Requirements of Settlement-Free Peering Policies

Ali Nikkhah
Department of Computer Science
University of California
Irvine, USA
ali.nikkhah@uci.edu

Scott Jordan
Department of Computer Science
University of California
Irvine, USA
sjordan@uci.edu

***Abstract*—Peering between two networks may be either settlement-free or paid. In order to qualify for settlement-free peering, large Internet Service Providers (ISPs) require that peers meet certain requirements. It is widely perceived that these requirements represent the conditions under which the two peering networks perceive a roughly equal exchange of value. However, the academic literature has not yet shown the relationship between these settlement-free peering requirements and the value to each interconnecting network.**

We analyze the value to each network from the most common and important requirements. Large ISPs often require potential settlement-free peers to interconnect at a minimum of 6-8 locations from a predetermined list. We find that there is a substantial benefit from this requirement to the ISP, but little incremental benefit from a larger number of interconnection points. Large ISPs often require that the ratio of incoming traffic to outgoing traffic remain below approximately 2:1. We find that this requirement ensures a roughly equal exchange of value.

***Index Terms*—Net Neutrality, Interconnection, Peering Policies**

I. Introduction

An Internet Service Provider (ISP) provides the capability to transmit data to and receive data from all or substantially all Internet endpoints. In order to provide this Internet access service, an ISP must make arrangements with other networks to interconnect and exchange traffic. An interconnection arrangement is for *peering* if and only if each network agrees to accept and deliver traffic with destinations in its customer cone. Historically, peering was principally used by tier-1 networks. Peering may be either paid (i.e., one interconnecting network pays the other) or settlement-free (i.e., without payment by either interconnecting network to the other). Large ISPs often require that peers meet certain requirements, including a specified minimum number of interconnection points and a traffic ratio less than 2:1. The conventional wisdom is that these requirements are related to the perception of roughly equal value from the peering arrangement, but the academic literature has not yet established such a relationship.

We studied the settlement-free peering policies of the ten largest ISPs in the United States. Table I summarizes the most relevant requirements of these policies. An estimate of the number of subscribers of each ISP in 2021 [1] is given, as settlement-free peering policies differ with the number of subscribers. An ISP's predisposition towards or against peering, as noted by PeeringDB [2], is given. We also show the minimum number of Internet exchange points (IXPs) required. The four largest ISPs each require interconnection at a minimum of 4 to 8 IXPs from specified lists; we show the length of the list in the table. They also require that incoming and outgoing traffic be roughly balanced. The next six largest ISPs require interconnection at a specified minimum number of interconnection points (but often less than 4), and may or may not require roughly balanced traffic. We henceforth focus on the settlement-free peering requirements of the four largest ISPs.

TABLE I: Settlement-free peering requirements

ISP	Subscribers	Inclination	# IXPs for Peering	# IXPs in ISP List	Traffic Ratio
Comcast	31,901,000	Selective	4	12	Balanced
Charter	30,089,000	Selective	6-8	15	-
AT&T	15,504,000	Selective	6	12	2:1
Verizon	7,365,000	Restrictive	8	-	1.8:1
Cox	5,530,000	Selective	2	15	Balanced
CenturyLink	4,519,000	Selective	6	10	1.5:1
Altice	4,386,200	Selective	2	-	1.8:1
Frontier	2,799,000	Selective	3	6	-
Mediacom	1,463,000	Open	-	5	-
TDS Telecom	526,000	Open	-	9	-

The academic literature provides little insight into why large ISPs impose these settlement-free peering requirements or how these requirements are related to either the ISP's network cost or its perception of value. Although there are many papers that consider various aspects of peering, there are few that analyze settlement-free peering requirements, and fewer yet that attempt to relate these requirements to the value to each interconnecting network of the peering agreement. Lodhi et al. [3] studied PeeringDB data. They found that the volume of traffic that an ISP carries on its network is positively correlated with the number of IXPs at which it interconnects, and that ISPs with large traffic volumes and a large number of subscribers are more likely to have a selective or restrictive peering inclination. However, peering inclination is a coarse measure, and they do not analyze the particular requirements in settlement-free peering policies (e.g., the minimum number of IXPs or maximum traffic ratio). We have not found any academic papers that do. The closest may be Johari and

This material is based upon work supported by the National Science Foundation under Grant No. 1812426

Tsitsiklis [4], which discusses the selection of IXPs in a few networks with idealized and regular topologies. There are some papers that discuss the presence of traffic ratio requirements [5], [6], and that suggest these requirements are related to ISP network costs. However, they do not analyze the effect of the traffic ratio upon costs, and thus are not concerned with relating traffic ratio requirements to the value to the ISP of the peering agreement. Indeed, there are some papers that are skeptical that traffic ratios relate to the benefits to each interconnecting party [7]–[9].

The focus of this paper is to relate the settlement-free peering requirements of large ISPs to the value the peering arrangement brings to the ISP. In section II, we develop a model that will enable us to examine the effect of the number and location of IXPs and the traffic ratio. In Section III, we determine the average distances on each portion of an ISP's network. We then model the average traffic-sensitive cost associated with carrying the traffic over these average distances. In section IV, we analyze settlement-free peering requirements about the number and location of IXPs. Large ISPs require interconnection at a minimum of 4 to 8 interconnection points from specified lists. We find that the ISP's traffic-sensitive cost is a uni-modal function of the number of interconnection points, and estimate that it is minimized with 8 IXPs. We also observe that there may be little value in requiring interconnection at more than 6 IXPs. In section V, we analyze settlement-free peering requirements on traffic ratios. Large ISPs require that the traffic ratio not exceed a specified threshold. We show that for traffic ratios above 2:1, the ISP's traffic-sensitive cost increases as the number of IXPs increases, and it is rational for the ISP to not agree to settlement-free peering. When traffic ratios are at or below 2:1, we also estimate that requiring interconnection at more than 8 IXPs is of little incremental value.

II. Model

A. Topology

Our model of the topology of an ISP's network consists of the ISP's service territory, the location of IXPs, and the segments of its network.

While most ISPs do not offer residential broadband Internet access service over the entire contiguous United States, we see little in their settlement-free peering policies that are specific to their service territory, other than that a subset of the IXPs at which they peer are concentrated near their service territory. Thus, we focus on an ISP whose service territory covers the contiguous United States, which we represent using the set of actual longitudes and latitudes of the contiguous United States.

We consider M IXPs at which an ISP and an interconnecting network may agree to peer. In numerical results below, we use the $M = 12$ largest IXPs in the United States, located at Ashburn, Chicago, Dallas, San Jose, Los Angeles, New York, Seattle, Miami, Atlanta, Denver, Boston, and Minneapolis [2]. Denote the coordinates (in longitude and latitude) of these M IXPs by $IXP(i)$, and the set of the locations of these IXPs by I^M. The four largest ISPs each interconnect at a minimum

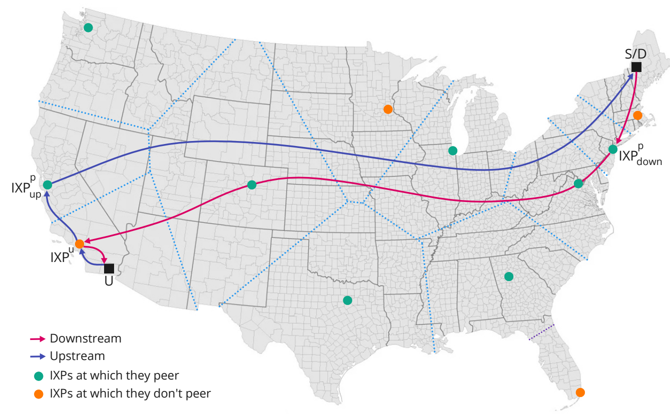

Fig. 1: Topology of an ISP's network

of 9 of these 12 IXPs, but for simplicity, we assume they each interconnect at all 12. An ISP and an interconnecting network often agree to interconnect at a smaller number $N < M$ of IXPs. Denote the set of N IXPs at which they agree to interconnect by $l^N \subseteq \{1, ..., M\}$, and the set of locations of these IXPs by $I^N = \{IXP(i), i \in l^N\} \subseteq I^M$.

We model an ISP's network as partitioned into a single backbone network, multiple middle mile networks, and multiple access networks. We model each access network as spanning a single U.S. county. While we recognize that the geographical sizes and topologies of access networks differ widely, this assumption will not significantly affect the results in this paper, since settlement-free peering policies depend more critically on the number of interconnection points than on the topologies of access networks. Denote the geographical region of access network j by $Access(j)$, and the location of the geographical center of access network j by $A(j)$. In the numerical results below, we assign these locations using the actual longitudes and latitudes of the center of each county in the contiguous United States.

Consider an ISP and an interconnecting network that agree to interconnect at the N IXPs in l^N. Denote by $R^N(IXP(i))$ the geographical region that consists of the union of access networks for which the closest IXP in l^N is IXP i, namely

$$R^N(IXP(i)) = \bigcup_{\substack{j \,|\, \|A(j)-IXP(i)\| \leq \\ \|A(j)-IXP(i')\| \,\forall i' \in l^N}} Access(j) \qquad (1)$$

Figure 1 approximately illustrates these regions when an ISP and an interconnecting network agree to interconnect at all 12 IXPs. (We will discuss the case when $N = 8$ below.)

B. Traffic Matrices

The locations of end users of the ISP are represented by a probability distribution over the ISP's service territory. We decompose this distribution into (a) a distribution of the number of end users in each access network and (b) for each access network, the distribution of end users within the access network.

Denote the probability that an end user resides within access network j by $P(j)$. We assume that end users are distributed across access networks according to the population of the county associated with the access network, which we denote by p_j and assign using U.S. census data [10], and we denote the population of the contiguous United States by $p = \sum_j p(j)$. It follows that $P(j) = p_j/p$. We further assume that end users are uniformly distributed within each access network, and we denote the size of county j by s_j, which we assign using the U.S. Gazetteer [11].

We focus first on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. Denote the source's location by S and the end user's location by U. We assume that S and U are independent and that the marginal distributions of S and U are given by the joint distribution of the population with each access network and the distribution of end users within each access network.

We assume that both networks use hot potato routing. Along the route S to U, denote the location of the IXP at which downstream traffic enters the ISP's network by IXP^p_{down} and the location of the IXP closest to the end user by IXP^u. For example, suppose S is in Maine and U is in Imperial county, California. Then, as illustrated in Figure 1, IXP^p_{down} might be in New York (if the two networks do not agree to peer in Boston) and IXP^u is in Los Angeles.

The ISP carries traffic on only a subset of the route from S to U. It carries traffic on its backbone from IXP^p_{down} to IXP^u, and it carries traffic on a middle mile network and an access network from IXP^u to U. The subset of the route that is on the ISP's network thus depends on the joint distribution of (IXP^p_{down}, IXP^u, U).

The access network on which U resides is distributed according to $\{P(j)\}$, and U is uniformly distributed within the access network. The IXP closest to the end user is a deterministic function of U, namely $IXP^u = (g' | U \in R^M(g'))$. The IXP at which downstream traffic enters the ISP's network (IXP^p_{down}) is independent of the end user, and it is the IXP closest to the source among the IXPs at which they agree to peer, i.e. $IXP^p_{down} = (g | S \in R^N(g))$. Since end users are assumed to be distributed according to U.S. county population statistics:

$$P(IXP^p_{down} = g) = \frac{1}{p} \sum_{Access(j) \subset R^N(g)} p(j) \tag{2}$$

III. Traffic-sensitive costs

In this section, we first determine the distances that an ISP carries traffic on each portion of its network. We next calculate the average distance using traffic matrices. Finally, we model the average traffic-sensitive cost associated with carrying the traffic over these average distances.

For downstream traffic, the distance from S to U on the ISP's backbone network is a function of the location of the IXP at which downstream traffic enters the ISP's network (IXP^p_{down}) and the location of the IXP closest to the end user (IXP^u). Denote the distance on the ISP's backbone network between these two IXPs by $D^b(IXP^p_{down}, IXP^u) = \|IXP^p_{down} - IXP^u\|$.

The distance from S to U on the ISP's middle mile network is a function of the location of the IXP closest to the end user (IXP^u) and the location of the access network on which U resides. Denote the distance on the ISP's middle mile network between these two locations by $D^m(IXP^u, U) = \|IXP^u - A(j)\|$, where $U \in R^M(IXP^u)$ and $j \,|\, (U \in Access(j))$.

The distance from S to U on the ISP's access network is a function of the location of the end user. Denote the distance on the ISP's access network by $D^a(U) = \|A(j) - U\|, j \,|\, (U \in Access(j))$. The distance can be determined by the location of U within the access network.

We now calculate the average distances over the ISP's backbone, middle mile, and access networks for both downstream and upstream traffic. The IXP at which downstream traffic enters the ISP's network (IXP^p_{down}) is independent of the end user and thus independent of the IXP closest to the end user (IXP^u). The IXP at which downstream traffic enters the ISP's network depends on the IXPs at which they agree to interconnect. Consider an ISP and an interconnecting network that agree to interconnect at the N IXPs in I^N. The average distance of downstream traffic on the ISP's backbone network is:

$$\begin{aligned} ED^b_{down} = & \sum_{g \in I^N} \sum_{g' \in I^M} \\ & D^b(g, g') P(IXP^p_{down} = g) P(IXP^u = g') \end{aligned} \tag{3}$$

The probability distribution of IXP^p_{down} was given in (2). The probability distribution of IXP^u can be similarly represented as:

$$P(IXP^u = g') = \frac{1}{p} \sum_{Access(j) \subset R^M(g')} p(j) \tag{4}$$

There is also upstream traffic. For upstream traffic using hot potato routing, the IXP at which the traffic exits the ISP's network is the IXP closest to the end user at which they agree to peer, i.e. IXP^p_{up}. For example, suppose the end user is in Imperial county, California, and the destination D is in Maine. Then, as illustrated in Figure 1, IXP^u is in Los Angeles and IXP^p_{up} might be in San Jose (if the two networks do not agree to peer in Los Angeles). The ISP might still carry traffic across a portion of its backbone, namely from IXP^u to IXP^p_{up}, and the average such distance is:

$$ED^b_{up} = \sum_{g' \in I^M} D^b(g \mid g' \in R^N(g), g') P(IXP^u = g') \tag{5}$$

The distance on the ISP's middle mile network is a function of (IXP^u, U). It is the same for downstream and upstream. The average distance is:

$$ED^m = \sum_{g' \in I^M} \sum_{A(j) \subset R^M(g')} D^m(g', A(j)) P(j) \tag{6}$$

The distance on the ISP's access network is a function of U. It is the same for downstream and upstream. Since end users

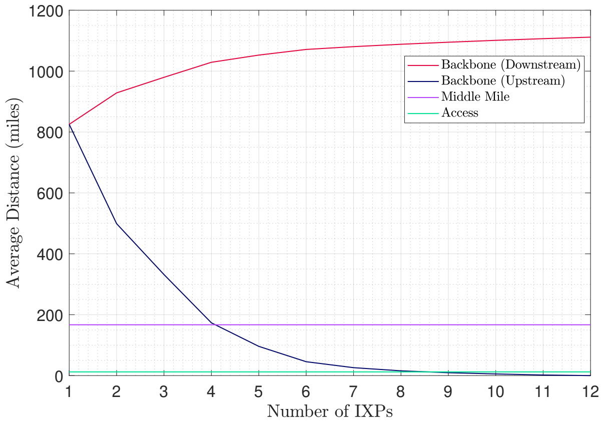

Fig. 2: Average distances over an ISP's network

are uniformly distributed within each access network, but not between access networks, the average distance is:

$$ED^a = \sum_j \frac{p_j}{ps_j} \int_{U \in Access(j)} D^a(U) \tag{7}$$

Figure 2 shows the average distances over the ISP's backbone, middle mile, and access networks for both downstream and upstream traffic, as a function of the number of IXPs at which the interconnecting networks agree to peer. The average distance that downstream traffic is carried over the ISP's backbone is increasing and concave with the number of IXPs at which they agree to peer, since increasing this number of IXPs results in the IXP at which downstream traffic enters the ISP's network becoming further from U. However, the average distance that upstream traffic is carried over the ISP's backbone is decreasing and convex, since increasing this number of IXPs results in the IXP at which traffic exits the ISP's network becoming closer to the ISP's customer. The average distances over the ISP's middle mile and access networks are functions of (U, IXP^u). The number of IXPs at which the two parties agree to peer does not affect the location of the IXP closest to the end user (IXP^u), and thus these average distances are independent of the number of IXPs at which they agree to peer.

We now turn to modeling the average traffic-sensitive cost associated with carrying the traffic over these average distances. We only consider here traffic-sensitive costs, because non-traffic-sensitive costs do not vary significantly with the number of IXPs or the traffic ratio.

Traffic-sensitive costs are a function of both distance and traffic volume. We model traffic-sensitive costs as linearly proportional to the average distance over which the traffic is carried on each portion of the ISP's network [12], and linearly proportional to the average volume of traffic that an ISP carries on each portion of its network. However, we model the cost per unit distance and per unit volume differently on the backbone network, the middle mile networks, and the access networks. Denote the cost per unit distance and per unit volume in the backbone network by c^b, in the middle mile networks by c^m, and in the access network by c^a. Denote the volume of traffic by V. The ISP's traffic-sensitive cost is thus $V\left(c^b ED^b + c^m ED^m + c^a ED^a\right)$.

Given a source-destination traffic matrix, the average distance across the ISP's access networks in (7) is independent of the number of IXPs at which the interconnecting networks peer. In addition, the average distance across the ISP's middle mile networks in (6) is constant, once we fix $M = 12$. The variable portion of the ISP's traffic-sensitive cost is thus only $C = c^b ED^b V$.

IV. Number and list of IXPs

With our model in place, we now turn to analyzing the effect on an ISP's traffic-sensitive costs of the number of IXPs at which peering occurs and of the traffic ratio. We are in particular interested in explaining the settlement-free peering policies of large ISPs. As shown in Table I, large ISPs require other parties who wish to have settlement-free peering to interconnect at a minimum specified number of IXPs. For Comcast, Charter, AT&T, and Verizon, this minimum is between 4 and 8. In addition, large ISPs often specify a list of eligible IXPs that this minimum must be chosen from. The academic literature provides little insight into why large ISPs require interconnection at a minimum specified number of IXPs, nor why they require that they be selected from a list of eligible IXPs.

The variable traffic-sensitive cost of downstream traffic is:

$$C_{down} = c^b V_{down} ED^b_{down} \tag{8}$$

where ED^b_{down} is given in (3).

The variable traffic-sensitive cost of upstream traffic is:

$$C_{up} = c^b V_{up} ED^b_{up} \tag{9}$$

where ED^b_{up} is given in (5).

Figure 3 shows the effect of the number of interconnection points at which the two parties agree to peer on the variable traffic-sensitive costs of both downstream and upstream traffic. (The costs in the figure are normalized by the cost per unit distance and per unit volume, and by the combined downstream and upstream traffic volume.) The costs are a function of $\|IXP^p_{down} - IXP^u\|$ and $\|IXP^p_{up} - IXP^u\|$.

IXP^u is independent of the number of IXPs at which they peer, since the IXP closest to the end user is a deterministic function of U. However, IXP^p_{down} does depend on the number of IXPs at which they peer. The IXP at which traffic enters the ISP's network is the IXP closest to the source among the IXPs at which they agree to peer, i.e. $IXP^p_{down} = (g| S \in R^N(g))$, and thus as the number of IXPs at which they peer increases, IXP^p_{down} moves farther from IXP^u and $\|IXP^p_{down} - IXP^u\|$ increases. Thus, the variable traffic-sensitive downstream cost increases. The variable traffic-sensitive downstream cost is concave, because the incremental distance from IXP^p_{down} to IXP^u associated with adding another IXP decreases, namely there are decreasing returns.

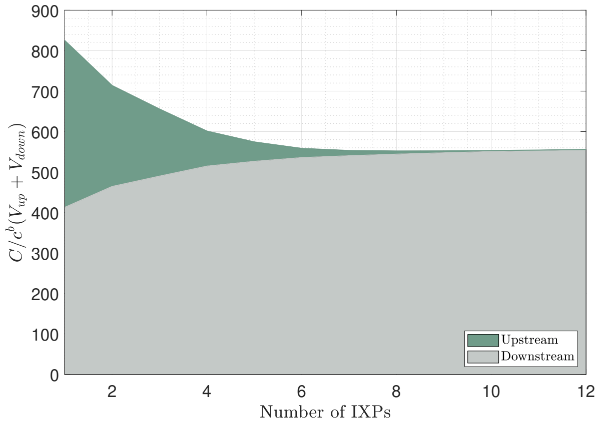

Fig. 3: Traffic-sensitive costs

However, at the same time, the ISP exchanges upstream traffic at the IXP closest to the end user at which they agree to peer, i.e. $IXP^p_{up} = (g|U \in R^N(g))$, and thus as the number of IXPs at which they peer increases, IXP^p_{up} moves closer to IXP^u and $\|IXP^p_{up} - IXP^u\|$ decreases. Therefore, the variable traffic-sensitive upstream cost decreases. The variable traffic-sensitive upstream cost is convex, because the incremental distance from IXP^u to IXP^p_{up} associated with adding another IXP decreases, namely there are decreasing returns.

When the traffic ratio is 1, the decrease in the variable traffic-sensitive upstream cost exceeds the increase in the variable traffic-sensitive downstream cost. In the upstream route, IXP^p_{up} is the closest IXP to IXP^u, whereas, in the downstream route, IXP^p_{down} could be any IXP (including the closest or farthest IXP from IXP^u). Thus, as the number of IXPs increases, the absolute value of the slope of $\|IXP^p_{up} - IXP^u\|$ in the upstream route is higher than the slope of $\|IXP^p_{down} - IXP^u\|$ in the downstream route. It follows that the total cost decreases.

We observe that, when the traffic ratio is 1, the variable traffic-sensitive cost is uni-modal with a minimum at $N = 8$. We also observe that there is less than a 2% difference in the cost between $N = 6$ and $N = 8$, so this indicates there may be little value in requiring interconnection at more than 6 IXPs. Indeed, Charter, AT&T, and Verizon each require 6-8 IXPs for settlement-free peering.

We also wish to examine why ISPs require that the IXPs at which the two parties peer be selected from a specified list. To answer this question, our model selects the N IXPs at which to peer, from the list of $M = 12$ IXPs, so as to minimize its variable traffic-sensitive cost:

$$I^N = \arg\min_{I^N} c^b \left(V_{down} ED^b_{down} + V_{up} ED^b_{up}\right) \tag{10}$$

The cost is typically minimized by selecting IXPs that span the country, so that the average distances the ISP carries traffic across its backbone are relatively small. Furthermore, when selecting a moderate or large number of IXPs, the cost is typically minimized by selecting more IXPs near where there are higher populations. Indeed, Comcast not only requires that potential settlement-free peers agree to peer at a minimum of 4 IXPs from Comcast's list of IXPs, it also requires that at least 1 of these 4 be on the west coast, that at least 1 be on the east coast, and that at least 1 be in a central region. For $N = 4$, our model chooses 1 on the west coast (Los Angeles), 2 on the east coast (Ashburn and Atlanta), and 1 in the middle (Chicago). All 4 of these cities are on Comcast's list. Similarly, Charter not only requires that potential settlement-free peers agree to peer at a minimum of 6-8 IXPs from Charter's list of IXPs, it also requires that at least 2 of these be in an eastern region, at least 2 be in a western region, and at least 2 be in a central region. For $N = 8$, our model chooses 4 on the east coast (Ashburn, New York, Miami, and Atlanta), 2 on the west coast (Los Angeles and Seattle), and 2 in the middle (Chicago and Dallas). All 8 of these cities are on Charter's list.

Our model thus not only explains why large ISPs require settlement-free peers to meet at a minimum of 4-8 IXPs, it also explains why these IXPs are geographically distributed across the country. Furthermore, it also predicts that more will typically be on the east coast, due to its greater population, than on the west coast or in the middle.

V. Traffic ratio

In this section, we examine the effect of the traffic ratio on variable traffic-sensitive costs. Two networks will agree to settlement-free peering if and only if the arrangement is superior for both parties compared to alternative arrangements including paid peering and transit. The conventional wisdom is that settlement-free peering thus occurs if and only if the two parties perceive that they are gaining an approximately equal value from the arrangement. Furthermore, the conventional wisdom, when the two parties are both tier-1 networks, is that the perceived value is related to the size of each network, the number of customers of each party, and the ratio of traffic exchanged in each direction.

Indeed, the settlement-free peering policies of large ISPs often place limits on the ratio of downstream traffic to upstream traffic. AT&T's settlement-free peering policy requires that this traffic ratio not exceed 2:1, and Verizon's settlement-free peering policy requires that this traffic ratio not exceed 1.8:1.

We use our models to investigate the effect of the traffic ratio on the value to each interconnecting party, when the two parties are both ISPs. We use the variable traffic-sensitive cost as a proxy for value. Denote the ratio of downstream traffic to upstream traffic by $r = \frac{V_{down}}{V_{up}}$.

The variable traffic-sensitive cost (C) was plotted in Figure 3 for a traffic ratio of 1. For general traffic ratios, the variable traffic-sensitive cost is:

$$C = c^b (V_{down} + V_{up}) \frac{r ED^b_{down} + ED^b_{up}}{r + 1} \tag{11}$$

Figure 4 shows the effect of the traffic ratio on the variable traffic-sensitive cost. For traffic ratios at or below 2, the

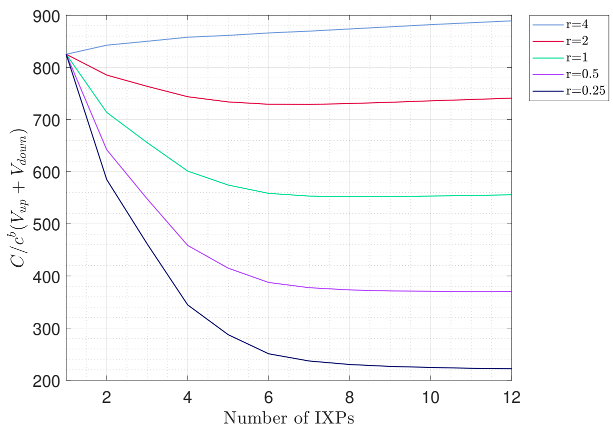

Fig. 4: Effect of the traffic ratio on the variable traffic-sensitive cost using hot potato routing

decrease in the upstream cost with the number of IXPs dominates the corresponding increase in the downstream cost, since the decrease in upstream cost due to *all* traffic exiting the ISP's network at a closer IXP outweighs the relatively small increase in downstream cost due to *some* traffic entering the ISP's network at a further away IXP. Recall that when the traffic ratio is 1, there is less than a 2% difference in the cost between $N = 6$ and the N at which cost is minimized, so there may be little value in requiring interconnection at more than 6 IXPs. We now find that when the traffic ratio is 0.5, there is less than a 2% difference in the cost between $N = 7$ and the N at which cost is minimized, and when the traffic ratio is 2, there is less than a 2% difference in the cost between $N = 4$ and the N at which cost is minimized.

In contrast, when the traffic ratio is 4, the increase in the downstream cost dominates the decrease in the upstream cost, since the downstream traffic volume is 4 times higher than the upstream traffic volume. As a result, the total cost increases with the number of IXPs, and thus it is no longer rational for the ISP to agree to settlement-free peering.

In conclusion, the traffic ratios at which an ISP will perceive approximately equal value from peering depends on the difference in value it is willing to accept, and the alternatives it has to deliver and receive traffic. However, we would expect the maximum acceptable traffic ratio to be 2:1 or less. Indeed, we observe that amongst the four largest ISPs, one specifies a maximum traffic ratio of 2:1, one specifies a maximum traffic ratio of 1.8:1, and the other two do not indicate specific thresholds but instead require a "general balance" of traffic. In addition, we observe that for traffic ratios at or below 2:1, it remains rational to require interconnection at a minimum of 6-7 IXPs.

VI. Conclusion

In order to explain the common settlement-free peering requirements of large ISPs, we examined the effect of the minimum number of interconnection points, the locations of these interconnection points, and traffic ratios on an ISP's variable traffic-sensitive costs. The four largest ISPs require peering at a minimum of 4 to 8 IXPs. Most of the large ISPs require interconnection at multiple interconnection points from specified lists. Most also require that the traffic ratio remains below a specified threshold, or that traffic should be generally balanced.

We find that when the traffic ratio is 1, the variable traffic-sensitive cost is uni-modal, and estimate that it is minimized with 8 IXPs. In our model, there is less than a 2% difference in the cost between 6 and 8 IXPs, which indicates there may be little value in requiring interconnection at more than 6 IXPs. We also show that the ISP's cost is typically minimized by selecting interconnection points that span the country and are near population centers.

We find that for traffic ratios at or below 2, the decrease in the upstream cost with the number of IXPs dominates the corresponding increase in the downstream cost, and thus interconnecting at 6 to 8 IXPs results in close to a minimum total cost. Requiring interconnection at more than 8 interconnection points is of little incremental value. In contrast, when the traffic ratio is above 2, the variation in the downstream cost with the number of IXPs dominates. As a result, the total cost increases with the number of IXPs, and thus it is no longer rational for the ISP to agree to settlement-free peering.

References

[1] Leichtman Research Group, Inc., "About 2,950,000 Added Broadband From Top Providers In 2021," *Retrieved from https://www.leichtmanresearch.com/about-2950000-added-broadband-from-top-providers-in-2021/. Accessed March 15, 2022.*

[2] PeeringDB, "The Interconnection Database," *Retrieved from https://www.peeringdb.com/. Accessed February 6, 2022.*

[3] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, and K. Claffy, "Using PeeringDB to understand the peering ecosystem," *Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.

[4] R. Johari and J. N. Tsitsiklis, "Routing and peering in a competitive Internet," in *Proceedings of the IEEE Conference on Decision and Control*, vol. 2, pp. 1556–1561, 2004.

[5] P. Faratin, D. D. Clark, S. Bauer, and W. Lehr, "Complexity of Internet interconnections: Technology, incentives and implications for policy," The Research Conference on Communications, Information, and Internet Policy, 2007.

[6] A. Dhamdhere, C. Dovrolis, and P. Francois, "A value-based framework for Internet peering agreements," in *International Teletraffic Congress*, pp. 1–8, 2010.

[7] D. Zarchy, A. Dhamdhere, C. Dovrolis, and M. Schapira, "Nash-peering: A new techno-economic framework for Internet interconnections," in *IEEE International Conference on Computer Communications*, pp. 403–408, 2018.

[8] D. D. Clark, W. Lehr, and S. Bauer, "Interconnection in the Internet: the policy challenge," The Research Conference on Communications, Information, and Internet Policy, 2011.

[9] R. A. Steenbergen, "A guide to peering on the Internet," in *NANOG 51*, 2011.

[10] United States Census Bureau, "Vintage County Population Estimates Totals," *Retrieved from https://www.census.gov/programs-surveys/popest/technical-documentation/research/evaluation-estimates/2020-evaluation-estimates/2010s-counties-total.html. Accessed January 26, 2022.*

[11] United States Census Bureau, "The U.S. Gazetteer Files," *Retrieved from https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.2021.html. Accessed January 26, 2022.*

[12] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani, "How many tiers? pricing in the Internet transit market," in *SIGCOMM*, pp. 194–205, 2011.