
IXP	Internet Exchange Point
PoP	Point of Presence (usually an IXP)
PPP	Possible Peering Points
PPC	Possible Peering Contracts
APC	Acceptable Peering Contracts
TM	Traffic Matrix
GEO-PP	GEOgraphic Peering Prediction
RCA-PP	Routing Cost Aware Peering Prediction

II. Peering Background and Motivation

ISPs have the option to establish bi-lateral peering relations at either private or commercial Internet eXchange Points (IXPs). They can install dedicated physical links which allows them to exchange a higher volume of traffic [31]. Similarly, ISPs can also form multi-lateral peering relations, commonly known as public peering, with ISPs in a shared environment within an IXP. ISPs can connect to a central route-server (RS) where multiple peering relations are entertained simultaneously. In either case, ISPs carry their own traffic to a common PoP [13]. In this paper, our main focus is on bi-lateral agreements in such public peering settings.

A. History of Peering: An Evolutionary Tale

Fig. 2 portrays the decades long evolution in the principles practiced among peering entities. In the legacy model, content providers (CPs) and access ISPs were horizontally separated and dependent on transit service from upstream providers who peered only with ISPs at a similar tier. Around 2002, CPs began to rely on content delivery networks (CDNs) [14] and over time, CPs established their own large fiber networks. During 2009 [15], they started bypassing both transit providers and CDNs to peer directly with access ISPs following a "Donut Peering" model. Traffic ratio between CPs and access ISPs was noticeably uneven, and for CPs, it was more beneficial to put caches directly inside access ISPs [16]. Generic peering policy was not enough, and access ISPs introduced the term "Paid Peering" to charge CPs [17]. Arguably, peering is a key ingredient for the Internet ecosystem, and having a more automated peering process supported by meta-peering tools will allow ISPs to evolve.

B. Why Automated Peering?

Peering is a handshake, but finding the right ISP is difficult. To ease the process, network admins typically meet other ISPs' representatives in-person at informal events such as NANOG, Global Peering Forum, or in CEE Peering Days [32]. After discussing the traffic volume to see if peering would generate enough savings for both the entities, agreements and peering policies are negotiated. Sometimes, conflicts happen due to lack of prior knowledge about other ISP's traffic amount or strategic business tussles. This may lead to disputes or de-peering, e.g., Cogent de-peered Level3 (2003), AOL (2002), and Telia (2008) due to imbalanced traffic ratio, and Sprint (2008) for not respecting the exchange criteria [33].

Deficiency in number of peer selection algorithm case-studies implies less preparation of ISP admins for the opponent ISP traffic behavior. An ISP has the complete knowledge of its own traffic matrix and router-level network topology. But, it can only have a rough idea about which ISP is sending the maximum traffic towards itself and attracts most of its customer traffic. The key question is, whether they can estimate each others' traffic after peering? To avoid future disputes, ISPs typically undergo a temporary "trial peering" periods of several weeks to determine the exchanged traffic amount before provisioning the long-term peering session [34].

Optimal peer selection is a hard problem. Game-theoretic approaches focus mostly on economic analysis by considering both routing and congestion cost [25] to study the capacity and pricing decisions made by service providers [35]. Earlier works [36], [37] focused on formulating an optimal peering problem to determine the maximum peering points along with their strategic placement or a negotiation-based platform for ISPs to jointly determine routing path for traffic exchange. While goal of these studies was to minimize the interconnection cost without compromising the service quality and understand the Internet-wide negotiation mechanisms, our work differs from earlier research as we focus primarily on automating the peer selection process and suggesting possible PoPs according to pre-defined ISP specific criteria.

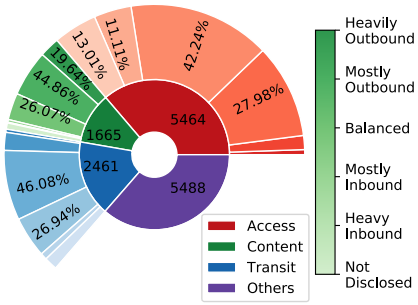

Fig. 6: Inbound vs. outbound traffic ratio from PeeringDB

	HI	MI	B	MO	HO
HI	0.8	2.3	2.6	1.1	0.4
MI	2.9	8.7	11.4	4.6	1.5
B	3.4	14.1	19.2	7.9	2.4
MO	1.0	3.7	5.3	2.2	0.6
HO	0.4	1.2	1.6	0.6	0.1

TABLE II: Peering pairs traffic ratio type percentage according to CAIDA

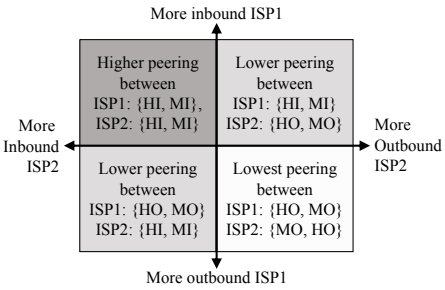

Fig. 7: Peering pairs traffic ratio (based on Table II)

III. Peering Process and Current Scenario

Fig. 3 shows three combinations of PoP locations for two ISPs where they can peer. *A* and *B* represents when ISPs are not located in the same PoP, but willing to peer and essentially agree on the closest place from both. There is only one way of peering in such cases. For *C*, ISPs overlap and there are at least two common PoPs between them, so they can either exchange traffic at all of them (*case 1*) or at only one location (*cases 2* and *3*). ISPs can also peer without being physically present in an IXP and/or connect through a third-party reseller through *Remote Peering* (RP) [29]. Despite being an option in practice, we do not consider RP in our model as they are opaque and controversial in terms of its performance benefits.

As mentioned earlier, an IXP hosts ISPs inside a physical location (PoP) equipped with multiple Ethernet switches (to perform layer 2 functionality) so that they can route traffic among themselves. An ISP needs to bring its traffic to the PoP, purchase port capacity from IXP to connect its router to the interfaces of 1/10/100 GbE switch, and pay for the colocation costs (e.g., electricity bill, cooling fee, security and others) [38]. The IXP allocates Meet-Me Room (MMR) for each of the participating ISPs where it can land its fiber and house all of its routers to interconnect with other ISPs. Fig. 4 portraits a high-level representative architectural overview of an IXP and the associated costs that an ISP has to pay. An IXP also hosts direct connectivity between two ISPs.

Inspired by Norton's *Peering Playbook* [5], we break-down the entire peering process into four phases (see Fig. 5) and restrict our focus specifically on the automation effort undertaken in each phases. As mentioned in Section I, *meta-peering* encompasses everything including the tools used by ISPs, algorithms developed by researchers in academia and associated companies who are contributing to this area and tackling the issues related to automation.

A. Pre-Peering Phase: Key Peering Metrics

From a purely economic perspective, if the peering cost (including the connectivity and the maintenance cost) is less than the transit cost, both the *requester* and the *candidate* ISPs will be interested to form the peering relationship.

More Control: Regardless of economic benefits, an ISP may be interested in peering with more ISPs to gain control over its traffic and influence route path selection instead of letting someone else (upper transit provider) to treat it as hot-potato. Thus, a requester ISP always looks for such candidate ISPs to peer who has significant presence in some other area and is willing to deliver requester ISPs' traffic there while keeping their individual traffic local. Also to avoid the "tromboning" effect and to reduce the latency, requester ISP prefers peering over transit. Each time an ISP peers with someone new, the congestion reduces, reliability increases, and therefore, the end-to-end service quality for the users is improved [39].

Traffic Ratio: A useful metric to consider for identifying potential peers is the balance of inbound vs. outbound traffic for an ISP. Fig. 6 presents an interesting overview of the traffic ratios of all the ISPs present in PeeringDB. CPs are mainly interested in producing and disseminating the content, as a result, 90% of them are outbound or balanced in nature; while access ISPs care about the end-user connectivity, thus, are mostly (82% of them) balanced to heavily inbound. On the other hand, transit ISPs serve in the middle to connect CPs and access ISPs. Almost half of the transit ISPs' traffic ratio is balanced. Using this information, a requester ISP, depending on its business strategies, may identify a peering candidate.

Based on CAIDA's inter-ISP relationship information, Table II shows the percentage of peering ISP pairs that have different traffic ratio type. Balanced ISPs are the most popular for peering and the maximum number of peering happens if both of them are balanced (19.149%). If either of them is balanced, there is a good chance of them being involved in peering. We populate a peering possibility quadrant in Fig. 7 based on ISPs traffic nature. Here, ISP1 (or ISP2) can be either *requester* or *candidate*, it does not impact the conjecture.

PoP Frequency: Having more PoPs attracts more ISPs which are interested in expanding their network connectivity. Therefore, PoP frequency is also a key metric for identifying potential peers. Fig. 8a shows the CDF of ISPs' PoP frequency. Access ISPs' target is the end customers and they operate in specific regions, which is why they have fewer PoPs. Transit ISPs form the backbone of the Internet: they lay out fiber across the country and establish PoPs in different locations to provide transit services to other ISPs. They usually have more PoPs then any other ISPs. For instance, the maximum number of PoPs for a particular ISP from access, content, and transit categories are 59, 324, and 380, respectively. However, one notable observation from our analysis is that by putting caches directly inside an access ISP or purchasing racks from datacenters, CPs are spreading their footprint and increasing the number of PoPs. In addition to CP, transit and access ISPs, some institutions also identify themselves as as educational,

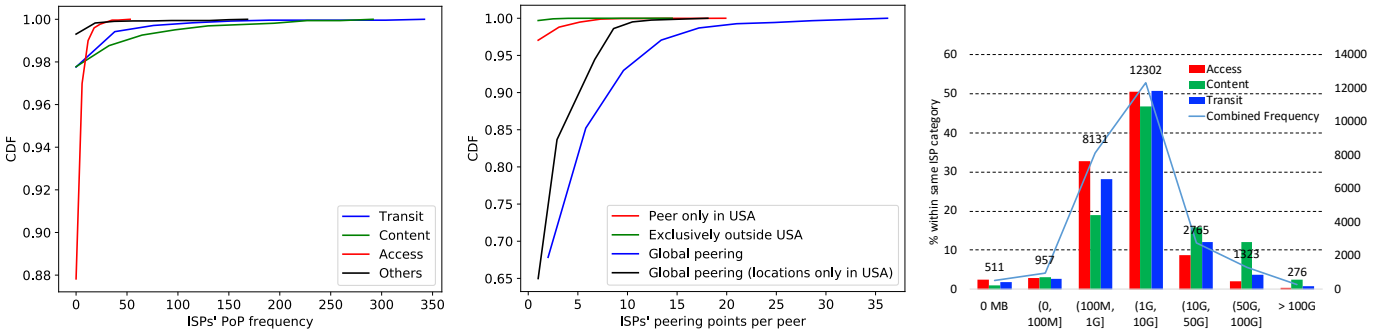

(a) CDF of ISPs' PoP frequency (b) CDF of ISPs' peering points per peer (c) Port Capacity of ISPs*

Fig. 8: ISPs' PoP analysis according to PeeringDB (* one ISP may have multiple ports [40] with different speeds).

government or similar non-profit organizations. We mark them as *others*, but exclude them from our study.

Fig. 8b shows the CDF of peering PoP frequency for ISP pairs using CAIDA AS Relationship – with geographic annotations dataset [41]. According to the figure, more than 95% of the ISPs, which peer globally, use less than 15 PoPs located inside the US for exchanging data with their peers.

Customer Cone Size: Traffic volume and advertised IP address space are vital for deciding the peering partners. A requester ISP, having a large customer cone tries to peer with ISPs with higher traffic volume and larger address space. Earlier work [42] shows a strong correlation between the advertised prefix count and traffic volume for both the access and transit ISPs except for CPs – since they do not serve the end-users directly, not much address space is needed for them. This validates our use of publicly available BGP-advertised address space for estimating network's traffic volume.

Coverage Area: The internal topology of an ISP is proprietary information, and ISPs are restrictive when it comes to publishing it. As a result, identifying the exact coverage of an ISP is difficult, and we have to rely on publicly available information about their PoP locations[2]. Large interconnection hubs (i.e., IXPs) are usually located in bigger business cities or near the trans-oceanic cable landing points [14]; thus, PoPs may reflect the business interest of an ISP. If an ISP has more PoPs in a certain area, that means its customer cone is concentrated around that region [43]. Further, to support a higher volume of customers, it increases its presence in nearest IXPs or private facilities, which ensures higher port capacity [40] for traffic exchange, and improves infrastructure [14]. On the other hand, ISPs, thriving for country-wise reachability, mostly have a presence in coast-to-coast (in the US case), or those who want to expand beyond a specific country and ensure global connectivity, they have to have a presence in different countries. To gain the benefits of more extensive coverage, ISPs peer. In this regard, geographical footprint coverage acts as one of the deciding factors to consider, and hence, a larger overlap between two ISPs will reduce the likelihood of them peering.

Overall, understanding and evaluating these metrics plays a critical role in developing robust peering strategies with regards to whom ISPs peer with and what locations do they peer in. This can vary significantly from ISP to ISP. Larger ISPs may choose to prioritize metrics like traffic ratio and coverage overlap to maintain their extensive reach and balance exchange traffic over their massive network backbones. On the other hand, smaller ISPs will lean more towards POP frequency and customer cone size to enhance local connectivity and growth potential.

[2]Router level information of ISPs are available from CAIDA, however, the locations are mostly approximate and usually converges around the PoP locations.

TABLE III: Frequently asked requirements by ISPs

Requirements	*% of all ISPs*
24*7*365 Support	83.33
Minimum traffic volume (inbound or outbound)	69.44
No static route/ default route	66.67
Do not announce third party routes. (Only self customer cone)	66.67
Consistent route announcements in all inter-connecting locations.	63.89
Minimum geographic presence (peering at least in PoP count)	58.33
Interconnection speed at each point	52.78
Provide security; handle DDoS and abuse	44.44
Traffic ratio (in-bound: out-bound)	44.44
Routes registered in IRR, RIPE, ARIN	33.33
Existing transit customer can not peer	33.33

B. Peer Selection Phase

ISPs typically outline their criteria for peering and describe the general guidelines. To sort out the most common prerequisites, we considered top 50 ASes according to CAIDA's AS-Rank [44] and utilized PeeringDB categorization. We gathered peering policies from each ISP's website and classified them as interconnection requirements, routing requirements, and general conditions. Some ISPs do not publish their policies so we could not find them while some ISPs have multiple entries in the CAIDA's top 50 AS list. Furthermore, no CPs were in the top 50, so we considered some prominent CPs varying from gaming industry (Blizzard), social media (Facebook) to CDNs (Limelight). We accumulated peering criteria for 36 ISPs from around the globe.

Table III lists the most popular requirements and shows how many ISPs ask for them. Most of the ISPs want to ensure that

the requester operates a 24*7*365 Network Operations Center (NOC) and provides escalation path for resolving networking issues. Apart from the listed items in the table, other trendy desiderata are: omnipresent geographic footprint, adequate backbone capacity, essential peering port size in PoPs (see Fig. 8c), maintaining financial stability, and redundancy of the requester network. Candidate ISP also does not peer with its direct customer or a customer of any of its peers.

Most of the candidate ISPs look for a similar sized requester ISP who runs backbone with at least 10 Gbps capacity or has at least 50% of candidate network's capacity [45]. Large ISPs like Telia, Orange, or Liberty ask for 100 Gbps backbone, while ISPs like Frontier and Blizzard do not set the amount but make sure that the requester can handle the projected load without any congestion. Some ISPs expect requester to announce an aggregated subnet size of /24 or greater (e.g., Time Warner and Zayo), some are more specific requiring /30 or above (e.g., Google), while some ISPs (e.g., Swisscom) are more reluctant and ask for address range of at least /11. Although it may vary, having a wide geographic presence is desirable from a requester ISP. Depending on the nature of the candidate, it may require at least 2 (Cox) to 8 (Charter, GTT) locations for peering in the US (Telstra requires presence in 8 Australian cities), 2 (AT&T) to 5 (Orange) countries in Europe or have 50% of the candidate's (e.g. Verizon) geo-coverage.

C. Establishing BGP Session Phase

Business relationships (direct customer/provider/peer), as well as the intention to limit the routing table size for scalability, and to gain control over in/out-bound traffic (by implementing MED or LocalPref) play vital roles in setting up a BGP session [46]. Erroneous manual configurations often lead to instability, generate excessive misconfigured route announcements, and cause unintentional blackhole routes.

Incidents like inadvertent prefix leak causing service outage for Google and Cloudflare customers in 2018 [47], or sending the entire Japan's Internet into dark for more than an hour in 2017 [48] amplifies the importance of meticulous BGP configuration. To prevent such occurrences from happening and reduce exchanges among ISP admins each time either ISP expands its geographic presence by joining in a new IXP collocation center, first-ever "peering-over github" network [19] has been introduced. Since most of the ISPs keep their information updated in PeeringDB [42], *Coloclue* leverages these information to find out the common IXPs, calculate the max-prefix and establishes BGP sessions. This is, by far, the only automation effort towards setting new or updating the existing BGP sessions between two networks.

D. Post-Peering (Monitoring) Phase

Once the BGP session is established among the neighbors, an ISP keeps monitoring all the remote BGP services with its neighbors for ensuring the least amount of BGP outage or black hole, and it compares the aggregate traffic in both direction so that the measured traffic-ratio does not violate the agreement. To automate the process, ISPs can either set up their internal monitoring system (BGPStream [49]) or can utilize external services like ThousandEyes [50], or NLNOG RING [51].

Among other BGP manipulators, Noction Intelligent Routing Platform (Noction IRP) [52] integrates intelligence to the routing decisions. By actively probing remote prefixes for packet loss, latency, throughput, and long-term reliability, Noction IRP optimizes the performance of the routes and is able to bypass congestion and outages. With constant monitoring, it can automatically alert the network admin for various types of errors instantaneously.

IV. Automating Peer Selection: Meta-Peering

The case studies in the previous section (especially in Section III-B and III-A) motivates us to look for a framework that can perform automated peer selection. Having a framework that is fit for all ISPs may be impossible to find; however, we perform a significant amount of sanity checks on different performance metrics using our framework to demonstrate the huge potential it has. In this section, we discuss the framework in general, and Section V covers the meta-peering methodologies proposed.

A. The Meta-Peering Framework

In general, ISP admins possess traffic statistics and detailed intelligence about their own network, but have limited data about their competitor ISPs while making the peering decisions. Our framework leverages the publicly available data and, based on the internal policy of the requester ISP, produces a guideline for peering contracts. The algorithm contains a heuristic function that runs for both ISPs independently and enumerates the possible options separately. Comparing these two lists, the algorithm generates a final list of peering points.

Fig. 9 presents an overview of our proposed framework. Considering the PoP locations, traffic matrices, port capacities in IXPs, or private facilities where both the requester and the candidate have their presence, the framework suggests to the network admin whether the candidate may agree with a particular peering offer or not. Utilizing this tool, the network admin can also come up with a list of potential ISPs; select the appropriate ones, and identify possible peering contract offers to these potential peers. In the figure, each gray box represents an autonomous module of the framework, which we shall discuss later. We adopt such a modular approach to support future extensions by adding newer components. The algorithm runs from one ISP's perspective and simulates the counterpart from publicly available information. Hence, the heuristic function requires limited shared data and can run independently without breaching any privacy rules.

B. Terminology

We categorize the input data for the framework as following:

Known data: Population at PoP locations, requester's own Traffic Matrix (TM), port capacity at PoPs.

Estimated data: Candidate's TM. We call it Estimated Traffic Matrix (ETM).

Outputs: Possible Peering Points (PPPs), Possible Peering Contracts (PPCs), Acceptable Peering Contracts (APCs).

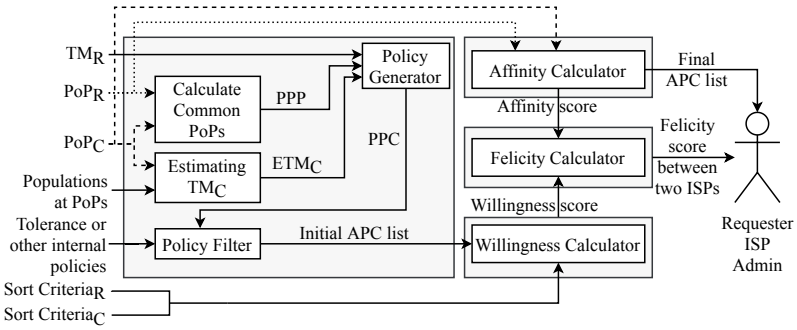

Fig. 9: Automatic peer selection framework

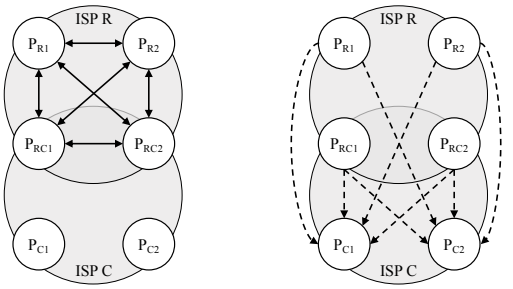

(a) From ISP R to ISP R

(b) From ISP R to ISP C

Fig. 10: Traffic exchange between ISPs

1) Possible Peering Points (PPPs): We use PeerinDB to compile the PoP list and identify common PoPs for both requester and candidate ISPs. Traffic can be exchanged at any combination of these common PoPs, except for the empty set. We call the list of all PoP combinations as Possible Peering Points (PPPs). Let us denote the total number of common PoP locations ($|PPP|$) as r.

2) Possible Peering Contracts (PPCs): PPPs give only the set of locations for possible peerings. After that, we feed TM and ETM to the *Policy generator* for generating the list of Possible Peering Contracts (PPCs). We use TM to compute the traffic flow at particular PoPs. Let us denote the list of PPC as $\mathcal{PPC}$.

$\mathcal{PPC}$ is the collection of all sets of PPP locations where the two ISPs may establish peering. For example, if there are three common PoPs between two ISPs, there would be seven different PPPs. Assuming the common cities are A, B, and C for both ISPs, PPPs in this case are: peering at all three cities (A-B-C), at different combinations of two locations (one of A-B, A-C, or B-C), or peering at a single location (either A, B, or C). Since *no peering* is not a valid option here, we discard that. The more the common locations, the more PPPs will be generated. Hence, for r number of common PoP count, the PPC count will be:

$$|\mathcal{PPC}| = 2^r - 1. \tag{1}$$

3) Acceptable Peering Contracts (APCs): With all the PPCs being populated, our algorithm sorts them according to the requestor ISP's internal policies and sorting strategy. *Policy filter* eliminates some impermissible options from the list at this stage if they do not qualify. We refer to these selected contracts as APCs and identify the list as $\mathcal{APC}$ with $|\mathcal{APC}| = z$. Note that $\mathcal{APC} \subseteq \mathcal{PPC}$, and $z \leq 2^r - 1$.

C. ISP Network Model

Analyzing a candidate network to come up with the traffic volume, routing topology, and link capacities is most critical when deciding the PPCs. We have used TMs for estimating the traffic amount between every possible origin-destination (OD) node pair, in this case, ingress and egress points of the network to model the candidate's traffic volume. Existing TM calculation techniques rely either on statistical estimation based on routing matrix inference from SNMP link counts [53] or measuring OD flows for a certain period [54] . Although achieving a proper balance between these methodologies of measurement, inference, and TM modeling has been heavily investigated [55] , we have taken a simplistic but most common way to model TM of an ISP network.

We collect the population and router data for each PoP of the ISP(s) to estimate their TM (ISP pair wise) using the *gravity model*. An extensive study on the gravity model shows that it can be used for practical traffic matrix synthesis [56]. Moreover, it was found that a tomo-gravity model with a relativity factor in the gravity model can attain a better prediction of the traffic in the TM [57]. The gravity model contains mass and distance in the equation and brings the metaphor of physical gravity. Based on previous study [55], we assume the traffic amount between two PoPs should be proportional to the population of these PoPs. With this rationale, we use the following derivation where $Force$, m, a, G represent the gravitational force, mass of the items, acceleration due to attraction, and gravitational constant:

$$Force = G * \frac{m_1 * m_2}{d^2} \tag{2a}$$

$$Force = m_1 * a_1 \qquad Force = m_2 * a_2 \tag{2b}$$

$$\Rightarrow a_1 = G * \frac{m_2}{d^2} \qquad \Rightarrow a_2 = G * \frac{m_1}{d^2} \tag{2c}$$

$$\Rightarrow T_{1,2} = G * \frac{m_2}{d^2} \qquad \Rightarrow T_{2,1} = G * \frac{m_1}{d^2} \tag{2d}$$

where a_1 represents the acceleration of m_1 towards m_2. From ISP's business context, for some ISP R the value of m_1 should be proportional to the router density of ISP R and the total population in that area (PoP_1). Hence, if the number of routers of ISP R at PoP_1 is $R_{R,1}$ and the population is p_1, then $m_1 = \frac{R_{R,1}}{\sum_k R_{k,1}} * p_1$, where $\sum_k R_{k,1}$ represents the sum of all routers corresponding to all the ISPs present at PoP_1. Let $T_{1,2}$ ($T_{2,1}$) represents the amount of traffic flow from PoP_1 to PoP_2 (PoP_2 to PoP_1, respectively), and d is the distance between these two locations (cities). As everyone in a certain location (city or state) may not have a connection to the Internet, we define G as the usage factor by combining the Internet penetration percentage in a specific state [58] and per-person Internet usage (calculated globally [59]). We consider statewide Internet penetration for the US only, but it can be adapted to any specific country. Let s_d be the Internet penetration percentage in the destination states (for the US regional peering) or countries (for global peering) respectively,

and u be the per-person usage percentage. We express G as:

$$G = s_d * u. \tag{3}$$

The value of G can also be extended to consider the country-wise per-person network-connected device count, max peak-hour behavior and other factors to achieve further fine-tuning. Since, traffic from one PoP to another PoP depends on the router distribution of the ISPs, hence to incorporate that in the notation, we denote $T_{1,2}^{R\to C}$ as the traffic amount that ISP R is sending from PoP_1 to ISP C at PoP_2. Assuming s_1 and s_2 to be the Internet penetration percentage of the two locations where PoP1 and PoP2 are located (PoPs can be in the same locations as well, then $s_1 = s_2$), we can rewrite Eq. 2 as:

$$\begin{aligned} T_{1,2}^{R\to C} &= s_2 * u * \frac{p_2}{d^2} * \frac{R_{C,2}}{\sum_k R_{k,2}}, \\ \text{and } T_{2,1}^{C\to R} &= s_1 * u * \frac{p_1}{d^2} * \frac{R_{R,1}}{\sum_k R_{k,1}}. \end{aligned} \tag{4}$$

The dimension of the traffic matrix depends on the number of locations where the ISPs R and C are present. If N_{PoP}^R and N_{PoP}^C represents the total number of PoPs that R and C has respectively, then the TM and ETM both will have a dimension of $N_{PoP}^R \times N_{PoP}^C$. In a real-world scenario, the requester should possess its own TM and only estimate the candidate's network, but here, we relied on the above model for estimating both the requester and the candidate's TM as we did not have TM information available for any ISP. However, in the following, we still call requester traffic as TM and candidate's traffic as ETM.

D. Offloaded Traffic Estimation

During our study, we found only *Microsoft* to explicitly mention that they would try to carry the data through their network to a PoP nearest to the user location. They also expect the requester ISPs to announce their entire and consistent set of prefixes in all the PoPs, so Microsoft can take advantage of choosing the exit point. In contrast, many ISPs do not follow this practice and simply implement the *hot potato* technique. While calculating the ETM of the candidate, we assumed that the traffic would enter or exit with proper discretion, and, that the entire volume would be proportionately distributed to the port capacities between all the common PoPs. As such, we did not consider the 'closer to geo-location' phenomenon here. We used unit values while calculating data flow for both the ISPs. Further, instead of simply averaging the total traffic and distributing it, we utilized the port capacity of each PoP as the weight factor to calculate the weighted average. This gave us a better approximation of ETM for ISPs.

Consider Fig. 10, where requester (R) and candidate (C) both have 4 PoPs individually and P_{RC1}, P_{RC2} are the common exchange points between them. Traffic from R can go to any PoP including P_{C1}, P_{C2} via P_{RC1} or P_{RC2} if they agree to peer. From R's point of view, outgoing traffic from P_{R_1} is

$$\sum_{j=1}^{n-r} T_{P_{R1},P_{Cj}}^{R\to C}, \tag{5}$$

where n is the number of PoPs C has, while r is the number of common PoPs. These traffic will go via P_{RC1} or P_{RC2}, and based on port capacities, we distribute the traffic among these exit points. Considering v_{RCk} to be the port capacity of k-th common PoP (P_{RCk}), we formulate the traffic that R will offload to C via P_{RCk} as:

$$T_k^{R\to C} = \frac{v_{RCk} * \sum_{i=1}^{m}\sum_{j=1}^{n-r} T_{P_{Ri},P_{Cj}}^{R\to C}}{\sum_{k=1}^{r} v_{RC_k}}. \tag{6}$$

In general, ISPs do not exhaust the total capacity of a single port, rather they limit the utilization to a certain threshold. Should the need arise, additional port is either purchased from the IXP or new port is activated if ISP owns the switch. Previous studies [60] found that 95th percentile average utilization varies from 36% to maximum utilization of 50% at peak hours. There were incidents when the utilization reached 90%, but those accounted for less than 10% of the cases [61].

E. Meta-peering as an Optimization Problem

From the discussion (in this section and in Section III) made on different metrics related to the decision-making of peering between two ISPs, we can argue that the peering decision problem can be formulated as an optimization problem. Intuitively, the decision of two ISPs to peer is related to the individual willingness of each ISP, and the longevity of that relationship depends on some stability metric. Overall, from the perspective of an ISP, willingness to peer consists of the traffic exchange benefit, internal routing cost reduction, etc. On the other hand, the stability of the peering relationship depends on how these ISPs compare, i.e., in terms of size, traffic ratio between them, internal routing cost ratio, etc. The methodologies presented in Section V focus on determining the willingness and stability of any ISP pair for different *peering contracts* (where to peer, what traffic exchange ratio to follow) and finding the optimum contract. If the optimum contract is good enough for both parties (ISPs), they should decide to peer, otherwise, a peering relationship should not be formed. Thus, for the peering contracts, the optimization problem is to find the peering contract PC_{OPT} that maximizes some function F (the *Felicity* score) that has the willingness and stability metric as its argument for some ISP pair;

$$PC_{OPT} = \arg\max_{i\in\mathcal{APC}} F\big(W_i^{R,C}, S_i^{R,C}\big), \tag{7}$$

where $W_i^{R,C}$ and $S_i^{R,C}$ respectively represent the willingness metric (value) and stability metric (value) for ISP pair (R, C) when using the i^{th} contract from the $\mathcal{APC}$ list. The *Felicity* scores calculated for different contracts are compared with a threshold value to check if the ISP pair (R, C) should peer or not, detailed later in Section V-D. Also, if the ISP pairs decide to peer, then the contract with the highest felicity score provides us with the peering locations (PoPs).

V. Meta-peering Methodologies

Earlier study [36] found that peering point placement problems under traffic cost constraints are NP-complete. Our framework solves the same peering (point) placement issue,

TABLE IV: Summary of Commonly Used Notations

Notation	*Description*
$T_{1,2}^{R\rightarrow C}$	Amount of traffic that ISP R is sending from location 1 to ISP C at location 2
N_{PoP}^{R}	Total Number of PoP(s) of R
N_{PoP}^{C}	Total Number of PoP(s) of C
N_{Rtr}^{R}	Number of locations where R has router(s)
N_{Rtr}^{C}	Number of locations where C has router(s)
$W_i^{R\rightarrow C}$	*Willingness* of R towards C (using contract i)
$W_i^{C\rightarrow R}$	*Willingness* of C towards R (using contract i)
$W_i^{R,C}$	Pairwise *Willingness* between R and C (using contract i)
α_R	*Affinity* score of R (GEO-PP)
α_C	*Affinity* score of C (GEO-PP)
$\alpha_{R,C}$	Pairwise *Affinity* score of R & C (GEO-PP)
$\xi_{R,C}$	Pairwise *Similarity* score of R & C (GEO-PP)
$S_i^{R\rightarrow C}$	*Stability* of R towards C (RCA-PP)
$S_i^{C\rightarrow R}$	*Stability* of C towards R (RCA-PP)
$S_i^{R,C}$	Pairwise *Stability* between R and C (using contract i)
$F_i^{R,C}$	*Felicity* score of R & C using contract i
τ	Threshold value to compare with Felicity score
β, γ, δ	Constants to calculate Felicity score
$F_{R,C}$	*Overall Felicity score* of R & C (is compared with T for peering decision)

using one of the two meta-peering techniques proposed in this work. At first, we introduce the *GEOgraphic Peering Prediction* (GEO-PP) method, which performs meta-peering by focusing on the geographic overlap between two ISPs.

Then, we discuss another meta-peering methodology, called *Routing Cost Aware Peering Prediction* (RCA-PP), that analyzes the peering relationship from the perspective of internal routing cost. Table IV provides the summary of commonly used notations for our proposed methodologies.

A. Methodology I: GEO-PP

The operational framework of GEO-PP is depicted in Fig. 9. It uses the TM of the requester ISP, ETM of the candidate ISP, PoPs of both of them, population data for ingress and egress cities, and tolerance information as parameters to generate the $\mathcal{APC}$. The data regarding PoP information is obtained from PeeringDB, while population data is gathered from a population database. This initial APC list is then forwarded to the *willingness calculator* where they are sorted and ranked from both ISPs standpoint ($\mathcal{APC}^R$ for requester, $\mathcal{APC}^C$ for candidate), based on some criteria (discussed below). These ranks are then used to calculate the willingness of peering, we call them APC *willingness* scores. The final goal of the *willingness calculator* is to obtain the optimum APC list, $\mathcal{APC}^*$, where the individual APCs are ordered in such a way that maximizes the overall combined *willingness* score of the ISP pair, hence preferable for both. Next, the framework calculates the stability of possible peering relation using two additional terms, a) *affinity* score, a measure of the overlap between the coverage of the two ISPs, and b) *similarity* score, a measure of the size similarity of the two ISPs. The affinity is calculated using the PoP data from PeeringDB, while the similarity score is measured using the CAIDA customer cone data. These three metrics are then combined to obtain the felicity score, which ultimately predicts whether the two ISPs should peer or not.

1) Peering Willingness: Referring to Fig. 9, the *willingness calculator* receives the preliminary APC list, which is then ranked based on each ISP's own sorting criteria to formulate $\mathcal{APC}^R$ and $\mathcal{APC}^C$. As a measure of tolerance, we utilize the ratio of outbound to inbound traffic amount to prepare the $\mathcal{APC}^R$ or $\mathcal{APC}^C$ (while eliminating PPCs). These ratios are essentially peering policy decisions ISPs make and vary from 1:1.5 (*Telstra, CenturyLink*) to 1:3 (*Liberty Global*), with the most popular being 1:2, maintained by *Zayo, AT&T, GTT Communication* and others. The sorting criteria of APC from the perspective of the ISPs can be one of the following three options: a) **Own:** maximize the requester ISP's outbound traffic towards the candidate ISP, regardless of how much traffic it receives from the candidate, b) **Diff:** minimize the absolute difference between in/out-bound traffic for the requester, or c) **Ratio:** choose peers with lower in/out-bound traffic ratio. These three criteria are the most common metrics ISPs consider when they are trying to decide to peer or not [62]. Additionally, based on the observation from Fig. 8b, we set 15 as the max common PoP count for any ISP pair.

To quantify the 'robustness' of an APC, we take these lists, and, for each APC, $\mathcal{APC}_i$, we calculate the difference of ISPs' preferences, the rank of $\mathcal{APC}_i$ in $\mathcal{APC}^R$ and $\mathcal{APC}^C$, and take the square of it for a positive value. We then normalize this value by the square of the maximum difference regardless of any specific APC to include the worst-case scenario when an APC is most preferred by the requester ISP but is least preferred by the candidate. In most cases, $\mathcal{APC}^R$ and $\mathcal{APC}^C$ will contain the same APCs but in different order of preference. In case of a scenario when $\mathcal{APC}^R$ and $\mathcal{APC}^C$ do not include the same items, the rank of the missing $\mathcal{APC}_i$ in the counterpart's list is set to infinity so that it is preferred the least, while making sure the list contains the same items. Let, $\mathfrak{R}_i^R$ and $\mathfrak{R}_i^C$ be the rank of a particular $\mathcal{APC}_i$ in $\mathcal{APC}^R$ and $\mathcal{APC}^C$, respectively. We calculate the individual willingness scores of R, $W_i^{R\rightarrow C}$, and C, $W_i^{C\rightarrow R}$, for that particular contract as follows:

$$W_i^{R\rightarrow C} = 1 - \frac{\mathfrak{R}_i^R - 1}{|\mathcal{APC}|}, \qquad W_i^{C\rightarrow R} = 1 - \frac{\mathfrak{R}_i^C - 1}{|\mathcal{APC}|}. \quad (8)$$

The combined willingness score for a particular $\mathcal{APC}_i$ is calculated as the geometric mean of the individual scores as follows:

$$W_i^{R,C} = \sqrt{W_i^{R\rightarrow C} * W_i^{C\rightarrow R}}\,. \quad (9)$$

2) Peering Stability: GEO-PP quantifies stability of the peering relationship between two ISPs based on how 'affine' and similar they are to each other. For this, we compute two corresponding scores as we detail next.

Affinity Score: An ISP will be interested in peering with another ISP if the relationship would expand its coverage

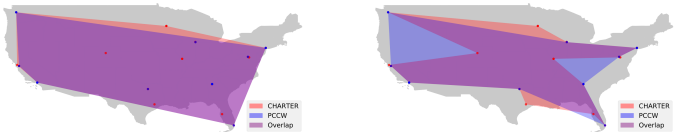

(a) Convex-Hull overlapping (b) Polygon overlapping

Fig. 11: Inter-ISP overlap

area; otherwise, there may not be enough incentive to peer with someone who is covering the same locations or has a smaller coverage area. We call this interest as *affinity* score of an ISP to peer with another ISP. To represent the coverage area of an ISP, we initially calculated the *convex-hull* using all of its PoPs. But, we observed that a bigger convex-hull may unintentionally cover a smaller ISP's coverage area and reduce the affinity score. So, instead of convex-hull, we prefer a regular polygon to represent an ISP's coverage area. Fig. 11a and Fig. 11b present the comparison between the coverage areas of two ISPs using the two methods. Furthermore, due to uneven distribution of population, a larger area coverage does not necessarily indicate that an ISP serves a wide customer pool in different regions. So to represent the coverage area in a more pragmatic way, we implement the flood fill algorithm to calculate the total population in the covered region [63]. We convert the entire coverage area into a grid of five-square miles cells, and estimate the total population using the Gridded Population of the World (GPW) [64]. After this, we calculate the affinity scores α_R and α_C, respectively, for the requester and candidate ISPs as follows. Let A_R and A_C be the population in coverage areas of the requester and the candidate, respectively, and A_o be the overlapped area's population. We express the affinity scores based on the overlap as:

$$\alpha_R = \frac{A_C - A_o}{A_R \cup A_C} = \frac{A_C - A_o}{(A_R - A_o) + (A_C - A_o) + A_o}, \quad (10)$$

$$\alpha_C = \frac{A_R - A_o}{(A_R - A_o) + (A_C - A_o) + A_o}. \quad (11)$$

Similar to the combined willingness score, we use geometric mean to calculate the combined affinity score:

$$\alpha_{R,C} = \sqrt{\alpha_R * \alpha_C}. \quad (12)$$

Similarity Score: An interesting feature about the affinity score is that the measure of affinity is relative to the size of the ISPs. Considering one large and one small ISP, the same amount of increase in coverage does not equate to equal increase in both their affinity scores. It is because the affinity basically refers to the percentage of coverage increase, and for the larger ISP (in this case), the change in affinity score will be less compared to the smaller ISP. So, it makes sense for the ISPs to peer with other ISPs of similar sizes. To test that statement, we introduce a metric called similarity score (ξ), which essentially measures the level of similarity between ISPs. To quantify it, we use the CAIDA AS-Rank API [44] to accumulate all of the announced prefixes and covered address spaces to get the total numbers, and come up with three types of similarity scores, i.e., similarity based on PoP (ξ_{pop}), similarity based on address (ξ_{add}), and similarity based on prefix (ξ_{pre}). These scores were calculated in an identical manner: as a ratio of number of PoPs, IP addresses, and IP prefixes respectively of the two potential peering ISP pairs, with the larger value at the denominator. This can be expressed as follows:

$$\xi_k^{R,C} = \frac{\min(n_R, n_C)_k}{\max(n_R, n_C)_k}; k = \{pop, add, pre\}. \quad (13)$$

where n_R and n_C are the number of elements of ISPs R and C respectively, based on the score criteria, i.e., PoPs, addresses, and prefixes. Higher the similarity score, the more similar ISPs are. The overall similarity score of the pair will be the average of the three types of the similarity scores as follows:

$$\xi_{R,C} = \frac{1}{3}\left(\xi_{pop}^{R,C} + \xi_{add}^{R,C} + \xi_{pre}^{R,C}\right). \quad (14)$$

3) Felicity in Peering: Using the willingness (Eq. 9), affinity (Eq. 12) and similarity (Eq. 14) scores, we compute the felicity score of $\mathcal{APC}_i$, which is a measure of the merit of that particular contract and plays a crucial role in the prediction of whether the two ISPs would be peering or not, and is computed as a geometric mean of the above scores. We propose two different felicity scores, one using only the affinity score as a measure of peering stability, while the other one uses both the affinity and similarity scores. This was done to evaluate the impact of the different stability parameters on the felicity metric. For the ISP pairing, (R, C), the felicity score for contract $i \in \mathcal{APC}$ is calculated in two difference ways, as follows:

$$F_i(W_i^{R,C}, S_i^{R,C}) = \left[(W_i^{R,C})^\beta * (\alpha_{R,C})^\gamma\right]^{\frac{1}{\beta+\gamma}} \quad \text{with } S_i^{R,C} = \alpha_{R,C}, \text{ and} \quad (15)$$

$$F_i(W_i^{R,C}, S_i^{R,C}) = \left[(W_i^{R,C})^\beta * (\alpha_{R,C})^\gamma * (\xi_{R,C})^\delta\right]^{\frac{1}{\beta+\gamma+\delta}} \quad \text{with } S_i^{R,C} = \{\alpha_{R,C}, \xi_{R,C}\} \quad (16)$$

where $F_i(.)$ is simply the geometric mean of input parameters using constant exponents β, γ, and δ. The geometric mean of the willingness, $W_i^{R,C}$, and the stability, $S_i^{R,C}$, implies that peering contracts with high willingness and high stability *at the same time* will be preferred. Further, in Eq. 16, the stability $S_i^{R,C}$ is a tuple of the affinity and similarity scores, emphasizing that the ISPs are more likely to peer if they both have low coverage overlap and similarity in size.

Now, the $\mathcal{APC}_i$ with the highest felicity may not always be the most preferable contract for an ISP. For example, it may propose a peering location where the ISP already has good coverage and thus provide an insufficient incentive for that ISP to peer. A better alternative would to be take into account, the felicity of all the contracts within the APC list and come up with a combined felicity score. That way, ISPs will have more options to choose from and select contracts that may not necessarily have the highest individual felicity score, but more in line with their operational requirements. To formulate this combined felicity score, the main task would be to maximize the willingness score by obtaining an optimum APC list, $\mathcal{APC}^*$, as the affinity and felicity scores remain constant for all $i \in \mathcal{APC}$.

Optimal APC List Formulation: Finding the optimal APC list refers to the ordering of each $\mathcal{APC}_i$ within the list in such a way that maximizes the overall willingness score of the list and accurately represents the benefit of peering to both ISPs. Let j_i be the rank or order of $\mathcal{APC}_i$ and is expressed as $j_i = \Re(\mathcal{APC}_i)$ where $j_i \in \mathcal{J} = 1 \dots z$. Based on that, obtaining the optimal APC list, $\mathcal{APC}^*$, can be expressed as a maximization problem for the requester ISP as follows:

$$\mathcal{J}^* = \arg\max_{\mathcal{J}} \sum_{i=1}^{z} W_i^{R,C} * W_i^R, \tag{17}$$

$$\textbf{s.t. } \Re(APC_1) < \Re(APC_2) < ... < \Re(APC_z) \tag{18}$$

$$\text{and } z \leq 2^r - 1. \tag{19}$$

The outcome of this optimization is the rank order of all the $\mathcal{APC}_i$ for the optimally ordered APC set, $\mathcal{APC}^*$ (i.e., $\mathcal{J}^*$) and occurs when the ordering of APCs in the solution set is monotonically decreasing in terms of the two ISPs' combined willingness. A higher combined willingness score ($W_i^{R,C}$) may not always reflect the best APC choice for the requester based on two hypotheses: 1) the requester will make the peering attempt, and 2) a rational ISP will always prefer an APC with maximum benefit. So, Eq. 17 also considers the willingness score of the requester (W_i^R) for that specific $\mathcal{APC}_i$ and the constraint in Eq. 18 assures that the optimal ordering requirement is satisfied. Using this optimal $\mathcal{APC}^*$, the combined willingness score of the ISP pairing can be calculated as the average of all the individual willingness scores within the list.

Finally, the combined felicity score for the pairing using the overall willingness score is formulated as follows:

$$F_{R,C} = \max_{i \in \mathcal{APC}^*} F_i \left(\frac{\sum_{i \in \mathcal{APC}^*} W_i^{R,C}}{|\mathcal{APC}|}, S_i^{R,C} \right). \tag{20}$$

Here, the first argument of $F_i(.)$, $\frac{\sum_{i \in \mathcal{APC}^*} W_i^{R,C}}{|\mathcal{APC}|}$, is the overall willingness score for the ISP pairing (R, C). Furhter, $S_i^{R,C}$ is fixed for all contracts $i \in \mathcal{APC}^*$ as the stability parameters remain constant regardless of the changes in the APC list. Using the combined felicity, the model predicts whether ISPs pair will peer or not by comparing it with a threshold, τ. If $F^{R,C} \geq \tau$, the ISPs are predicted to peer, while not if $F^{R,C} < \tau$.

B. Methodology II: RCA-PP

For this specific methodology, RCA-PP, we extend the definition of Possible Peering Points ($PPPs$) slightly. In GEO-PP the $PPPs$ only included the common PoPs, whereas for this case we also include common router locations of the ISP pairs. Similarly, we extend the notion of Possible Peering Contracts ($PPCs$) along with the Acceptable Peering Contracts ($APCs$) in the same way. Moreover, due to this modification, the traffic matrix for this method is also different than GEO-PP. Let us denote N_{Rtr}^R and N_{Rtr}^C as the locations where R and C has routers (due to the nature of PoP locations, each PoP locations also has at least a router from its corresponding ISPs). The traffic matrix for RCA-PP has a dimension of $N_{Rtr}^R \times N_{Rtr}^C$.

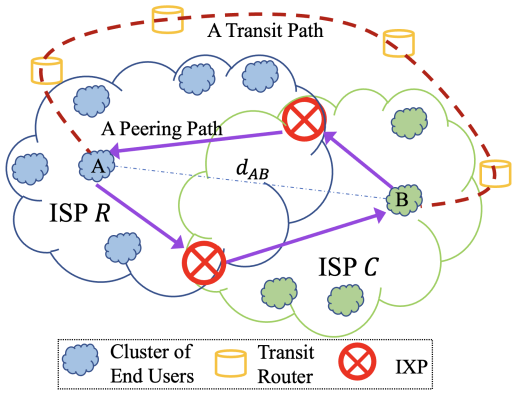

Fig. 12: Sample peering and transit path for an ISP pair.

1) Cost of Traffic: The cost of peering traffic depends on the perspective of ISPs. When an ISP carries traffic through its own network, the ISP tries to minimize the increased congestion (cost) along with the cable infrastructure (cost) to facilitate this data transmission. On the other hand, ISPs can decide to send traffic using transit services by paying a transit fee (set by the transit provider) as the cost of traffic. Overall, it is hard to put a dollar cost on the exchanged traffic, but in general, the cost is an increasing function of the length the traffic needs to traverse [5].

For illustrative purposes, let us consider the scenario of Fig. 12. In this scenario, ISPs R and C are peering at two locations (IXPs), indicated by red crossed circles. These locations are connected through internal routing paths (depicted in purple). Alternatively, the ISPs have the option to exchange traffic through a transit service, traversing a path marked in maroon dashed lines that involves the transit provider's routers (depicted as orange cylinders). As the transit path often exceeds the shortest available route and, at times, even surpasses the peering path in length, a parameter called the *path stretch factor* is employed to calculate the effective distance that traffic might need to cover via a transit path. A detailed discussion on the path stretch factor is provided in the subsequent section. Now, let us assume that certain traffic needs to journey from point A (belonging to ISP R) to point B (belonging to ISP C), and d_{AB} represents the distance between these two points (refer to Fig. 12). Additionally, let d_I denote the distance covered by the traffic while navigating internal routes (say d_r) within ISP R up to the peering location. We can estimate the internal routing cost of the ISP as

$$C_I = a_I \times \sum_r d_r = a_I \times d_I, \tag{21}$$

where a_I is a constant and $\sum_r d_r = d_I$ is summation of all the internal routes that the traffic need to traverse to reach the peering location. Similarly, the transit cost (C_T) can be expressed as

$$C_T = a_T \times d_{AB} \times f, \tag{22}$$

where f signifies the path stretch factor and a_T is a constant.

2) Path Stretch Factor: When data flows on the Internet, the route does not consistently follow the most direct path for various reasons. The path stretch factor, f, between two

endpoints, A and B, can be broadly characterized as the degree to which the route taken from A to B exceeds the length of the shortest available path. This can be quantified as the ratio of the actual path length to the shortest path length. One illustrative example is given in Fig. 12, where traffic is compelled to take a more extended route due to some constraints, e.g., policy considerations. For RCA-PP, determining a good and representative path stretch, f, for calculating the transit cost in (22) is necessary.

Several research studies have assessed the router-level hop path stretch factor in the context of the Internet. Tangmunarunkit *et al.* were among the pioneers, conducting one of the initial studies [65] to calculate the router-level path stretch resulting from routing policies. Building upon this work, Gao *et al.* expanded the scope to encompass AS-level hop stretch [66]. While both studies delineate trends in path stretch, they refrain from presenting global averages. Similarly, Mühlbauer *et al.* calculated the path stretch factor on a large-scale simulated AS-level graph constructed with data from CAIDA [67]. Their computations showed observed path stretch of 1.3, 2.1, and 2.9 in terms of AS-level hops, router-level hops, and geographical distances, respectively. Based on these studies, we chose $f = 2$ for our study.

3) Peering Willingness: The *peering willingness* of an ISP R to peer with another ISP C using some contract $i(\in \mathcal{APC})$ is denoted by $W_i^{R\to C}$, and is defined as the ratio of transit cost to peering cost from the perspective of the ISP R. Hence, we have

$$W_i^{R\to C} = \frac{\sum_t C_T(t,R,C)}{\sum_t C_I(i,t,R,C)} = \frac{\sum_t a_T * d_{(AB)_t} * f}{\sum_t a_I * d_{I_t}(i)}\ , \quad (23)$$

where $C_T(t,R,C)$ and $C_I(i,t,R,C)$ are the transit and internal routing costs for some traffic flow t that ISP R needs to send to C, and a sum over all the traffic t gives the total cost. Also, $d_{(AB)_t}$ and $d_{I_t}(i)$ respectively represent the geographical and internal routing distance that traffic t needs to travel in transit and peering scenarios. Higher value of $W_i^{R\to C}$ means higher inclination of ISP R to engage in peering with ISP C using peering contract i. The *pairwise peering willingness* between ISPs R and C for contract i is formulated as follows:

$$W_i^{R,C} = \sqrt{W_i^{R\to C} \times W_i^{C\to R}}, \quad (24)$$

and is the representation of the overall peering inclination of the ISP pair (R,C).

4) Peering Stability: We define *Peering Stability* for an ISP R peering with ISP C with a contract $i \in \mathcal{APC}$, denoted as $(S_i^{R\to C})$, as the ratio of the minimum attainable cost using any of the contracts from $\mathcal{APC}$ to the cost of using contract i. Thus, we have,

$$S_i^{R\to C} = \frac{\min_{\tilde{i}\in\mathcal{APC}} \sum_t C_I(\tilde{i},t,R,C)}{\sum_t C_I(i,t,R,C)}, \quad (25)$$

where $\min_{\tilde{i}\in\mathcal{APC}} \sum_t C_I(\tilde{i},t,R,C)$ is the minimum cost that ISP R would have incurred using various peering contracts from $\mathcal{APC}$, and $\sum_t C_I(i,t,R,C,)$ is the cost when using the specific contract i. From the perspective of both ISPs, if the selected solution ($i \in \mathcal{APC}$) doesn't deviate significantly from the minimum achievable solution for both ISPs, then they are likely to establish a peering agreement. Therefore, in a manner similar to GEO-PP's, we define the peering stability for an ISP pair (R,C) for some contract i as,

$$S_i^{R,C} = \sqrt{S_i^{R\to C} \times S_i^{C\to R}}\ , \quad (26)$$

which offers a quantitative measure of the overall stability of the peering relationship.

C. Felicity in Peering

Finally, the $W_i^{R,C}$ and $S_i^{R,C}$ values from the perspective of both ISPs are used to find *the optimum contract* given by Eq. 7 that have the highest felicity score. To calculate the felicity score we use a very similar equation as used in Eq. 15, and define it in the RCA-PP method for ISP pair (R,C) with contract $i \in \mathcal{APC}$ as;

$$F_i^{R,C} = \left[(W_i^{R,C})^\beta \times (S_i^{R,C})^\gamma\right]^{\frac{1}{\beta+\gamma}}, \quad (27)$$

where β and γ are constants but can have very different values than in Eq. 15. Lastly, the (final) felicity score that is compared with the threshold τ for peering decision is given by

$$F_{R,C} = \max_{i\in APC} F_i^{R,C}. \quad (28)$$

D. Peering Decision with Meta-peering

GEO-PP and RCA-PP are meta-peering methodologies to aid us in a two-step decision-making of 1) whether two ISPs should peer, and 2) where they should peer (if they decide to peer).

1) Peering Partner Choice: The felicity score indicates the overall quality of the peering relation between ISPs . The higher the felicity score, the better matched the ISP pairs for peering. In the GEO-PP framework, the *felicity score* between two ISPs is calculated using either Eq. 15 or 16. We set a tunable threshold value, τ, on the felicity score, which helps determine the peering decision; felicity higher than τ means the ISP pair should peer and vice versa. Similarly, for the case of RCA-PP, we use Eq. 27 to calculate the felicity score of some ISP pair for all possible peering contracts and use the maximum value to decide if the pair should peer or not.

2) Peering Location Selection: If the ISP pairs decide to peer, the next question is which location(s) to peer. For the GEO-PP method, the optimal APC list, $\mathcal{APC}^*$ is calculated by solving the maximization problem given by Eq. 17. The peering locations are then obtained from the optimized list ($\mathcal{APC}^*$). On the other hand, for RCA-PP, the contract with the highest felicity score provides the PoP location(s) where the ISP pair should peer to have the higher willingness and stability.

VI. Evaluation

The methodology of GEO-PP and RCA-PP are developed using robust arguments backed up by statistical data. In this section, we demonstrate the effectiveness of these two methods by comparing and analyzing their outputs. Initially, we discuss the significance of each of the metrics that are being calculated in the proposed methods using real-life examples

(Sections VI-C1 and VI-C2). Later, we check the *balanced* accuracy of the peering prediction using both methods. Finally, we perform a comparison of the total internal routing costs using our proposed methods and the original PoP locations. Both (proposed) methodologies showed promising results in predicting peering partners and locations.

A. Datasets

To check the performance GEO-PP and RCA-PP, we used publicly available data from CAIDA [67], PeeringDB [62], and CEDAC (NASA) [64]. From CAIDA, we accessed data about routers (i.e., latitude, longitude, and IP addresses), router-to-ASN assignments, and ASN peering relationships. Autonomous system number or ASN is a unique number assigned to each Autonomous Systems (ASes). From PeeringDB, we extracted information on the ASNs, e.g., PoPs and port capacities purchased at IXPs. Interestingly, CAIDA also stores the PeeringDB databases and is usually up-to-date, and one can access the PeeringDB data from CAIDA as well. Lastly, we used [64] to get population distribution data given for different grid sizes.

ASNs or ISPs: ISPs are the organizations that provide different services, i.e., Access, Content, or Transit. Larger ISPs can own different types of ASNs under their name. Hence, it is not straightforward to tag an ISP with any specific service. On the other hand, ASNs always follow a single role, and we can get the ASN type for any ASN from PeeringDB. Furthermore, peering policies are ultimately implemented among ASes even though the contracts are made among ISPs. Due to these reasons, in our simulation analysis, we used ASNs instead of an ISP or the whole organization.

ASNs Analyzed: To check the performance of our proposed methods, we mainly focused on two sets of ASNs. The first set focuses on the ASNs that are mostly active in the US, and majority of their PoPs are inside the mainland US. In the second set, the ASNs are filtered from the global list of ASNs such that each ASN has a similar number of peering to non-peering relationship information in CAIDA. In other words, the second set focuses on maintaining a balanced peering relationship data. To give numerical values, there are 1,324 ASN pairs in the first set of ASNs and 1,528 ASN pairs in the second set of ASNs. In the first set, 1,196 out of 1,324 pairs have a peering relationship and the rest have a non-peering relationship. Hence, the peering to non-peering data is heavily imbalanced with a ratio of $90:10$. Overall, the reasons behind this imbalanced data are: i) the original dataset from CAIDA is also imbalanced (with a 70% of peeering data to 30% of non-peering data), and ii) the peering information in CAIDA on the large ASNs active in the US are heavily biased as well (with much more data on peering ISP pairs). On the other hand, the second set of ASNs has a much better data balance. 709 out of 1,528 pairs of the second set have peering relationship and the rest of the pairs have non-peering relationship, i.e., a ratio of $47:53$.

Ground Truth: To evaluate the performance of our method in predicting peering partners and locations, we took aid from the CAIDA and PeeringDB data. CAIDA maintains a dataset of ASN peering relationships [68], where it stores peering relationships between different ASN pairs. CAIDA uses its methodology, focused on how traffic flows through the network, to identify peering ASNs, and may not be 100% accurate. However, without any alternate source, we assume this to be the ground truth. Our idea is to check if our methodologies, which only use publicly available data, can predict these relationships. On the other hand, to get the ground truth on the peering locations for any ASN pair, we assumed that if CAIDA says that two ASNs are peering, and if those two ASNs are present at some location (in a facility or an IXP), we assumed that location to be a peering location for that ISP pair, and denote that as the *original peering location*.

B. Evaluation Procedure

1) Heuristic Method: To evaluate the performance of GEO-PP using the Heuristic method, we use Eqs. 15 and 16 to find the felicity scores for individual ASN pairs. When calculating these scores we use Grid Search [69] to get the best values for different parameters (i.e., β, γ, and δ) and threshold, τ, that gives the highest accuracy with respect to the ground truth obtained from CAIDA. Similarly, for RCA-PP, we follow a similar approach with Eq. 27 and optimize the parameters to attain the best accuracy. On the other hand, the predicted peering locations are given directly by the contract $i \in APC$ that has the highest felicity score (Fig. 13).

2) Computational Complexity: The computational complexity of GEO-PP is primarily dependent on obtaining the optimal APC list, $\mathcal{APC}^*$, i.e., peering willingness, as the stability components, affinity and similarity, remain constant for all individual APCs. The main computational steps for determining willingness consist of 1) estimating the TM of both requester (R) and candidate (C) ISPs, 2) calculating the individual and combined willingness scores and 3) sorting them based on that. The cost related to the computation of TM will be dependent on the number of PoPs of both R and C (let, denoted as N^R_{PoP} and N^C_{PoP} respectively) and the complexity can be expressed as $\mathcal{O}(N^R_{PoP}N^C_{PoP})$. For an APC list with length z, the complexity of obtaining the individual and combined willingness scores is $\mathcal{O}(z)$ and the complexity of sorting the list according to the scores is $\mathcal{O}(z \log z)$. Now according to the discussion in Section III, the maximum value of z is, $z = 2^r - 1$, where, r is the number of common PoPs between the ISPs. So, the worst case complexity of GEO-PP is $\mathcal{O}(N^R_{PoP}N^C_{PoP}+(2^r-1)\log(2^r-1)) \equiv \mathcal{O}(N^R_{PoP}N^C_{PoP}+r2^r)$. This will be heavily influenced by the value of r. For a lower number of common PoPs, the first term, $N^R_{PoP}N^C_{PoP}$, will play the determining role, whereas for a higher number of common PoPs, the second term, $r2^r$, will be the defining factor for the complexity.

RCA-PP uses a different traffic matrix compared to GEO-PP, and the complexity of that calculation is $\mathcal{O}(N^R_{Rtr} * N^C_{Rtr})$. On the other hand, the calculation of willingness and stability is done for each contract $i \in \mathcal{APC}$. For $APC\ i$, the internal routing cost is calculated in two different ways: 1) when both ASNs are selfish and try to minimize their internal traffic routing, and ii) when ASNs cooperate to minimize the total

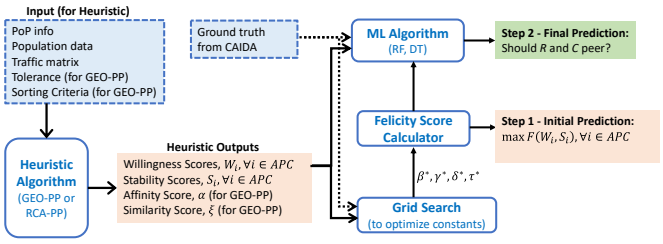

Fig. 13: Evaluation workflow of GEO-PP and RCA-PP.

routing cost, i.e., minimize the system cost. Hence, if the peering contract i has $N_{PoP}(i)$ number of PoP location, then RCA-PP needs to make $2 \times N^R_{Rtr} \times N^C_{Rtr} \times N_{PoP}(i)$ decision making. Thus for a total of $|APC| = z$ number of contracts, the computational complexity is $2 \times N_R \times N_C \times N_{PoP}(i) \times z$. So overall the computational complexity for the RCA-PP method is $\mathcal{O}\Big(N^R_{Rtr}N^C_{Rtr} + N^R_{Rtr}N^C_{Rtr}N_{PoP}(i)z\Big)$, and that can be upper bounded by the value of $\mathcal{O}\Big(N^R_{Rtr}N^C_{Rtr}\big(1 + r(2^r - 1)\big)\Big) \equiv \mathcal{O}\Big(N^R_{Rtr}N^C_{Rtr}r2^r\Big)$.

3) Machine Learning: The felicity score calculations for both GEO-PP and RCA-PP were done by a geometric mean function on willingness and stability metrics. To check if we can improve the peering partner prediction accuracy, we use machine learning method on top of our willingness and stability metric as an alternate solution to the felicity function. For this prediction method, the willingness and stability metrics are used as the input for the machine learning (ML) model (i.e., Random Forest) to perform peering partner prediction as a binary classification (Fig. 13).

C. Analysis of GEO-PP and RCA-PP metrics

1) GEO-PP: A higher felicity score in GEO-PP indicates a better matching between the ASNs. This can be easily understood by looking into the constituent components that make up the felicity score: the willingness, affinity, and similarity scores. A higher felicity score indicates 1) a selection of APCs that offer more beneficial traffic exchange agreements and alignment of routing policies, i.e., higher willingness scores, 2) a more significant increase of the total coverage area, i.e., higher affinity scores, and/or 3) a higher similarity (in terms of IP address and PoP counts) between the ASNs. Thus, ISPs will prefer peering agreements that offer better felicity scores due to the increase in the (inherent) utility associated with it. So, by setting a threshold on the felicity score, we can get a rough estimation of whether two potential ISP pair will peer or not (peer if the felicity score is higher than the threshold).

To check the individual relationship of the willingness ($W^{R,C} = \frac{\sum_{i \in \mathcal{APC}^*} W_i^{R,C}}{|\mathcal{APC}|}$) and stability metrics ($\alpha_{R,C}$ and $\xi_{R,C}$) to the felicity score, we plot them against each other using a threshold value of, $\tau = 0.1$ and obtain trend lines for each case as depicted in Fig. 14. The trend-lines exhibit similar characteristics in both datasets, so we only report the results for the balanced dataset here. Both willingness and stability show positive trends with respect to the felicity score, with the stability metrics showing a stronger trend. This can be attributed to the much narrower spread of the willingness score, as for majority of the pairs, the peering willingness was between 0.5 (50%) to 0.7 (70%). This also indicates that, without considering the peering stability, most of the ISP pairs would prefer to peer than not. As for the peering stability, the combined affinity and similarity scores showed stronger trends compared to only using affinity. This makes sense as discussed earlier in V-A1 that the affinity score is relative to the size of the ISPs and an equal amount of increase in coverage does not necessarily indicate an equal increase of affinity scores i.e. peering stability.

The relation of the similarity metrics to the felicity score is illustrated in Fig. 15. All three of them show positive trends with $\xi^{R,C}_{add}$ and $\xi^{R,C}_{pre}$ exhibiting stronger relations. This indicates that the more similar the ASNs are in these two aspects, the higher the felicity score (more likely the ISPs are to peer). As for ξ_{pop}, it shows a marginal positive trend. This can be attributed to the fact that the datasets appear to be heavily skewed in terms of the number of PoPs belonging to the ASNs, meaning most of them are similar in terms of PoP count, which explains why there was not much of a positive trend, and using a more diverse dataset in terms of PoP count should alleviate this issue.

2) RCA-PP: The calculation of willingness and stability for this method is computationally very expensive when done with higher granularity. We used the methodology described in Section V-B and V-D to calculate the $W_i^{R,C}$ and $S_i^{R,C}$ for all contracts $i \in \mathcal{APC}$ for 73 ISP pairs using a gridded map of the US. In the simulation, instead of using all possible contracts ($\mathcal{APC}$), we used a greedy method to reduce the number of contracts. The greedy algorithm used is an iterative algorithm that finds the best PoP location given a set of PoP locations already chosen, and adds the new PoP location to the previous set of PoPs. We specifically considered three different types of ASN pairs, i.e., Access-Access, Access-Content, and Content-Content.

From Fig. 16, we see that with the increase of peering points, the peering willingness of Access-Content and Content-Content type ASN pairs increases rapidly and reaches around 80% and 70% respectively. The results for Access-Access type ISPs are not provided because it took tremendous computational power to analyze those with the proposed methodology, and for the few cases it was analyzed, the pairs showed very low peering willingness. On the other hand, peering stability ($S_i^{R,C}$) of different types of ASN pairs showed a gradually decreasing trend with the increase of peering locations. The main reason behind this was that with more options to exchange traffic, both ASNs tried to minimize their own cost, thus decreasing the stability of the relationship. However, even with the selfish decision-making of the ASNs, the stability values remained quite high, with an average of 97% for Access-Content pairs and 89% for Content-Content pairs (Fig. 17). In the figure, the shaded region shows the standard deviation of the $S_i^{R,C}$ in the lower limit (since the average is really close to 100%, the upper limit is not shown here).

While peering willingness serves as a useful indicator for assessing mutual incentives for peering, it doesn't distinguish between public and private peering. In instances where an ISP

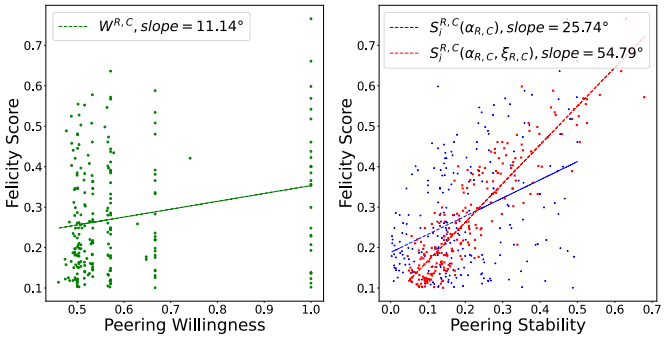

Fig. 14: Analysis of Peering Willingness and Stability

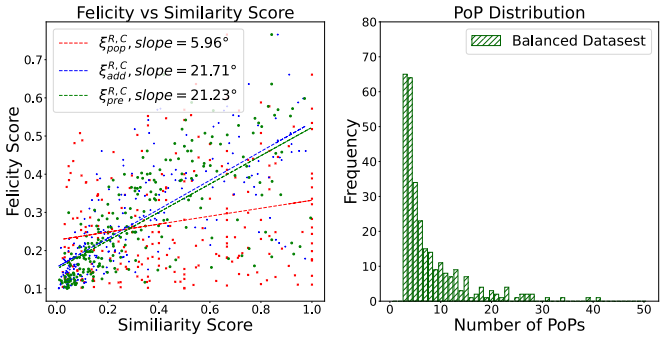

Fig. 15: Analysis of Similarity Metrics

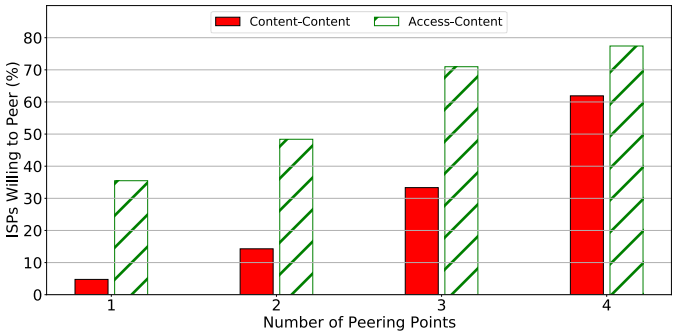

Fig. 16: Peering Willingness (RCA-PP) - multiple peering points.

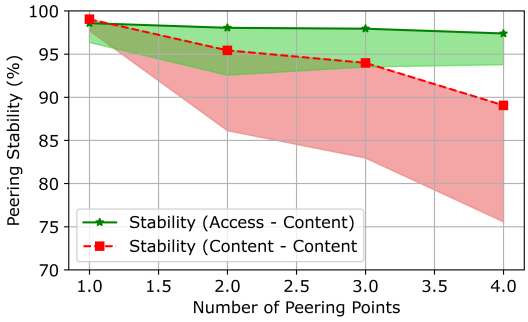

Fig. 17: Peering Stability (RCA-PP) - multiple peering points.

pair (R, C) has a similar volume of traffic to exchange, resulting in comparable peering (and transit) costs, and $W_i^{R,C} > 1$, they are likely to opt for public peering. Conversely, for highly asymmetric traffic, typical in Access-Content ISP pairs, if the peering costs are also highly asymmetric, a high $W_i^{R,C}$ might only lead to a preference for private peering, involving some monetary exchange.

In summary, the Access-Content type ISP pair showed high peering willingness and stability for a range of peering location(s), thus making it the ideal pair to have a peering relationship. Moreover, although the $A-A$ type ISP pair showed the highest stability for single-point peering, it exhibited low peering willingness.

D. Comparison of GEO-PP and RCA-PP

1) Peering Prediction Accuracy: The accuracy of our proposed methodologies in predicting whether two ASNs should peer or not is discussed in this section. Since we have a dataset

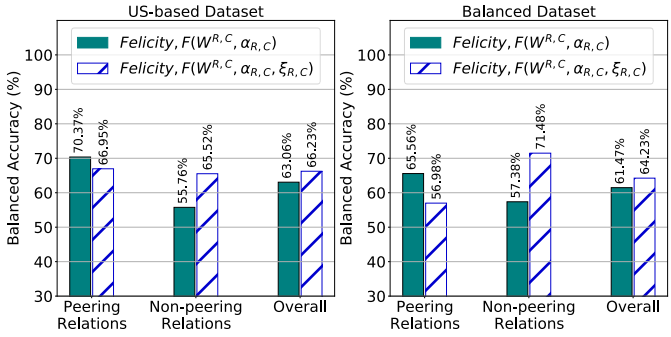

Fig. 18: Performance of GEO-PP (with heuristic only, no ML)

that is really imbalanced, all our performance calculation is done with balanced accuracy instead of just accuracy. Balanced accuracy is defined as the average of each individual class's accuracy, and is a better representation of performance.

GEO-PP: The prediction process in GEO-PP follows Fig. 13 and operates in two main steps, 1) obtaining an initial prediction through the felicity score, and 2) fine-tuning the predictions using a simple machine learning approach. For the felicity calculation, we considered both Eqs. (15) and (16). To obtain the optimum value of the weights, we used the grid search method [69], and the resultant values were, $\beta = 0.36, 0.30, \gamma = 0.16, 0.20, \delta = 0.48, 0.40$ and threshold, $\tau = 0.2, 0.2$ for the US-based and balanced datasets respectively. The results of step 1 are depicted in Fig. 18. The felicity score with the willingness and affinity score performs well in predicting peering relations in both the datasets. However, including the similarity score when calculating the felicity score showed considerable improvement in the prediction accuracy of non-peering relations. Overall, the felicity score using the combined affinity and similarity scores as a stability metric (Eq. 16) outperforms the other one with a modest balanced accuracy of 64% and 66% respectively.

The willingness and stability metrics ($W^{R,C} = \frac{\sum_{i \in \mathcal{APC}^*} W_i^{R,C}}{|\mathcal{APC}|}$, $\alpha_{R,C}$, and $\xi_{R,C}$), as well as the optimized felicity score using both affinity and similarity as the stability metric (Eq. 20), are then fed into a machine learning (ML) algorithm as training data, for which we consider two approaches in GEO-PP: Random Forest Classifier and Decision Tree Classifier. For the willingness score, we use the average value of the scores in all three sorting criteria

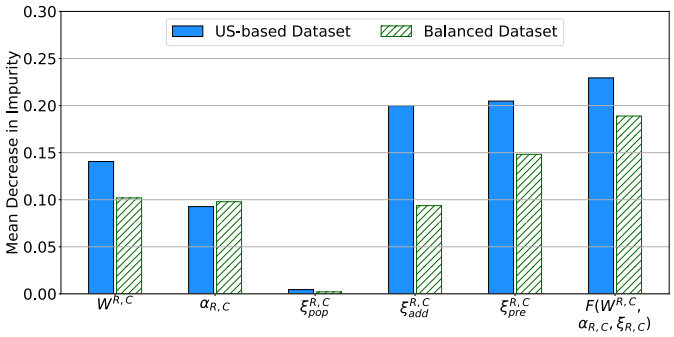

Fig. 19: Feature Importance: GEO-PP

for each ISP pair, as 1) for 87% of the pairs, the scores are within 0.1 of each other and 2) RCA-PP does not use similar sorting criteria. The feature importance of these parameters is depicted in Fig. 19. The felicity score exhibits the strongest impact with regards to the final prediction followed by the similarity metric $\xi^{R,C}_{pre}$. The similarity metric $\xi^{R,C}_{add}$, willingness and affinity come after them and show similar level of impact, particularly in the balanced dataset, with $\xi^{R,C}_{pop}$ showing the least importance.

The output from step 2 is illustrated in Fig. 20 and 21. The model performed quite well with the balanced dataset reaching a peak balanced accuracy of 79.9% and 78.5% with the two ML approaches (Fig. 21), but struggles with the US-based dataset, only reaching a modest accuracy of 66% and 64% respectively (Fig. 20), which is just a marginal improvement from the heuristic approach of step 1. The reason behind it can be directly attributed to the skewness of the US-based data with regards to the CAIDA ground truth. As mentioned earlier, the US-based data has a 90-10 split in terms of peering and non-peering ASN pairs, which reflects why our accuracy is quite average, as class imbalance or bias in training data can cause the model to produce erroneous predictions by exhibiting the same bias in the decision-making process. So although the model can reach accuracies of up to 98% when predicting peering relations in that dataset, it fails to replicate that for non-peering relations, dropping the balanced accuracy. That's why it has to be noted that, we can improve the performance of our model in terms of predicting peering relations in both datasets by lowering the threshold, τ, to less than the optimal value (0.2) found for the balanced accuracy, and reaching an accuracy of 90% for both. However, this will fail to predict majority of the non-peering relations correctly, lowering the balanced accuracy, which is why we used balanced accuracy instead of a regular accuracy calculation.

RCA-PP: RCA-PP mainly generates two metrics by analyzing the router distribution of the ISP pairs, which are $W^{R,C}_i$ and $S^{R,C}_i$. Moreover, these two metrics are generated from the individual peering willingness and stability of each ASN in the ASN pairs (i.e., $W^{R\rightarrow C}_i, W^{C\rightarrow R}_i, S^{R\rightarrow C}_i$, and $S^{C\rightarrow R}_i$). To check if these metrics can be used to predict whether two ASNs should peer or not, we perform a similar study as was done for GEO-PP (previous subsection) with more than 3,000 ISP pairs. Due to the requirement of high computational power to run RCA-PP, we used a modified version of RCA-PP when generating these metrics. In the modified version, we assumed that the traffic matrices are only confined to the largest cities (42 densely populated area) and the PoP locations were also confined to the largest 148 IXPs in the US. We use the output of RCA-PP as the input to ML-based Random Forest (RF) and Decion Tree (DT) models, and the output of the models are binary classification indicating if the ASN pair should peer or not.

The output from the RF and DT models are shown in Figs. 22 and 23 for the sets of the US-based ISPs and the balanced dataset respectively. For the US-based dataset (Fig. 22), we see that with less training data, the models were biased and has very high (or low) accuracy when predicting peering (or non-peering) relationship. However, with the increment of training data, the peering prediction accuracy decreased slowly while the non-peering prediction accuracy improved faster, thus helping the overall balanced accuracy to go up. Overall, on average, the balanced accuracy improved with training data; however, the highlighted area (in orange), which portrays the error interval, started to expand when training data increases to more than 50%. From this observation, it is likely that having training data of around 50% may be the optimum choice to train the ML models. Fig. 23 shows the same set of results for the balanced ISP dataset. Overall, the prediction performance (balanced accuracy) shows a similar trend to that we had for the US-based ISPs set. However, there are two main differences; firstly, the accuracy is modestly high with low training data, and with the increase of training data, the models do not gain much leverage, and the final accuracy becomes around 66% with 90% training data. Secondly, both class accuracies overall showed a steady or increasing accuracy value, whereas, for the US-based ISPs, the peering relations accuracy decreased (green star) to compensate for the increase in accuracy for the non-peering relationship (blue square). For this set, similar to the US-based ISP set, 50% training data seems to have better error intervals while ensuring close to the best average accuracy achievable.

Since RCA-PP focuses on the internal routing cost, and although it is good for understanding the trends in different types of ISP pairs' peering relationships, it could not perform well in predicting when two ISPs should peer. This observation also tells us that when deciding if two ASNs should peer, they focus on more than just the internal routing costs. To find the importance of the $W^{R,C}_i$ and $S^{R,C}_i$ on decision-making, we looked into the feature importance of these two metrics. Fig. 24 shows the feature importance of all the metrics when computed individually from one ASN's perspective and when the metrics are calculated pairwise. From the figure, it is evident that $S^{R,C}_i$ does not have much effect on the peering decision making of two ASNs. On the other hand, $W^{R,C}_i$ has significant importance on the peering relationship of two ASNs.

Overall Comparison: To test the effectiveness of the proposed framework (both heuristic and ML), we compare our algorithms with other state-of-the-art approaches from the literature. To the best of our knowledge, there is no other prior approach that aims to automate the peering decision process.

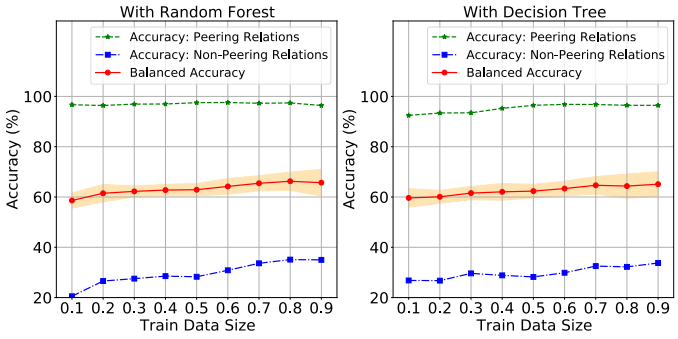

Fig. 20: Performance of GEO-PP (US-based ISPs).

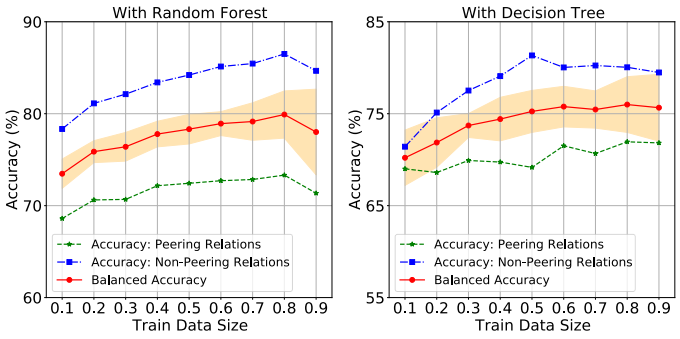

Fig. 21: Performance of GEO-PP (Balanced Dataset).

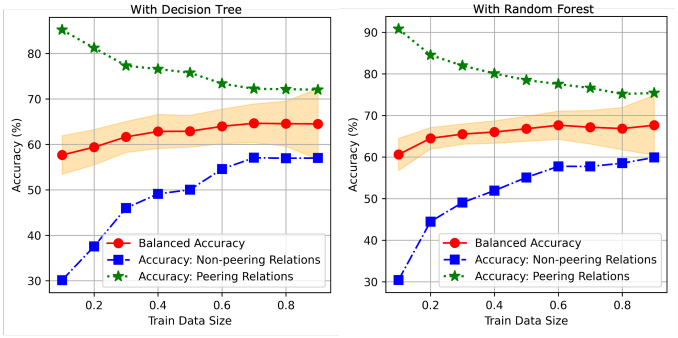

Fig. 22: Performance of RCA-PP (US based ISPs).

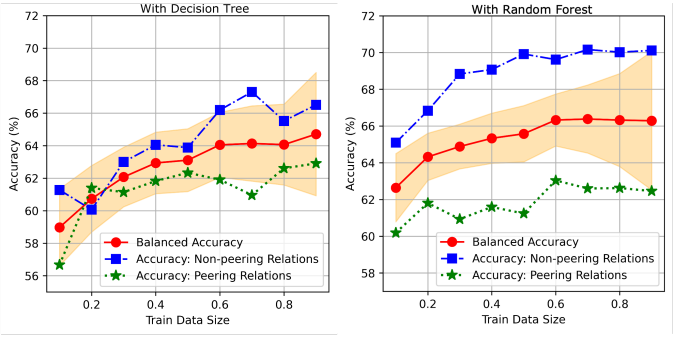

Fig. 23: Performance of RCA-PP (balanced dataset).

Therefore, we develop a baseline method to determine if two ISPs will peer using only the similarity metric. It is based on the idea that ISPs of similar sizes are more likely to peer than those with size discrepancies. This holds true in real-life examples as Tier 1 ISPs have always peered with each other in a settlement-free manner, while being much more selective when peering with providers of lower tiers, with a similar trend followed by Tier 2 and 3 ISPs [5]. The reasoning here is that ISPs will get into peering agreements only if it allows a roughly equal exchange of benefits, which is possible when they have similar levels of networking infrastructure (i.e. size) [70]. To quantify this concept, we take the weighted geometric mean of all three similarity score parameters and compare it with a threshold to identify whether or not the ISPs will peer. Similar to our heuristics, we use the grid search method to determine the values of the weights and threshold that will yield the highest accuracy. For the ML approach, we compare our models with the work in [71] which uses all available data parameters from PeeringDB to formulate 18 different features and predict peering relations through the Random Forest algorithm.

The results are shown in Fig. 25. In case of the heuristic approaches, GEO-PP outperforms both RCA-PP and the similarity baseline with balanced accuracies of 66% and 64% in the US-based and balanced dataset respectively, using the felicity scores given by Eq. 16 (compared to 53% and 53.5% for RCA-PP, 62% and 60% for similarity baseline). When using the ML-based approach, both our algorithms perform much better with accuracies of 66% and 80% for GEO-PP and 66% and 67% for RCA-PP in US-based and balanced datasets respectively. So GEO-PP outperforms RCA-PP in all cases. This can be attributed to two factors. Firstly, there is no

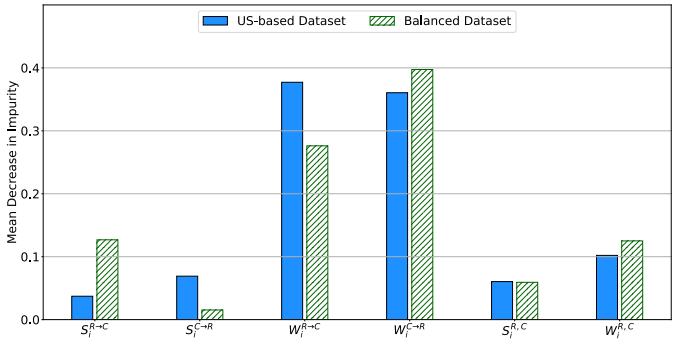

Fig. 24: Feature Importance - RCA-PP

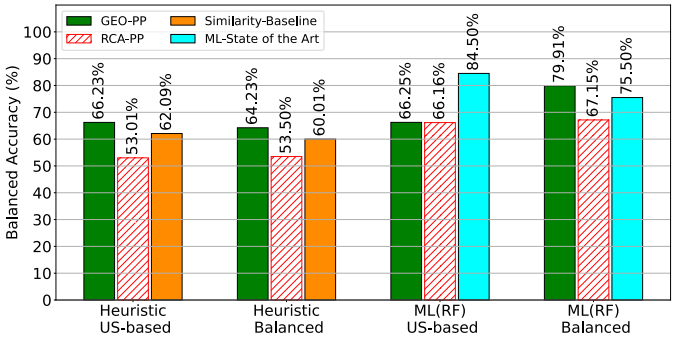

Fig. 25: Overall Performance Comparison of GEO-PP (with Heuristic and Machine Learning)

metric in RCA-PP that represents the similarity between two ASNs. From our observation of the feature importance, it is evident that the similarity between ISPs (in size, geographical coverage, or IP address) and possible extension of the service area (affinity score) are crucial for any ISP pair when deciding to pair up for peering. Secondly, the $W_i^{R,C}$ and $S_i^{R,C}$ metrics

are calculated with the modified RCA-PP method to reduce computational complexity, which may have hindered the actual potential of the RCA-PP method. As for their comparison with the state-of-the-approach in [71], GEO-PP outperforms it in the balanced dataset but lags behind in the US-based case. This can be attributed to the following: 1) ISPs usually require potential partners to peer at a minimum number of locations from a predefined list of IXPs, and 2) the requirement is more prevalent in ISPs operating in the US compared to those with global scope [70]. This data is not publicly available for vast majority of providers so we could not include it in our methodologies. PeeringDB includes a naive representation of this requirement through a binary parameter without listing the actual values which is considered as a feature in [71]. So due to a combination of this, GEO-PP loses out in the US-based dataset but outperforms the state-of-the-art on the balanced dataset. The reason we did not include that specific parameter from PeeringDB in our model is that as it is a naive binary implementation (only specifies if a particular ISP requires multiple locations to peer, without disclosing the actual value), it would not generalize well as evident by the reduction in accuracy of [71] when moving from the biased US-based dataset to the neutral balanced dataset. Another contributing factor to that is the number of features in [71] is 18 compared to 6 and 4 for GEO-PP and RCA-PP respectively, reducing the likelihood of over-fitting in the proposed approaches. This indicates that our model is more robust in situations where there is a balance between the peering and non-peering relations in the training data. Additionally, the approach in [71] cannot predict peering locations that both of our algorithms can, which is another advantage of our framework.

2) Internal Routing Cost: We calculated the internal routing costs using the PoP locations suggested by our proposed methods and the original PoP locations (extracted from PeeringDB using the method discussed in Section VI-A). Fig. 26 shows the predicted and the original PoP locations for the ASN-11686 and ASN-21928 pair (as an example). In the figure, the stars (in black line) point to the original PoP locations of this ASN pair (In total 9 locations). The green squares and red circles depict the predicted PoP locations using GEO-PP and RCA-PP respectively. GEO-PP predicts PoP locations from the locations where the ASN pairs are already peering; thus, the prediction is always a smaller subset of the original PoP locations (as can be seen in the figure). On contrary, RCA-PP can suggest peering locations outside the original PoP locations if there are locations that have routers from both ASNs. However, RCA-PP can only suggest up to four peering locations (an upper bound set by us to control computational complexity). The main takeaway from the above discussion is that in most cases, both our methods only suggest a few locations for peering, whereas, in actuality there could be many common PoP locations for the respective ISP pair.

Fig. 27 shows the cost comparison analysis for the proposed methods. The figure is a histogram where the (total internal routing) costs for each ASN pairs for the two methods are normalized by dividing their respective values by the original internal routing cost (using the original PoP locations) calculated by the method discussed in Section V. From the figure,

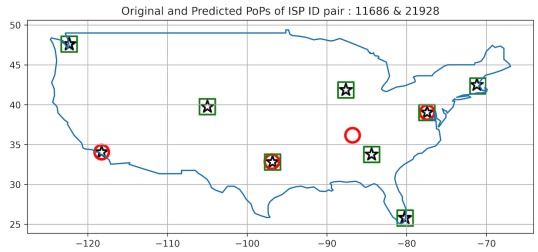

Fig. 26: Predicted PoPs (Green: GEO-PP, Red: RCA-PP) with inferred PoPs from PeeringDB (Black)

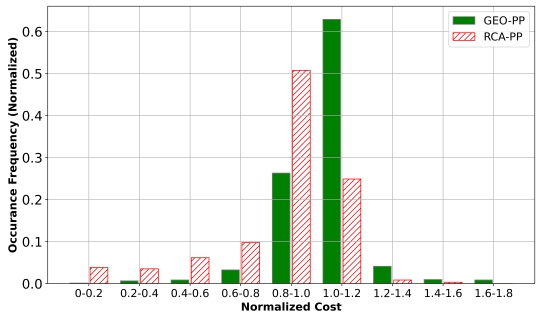

Fig. 27: Cost Comparison of our methods with original

it is evident that the proposed methods with their proposed PoP locations have internal routing costs less than or equal to the original cost calculated with the original PoP locations for the majority of the cases. Moreover, we see that almost always the internal routing cost of RCA-PP is smaller than the original cost, suggesting the achievement of the objective of the method, which is minimizing the cost. Furthermore, although GEO-PP has some occurrences that have cost greater than the original, that cost is still within (20%) of the original cost (having a bar in the range 1-1.2).

The comparison between the original (ground truth) and the predicted PoP locations is done in terms of their total number, and is shown in Fig. 28. Similar to Fig. 27, the total number of predicted PoP locations is divided by the original number of PoP locations to achieve a normalized PoP value. Hence, having a value less than or equal to one means the predicted number of PoPs is less than or equal to the original number of PoPs. From the figure, we see that GEO-PP almost always predicts PoP locations less than the number of original PoP(s), whereas RCA-PP in 20% cases (Fig. 28) predicts more PoPs than the original. This increased number of PoP prediction is because RCA-PP has the freedom to choose locations where just the routers of the ASN are present. Also, this increased PoP locations helped RCA-PP to decrease the internal routing cost, which we observed in Fig. 27.

VII. Summary and Conclusion

We introduced "meta-peering" as a combined effort towards automating the entire peering process among ISPs. As part of the automation process, we focused on the peer selection technique and formulated the peer selection sub-process as an optimization problem. Using PeeringDB and CAIDA datasets, we estimated the traffic matrix of an ISP, identified its PoPs,

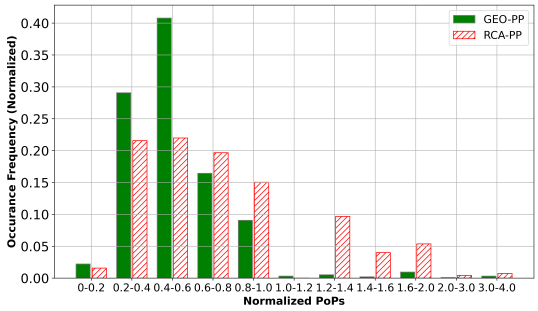

Fig. 28: Comparison of (Suggested Peering Points to Original PoP(s) - our methods to original

and then described a framework to suggest the best candidates for a requester ISP along with its best peering locations. We introduced the concept of felicity score to represent the interest of peering between an ISP pair. We found that ISPs mostly (more than half of them) prefer to offload as much traffic as they can, and similarity between two ISPs play a major role in the decision making of peering partner(s). Our felicity score calculations warrant further investigation and feedback from the ISP community to establish more precise and stable metric sets for peer selection. Moreover, we provide a web-service with basic features for test purpose.

We have proposed two meta-peering methods, namely GEO-PP and RCA-PP to select the peering partner and peering point(s). When making a peering decision, i.e., whether two ISPs should peer, GEO-PP focuses more on the geographical overlap (quantifying the business competition) and size similarity between the two ISPs. On the other hand, RCA-PP focuses on the internal routing cost from the perspective of the ISPs. Our study confirmed that the similarity between two ISPs is crucial when deciding peering partner. The willingness factors calculated by both methods showed promising significance, followed by the geographical overlap between ISPs. However, the stability metric proposed by RCA-PP, which focuses on the deviation of cost from the minimum cost attainable by the ISP pairs, did not have much significance in peer selection decisions. The peering partner prediction accuracy of both methods improved when the willingness and stability metrics (calculated by the proposed methods) were used as an input to an ML algorithm, and the output of the ML algorithm was to predict if the ISP pair should peer or not. Overall, GEO-PP and RCA-PP predicted approximately 74% and 67% of the peering relationships (from CAIDA's determination) correctly (considering both datasets) when aided by ML. Our analysis also showed that both the proposed methods are very good at predicting peering locations. In most cases, the predicted PoP locations resulted in a decrease in internal routing cost while ensuring the total number of predicted PoP count is not greater than the actual PoP count found from PeeringDB.

Our work can be extended on two primary directives. The first directive involves the amalgamation of the proposed heuristic methodologies to check whether this integration enhances overall performance. Additionally, further exploration of various peering metrics outlined in Section III-A and their incorporation into our proposed heuristics is possible. Lastly, the results of the RCA-PP were obtained from a modified version of RCA-PP (due to computational complexity), and further experiments can be done on that to check its actual performance. The second directive centers on developing data-driven peering models. This choice is motivated by two principal factors: 1) the abundant availability of publicly accessible data on ISPs, and 2) some preliminary work has shown promising results [71]. Hence, ML-based models where the input encompasses either complete or partial data on ISPs, and the output is peer selection, could be a plausible objective using the extensive measurement datasets.

ACKNOWLEDGMENTS

This work is supported in part by NSF awards 1814086 and 1816396.

REFERENCES

[1] P. K. Dey, S. Mustafa, and M. Yuksel, "Meta-Peering: towards automated ISP peer selection," in *Proceedings of the Applied Networking Research Workshop*, 2021, pp. 8–14.

[2] M. I. I. Alam, S. Mustafa, K. Kar, and M. Yuksel, "Modeling and Automating ISP Peering Decision Process: Willingness and Stability," in *Proceedings of IEEE International Conference on Communications (ICC)*, May 2022, pp. 371–376.

[3] P. Maigron, "World - Autonomous System Number statistics - Sorted by number," Accessed: Sep 03, 2024. [Online]. Available: https://www-public.imtbs-tsp.eu/~maigron/rir-stats/rir-delegations/world/world-asn-by-number.html

[4] D. Meyer, "Management of ISPs, Peering, China Firewall Cause Big Challenges for SD-WAN," SDX-Central, Accessed: Sep 03, 2024. [Online]. Available: https://www.sdxcentral.com/articles/news/management-of-isps-peering-china-firewall-remain-big-challenges-for-sd-wan/2018/08/

[5] W. B. Norton, *The Internet peering playbook: connecting to the core of the Internet*. DrPeering Press, 2014.

[6] A. Ahmed, Z. Shafiq, H. Bedi, and A. Khakpour, "Peering vs. transit: Performance comparison of peering and transit interconnections," in *2017 IEEE 25th International Conference on Network Protocols (ICNP)*. IEEE, 2017, pp. 1–10.

[7] A. Formoso, J. Chavula, A. Phokeer, A. Sathiaseelan, and G. Tyson, "Deep Diving into Africa's Inter-Country Latencies," in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, 2018, pp. 2231–2239.

[8] E. Carisimo, J. Del Fiore, D. Dujovne, C. Pelsser, and J. Alvarez-Hamelin, "A first look at the Latin American IXPs," *ACM SIGCOMM Computer Communication Review*, vol. 50, pp. 18–24, 03 2020.

[9] K. Zarifis, T. Flach, S. Nori, D. Choffnes, R. Govindan, E. Katz-Bassett, Z. M. Mao, and M. Welsh, "Diagnosing path inflation of mobile client traffic," in *International Conference on Passive and Active Network Measurement*. Springer, 2014, pp. 23–33.

[10] K. Hinton, J. Baliga, M. Feng, R. Ayre, and R. S. Tucker, "Power consumption and energy efficiency in the internet," *IEEE Network*, vol. 25, no. 2, pp. 6–12, 2011.

[11] C. Ge, Z. Sun, N. Wang, K. Xu, and J. Wu, "Energy Management in Cross-Domain Content Delivery Networks: A Theoretical Perspective," *IEEE Transactions on Network and Service Management*, vol. 11, no. 3, pp. 264–277, 2014.

[12] A. Nikkhah and S. Jordan, "Towards Equitable Peering: A Proposal for a Fair Peering Fee Between ISPs and Content Providers," *IEEE Transactions on Network and Service Management*, pp. 1617–1633, 2023.

[13] I. van Beijnum, "Transit vs peering: what makes sense when?" Packet Pushers, Jan. 2015, Accessed: Sep 02, 2024. [Online]. Available: https://packetpushers.net/transit-vs-peering-makes-sense/

[14] V. Stocker, G. Smaragdakis, W. Lehr, and S. Bauer, "Content may be king, but (peering) location matters: A progress report on the evolution of content delivery in the internet," 2016.

[15] "The 'Donut Peering' Model: Optimizing IP Transit for Online Video," MZIMA White paper, Sep. 2009, Accessed: Sep 02, 2024. [Online]. Available: http://tnarg.org/pdf/donut_peering.pdf

[16] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the expansion of Google's serving infrastructure," in *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 2013, pp. 313–326.

[17] R. T. Ma, "Pay or Perish: The Economics of Premium Peering," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 2, pp. 353–366, 2017.

[18] B. Schlinker, T. Arnold, I. Cunha, and E. Katz-Bassett, "PEERING: Virtualizing BGP at the Edge for Research," in *Proceedings of ACM International Conference on Emerging Networking Experiments And Technologies (CoNEXT)*, 2019, p. 51–67.

[19] J. Snijders, "Automated Peering Operations," NANOG 70, 2017.

[20] Z. Liu, A. Manousis, G. Vorsanger, V. Sekar, and V. Braverman, "One Sketch to Rule Them All: Rethinking Network Flow Monitoring with UnivMon," in *Proceedings of the 2016 ACM SIGCOMM Conference*, 2016, pp. 101–114.

[21] B. Li, J. Springer, G. Bebis, and M. H. Gunes, "A survey of network flow applications," *Journal of Network and Computer Applications*, vol. 36, no. 2, pp. 567–581, 2013.

[22] B. Woodcock and M. Frigino, "2016 Survey of Internet Carrier Interconnection Agreements," Packet Clearing House (Survey), Nov. 2016, Accessed: Sep 10, 2024. [Online]. Available: https://www.pch.net/resources/Papers/peering-survey/PCH-Peering-Survey-2016/PCH-Peering-Survey-2016.pdf

[23] "BIG-IP Link Controller," f5 Datasheet, Accessed: Sep 10, 2024. [Online]. Available: https://www.f5.com/pdf/data-sheet/big-ip-platforms-datasheet.pdf

[24] Y. Sverdlik, "Google and level 3 interconnect network backbones," DataCenter Knowledge, Feb. 2016, Accessed: Sep 10, 2024. [Online]. Available: https://www.datacenterknowledge.com/networking/google-and-level-3-interconnect-network-backbones

[25] S. Secci, J.-L. Rougier, A. Pattavina, F. Patrone, and G. Maier, "Peering equilibrium multipath routing: a game theory framework for internet peering settlements," *IEEE/ACM Transactions on Networking*, vol. 19, no. 2, pp. 419–432, 2010.

[26] M. I. I. Alam, E. Anshelevich, K. Kar, and M. Yuksel, "Pricing for Efficient Traffic Exchange at IXPs," *IEEE/ACM Transactions on Networking*, no. 01, pp. 1–16, 2023.

[27] C.-H. Hsu and M. Hefeeda, "ISP-Friendly Peer Matching without ISP Collaboration," in *Proceedings of the 2008 ACM CoNEXT Conference*, 2008, pp. 1–6.

[28] R. T. Ma, D.-m. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement," in *Proceedings of the 3rd international workshop on Economics of networked systems*, 2008, pp. 61–66.

[29] 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*, 2018, pp. 265–278.

[30] J. Ramseyer and J. Heichman, "Peering automation at Facebook," Facebook, May 2021, Accessed: Sep 10, 2024. [Online]. Available: https://engineering.fb.com/2021/05/20/networking-traffic/peering-automation

[31] V. Giotsas, G. Smaragdakis, B. Huffaker, M. Luckie *et al.*, "Mapping Peering Interconnections to a Facility," in *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*. ACM, 2015, p. 37.

[32] D. Temkin, "The Real Cost of Public IXPs," NANOG 67, Jun. 2016, Accessed: Sep 10, 2024. [Online]. Available: https://archive.nanog.org/sites/default/files/Temkin_The_Real_Cost.pdf

[33] S. Bafna, A. Pandey, and K. Verma, "Anatomy of the Internet Peering Disputes," *arXiv preprint arXiv:1409.6526*, 2014.

[34] T. W. Cable, "Ipv4 and ipv6 settlement-free peering policy." [Online]. Available: https://help.twcable.com/twc_settlement_free_peering_policy.html

[35] N. Shetty, G. Schwartz, and J. Walrand, "Internet QoS and Regulations," *IEEE/ACM Transactions on Networking (TON)*, vol. 18, no. 6, pp. 1725–1737, 2010.

[36] R. Johari and J. N. Tsitsiklis, "Routing and Peering in a Competitive Internet," in *2004 43rd IEEE CDC (IEEE Cat. No. 04CH37601)*, vol. 2. IEEE, 2004, pp. 1556–1561.

[37] R. Mahajan, D. Wetherall, and T. Anderson, "Negotiation-Based Routing Between Neighboring ISPs," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation-Volume 2*. USENIX Association, 2005, pp. 29–42.

[38] D. Arena, "WORKSHOP: STARTING AN IXP," RIPE Regional Meeting, Dubrovnik (HR), Aug. 2011. [Online]. Available: https://www.ripe.net/participate/meetings/regional-meetings/see/dubrovnik-2011/presentations/IXP%20Workshop%20Part%20I%20-%20Daniele%20Arena.pdf

[39] ThousandEyes, "ISP Peering & Settlement-Free Peering," Accessed: Sep 10, 2024. [Online]. Available: https://www.thousandeyes.com/learning/techtorials/isp-peering

[40] J. C. Cardona Restrepo and R. Stanojevic, "A History of an Internet eXchange Point," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 2, pp. 58–64, 2012.

[41] "AS Relationships – with geographic annotations," CAIDA, Accessed: Sep 10, 2024. [Online]. Available: http://www.caida.org/data/as-relationships-geo/

[42] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis *et al.*, "Using PeeringDB to Understand the Internet Peering Ecosystem," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.

[43] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig, "Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN," *ACM SIGCOMM Computer Communication Review*, vol. 48, no. 1, pp. 28–34, 2018.

[44] "CAIDA AS Rank," CAIDA, 2019, Accessed: Sep 10, 2024. [Online]. Available: https://api.asrank.caida.org/v2/docs

[45] Swisscom, "Swisscom Peering Policy IP-Plus (AS 3303)," May 2018, Accessed: Sep 10, 2024. [Online]. Available: https://www.swisscom.ch/content/dam/swisscom/de/biz/wholesale/ott/20181105_swisscom-peering-policy.pdf

[46] M. Caesar and J. Rexford, "BGP Routing Policies in ISP Networks," *IEEE network*, vol. 19, no. 6, pp. 5–11, 2005.

[47] D. Goodin, "Google goes down after major BGP mishap routes traffic through China," Ars Technica, Nov. 2018, Accessed: Sep 10, 2024. [Online]. Available: https://arstechnica.com/information-technology/2018/11/major-bgp-mishap-takes-down-google-as-traffic-improperly-travels-to-china/

[48] R. Chirgwin, "Google routing blunder sent Japan's Internet dark on Friday," The Register, Aug. 2017, Accessed: Sep 10, 2024. [Online]. Available: https://www.theregister.co.uk/2017/08/27/google_routing_blunder_sent_japans_internet_dark/

[49] C. Orsini, A. King, D. Giordano, V. Giotsas, and A. Dainotti, "BGPStream: A Software Framework for Live and Historical BGP Data Analysis," in *Proceedings of the 2016 Internet Measurement Conference*, 2016, pp. 429–444.

[50] Cisco, "ThousandEyes: BGP Route Network Monitoring Solution," June 2024, Accessed: Sep 10, 2024. [Online]. Available: "https://www.thousandeyes.com/solutions/bgp-and-route-monitoring"

[51] J. Snijders, "Nlnog ring," Accessed: June 2024. [Online]. Available: https://ring.nlnog.net

[52] "Intelligent Routing Platform," Noction Network Intelligence, White paper, Accessed: Sep 10, 2024. [Online]. Available: https://www.noction.com/resource_center/irp_white_paper

[53] Z. Hu, Y. Qiao, and J. Luo, "Coarse-grained traffic matrix estimation for data center networks," *Computer Communications*, vol. 56, pp. 25–34, 2015.

[54] X. Ros-Roca, L. Montero, J. Barceló, K. Nökel, and G. Gentile, "A practical approach to assignment-free Dynamic Origin–Destination Matrix Estimation problem," *Transportation Research Part C: Emerging Technologies*, vol. 134, p. 103477, 2022.

[55] K. Swetha, U. Prabu, G. Angel, and Y. Lahari, "Traffic Matrix Estimation Techniques-A Survey on Current Practices," in *2023 International Conference on Sustainable Computing and Data Communication Systems (ICSCDS)*. IEEE, 2023, pp. 661–668.

[56] P. Tune, M. Roughan, and C. Wiren, "Hierarchical Traffic Matrices: Axiomatic Foundations to Practical Traffic Matrix Synthesis," in *2018 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2018, pp. 1591–1600.

[57] H. Zhou, L. Tan, F. Ge, and S. Chan, "Traffic matrix estimation: Advanced-Tomogravity method based on a precise gravity model," *International Journal of Communication Systems*, vol. 28, no. 10, pp. 1709–1728, 2015.

[58] "Internet usage penetration in the United States in November 2021, by state," Statista, 2021, Accessed: Sep 10, 2024. [Online]. Available: https://www.statista.com/statistics/184691/internet-usage-in-the-us-by-state/

[59] "Cisco Visual Networking Index: Forecast and Trends, 2017–2022 White Paper," CISCO, Feb. 2019, Accessed: Sep 10, 2024. [Online].

Available: https://cloud.report/Resources/Whitepapers/eea79d9b-9fe3-4018-86c6-3d1df813d3b8_white-paper-c11-741490.pdf
[60] N. Feamster, "Revealing utilization at internet interconnection points," *arXiv:1603.03656*, 2016.
[61] BEREC, "Draft berec report on ip-interconnection practices in the context of net neutrality," BEREC, Tech. Rep., 2017, Accessed: Sep 10, 2024. [Online]. Available: https://www.berec.europa.eu/sites/default/files/files/document_register_store/2012/12/BoR_%2812%29_130__IP_IC_Assessment_NN_Report_publication2.pdf
[62] PeeringDB, "PeeringDB Website," Accessed: Sep 10, 2024. [Online]. Available: https://www.peeringdb.com
[63] V. Kshirsagar, "Flood Fill Algorithm Explained (with C++ & Python code)," 2023. [Online]. Available: https://favtutor.com/blogs/flood-fill-algorithm/
[64] C. for International Earth Science Information Network CIESIN Columbia University, "Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11," NASA Socioeconomic Data and Applications Center (SEDAC), 2018, Accessed: Sep 10, 2024. [Online]. Available: https://doi.org/10.7927/H4JW8BX5
[65] H. Tangmunarunkit, R. Govindan, S. Shenker, and D. Estrin, "The Impact of Routing Policy on Internet Paths," in *Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No. 01CH37213)*, vol. 2. IEEE, 2001, pp. 736–742.
[66] L. Gao and F. Wang, "The Extent of AS path Inflation by Routing Policies," in *Global Telecommunications Conference, 2002. GLOBECOM'02. IEEE*, vol. 3. IEEE, 2002, pp. 2180–2184.
[67] CAIDA, "CAIDA Data," Accessed: Sep 10, 2024. [Online]. Available: http://caida.org/data
[68] "The CAIDA AS Relationships Dataset," June 2024, https://www.caida.org/catalog/datasets/as-relationships/.
[69] P. Liashchynskyi and P. Liashchynskyi, "Grid Search, Random Search, Genetic Algorithm: A Big Comparison for NAS," *arXiv preprint arXiv:1912.06059*, 2019.
[70] A. Nikkhah and S. Jordan, "Analysis of the Requirements of Settlement-Free Interconnection Policies," *IEEE Transactions on Network and Service Management*, vol. 20, no. 4, pp. 4028–4046, 2023.
[71] S. Mustafa, P. K. Dey, and M. Yuksel, "Peer Me Maybe?: A Data-Centric Approach to ISP Peer Selection," in *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 2022, pp. 1–9.

Md Ibrahim Ibne Alam received his B.Sc. and M.Sc. degree in Electrical and Electronics Engineering from Bangladesh University of Engineering & Technology (BUET), Dhaka, Bangladesh in 2014 and 2017 respectively. He is currently pursuing his Ph.D. degree in Electrical Engineering at Rensselaer Polytechnic Institute. His research interests include Game theory in networked systems, peering process between ISPs, pricing policy at IXPs, optimization of MAC protocol, and Machine Learning with a focus on Federated Learning and NLP.

Anindo Mahmood is a Ph.D. student at the Department of ECE in the University of Central Florida. He received his M.S. degree in Electrical Engineering from the University of Texas Rio Grande Valley in 2022. His research interests include peering in computer networks, machine learning techniques for wireless communication and spectrum sharing.

Prasun Kanti Dey received his B.Sc. degree in Computer Science and Engineering (CSE) from Bangladesh University of Engineering and Technology, Bangladesh. He received his M.Sc. in CSE from University of Nevada- Reno, Reno, NV in 2016 and Ph.D. in Computer Engineering from University of Central Florida in 2019. He is currently with MathWorks Inc. as a Senior Software Engineer. His research interests include network management and security, SDN and distributed systems, network measurement and performance analysis, and network economics. He is a member of both ACM and IEEE.

Murat Yuksel is a Professor at the ECE Department of the University of Central Florida (UCF), Orlando, FL, and a Visiting Scientist at MIT Lincoln Labs. He served as the Interim Chair of ECE at UCF from 2021 to 2022. Prior to UCF, he was a faculty member at the CSE Department of the University of Nevada, Reno, NV. He received his B.S. degree in computer engineering from Ege University, Izmir, Turkey in 1996, and M.S. and Ph.D. degrees in computer science from Rensselaer Polytechnic Institute, Troy, NY, in 1999 and 2002, respectively. His research interests are in the areas of networked, wireless, and computer systems with a recent focus on wireless systems, optical wireless, spectrum sharing, network economics, network architectures, and network management. He has been on the editorial boards of Computer Networks, IEEE Transactions on Communications, IEEE Transactions on Machine Learning in Communications and Networking, and IEEE Networking Letters. He has published more than 200 papers at peer-reviewed journals and conferences, and is a co-recipient of five Best Paper, one Best Paper Runner-up, and one Best Demo Awards. He is a senior member of IEEE and ACM.

Koushik Kar received his Ph.D. in electrical and computer engineering from the University of Maryland at College Park in 2002 and has been a faculty member at Rensselaer Polytechnic Institute since then. He has held short-term visiting researcher appointments at Bell Laboratories and IBM Research. His primary research expertise is in developing and analyzing low-complexity and decentralized optimization algorithms for communication networks and other networked systems. Dr. Kar received the CAREER Award from the National Science Foundation in 2005, and won multiple best paper awards. He has served on the editorial board of journals such as IEEE/ACM Transactions on Networking and IEEE Transactions on Mobile Computing, and has been a Technical Program Committee Co-Chair for international conferences such as ACM MOBIHOC 2016 and IEEE LANMAN 2020.