

Complexities in Internet Peering: Understanding the “Black” in the “Black Art”

Aemen Lodhi
Georgia Institute of Technology
aemen.lodhi@gatech.edu

Nikolaos Laoutaris
Telefonica Research
nikos@tid.es

Amogh Dhamdhare
CAIDA, UCSD
amogh@caida.org

Constantine Dovrolis
Georgia Institute of Technology
constantine@gatech.edu

Abstract—Peering in the Internet interdomain network has long been considered a “black art”, understood in-depth only by a select few peering experts while the majority of the network operator community only scratches the surface employing conventional rules-of-thumb to form peering links through ad hoc personal interactions. Why is peering considered a black art? What are the main sources of complexity in identifying potential peers, negotiating a stable peering relationship, and utility optimization through peering? How do contemporary operational practices approach these problems? In this work we address these questions for Tier-2 Network Service Providers. We identify and explore three major sources of complexity in peering: (a) inability to predict traffic flows prior to link formation (b) inability to predict economic utility owing to a complex transit and peering pricing structure (c) computational infeasibility of identifying the optimal set of peers because of the network structure. We show that framing optimal peer selection as a formal optimization problem and solving it is rendered infeasible by the nature of these problems. Our results for traffic complexity show that 15% NSPs lose some fraction of customer traffic after peering. Additionally, our results for economic complexity show that 15% NSPs lose utility after peering, approximately, 50% NSPs end up with higher cumulative costs with peering than transit only, and only 10% NSPs get paid-peering customers.

Keywords: Internet, Autonomous System interconnections, settlement-free, paid peering, IXPs, economic utility

I. INTRODUCTION

The Internet interdomain network is a complex network of approximately 50,000 Autonomous Systems (ASes) which interconnect with one another through transit (customer-provider) or peering links. Peering links fall in one of two categories: (a) settlement-free and (b) paid. A recent study showed the presence of a rich peering fabric at a major European IXP, with more than 67% of all possible links formed between 400 member ASes [1]. Moreover, the importance of peering has grown in view of paid-peering which has risen as one of the main instruments for catching up with asymmetric traffic due to CDNs, online video traffic, etc. However, despite its widespread adoption, peering has been considered a “black art” [2]. The findings of a recent survey showed that 99.5% peering relationships were formed through ad hoc personal interactions, without any formal economic analysis or agreements [3].

Following the recent Level3-Comcast peering dispute an intense public debate started around peering [4]. It touched upon nearly all aspects of network interconnections including pricing, traffic ratios, costs, performance, network neutrality, the power of access ISPs, regulation, etc. However many fundamental questions are still unanswered: What makes peering so complex that it is understood by a small community of peering coordinators only? What are the main sources of complexity in peering that force the majority of the peering community to resort to simple rules-of-thumb for link formation? Why has peering defied the development of a methodical quantitative approach? Is the general notion that peering saves costs, always valid? What limits the ability of peers to forecast the future of their peering relationships?

In this work we address these questions for a specific class of ASes, the *Tier-2 Network Service Providers (NSPs)*. We choose to focus on NSPs (or transit providers) because they appear in all three AS roles in the interdomain network, namely providers, customers and peers, simultaneously. Additionally, their “selective” peering policies are more complex than the simpler “open” and “restrictive” policies of stubs and Tier-1 providers respectively. We focus on three major sources of complexity which we evaluate separately:

- 1) Limited ability to determine and accurately predict traffic flows.
- 2) Limited ability to accurately forecast the effect of peering on utility owing to a complicated pricing structure.
- 3) Infeasibility of determining the optimal set of peers because of the combined effects of topology and routing policy.

We use computational modeling and simulations, with realistic parameters and interconnection strategies to study these problems. Our results for traffic complexity show that formation of new peering links may cause 15% of Tier-2 NSPs to lose the number of flows being routed through them while some Tier-2 NSPs can lose as much as 80% of their customer traffic because of a wrong peering decision. Furthermore, our results for economic complexity show that 10% NSPs actually lose utility after engaging in peering. We find that

the conventional wisdom that peering reduces transit costs is not always true as approximately 50% of the NSPs end up with higher cumulative costs. However, peering may also increase customer revenue. We show that an exhaustive analysis of a combination of peering links is necessary to ensure that the set of peers of the NSP yields optimal utility. However, such a scheme is practically infeasible because of a large number of potential peers.

The rest of this paper is structured as follows. We briefly discuss related work in section II and motivate our study by outlining contemporary peering practices in section III. In section IV, we explain our model and its parameterization. In sections V, VI and VII we discuss the three major complexities in peering. We conclude and outline future work in section VIII.

II. RELATED WORK

The complexity of the interdomain network, encompassing its topology, dynamics and economics, has been the subject of many recent works. Faratin et al. discuss the growing complexity of Internet interconnection and economics [6], [7]. Steenbergen and Norton discuss the operational practices of peering operators [8], [2].

Ager et al. show that a rich peering fabric exists at IXPs which generally eludes topological studies of the interdomain network [1]. Lodhi et al. describe the state of the peering ecosystem as captured by PeeringDB [9]. Discovering the AS-level topology has been an area of active research, yet peering links remain elusive [10], [11]. Gursun et al. study the potential of different techniques for interdomain traffic matrix completion [12].

There has been much prior work on economic analysis of peering in the interdomain network covering various aspects of the peering ecosystem including pricing, paid-peering, utility optimization heuristics, policy adoption, etc. An extensive line of research has taken an analytical approach to explore paid-peering. Shrimali et al. study linear pricing schemes for paid-peering between two providers [13]. Dhamdhare et al. propose a quantitative framework to determine the value of a peering link for both peers involved [14]. Ma et al. analyze the use of Shapley value for revenue distribution among peers [15]. There is also prior work on the game-theoretic analysis of settlement-free peering and transit vs. peering [16], [17], [18], [19], [20]. However, for reasons of mathematical tractability, these models often study networks with a small number of players. Therefore, many complexities that arise out of the interaction of a large number of autonomous agents do not appear in these works. Furthermore, they also ignore many real-world features of the interdomain network, e.g., a highly skewed traffic matrix, geographic co-location constraints, ratio-based peering policies, capacity constraints of IXP ports, complex non-linear pricing structure, etc. Another line of work has focused on models in which ASes select transit providers, settlement-free peers or peering policies based on economic factors and other constraints to optimize their utility [21], [22], [23], [24].

III. CONTEMPORARY PEERING PRACTICES

In this section we briefly describe how NSPs carry out peering in practice based on our discussions with network operators and peering policies published by NSPs.

A. Identification of potential peers

Identification of potential peers is the first step in peering. One of the ways that operators use to find potential peers is to analyze traffic flow data collected locally using NetFlow to identify other ASes with whom they exchange significant volume of traffic¹. Typically, this analysis only informs about traffic that is generated and consumed within the AS (and exchanged with the potential peer), and does not include traffic from the AS' customers (and potential peer's customers) that may be exchanged over the peering link after link formation.

B. Selective peering criteria

Since Tier-2 NSPs are in the transit business, they prefer other ASes to be their customers instead of peers. Thus, most of them adopt a *selective* criteria of some sort to deter peering applications by smaller ASes. For example, many large NSPs require their peers to be co-located at multiple geographic locations, maintain a lower bound on the traffic exchanged, 24×7 NOC, a lower bound on the capacity of the physical network, etc. NSPs almost always deny peering to their existing customers, while many NSPs also deny peering to their previous customers². Our background discussions with peering operators reveal that some large NSPs even deny peering to those ASes which are deemed potential customers.

C. Preventing asymmetric benefits

Peering is supposed to be mutually beneficial. Most ASes involved in peering would expect that the cost of peering would be borne equally by both parties. Furthermore, many ASes, typically large ISPs and NSPs, demand that the benefits derived by both parties in a peering relationship be roughly equal. For example, NSPs do not want their competitors to free-ride their networks through peering. NSPs use traffic exchanged over the peering link as a proxy for the benefits derived from a peering relationship. In order to limit the asymmetry of benefits, NSPs generally require that the ratio of inbound to outbound traffic be within certain bounds. If the bound is set to 1 then the traffic in both directions should be equal, if it is set to a value less than 1 then the inbound traffic should be less than the outbound traffic and vice versa. Analysis of peering policies published by large NSPs and our discussions with network operators reveal that contemporary values for this bound, in general, are between 2 and 3 [25], [26], [27]. This allows an NSP to form settlement-free relationships with most content providers but excludes a few "hyper-giants" [28] with whom the inbound traffic would be orders of magnitude greater than the outbound traffic.

¹What is deemed "significant" varies from AS to AS, but for typical Tier-2 NSPs it ranges anywhere from 1 to 7 Gbps.

²Previous customers are denied peering to dissuade existing customers from terminating their contracts and doing the same.

D. Paid peering

In general, NSPs form settlement-free peering relationships with all other ASes which satisfy the requirements of their peering policies, while offering paid peering relationships to those which do not do so. Paid peering is similar to a conventional customer-provider relationship, however, whereas a transit provider is responsible for providing a customer connectivity to the entire Internet, a paid-peering provider only offers access to its customers, in addition to itself. Correspondingly, the price for paid-peering is lower than conventional transit prices. Little public information is available about the modalities of paid-peering as well as prices for paid-peering, as details of most relationships are held private through non-disclosure agreements.

IV. MODEL

We describe our network formation and economic model in this section. We consider a population of N nodes, representing ASes. Each node seeks peers which conform to its peering policy. We describe different components of the model in detail as follows.

Link formation: In our model IXPs represent geographic locations. Each node is present at at least one IXP. A link can be formed between two nodes if they are co-located at an IXP. Nodes interconnect through two types of links: transit (or customer-provider) and peering links. A transit link between nodes i and j is denoted by L_{ij}^t and a peering link by L_{ij}^p .

Traffic matrix & routing: We denote the traffic matrix by T^3 . We denote the traffic generated at i and consumed at j by T_{ij} . We denote the traffic that is mutually exchanged between two nodes by $\hat{T}_{ij} = T_{ij} + T_{ji}$. If nodes i and j are connected through a link then the traffic that is sent from i to j is denoted by T'_{ij} . Note that $T_{ij} \leq T'_{ij}$ because T'_{ij} may also include traffic from the customers of i and j being exchanged over the link L_{ij} .

In our model, traffic is routed over the shortest path subject to two common policy constraints in the Internet [10] “prefer-customer-over-peer-over-provider links” and the “valley-free” routing property.

Transit Provider Assignment: Our model exogenously assigns each node with at least one transit provider with a few exceptions. We do not assign providers to Tier-1 nodes. The customer-provider relationships are fixed and do not change over the course of network formation.

The set of transit providers of j is denoted by $P(j)$ and the set of all customers of j are denoted by $C(j)$. A node which does not have any customers is designated a *stub* while a node which does not have a provider is designated a *Tier-1 NSP*; all other nodes are designated as *Tier-2 NSPs*. Our nodes of interest are invariably chosen from the set of Tier-2 NSPs.

Peering Policies: Nodes form peering relationships based on their peering policies which are assigned to them based on their status in the network hierarchy. We use the following peering policies in our model:

- 1) *Restrictive:* Nodes using this strategy do not peer with any other node unless required to ensure global reachability. Restrictive policy is assigned to Tier-1 NSPs.
- 2) *Open:* Nodes using this policy peer with all co-located nodes. All stubs use this peering policy.
- 3) *Selective:* Selective policy is used by Tier-2 NSPs. A node i using Selective peering policy will agree to peer with another node j if the traffic exchanged between them conforms to the following condition:

$$\frac{T'_{ji}}{T'_{ij}} \leq \sigma \quad (1)$$

where σ is a traffic ratio constraint which is uniform and constant across all nodes using Selective policy⁴. It is difficult to estimate T'_{ij} (and vice versa) without actually forming L_{ij}^p . Therefore, nodes use local traffic, T_{ij} , instead of actual traffic, T'_{ij} , to identify peers and verify policy constraints prior to link formation. All players of interest use Selective peering policy.

Since peering is a bilateral relationship both peers must conform to each other’s peering policies before a link is formed. The set of peers of a node i is denoted by $F(i)$.

Transit Cost and Revenue: Traffic exchanged over customer-provider links is metered. For the traffic $(T'_{ij} + T'_{ji})$ exchanged between i and its provider j over the link L_{ij}^t , i incurs a transit cost $TC(i, j)$ given by:

$$TC(i, j) = E^t(T'_{ij} + T'_{ji}) \times (T'_{ij} + T'_{ji}) \quad (2)$$

where $E(T'_{ij} + T'_{ji})$ is the price (\$/Mbps) for the corresponding traffic volume. The total transit cost of i is given by:

$$TC(i) = \sum_{j \in P(i)} TC(i, j) \quad (3)$$

The transit revenue, of i , $TR(i)$, is the sum of the costs incurred by the customers of i for their transit links with it.

Peering Costs at IXPs: Players utilize the *ports* at the IXPs to peer with one another and exchange their peering traffic. IXP costs are fixed monthly recurring costs based on the number and type of ports that a node has acquired at the IXP. When i peers with j , it checks if any of its existing IXP ports with leftover capacity $T'_{ij} + T'_{ji}$ and routes over that port, otherwise it acquires a new port of minimum size required to accommodate the traffic.

Let $p_c(i)$ be the number of ports of capacity c Mbps being utilized by i and $E^p(c)$ be the price of a port with capacity c . Then the total settlement-free peering cost of i , $IC(i)$ is given by:

$$IC(i) = \sum_c E^p(c) \times p_c(i) \quad (4)$$

⁴In reality, Selective peering policies have additional constraints, e.g., co-location at more than one location, minimum requirements for the volume of traffic exchanged, etc. However, we only use traffic ratios for simplicity.

³ T represents the 95th percentile traffic of the 5-min traffic distribution.

Paid Peering: There is no publicly available data about paid-peering prices. Therefore, we use anecdotal evidence gathered from peering coordinators and various online peering discussions [29] to set the paid-peering price for traffic volume t to one half that of the transit price for the same traffic volume. Finally, both paid-peers use a private interconnect at the IXP to exchange their traffic separate from settlement-free peering traffic. For simplicity, the private interconnect incurs a cost equal to the cost of the public interconnect with the same capacity. The paid-peering revenue and costs for node i are denoted by $PC(i)$ and $PR(i)$ respectively.

Analysis for a paid-peering relationship is carried out only if one of the peers does not satisfy the other’s peering policy requirements. For example, if j does not satisfy i ’s policy constraints, then i carries out a cost-benefit-analysis of accepting j as its paid-peering customer, based on traffic $T_{ij} + T_{ji}$. The cost-benefit-analysis involves calculating the effect of moving the traffic $T_{ij} + T_{ji}$ from the link on which it is currently being routed⁵ to the proposed link L_{ij}^p . i offers j to become a paid-peering customer if the cost-benefit-analysis indicates that acquiring j would increase its utility. j upon receiving the offer carries cost-benefit-analysis at its own end and accepts the offer if its analysis shows that its utility will also increase. On the other hand, if i does not satisfy j ’s peering policy, then an offer to become a paid-peering customer is made by j and a similar analysis is carried out, albeit with roles reversed. Thus, a paid-peering relationship is formed if and only if cost-benefit-analysis by both nodes shows that their utility will increase as a result of the relationship. Note that the actual cost may differ from the one estimated by cost-benefit-analysis because the actual traffic $(T'_{ij} + T'_{ji}) \geq (T_{ij} + T_{ji})$.

Utility: The utility of a player i is determined by its peering links, the traffic traversing those links and how this traffic is distributed over IXP ports. The utility of i is given by:

$$\pi_i = TR(i) + PR(i) - TC(i) - PC(i) - IC(i) \quad (5)$$

Note that although the transit prices and the traffic matrix T are constant, the underlying topology changes as i chooses different peers. Hence, the costs, revenues and utilities may change as the topology changes.

The objective of player i is to maximize its π_i through peering.

Network Formation: Starting from a random population, we create an initial topology by assigning each node (except Tier-1 nodes) with at least one transit provider using the constraints described in the model of Lodhi et al. [30]. It produces a network hierarchy similar to that of the Internet at scale, without any peering links. Tier-1 nodes are not assigned transit providers and they form a complete mesh of peering links among themselves similar to the Tier-1 ASes.

After the creation of initial topology, network formation proceeds in discrete iterations called *rounds*. In each round all Tier-2 NSPs play once, one at a time. When a Tier-2 NSP

⁵The current link carrying this traffic may be a transit link or an existing settlement-free or paid-peering link.

| Transit & Paid Peering | | | IXP Peering | |
|------------------------|-------------------------|------------------------------|------------------|------------------|
| Traffic t (Gbps) | Transit Price (\$/Mbps) | Paid Peering Price (\$/Mbps) | Port size (Gbps) | Price (\$/month) |
| $t < 1$ | 6 | 3 | 0.1 | 100 |
| $1 \leq t < 10$ | 4 | 2 | 1 | 800 |
| $10 \leq t < 100$ | 1 | 0.5 | 10 | 1700 |
| $t \geq 100$ | 0.4 | 0.2 | 100 | 7820 |

TABLE I: Transit and Peering Prices

i plays, it actively seeks to form peering relationships with other nodes, while the other nodes only evaluate incoming peering requests.

A. Parameterization

The values for our parameters are given in Tables I and II. Table I shows the median prices for different traffic ranges and port sizes at IXPs, reported by TeleGeography [31] and the websites of the following IXPs: AMS-IX [32], DEC-IX [33], LINX [34]. We consider high-end pricing at IXPs for the same port sizes so that the performance is comparable under transit and IXP peering. We ignore one-time fixed costs, e.g., initial IXP membership costs.

V. TRAFFIC UNCERTAINTY IN PEERING

In this section we evaluate the first source of complexity, i.e., the effect of imperfect traffic prediction prior to establishing a peering link. In section III-A we described that ASes typically employ NetFlow to identify its potential peers. We discuss each of these sources of uncertainty as follows.

A. Limited Information from NetFlow

Consider the network shown in figure 1 where a Tier-2 NSP i attempts to determine if j is a potential peer. In order to identify potential peers, a Tier-2 NSP i employs NetFlow to analyze the origin and destination of its local traffic, i.e., traffic generated within i and consumed at another co-located AS j and vice versa, i.e., \hat{T}_{ij} . i also estimates the ratio of inbound to outbound traffic T_{ji}/T_{ij} using this data. However, this analysis ignores the fact that x and y , which are customers of i and j respectively, may also exchange traffic over the proposed link L_{ij}^p . Thus, the maximum traffic that may be exchanged over L_{ij}^p is:

$$\max(T'_{ij} + T'_{ji}) = \hat{T}_{ij} + \hat{T}_{iy} + \hat{T}_{jx} + \hat{T}_{xy}$$

In this case, while NetFlow informs i that \hat{T}_{xy} flows through its network, i cannot determine if this traffic also flows through j . Even after employing other tools, e.g. *traceroute* and analysis of *BGP* announcements, i cannot be certain about the path taken by \hat{T}_{xy} because of asymmetric routing and multihoming in the interdomain network. This uncertainty may negatively affect the peering decisions of i as follows:

Premature rejection: Let

$$T_{ij} + T_{ji} \ll T'_{ij} + T'_{ji}$$

TABLE II: Input Parameters

| Parameter, Symbol, Description | Value | Explanation |
|--|-----------|---|
| Number of ASes N | 2000 | Time constraints for the simulation |
| Number of geographic locations G_{Max} | 100 | Based on approximate ratio of IXPs to peering networks in the Internet. PeeringDB ratio 18.55. Model ratio 20.0 [35] |
| Geographic expanse distribution | Zipf(1.6) | Based on data about number of participants at each IXP collected from PeeringDB [35]. IXP locations are randomly assigned to each node. |
| Maximum points-of-presence for an AS | 15 | |
| Generated traffic distribution | Zipf(1.2) | It generates a heavy-tailed distribution consistent with the behavior reported in [36], [37] & [28]. With this distribution, 0.1% of the ASes generate nearly 30% of the total traffic. |
| Consumed traffic distribution | Zipf(0.8) | Produces heavy-tailed distribution of incoming traffic, similar to internally measured traffic distribution at a large US public university. Estimated Comcast traffic [38]. Consumed traffic of a node is proportional to its points-of-presence, the rationale being that a node with large expanse will also have a large number of access customers. |
| Maximum consumed traffic | 8 Tbps | |
| Selective peering ratio σ | 3.0 | Peering policies of different NSPs, e.g., [25], [26], [27] |
| Multihoming Degree | 2 | Fixed for all non-Tier1 nodes for simplicity. [5] |




Fig. 1: Limited Information from NetFlow

Then i may assume that it does not exchange a significant volume of traffic with j and prematurely decide not initiate peering negotiations. Whereas, if it had accurate estimates of the traffic that would flow over the peering link, it would have moved down the check list of other peering policy requirements.

Premature acceptance: Let

$$\frac{T_{ji}}{T_{ij}} \leq \sigma, \quad \frac{T'_{ji}}{T'_{ij}} > \sigma$$

In this case, i may peer with j assuming that j satisfies its traffic ratio requirement. However, once the link is formed and traffic starts flowing over L_{ij}^p , i will determine that its peering requirements are not being met by j . Such situations, which arise out of inaccurate estimates of traffic prior to link formation, are one of the causes of peering conflicts in the real world.

B. Dynamic Routing

The interdomain network, constituting its physical structure and traffic flows, self-organizes itself through the collective actions of local and (in many cases) autonomous interactions of the ASes. This results in a complex network where peering links between different ASes may have non-local effects, i.e., they may affect the traffic flows over other ASes and links [21]. Once a peering link is formed by i , its customers (and those of the peer, if it has any) may update their routes given the changes in the network. Assuming default BGP

configuration to choose the shortest routes, these updates may cause customer traffic which was not previously routed through i to flow through i and vice versa. Both scenarios have a direct impact on the peering relationships of i and its utility. We illustrate both cases as follows.

Addition of traffic: Consider the network shown in figure 2. Prior to formation of L_{ij}^p , as shown in figure 2a, traffic \hat{T}_{xy} from its customer x bypasses i as route $x \rightarrow k \rightarrow Tier1 B \rightarrow j \rightarrow y$ is one hop short of the route $x \rightarrow i \rightarrow Tier1 A \rightarrow Tier1 B \rightarrow j \rightarrow y$. Since this traffic bypasses i , it has no way of measuring it. Additionally, x does not know the route taken by this traffic. However, once the peering link L_{ij}^p is formed, i offers a shorter route $x \rightarrow i \rightarrow j \rightarrow y$ than the one previously chosen by x . Hence, i may experience an upsurge in customer traffic, which it could not predict prior to the execution of its peering decision.

Reduction in traffic: Consider the network shown in figure 3. Prior to formation of L_{ij}^p , as shown in figure 3a, traffic \hat{T}_{xy} from its customer x flows through i taking route $x \rightarrow i \rightarrow Tier1 \rightarrow y$. However, once the peering link L_{ij}^p is formed, i no longer routes traffic through the $Tier1$ node because of “prefer-peer-over-provider” routing policy. This, however, increases the path length for \hat{T}_{xy} as it is routed over the path $x \rightarrow i \rightarrow j \rightarrow k \rightarrow y$. Therefore, x routes traffic away from i to its second provider j offering it a shorter path $x \rightarrow j \rightarrow k \rightarrow y$. Hence, i may experience a decline in customer traffic, which it could not predict prior to formation of the peering link. Note that these path-selection decisions may take place autonomously without any human intervention.

Although i can infer the topology of the interdomain network using different inference techniques, these techniques are limited in that they cannot accurately discover peering links and do not inform about the routes adopted by specific traffic flows [10], [11].

C. Computational Results

We illustrated the main sources of traffic uncertainty using simple examples in sections V-A and V-B. In this section we explore, through large-scale computational simulations based on our model of section IV, the extent to which these complexities manifest themselves in large-scale networks.

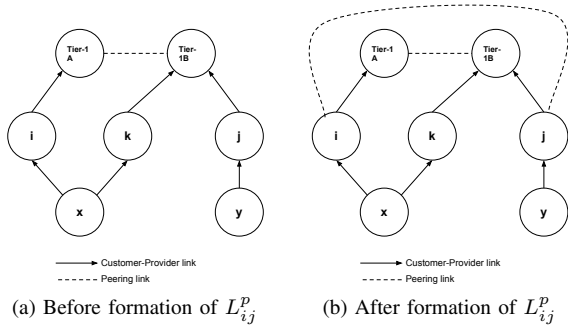


Fig. 2: Addition of traffic after formation of a peering link

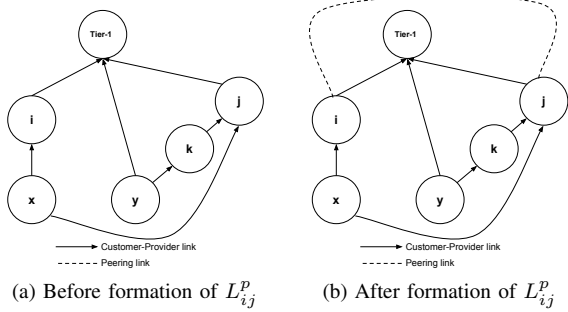


Fig. 3: Reduction in traffic after formation of a peering link

We carry out 1000 simulations of our model, each with a unique population and initial topology. As Tier-2 NSPs play during network formation, we measure the number of traffic flows and total traffic volume transiting each Tier-2 NSP i before and after it forms a peering relationship with another node j . We also measure the traffic T'_{ij} and T'_{ji} after L'_{ij} has been formed. Our objective is to determine the changes in traffic volume of i with each new peering link and the fraction of peering relationships which fall in the category of “premature acceptance” as described in section V-A.

We find that only 10% peering links fall in the category of “premature acceptance” by one of the peers, i.e., a posteriori analysis of the traffic on the peering link reveals that the traffic ratio is out of bounds for one of the peers. All peering links in this category are those which are formed between Tier-2 NSPs. If one of the peers has a major content provider and the other has a major access ISP as its customer, the traffic on link is likely to be asymmetric. However, such asymmetries are not detected during peer evaluation phase as players only take into account local traffic and ignore customer traffic.

Figure 4 shows the relative difference between the number of traffic flows and traffic volume before and after each peering link is formed by each Tier-2 NSP. We find that approximately 85% peering links result in an overall increase in the number of traffic flows and traffic volume transiting the player. However, because of the skewed nature of the traffic distribution, the addition or removal of a few traffic flows carrying traffic

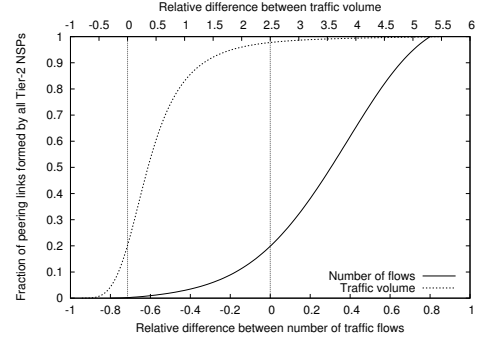


Fig. 4: Relative difference between number of traffic flows and volume before and after formation of a peering link

for major content or access providers can significantly affect the traffic volume transiting a node. Furthermore, changes in traffic volume have a direct bearing on the utility of the players.

Our analysis shows that the peering links contributing to a significant increase ($\geq 50\%$) in the number of traffic flows are the ones which are formed between two Tier-2 NSPs. These players have large number of customers which often find that a peering link between their providers offers them a shorter path to one another. Whereas, the most significant increase in traffic volume ($\geq 50\%$) arises from peering directly with major content and access providers. Interestingly, we find that the Tier-2 NSPs undergoing significant losses in traffic volume with peering are the ones which have major content or access providers as their customers. These large nodes are often multihomed to Tier-1 nodes providing them with short paths to the entire network. Any peering link formed by their Tier-2 providers that increases their path length by even a single hop can lead these large nodes to divert their traffic away from the Tier-2 provider.

These results imply that the scope of analysis for peering decisions should not be limited to local traffic only; instead it should also incorporate customer traffic data as much as possible. Furthermore, the identification and evaluation of peers using NetFlow is inherently inaccurate for Tier-2 NSPs and may result in premature rejection of peers, premature acceptance of peers which do not qualify and lead to conflicts, and even a negative impact on utility.

VI. ECONOMIC UNCERTAINTY IN PEERING

In this section, we evaluate the second obstruction to optimal peering choices, i.e., a complex transit and peering prices structure that has evolved as transit providers and IXPs try to lure customers towards themselves.

A. Analysis for Settlement-free Peering

Over a period of time simple rules-of-thumb have come in usage to decide the mode of traffic exchange, e.g., move as much traffic as possible to a settlement-free peering link to cut down costs.

Figure 7 shows that peering costs are much lower than transit costs for the same traffic volume. Hence, the first instinct of many operators is to offload as much traffic as possible on peering links. However, the analysis has to be carried out in totality because diverting traffic from a transit link to a peering link may also affect the transit price per unit traffic. We illustrate this complexity by the following example. Consider a Tier-2 NSP i with 10 Gbps upstream transit traffic. It would incur a monthly cost of \$10,000 in transit payments. Let us assume that i can divert as much as 50% of its traffic onto peering links at an IXP. With 5 Gbps transit and peering traffic each, the transit cost of i becomes \$20,000 with an additional IXP cost of \$1700, thus resulting in a total cost of \$21,700 - an increase of \$11,700 over the original cost. The transit cost increased dramatically because of the complex economies-of-scale engineered in the pricing structure.

Effect on utility components

The general notion is that settlement-free peering increases utility by lowering transit costs. In section V we showed that peering can change the traffic volume transiting through a network. We show that possible traffic variation and the complex pricing structure can affect all components of utility under settlement-free peering: transit and peering costs and transit revenues.

Computational Results: We simulate 1000 instances, each with a unique population and initial topology, of our model. In each simulation, we record the utility and its components of each Tier-2 NSP before and after it has committed its peering decisions.

Our analysis shows that 10% of players actually have their utility decreased after peering (90% C.I.). Furthermore, 1.5% of players have their utility decreased by more than 50%.

Figure 5 shows the change in transit revenue, transit costs and cumulative costs (sum of transit and peering costs) for players which undergo a decrease in utility. Although transit costs decrease for 80% of such players, yet the cumulative costs increase for 75% of them. Similarly, 34% of such players also face a loss in revenue as their customers divert traffic away from them after peering. Figure 6 shows similar analysis for players which increase their utility. Although cumulative costs increase for 50% of such players, yet they are able to register a net benefit through an increase in transit revenue as well. Thus, NSPs require more careful analysis of their utility before committing to large-scale peering.

B. Analysis for Paid-peering

Although paid-peering has been getting a lot of attention recently, our simulations show that only a small fraction of NSPs are able to form paid peering links. We attribute limited adoption of paid peering to the following two reasons:

- 1) On average, 70% potential peers of a Tier-2 NSP i satisfy its traffic ratio requirements. Hence, i cannot ask them to be its paid peering customers.
- 2) In approximately 50% of the paid-peering evaluations, on average, where a potential peer j does not satisfy the

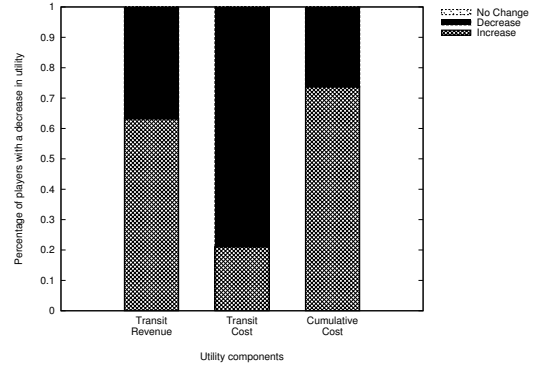


Fig. 5: Effect on utility components of Tier-2 NSPs with a decrease in utility after peering




Fig. 6: Effect on utility components of Tier-2 NSPs with an increase in utility after peering

traffic ratio of i , cost-benefit-analysis by j reveals that a mix of paid-peering and transit is more costly for j than its transit alone.

Similar to settlement-free peering, the cumulative costs of paid-peering and transit mix can exceed the costs under transit alone. Figure 8 shows the total cost of traffic exchange versus the fraction of total traffic that is diverted on a paid peering link. We find that although paid-peering is priced at half the transit price, yet adoption of paid peering favors only a small class of Tier-2 NSPs. We find that Tier-2 NSPs with the following characteristics generally benefit from being paid-peering providers:

- 1) Tier-2 NSPs with very large local traffic volume ($\geq 500Gbps$) which ensures that they can continue to use the same transit prices even after diverting a fraction of their traffic from transit to paid-peering links.
- 2) Tier-2 NSPs whose local inbound traffic is much greater than local outbound traffic which makes them attractive paid-peering providers.
- 3) Tier-2 NSPs which do not have large content providers as customers. Although having large content providers as customers is beneficial for transit revenue, yet large outbound traffic makes such transit providers unattractive for paid peering.

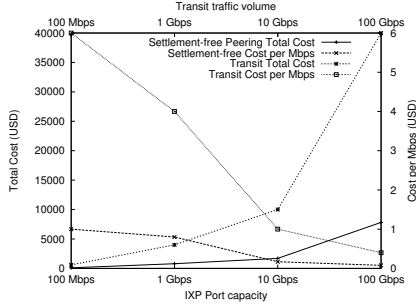


Fig. 7: Peering Cost vs. Transit Cost (monthly recurring)

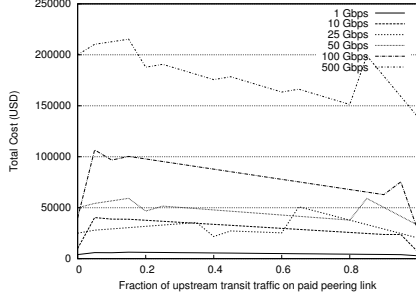


Fig. 8: Transit + Paid-Peering Cost vs. Fraction of transit traffic on a paid peering link (monthly recurring)

Furthermore, we find that having a large number of smaller customers does not turn a Tier-2 NSP into an attractive paid-peering provider because it generally results in balanced traffic ratios which favor settlement-free relationships.

These results imply that the scope of economic analysis for peering decisions should not be limited to reduction of transit costs only; instead it should also incorporate the effect of peering on customer revenue and transit costs.

VII. COMPLEXITY OF DETERMINING THE OPTIMAL SET OF PEERS

In this section, we evaluate if the myopic one-by-one analysis policy for peers described in section III is sufficient to get an optimal utility. Furthermore, we evaluate if it is feasible for an NSP to do a combinatorial optimization of its peering set instead of doing the one-by-one. For simplicity, we assume that NetFlow is able to provide accurate estimates of traffic that would be exchanged over any proposed link⁶.

A. Interdependence of Peering Links

We illustrate this complexity through a simple network in figure 9 where ASes i , j and k are co-located with one another. Let i and j be two Tier-2 NSPs where i actively seeks peers and j and k only respond to peering requests. Since k is a stub, it uses Open peering. Furthermore let $T_{ik} \gg T_{ij}$. i has a choice of four distinct *peering configurations* shown in the figure. Each configuration may incur different transit, paid-peering

⁶Suboptimal peering decisions would be even worse in the presence of limited information.

and IXP costs and yield different paid-peering revenues. i evaluates its utility under each configuration, beginning with configuration A shown in figure 9a.

In configuration A , i has no peers and incurs a steep transit cost for exchanging traffic through its upstream transit providers. Let π_i^A be the utility of i under configuration A .

In configuration B , shown in figure 9b, i evaluates peering with j . T_{ik} and T_{ki} will be routed through j under this configuration. Let

$$\frac{T_{ij} + T_{ik}}{T_{ji} + T_{ki}} > \sigma \quad (6)$$

Thus, j refuses settlement-free peering to i and instead offers i to become its paid-peering customer. i carries out cost-benefit-analysis for L_{ij}^p . Let i determine that $\pi_i^A > \pi_i^B$. Hence, i refuses to become paid-peering customer of j .

In configuration C , shown in figure 9c, i determines that

$$\frac{T_{ki}}{T_{ik}} \leq \sigma \quad (7)$$

Since k uses Open peering policy, it accepts peering with i . Let $\pi_i^C < \pi_i^A$ since i saves on transit costs under configuration C and peering costs are generally lower than transit costs. Hence, i peers with j .

Let i re-evaluate j in configuration D , shown in figure 9d. Now that traffic $T_{ik} + T_{ki}$ will not be routed over L_{ij}^p , the ratio computation in expression 6 no longer holds. Upon re-evaluation, j finds that:

$$\frac{T_{ij}}{T_{ji}} \leq \sigma \quad (8)$$

However, let i find that:

$$\frac{T_{ji}}{T_{ij}} > \sigma \quad (9)$$

Hence, i can acquire j as a paid-peering customer, increasing its utility. Thus, the evaluation of the four configurations reveals that:

$$\pi_i^B < \pi_i^A < \pi_i^C < \pi_i^D \quad (10)$$

Thus, i had to evaluate all possible combinations of peers to determine the optimal set of peers.

B. Infeasibility of Exhaustive Search

An NSP can use brute force approach to exhaustively search the space for all possible peer combinations, compute its utility for each combination and determine the one that gives it the optimal utility. Let K denote the number of peers and M the number of potential peers of i . Then the total number of combinations, Q , that i may need to evaluate is given by:

$$Q = \sum_{K=0}^M \binom{M}{K} \quad (11)$$

yielding a complexity of $O(2^M)$.

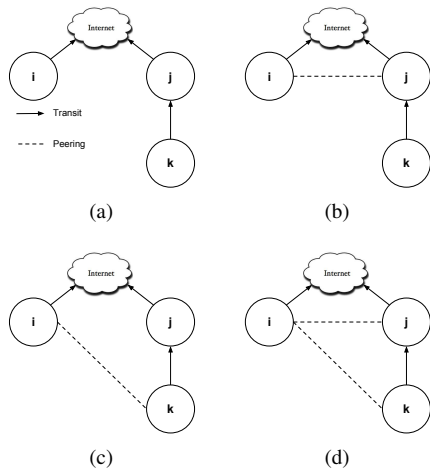


Fig. 9: Peering interdependencies

Our analysis shows that peering decisions may be intertwined with one another. Furthermore, it is infeasible even for a modest sized NSP to compute the optimal set of peers through exhaustive search.

VIII. CONCLUSIONS & FUTURE WORK

We have investigated three sources of complexity that preclude the adoption of a formal approach to optimal peer selection in the Internet interdomain network. We showed that limited traffic information provided by NetFlow, the complex nature of network topology combined with policy-based routing decisions limit the ability of a Tier-2 NSP from accurately forecasting the effect of its peering decisions.

This leaves us with a multitude of questions for future work. For example, what is the highest utility a typical Tier-2 NSP can expect to accumulate in the presence of these complexities? How close is the utility acquired from conventional rules-of-thumb peering to this theoretical upper bound? Are there practical ways to overcome these complexities? In other words, is there a case for smart peering?

REFERENCES

- [1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large European IXP," in *Proceedings of the ACM SIGCOMM 2012*.
- [2] W. Norton, *The Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2014.
- [3] "Survey of Characteristics of Internet Carrier Interconnection Agreements," <http://www.pch.net/resources/papers/peering-survey/PCH-Peering-Survey-2011.pdf>.
- [4] W. Norton, "Comcast Vs Level 3," http://drpeering.net/AskDrPeering/blog/articles/Upcoming/Entries/2011/2/10_Comcast_vs_Level_3.html, 2011.
- [5] A. Dhamdhere and C. Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem," *IEEE/ACM Transactions on Networking*, 2011.
- [6] P. Faratin, D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr, "Complexity of Internet interconnections: Technology, incentives and implications for policy," in *TPRC '07*, 2007.
- [7] P. Faratin, D. Clark, S. Bauer, W. Lehr, P. Gilmore, and A. Berger, "The Growing Complexity of Internet Interconnection," *Communications & Strategies*, vol. 1, no. 72, pp. 51–72, 4th quart 2008.
- [8] R. Steenbergen, "A Guide to Peering on the Internet," *NANOG 51*, 2011.
- [9] "Using PeeringDB to Understand the Peering Ecosystem, author=Lodhi, Aemen and Larson, Natalie and Dhamdhere, Amogh and Dovrolis, Constantine and others," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.
- [10] L. Gao, "On Inferring Autonomous System Relationships in the Internet," *IEEE/ACM Transactions on Networking*, vol. 9, pp. 733–745, 2000.
- [11] M. Luckie, B. Huffaker, k. claffy, A. Dhamdhere, and V. Giotsas, "AS Relationships, Customer Cones, and Validation," in *Proceeding of the ACM SIGCOMM IMC*, 2013.
- [12] G. Gürsun and M. Crovella, "On traffic matrix completion in the internet," in *Proceedings of the ACM SIGCOMM IMC*, 2012.
- [13] G. Shrimali and S. Kumar, "Paid Peering Among Internet Service Providers," in *Proceedings of the ACM GameNets*, 2006.
- [14] A. Dhamdhere, C. Dovrolis, and P. Francois, "A Value-based Framework for Internet Peering Agreements," in *International Teletraffic Congress*, Oct 2010.
- [15] R. T. Ma, D.-M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "Interconnecting Eyeballs to Content: A Shapley Value Perspective on ISP Peering and Settlement," in *Proceedings of the ACM NetEcon*, 2008.
- [16] N. Badasyan and S. Chakrabarti, "A Simple Game Theoretic Analysis of Peering and Transit Contracting Among Internet Access Providers," in *Telecommunications Policy Research Conference*, 2005.
- [17] S. Lippert and G. Spagnolo, "Internet Peering as a Network of Relations," *Telecommunications Policy*, vol. 32, pp. 33–49, February 2008.
- [18] P. Baake and T. Wichmann, "On the Economics of Internet Peering," *Netnomics*, vol. 1, pp. 89–105, 1999.
- [19] E. Anshelevich, B. Shepherd, and G. Wilfong, "Strategic Network Formation through Peering and Service Agreements," in *Proceedings of FOCS*, 2006.
- [20] E. Meiron, S. Mannor, A. Orda *et al.*, "Formation Games and the Internet's Structure," *arXiv preprint arXiv:1307.4102*, 2013.
- [21] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "Open Peering by Internet Transit Providers: Peer Preference or Peer Pressure," in *Proceedings of IEEE INFOCOM*, 2014.
- [22] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou, "Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet," in *Proceedings of ICALP*, 2002.
- [23] H. Chang, S. Jamin, and W. Willinger, "Internet Connectivity at the AS-level: An Optimization-Driven Modeling Approach," in *Proceedings of ACM SIGCOMM MoMeTools Workshop*, 2003.
- [24] —, "To Peer or Not to Peer: Modeling the Evolution of the Internets AS-level Topology," in *Proceedings of IEEE INFOCOM*, 2006.
- [25] "AT&T Peering Policy," <http://www.corp.att.com/peering/>.
- [26] "Comcast Peering Policy," <http://www.comcast.com/peering>.
- [27] "A Study of 28 Peering Policies," drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html.
- [28] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-domain Traffic," in *Proceedings of ACM SIGCOMM*, 2010.
- [29] "Oh, Elusive Paid Peering," http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2009/4/18_Oh_Elusive_Paid_Peering.html.
- [30] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An Agent-Based Model of Interdomain Network Formation, Traffic Flow and Economics," in *Proceedings of IEEE INFOCOM*, 2012.
- [31] TeleGeography, "IP Transit Pricing Service," <http://www.telegeography.com>, 2014.
- [32] AMS-IX, "Amsterdam Internet Exchange Pricing," <https://ams-ix.net/services-pricing/pricing>, 2014.
- [33] DEC-IX, "Frankfurt Internet Exchange Pricing," <http://www.de-cix.net/products-services/pricing/>, 2014.
- [34] LINX, "London Internet Exchange Pricing," <https://www.linx.net/service/servicefees.html>, 2014.
- [35] "PeeringDB," <http://www.peeringdb.com>, June 2014.
- [36] H. Chang, S. Jamin, Z. Mao, and W. Willinger, "An Empirical Approach to Modeling Inter-AS Traffic Matrices," in *Proceedings of ACM SIGCOMM IMC*, 2005.
- [37] A. Feldmann, N. Kammenhuber, O. Maennel, B. Maggs, R. De Prisco, and R. Sundaram, "A Methodology for Estimating Interdomain Web Traffic Demand," in *Proceedings of ACM SIGCOMM IMC*, 2004.
- [38] Rayburn, Dan, "Here is How The Comcast & Netflix Deal Is Structured, With Data & Numbers," <http://blog.streamingmedia.com/2014/02/heres-comcast-netflix-deal-structured-numbers.html>, February 2014.