

Meta-Peering: Towards Automated ISP Peer Selection

Prasun Kanti Dey
The MathWorks, Inc.
Massachusetts, USA
pdey@mathworks.com

Shahzeb Mustafa
University of Central Florida
Florida, USA
shahzeb.mustafa@knights.ucf.edu

Murat Yuksel
University of Central Florida
Florida, USA
murat.yuksel@ucf.edu

ABSTRACT

We introduce *meta-peering*, a term that encompasses the set of tools needed to ease and automate the Internet Service Provider (ISP) peering process; starting with identifying a list of ISPs that are likely to peer, generating respective BGP configurations, and monitoring these sessions for outages or peering agreement violations. In this paper, we describe how existing tools can be leveraged to implement *meta-peering* and focus on instrumenting the automation of peer selection process. Utilizing PeeringDB and CAIDA datasets to identify possible peering points for requester and candidate ISPs, we estimate candidate ISP's traffic matrix and consider ISPs' internal policies to generate acceptable peering contracts.

CCS CONCEPTS

- Networks → Network architectures; Programming interfaces; Network services; Network management;

KEYWORDS

Peering, Traffic Matrix, IXP, Interconnection, Automation

ACM Reference Format:

Prasun Kanti Dey, Shahzeb Mustafa, and Murat Yuksel. 2021. Meta-Peering: Towards Automated ISP Peer Selection. In *Applied Networking Research Workshop (ANRW '21), July 24–30, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3472305.3472325>

1 INTRODUCTION

With over 89K Autonomous Systems (ASes) [1] around the world, it is impossible for an individual AS to have global

This work was supported in part by U.S. National Science Foundation award number 1814086.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ANRW '21, July 24–30, 2021, Virtual Event, USA
© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8618-0/21/07...\$15.00
<https://doi.org/10.1145/3472305.3472325>

reachability. Organizations capable of maintaining enormous network backbone are extremely rare [2]. ISPs have to collaborate and establish *transit* or *peering* interconnections. Some may prefer *peering* over *transit* for better control on routing, low latency and reduced cost. Since traffic can reach the destination AS directly, the propagation delays for peering paths are often smaller [3] and it can help avoid extraneous traffic detours [4, 5].

Transit and peering, both come with the added complexity of achieving a stable state in accordance with inter-AS routing policies. Various tools have helped ISPs minimize human error by automating Border Gateway Protocol (BGP) sessions and network monitoring [6–9]. However, optimal selection of peers and Point-of-Presences (PoPs) is challenging to automate because (1) Peering is often an informal handshake [10], (2) Peering with multiple ISPs is a hassle even with modern switches [11, 12], (3) ISPs may choose sub-optimal PoPs in peering deals to minimize “bit-miles” [13], (4) Some ISPs are more *selective* than others in peering.

Potential peers need to be identified based on the estimated traffic, customer cone size and peering policy among other aspects. This paper focuses on answering a key question in the Internet peering: “How far can peer selection be automated?” We envision the ISP peering process to be well integrated and automated. We consider the entirety of every tool or algorithm needed for peering automation as *Meta-Peering*. This requires major innovations at different levels such as making 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. We focus on the first of many steps along this road: establishing a quantifiable system for optimal peer selection. Major contributions include:

- detailing the *meta-peering* concept and breaking down the peering activities into four phases;
- 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 [14] for access to recommended peering deals generated using our approach.

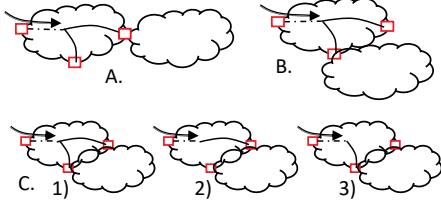


Figure 1: Possible PoP locations

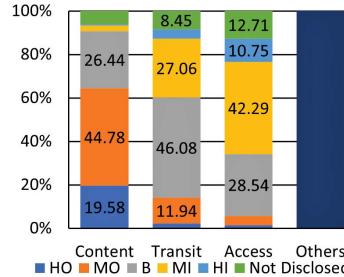


Figure 2: ISP traffic ratio

The paper is organized as follows: Sec. 2 discusses our motivation for automated peering and presents a *Meta-Peering* overview. Sec. 3 presents the *Meta-Peering* framework and formulates the problem of peer selection. Results are presented in Sec. 4, followed by summary and future directions.

2 MOTIVATION AND OVERVIEW

Dey et al. portray an approximate timeline of peering relationship evolution [15]. ISPs can choose either bi-lateral private peering using dedicated physical links or multi-lateral public peering using a Route-Server. In either case, network administrators typically negotiate peering deals at various events [16]. In some cases, ISPs start off with a “trial peering” period to avoid future tussles. Optimal peer selection is a hard problem and current methodologies are clearly inefficient. Earlier works [17, 18] formulate an optimal peering problem to determine the maximum peering points and their strategic placement or a negotiation-based platform for ISPs to determine routing path for traffic exchange. Such peering interconnections are facilitated by Internet eXchange Providers (IXPs) [19].

We focus on automating the peer selection process in bi-lateral agreements, and suggesting possible PoPs according to ISP specific criteria. Figure 1 shows three combinations of PoP locations for two ISPs where they can peer. *A* and *B* are the cases when ISPs are not located in the same PoP, but agree to peer on the closest place from both. 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 (*case 2* and *3*). ISPs can also peer without being physically present in an IXP in a *Remote Peering* (RP) [20] manner. Despite being an option in practice, we do not consider RP in our model as they are opaque and controversial in terms of their performance benefits.

2.1 Meta-Peering

Inspired by Norton’s *Peering Playbook* [21], we break down the entire peering process into four phases and focus specifically on the automation effort undertaken in each phase.

	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

HI: Heavily Inbound; MI: Medium Inbound; B: Balanced; MO: Medium Outbound; HO: Heavily Outbound

Table 1: Traffic ratios of peering ISPs (%)

2.1.1 Pre-Peering Phase: Key peering metrics. From an economic perspective, if peering can reduce operation cost for two ISPs, and if rerouting transit traffic through a peering channel is possible, they can be expected to peer.

More Control: An ISP may be interested in peering to get more control over its traffic and influence route path selection instead of letting someone transit provider to treat it as “hot potato”. Each time an ISP peers with someone new, the congestion reduces, reliability increases, and therefore, the end-to-end service quality for the users get improved [22].

Traffic Ratio: Figure 2 presents the traffic ratio distribution of all Access (5,207), Transit (2,413), and Content (1,619) ISPs from PeeringDB. CAIDA published the inter-ISP relationship information of 339K ISP pairs. Of them, 209K pairs were peering. Ignoring 57K pairs with undisclosed traffic ratios, Table 1 shows the rest of them as the percentage of peering ISP pairs, based on their traffic ratio type. It can be seen that some traffic ratios are more likely to peer than others, with Balanced-Balanced having the highest peering rate.

PoP Frequency: Having more PoPs attracts more ISPs that are interested in expanding their network. Access ISPs usually have a lower number of PoPs as they operate in smaller regions. Transit ISPs form the Internet backbone, with countrywide fiber network and a large number of PoPs. Content Providers (CPs) are also spreading their footprint by putting caches directly inside Access ISPs and Data-Centers [23].

Customer Cone Size: Traffic volume and advertised IP address space are vital for choosing 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 [24] shows a strong correlation between the advertised prefix count and traffic volume for Access and Transit ISPs. This validates our use of public BGP-advertised address space for estimating network’s traffic volume.

2.1.2 Peer Selection Phase: Many ISPs advertise their criteria and willingness for peering. We collected key requirements that are commonly expected from a peering candidate. Some trendy desiderata are: 24/7 Network Operations Center

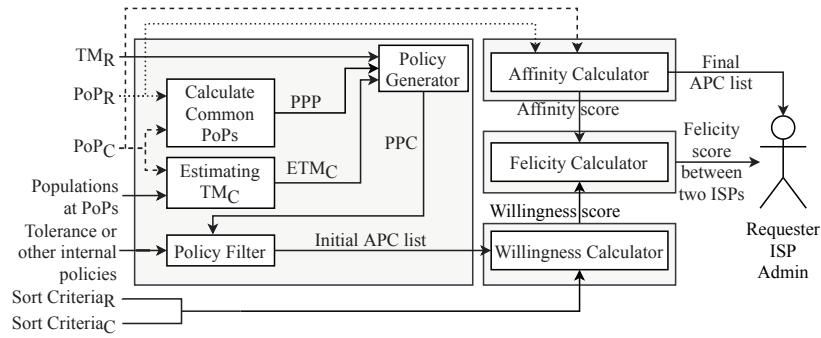


Figure 3: Automatic peer selection framework

(NOC), omnipresent geographic footprint, adequate backbone capacity, essential peering port size in PoPs maintaining financial stability, and redundancy of the requester network. Most ISPs look for a similar sized partners with capacity sufficient to handle the projected load.

2.1.3 Establishing BGP Session Phase: Business relations (customer/provider/peer), intention to limit the routing table size for scalability, control over in/out-bound traffic play vital roles in setting up a BGP session [25]. Erroneous configurations often lead to instability, misconfigured route announcements, and blackhole routes. Some recent incidents [26, 27] amplify the importance of meticulous BGP configuration. The first-ever “peering-over github”, *Coloclue*, [6] aims to prevent such occurrences and to make the peering process more dynamic. It identifies common IXPs, calculates the max-prefix and sets up BGP sessions. To the best of our knowledge, this is the only automation effort towards managing BGP sessions between two networks. Dynam-IX[28] and Route Bazaar [29] try to reduce the communication gap between two ISPs willing to peer and reduce the average time needed for the whole process.

2.1.4 Post-Peering (Monitoring) Phase. An ISP monitors all the remote BGP sessions with its neighbors for ensuring the least amount of BGP outage or black hole, and it compares both ways aggregate traffic so that the measured traffic ratio does not violate the agreement. To automate the process, ISPs either set up their internal monitoring system or purchase such services from several third party providers like BGPmon [30]. RING is another effort to enhance the BGP sessions monitoring process that provides a *friendly* access for a participating entity to all the other participants to view its own network from outside [31]. Finally, Noction Routing Platform [32] allows bypassing congestion and outages.

3 AUTOMATING PEER SELECTION

While making peering decisions, ISP admins have access to local traffic statistics, but have limited data about their

\mathcal{PPP}	Possible Peering Points
\mathcal{PPC}	Possible Peering Contracts
\mathcal{APC}	Acceptable Peering Contracts
\mathcal{APC}^*	Optimal \mathcal{APC}
$\mathfrak{R}(\mathcal{APC}_i)$	Rank of \mathcal{APC}_i
TM	Traffic Matrix
ETM	Estimated Traffic Matrix
$\bar{\omega}_{R_i}$	Willingness score for Requester
α_R	Affinity score for Requester
f_{RC}	Felicity score

Table 2: List of symbols

competitor ISPs. Intrigued by the case studies discussed in Section 2.1.2, we have developed an automated peer selection framework. We doubt there is a perfect model that fits for all ISPs, as such, we perform a careful sanity check to validate our observation before recommending the ISP peers.

3.1 Framework

Figure 3 presents an overview of our proposed framework that leverages publicly available data and produces a guideline of peering contracts based on requester ISP’s internal policy. Considering the PoP locations, traffic matrices, port capacities at common ISP locations, the framework suggests whether the candidate may agree with a particular peering offer or not. It can also be used to identify potential peers, and formulate respective peering contracts. The heuristic function requires limited shared data as it simulates the candidate’s specifications from publicly available information.

We use PeeringDB to identify common PoP locations among the requester and candidate ISPs, (Possible Peering Points, \mathcal{PPP}). Peering is possible at all combinations of these \mathcal{PPP} . Therefore, the number of Possible Peering Contracts (\mathcal{PPC}) between two ISPs generated by *Policy Generator* is $2^{|\mathcal{PPP}|} - 1$. The algorithm sorts \mathcal{PPC} according to the requester’s internal policies and sorting strategy. *Policy Filter* generates the Acceptable Peering Contracts (\mathcal{APC}) by eliminating impermissible options from the list if they do not qualify. The sorting criteria can be **Own** (maximize outbound traffic towards the candidate), **Diff** (minimize the absolute difference between in vs. outbound traffic), or **Ratio** (choose peers with lower in/out-bound traffic ratio).

The *Willingness Calculator* receives the sorted \mathcal{APC} , ranks them from the *most* preferred to the *least*, and normalizes the rank values from 1 to 0. We call these scores as *willingness scores*. The \mathcal{APC} for both candidate and requester ISPs are sorted again according to the willingness scores. Using \mathcal{APC}_R (Requester) and \mathcal{APC}_C (Candidate), the *Willingness Calculator* finds the optimal \mathcal{APC}^* , which is preferable for both, and also calculates the combined *willingness score* for

every \mathcal{APC}_i^* . Next, the framework considers the overlap between ISPs' coverage areas to calculate the *affinity* score. Utilizing these two metrics, the *felicity* score is computed. This value represents the ultimate likeliness of peering between two ISPs and is displayed to the network admin.

3.2 Methodology

3.2.1 Modeling ISP Network & TM. Earlier study [16] found that “peering point placement problems” under traffic cost constraints are NP-complete. As such, we use heuristics and simplify the APC generation process to solve the same issue as part of generating each APC. We use the population data for each PoP location and use the *Gravity Model* [33] to compute Traffic Matrix (TM) and Estimated TM (ETM). *Policy Generator* then measures the traffic flow between two ISPs at particular PoPs for each of the \mathcal{PPC} . We use TMs for estimating the traffic amount between every possible origin-destination (OD) node pairs to model candidate traffic volume. In a real world scenario, requester ISP will only have to estimate the candidate ISPs TM. However, we had to estimate both ISPs' TMs due to lack of access to any internal data. In that regard, we used port capacities at each PoP location to proportionately distribute the offloaded traffic for each ISP. This allowed us to simulate the fact that ISPs tend to keep port usage under 50% capacity [34, 35] and also gave a better TM approximation.

3.2.2 Willingness Score. In most of the cases, \mathcal{APC}_R and \mathcal{APC}_C will contain the same APCs but in different order of preference. In case of an unusual scenario where \mathcal{APC}_R and \mathcal{APC}_C do not include the same items, we set an infinite as rank value for the missing APC in the counterpart's \mathcal{APC} to make sure the list contains the same items but is preferred the least. Let, \mathfrak{R}_{R_i} and \mathfrak{R}_{C_i} be the rank of a particular APC_i in \mathcal{APC}_R and \mathcal{APC}_C . We calculate the individual ISP's willingness scores $\bar{\omega}_{R_i}$, and $\bar{\omega}_{C_i}$ as following:

$$\bar{\omega}_{R_i} = 1 - \frac{\mathfrak{R}_{R_i} - 1}{|\mathcal{APC}|}, \quad \bar{\omega}_{C_i} = 1 - \frac{\mathfrak{R}_{C_i} - 1}{|\mathcal{APC}|} \quad (1)$$

We want both of these values to be closer to each other so we express the combined willingness score, $\bar{\omega}_{RC_i}$ for APC_i as the geometric mean of $\bar{\omega}_{R_i}$ and $\bar{\omega}_{C_i}$. Finally, \mathcal{J}^* stores the orders of the APCs that are preferred by both but prioritizes the requester's preference. Here, $j_i = \mathfrak{R}(\mathcal{APC}_i^*)$ and $j_i \in \mathcal{J} = 1 \dots 2^{|\mathcal{PPC}|} - 1$ where j_1 ($\mathfrak{R} = 1$) means the best choice for both ISPs, the next preference gets j_2 ($\mathfrak{R} = 2$), and the rest follow accordingly.

Using individual APC's willingness score for each item in \mathcal{APC} , we define ISP pair's combined willingness score:

$$\omega_{RC} = \frac{\sum_{i \in \mathcal{APC}} \bar{\omega}_{RC_i}}{|\mathcal{APC}|} \quad (2)$$

Table 3: Generating willingness scores

\mathcal{APC}_i	\mathfrak{R}_{R_i}	\mathfrak{R}_{C_i}	$\bar{\omega}_{R_i}$	$\bar{\omega}_{C_i}$	$\bar{\omega}_{RC_i}$	$\bar{\omega}_{RC_i} * \bar{\omega}_{R_i}$	\mathcal{J}^*
1	5	1	0.429	1.0	0.655	0.281	4
2	6	2	0.286	0.857	0.495	0.141	6
3	2	4	0.857	0.571	0.699	0.599	1
4	7	3	0.143	0.714	0.319	0.046	7
5	3	5	0.714	0.429	0.553	0.395	2
6	4	6	0.571	0.286	0.404	0.231	5
7	1	7	1.0	0.143	0.378	0.378	3

We present an optimal rank calculation procedure in Table 3. For each \mathcal{APC}_i , \mathfrak{R}_{R_i} and \mathfrak{R}_{C_i} show the requester and the candidate's ranking of that particular APC. $\bar{\omega}_{R_i}$, $\bar{\omega}_{C_i}$, and $\bar{\omega}_{RC_i}$ refer to their individual and combined willingness scores. After that, we utilize the requester's preference of a particular APC before ranking them and produce the \mathcal{APC}^* . For example, $\bar{\omega}_{RC_1} > \bar{\omega}_{RC_5}$. If we ignore the requester's preference, APC_1 will be suggested as better than APC_5 . But this is not the case here, as the requester placed this APC in 5-th position. If we consider the requester's individual preference of an APC along with the combined willingness, we identify a better deal for the requester. In the right-most column, we show the rank of the APCs in their optimal order. So, APC_3 is the best option for the requester, APC_5 comes next and so on. This is the order they will appear in \mathcal{APC}^* . Figure 4b shows the willingness scores for all APCs for Charter-PCCW using *own* sorting criteria.

3.2.3 Affinity Score. An ISP may be more interested in peering if the relationship would expand its coverage area. We call this interest as *affinity* score of an ISP pair. To represent the coverage area of an ISP, we used its PoPs to draw a polygon. Figure 4a shows the comparison between the coverage areas of two ISPs, Charter and PCCW. Since a larger area coverage does not necessarily mean that an ISP serves a wide customer pool because of the uneven population distribution, we calculate the total population in its covered region. 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) [36]. We express the affinity scores for requester (α_R) and candidate (α_C) based on the overlap as:

$$\alpha_R = \frac{A_C - A_o}{A_R \cup A_C} = \frac{A_C - A_o}{(A_R - A_o) + (A_C - A_o) + A_o} \quad (3)$$

$$\alpha_C = \frac{A_R - A_o}{(A_R - A_o) + (A_C - A_o) + A_o} \quad (4)$$

Where A_R and A_C are the population in coverage areas of R and C, respectively, and A_o is the overlapped area's population. Similar to the combined willingness score, we use geometric mean to calculate the combined affinity score:

$$\alpha_{RC} = \sqrt{\alpha_R * \alpha_C}. \quad (5)$$

3.2.4 Felicity Score. We take weighted geometric mean of *willingness* and *affinity* scores to generate the ultimate felicity score for the requester-candidate ISP pair as:

$$f_{RC} = (\omega_{RC}^\beta * \alpha_{RC}^\gamma)^{\frac{1}{\beta+\gamma}} \quad (6)$$

where β and γ are constants. This felicity scoring emphasizes that the ISPs would want to peer more if they have both low coverage overlap and high willingness for possible peering deals. Our framework computes *felicity scores* between two ISPs at a time. Requester ISP's network admin needs to run this framework several times to compare multiple ISPs and identify the best peer. They can then utilize the associated \mathcal{APC}^* to suggest the PoPs to the candidate ISPs for peering.

4 RESULTS

We test our proposed *Meta-Peering* framework on real-world ISP data in the context of USA (see App. A). As part of this initial study, we select 23 large-scale ISPs (according to CAIDA AS-Rank [37]) with coverage in the USA according to PeeringDB [38]. For population database, we used US Census Bureau [39] and GPW [36].

Most ISP pairs had a small number of common PoPs, and therefore both of their $|\mathcal{APC}|$ and $|\mathcal{PPC}|$ were small. 79% of the pairs (of the selected 23 ISPs) had ≤ 15 and 11% had 16–21 common PoPs. It is no surprise that *Google* and *HE* pair was leading the sequence with 51 common PoPs. As $|\mathcal{PPC}|$ grows exponentially, we limit the number of common PoPs for potential peering locations to only 15. Currently, the algorithm does not impose any stringent filtering except for *traffic ratio*. Interestingly enough, *Cogent* did not disclose its traffic ratio in PeeringDB. While our algorithm excludes ISPs from peering consideration if the traffic ratio is undisclosed but since *Cogent* is a tier-1 transit ISP, we assumed its traffic to be balanced and accorded full consideration for peering.

We have also developed a web-application to provide a basic experience of the overall service [14]. While still under active development, it showcases key *Meta-Peering* features as discussed in this paper. We plan to open-source our code base along with the all data that we have used for this analysis as this can motivate the community to contribute.

4.1 Candidate Recommendation

Of all the 506 possible pairs, we found (from PeeringDB) that 65 prefer *diff*, 439 prefer *own* and only 2 prefer *ratio* as their sorting criteria. As such, we use *own* criteria for all ISPs for recommending potential candidates. After candidate identification, \mathcal{APC} is compiled that gives a list of good peering locations in order of contract desirability. For example, Figure 4 shows the willingness score and the coverage

area, that is used to calculate the affinity score (that remains the same for every contract in \mathcal{APC}), of the Charter-PCCW pair. Our analysis found that *Charter* gets the highest felicity scores among all ISPs. Its possible candidates are *Cablevision*, *Cox*, and *CenturyLink*. For each of these pairs, we generate multiple APCs that include the best peering locations. As an example, the APC below shows that Charter and CenturyLink should be peering at the following locations:

APC CoreSite Denver, Equinix Dallas, Equinix Miami, Equinix Chicago, Equinix Ashburn Exchange, Equinix Los Angeles, Equinix San Jose

4.2 Validation

To validate our identification of peering results, we compared them to CAIDA data. In order to train our model to align well with CAIDA, we varied the weights (β and γ separately) in Eq. 6 to calculate the *felicity* score and used a threshold value. Selecting the right threshold value is essential as if it is too low, the False-Positive rate will be too high, and if it is too low, False-Negative rate will be too high. Using $\beta = 0.15$, $\gamma = 0.85$ and $threshold = 0.55$, we were able to successfully identify 85% of the peering pairs from CAIDA with a False-Negative percentage of 15%. Figure 5 shows a detailed distribution of our results in comparison to CAIDA. It would be interesting to see whether the pairs we suggested to peer end up actually peering in a few years.

The results show that our algorithm performs best in identifying peering possibilities in which a transit ISP was involved. The algorithm suggested 56 pairs for transit providers and 40 of them are actually peering. Similarly, out of the 24 existing peering deals between CPs and transit ISPs, our algorithm successfully identified 22 cases. CAIDA data shows that none of the CPs from our list are peering, and this aligns with the fact that they treat each other as rivals and have minimal incentive in peering. Although our algorithm successfully identifies 85% of actual peering relations, its performance is sub-par for suggesting Access-Content peering contracts. This may be due to the traffic flow disparity between heavily outbound CPs and heavily inbound access ISPs. However, further exploration of Access-Content peering is needed.

4.3 Utilizing the Holistic View

Throughout the analysis, we observed that the algorithm tries to increase the felicity threshold as much as possible while selecting a candidate. As we have already mentioned, we are using the threshold value of 0.55. In our algorithm, it is possible that we identify peering possibility between ISP A and ISP B when ISP A is the requester, but when ISP B becomes requester, ISP A may not consider ISP B as a good fit for peering. This happens due to the higher threshold value. To check if we can improve the overall peering suggestion

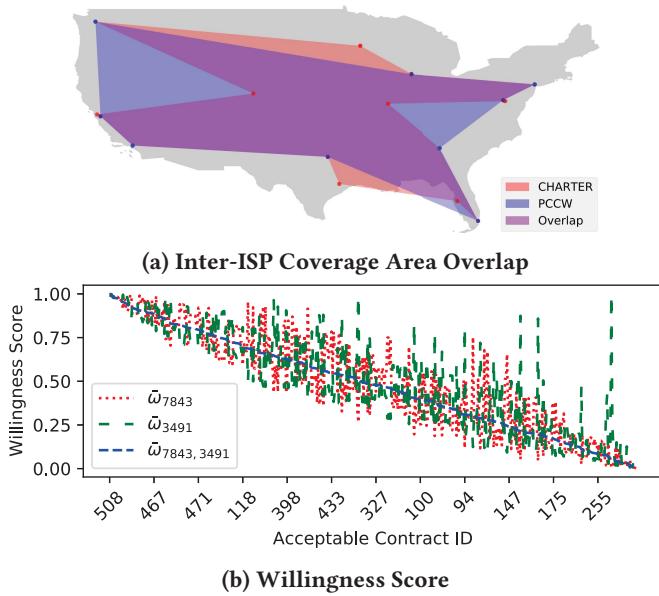


Figure 4: Charter-7843, PCCW-3491 Peering Results

utilizing a holistic view, we consider successful peering pair only when the felicity score between ISP A and ISP B is greater than the threshold and the felicity score of ISP B and ISP A is also greater than the threshold. Akin to our intuition, we identified that if we reset the weights differently ($\beta = 0.1$, $\gamma = 0.85$) and lower the threshold ($= 0.3$), our success rate (identifying already peering pairs from CAIDA) increases to 89%. This observation is particularly important since it shows that IXPs can play a vital role in future peering decisions as they have a bigger picture of the overall Internet traffic and can suggest good peering deals to ISPs. The validation with holistic view is also included in Figure 5.

5 SUMMARY

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 treated 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, 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 we could successfully identify 153 ISP pairs that are already peering according to CAIDA. We could not identify 27 of existing peering pairs, but that is mostly because some ISPs are big and cover the

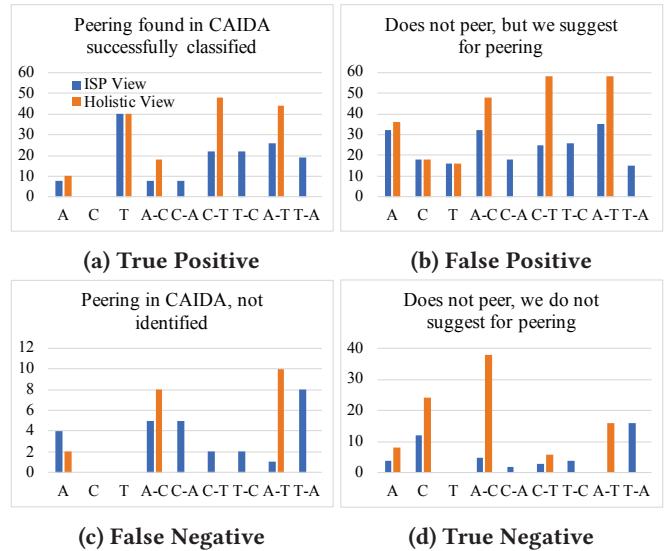


Figure 5: CAIDA validation and new recommendation comparison from individual ISP’s point-of-view vs. holistic view (Section 4.3).

entire area of the other ISP. Our felicity score calculations warrant further investigation and feedback from the ISP community to establish more precise and stable metric sets for peer selection. Finally, we provide a web service [14] with basic features of *Meta-Peering* for further testing and experimentation. We envision this service to become a tool for ISP network admins to use when making peering decisions.

The work we present here is at its infancy, developing a full-fledged prototype, implementing it within an ISP or an IXP may provide better insight and will improve the prediction. The framework considers geographic overlapping and traffic exchange willingness between two ISPs. But, it is easy to add newer modules such as cost benefit analysis and security overhead to extend the framework further. Besides, we only considered medium level US-based ISPs, as a result, higher ranked (i.e., CAIDA AS-Rank) ISPs such as Telia Company AB, or GTT Communications were not included. Expansion of current ISP set and their operational regions in the study is another vein for future work.

A LIST OF ISPS (ASN)

Access: Cable ONE-11492, Cablevision-6128, CenturyLink-209, Charter-7843, Comcast-7922, Cox-22773, TDS-4181, Windstream-7029
Content: Akamai-20940, Amazon-16509, Ebay-62955, Facebook-32934, Google-15169, Microsoft-8075, Netflix-2906
Transit: Columbus Networks-23520, Cogent-174, Hurricane Electric-6939, NTT-2914, PCCW-3491, Sprint-1239, Verizon-701, Zayo-6461

REFERENCES

[1] P. Maigron. (2019) World - Autonomous System Number statistics - Sorted by number. [Online]. Available: https://www-public.imsbs-tsp.eu/~maigron/RIR_Stats/RIR_Delegations/World/ASN-ByNb.html

[2] D. Meyer, "Management of ISPs, Peering, China Firewall Cause Big Challenges for SD-WAN," SDXCentral, Aug. 2018. [Online]. Available: <https://www.sdxcentral.com/articles/news/management-of-isps-peering-china-firewall-remain-big-challenges-for-sd-wan/2018/08/>

[3] 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.

[4] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett, "Peering at the internet's frontier: A first look at ISP interconnectivity in africa," in *International Conference on Passive and Active Network Measurement*. Springer, 2014, pp. 204–213.

[5] H. Galperín, "Localizing Internet infrastructure: Cooperative peering in Latin America," *Telematics and Informatics*, vol. 33, no. 2, pp. 631–640, 2016.

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

[7] "Transit vs peering: what makes sense when?" Noction, Nov. 2016. [Online]. Available: <https://www.notion.com/blog/using-peeringdb-to-set-up-your-internet-exchange-peering>

[8] E. B. Claise, "Cisco Systems NetFlow Services Export Version 9," Cisco Systems, Tech. Rep., 2004.

[9] 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.

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

[11] "BIG-IP Link Controller," f5 Datasheet. [Online]. Available: <https://www.f5.com/pdf/products/big-ip-link-controller-ds.pdf>

[12] R. Perkin, "Peering with Route Servers." [Online]. Available: https://www.rogerperkin.co.uk/bgp/peering-with-route-servers/?doing_wp_cron=1546013473.1083629131317138671875

[13] Y. Sverdlik, "Google and Level 3 Interconnect Network Backbones," DataCenter Knowledge, Feb. 2016. [Online]. Available: <https://www.pch.net/resources/Papers/peering-survey/PCH-Peering-Survey-2016/PCH-Peering-Survey-2016.pdf>

[14] "Meta-peering: Peering Query form." [Online]. Available: <https://prod-metapeering.herokuapp.com/>

[15] P. K. Dey and M. Yuksel, "Peering Among Content-Dominated Vertical ISPs," *IEEE Networking Letters*, vol. 1, no. 3, pp. 132–135, 2019.

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

[17] 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.

[18] 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.

[19] T. Böttger, G. Antichi, E. L. Fernandes, R. di Lallo, M. Bruyere, S. Uhlig, and I. Castro, "The Elusive Internet Flattening: 10 Years of IXP Growth," *CoRR*, vol. abs/1810.10963, 2018. [Online]. Available: <http://arxiv.org/abs/1810.10963>

[20] 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 IMC 2018*. ACM, 2018, pp. 265–278.

[21] W. B. Norton, *The Internet peering playbook: connecting to the core of the Internet*. DrPeering Press, 2011.

[22] ThousandEyes, "ISP Peering & Settlement-Free Peering." [Online]. Available: <https://www.thousandeyes.com/learning/techtorials/isp-peering>

[23] 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 IMC*. ACM, 2013, pp. 313–326.

[24] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis *et al.*, "Using peer-DB to understand the peering ecosystem," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.

[25] M. Caesar and J. Rexford, "BGP routing policies in ISP networks," *IEEE network*, vol. 19, no. 6, pp. 5–11, 2005.

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

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

[28] P. Marcos, M. Chiesa, L. Müller, P. Kathiravelu, C. Dietzel, M. Canini, and M. Barcellos, "Dynam-IX: A dynamic interconnection exchange," in *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*, 2018, pp. 228–240.

[29] I. Castro, A. Panda, B. Raghavan, S. Shenker, and S. Gorinsky, "Route bazaar: Automatic interdomain contract negotiation," in *15th Workshop on Hot Topics in Operating Systems (HotOS {XV})*, 2015.

[30] H. Yan, R. Oliveira, K. Burnett, D. Matthews, L. Zhang, and D. Massey, "BGPmon: A real-time, scalable, extensible monitoring system," in *2009 Cybersecurity Applications & Technology Conference for Homeland Security*. IEEE, 2009, pp. 212–223.

[31] J. Snijders, "NLNOG RING," 2010. [Online]. Available: <https://ring.nlrog.net>

[32] "Intelligent Routing Platform," Noction Network Intelligence, white paper. [Online]. Available: https://www.notion.com/resource_center/irp_white_paper

[33] D. Lam, D. C. Cox, and J. Widom, "Teletraffic modeling for personal communications services," *IEEE Communications Magazine*, vol. 35, no. 2, pp. 79–87, 1997.

[34] J. Brodkin, "Netflix packets being dropped every day because Verizon wants more money," Ars Technica, Feb. 2014. [Online]. Available: <https://arstechnica.com/information-technology/2014/02/netflix-packets-being-dropped-every-day-because-verizon-wants-more-money/>

[35] N. Feamster, "Revealing utilization at Internet interconnection points," in *The 44th Research Conference on Communication, Information and Internet Policy 2016*. TPRC, 2016.

[36] E. Doxsey-Whitfield, K. MacManus, S. B. Adamo, L. Pistolesi, J. Squires, O. Borkovska, and S. R. Baptista, "Taking advantage of the improved availability of census data: a first look at the gridded population of the world, version 4," *Papers in Applied Geography*, vol. 1, no. 3, pp. 226–234, 2015.

[37] "CAIDA AS RANK," CAIDA. [Online]. Available: <http://as-rank.caida.org/>

[38] "PeeringDB," 2004. [Online]. Available: <https://www.peeringdb.com>

[39] "Population Estimates APIs," US Census Bureau. [Online]. Available: <https://www.census.gov/data/developers/data-sets/popest-popproj/popest.html>