Are the Settlement-Free Peering Policy Requirements for ISPs and CDNs Based on Network Costs?

Ali Nikkhah[1], Scott Jordan[2]

Abstract

Peering is an interconnection arrangement between two networks for the purpose of exchanging traffic between these networks and their customers. Two networks will agree to settlement-free peering if this arrangement is superior for both parties compared to alternative arrangements including paid peering or transit. The conventional wisdom is that two networks agree to settlement-free peering if they receive an approximately equal value from the arrangement. Historically, settlement-free peering was only common amongst tier-1 networks, and these networks commonly require peering at a minimum specified number of interconnection points and only when the traffic ratio is within specified bounds. However, the academic literature does not explain how these requirements relate to the value to each network. More recently, settlement-free peering and paid peering have become common between ISPs and CDNs.

In this paper, we construct a network cost model to understand the rationality of common requirements on the number and location of interconnection points. We also wish to understand if it is rational to apply these requirements to interconnection between an ISP and a CDN. We construct a model of ISP traffic-sensitive network costs. We consider an ISP that offers service across the US. We parameterize the model using statistics about the population and locations of people in the contiguous US. We consider peering at the locations of the largest interconnection points in the US. We model traffic-sensitive network costs in the ISP's backbone network, middle-mile networks, and access networks. These costs are thus functions of routing policies, distances, and traffic volumes.

To qualify for settlement-free peering, large ISPs commonly require peering at a minimum of 4 to 8 mutually agreeable interconnection points. The academic literature provides little insight into this requirement or how it is related to cost. We show that the traffic-sensitive

[1]Ali Nikkhah is a Networked Systems Ph.D. student at the University of California, Irvine. Email: ali.nikkhah@uci.edu

[2]Scott Jordan is a Professor of Computer Science at the University of California, Irvine. He served as the Chief Technologist of the Federal Communications Commission during 2014-2016. Email: sjordan@uci.edu. Mailing address: 3214 Bren Hall, Department of Computer Science, University of California, Irvine, CA 92697-3435. Webpage: www.ics.uci.edu/~sjordan/

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

network cost decreases as the number of interconnection points increases, but with decreasing returns. The requirement to peer at 4 to 8 interconnection points is thus rational, and requiring interconnection at more than 8 points is of little value.

Finally, we turn to interconnection between an ISP and a CDN. Large ISPs often assert that CDNs should meet the same requirements on the number of interconnection points and traffic ratio to qualify for settlement-free peering. We show that if the CDN delivers traffic to the ISP locally, then a requirement to interconnect at a minimum number of interconnection points is rational, but a limit on the traffic ratio is not rational. We also show that if the CDN delivers traffic using hot potato routing, the ISP is unlikely to perceive sufficient value to offer settlement-free peering.

Keywords:
Broadband, Regulation, Net Neutrality, Interconnection, Peering

1. 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 *transit service* if and only if the transit provider agrees to accept and deliver traffic to and from the ISP, regardless of where this traffic is going to or from. In contrast, an interconnection arrangement is for *peering* if and only if each network agrees to accept and deliver traffic with destinations in its customer cone[3].

We focus on peering. 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). The conventional wisdom is that two Tier 1 networks agree to settlement-free peering if and only if the two networks perceive a roughly equal exchange of value from the peering arrangement. For example, if two Tier 1 networks are both ISPs with similar numbers of customers and similar size backbones, then they may perceive a roughly equal value from the exchange of traffic with destinations in their customer cones. Large ISPs often require that peers meet certain requirements, including a specified minimum number of interconnection points, a traffic ratio less than 2:1, and symmetric routing. The conventional wisdom is that these requirements are related to the perception of roughly equal value, but the academic literature has not yet established such a relationship.

More recently, it has become common for large ISPs and large content providers or content delivery networks (CDNs) to peer. However, there have often been disagreements between them over whether the peering arrangement should be settlement-free or paid. Large ISPs advertise the same settlement-free peering requirements for content providers as

[3] An ISP's customer cone consists of the union of its network, its customer's networks, its customer's customer's networks, etc.

for ISPs. However, large content providers do not satisfy requirements about traffic ratios, and often are more inclined to use non-symmetric routing. The academic literature has not lent much insight into when settlement-free peering between an ISP and a content provider is appropriate.

The focus of this paper is to relate the settlement-free peering requirements of large ISPs to the value the arrangement brings to the ISP. We represent value in terms of an ISP's traffic-sensitive costs. In particular, we examine settlement-free peering requirements on the minimum number of interconnection points, the locations of these interconnection points, and symmetric routing. We also wish to understand if it is rational to apply these settlement-free peering requirements to the interconnection between an ISP and a content provider.

We considered traffic ratios as another settlement-free peering requirement in [1]. We analyzed settlement-free peering requirements about traffic ratios between two ISPs using hot potato routing. Large ISPs require that the ratio of downstream to upstream traffic not exceed a specified threshold. We considered the conjecture that this requirement is related to a perception of roughly equal value. The traffic ratio determines the trade-off between the downstream and upstream costs. We showed that for traffic ratios above 2:1, the variation in the downstream cost with the number of IXPs dominates, and it is rational for the ISP not to agree to settlement-free peering. When traffic ratios are at or below 2:1, we estimated that requiring interconnection at more than 8 interconnection points is of little incremental value.

The paper is organized as follows. In Section 2, we summarize the settlement-free peering requirements of the ten largest ISPs in the United States. The four largest ISPs require interconnection at a minimum of 4 to 8 interconnection points from specified lists, that incoming and outgoing traffic be roughly balanced, and that the two parties use symmetric routing. 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.

In Section 3, we summarize the relevant research literature. Although a number of papers discuss settlement-free peering requirements, few analyze the relationship between these requirements and network costs. The academic literature thus provides little insight into why large ISPs impose settlement-free peering requirements or how these requirements are related to either the ISP's network cost or its perception of value. A few papers note that large ISPs interconnect at more interconnection points than do small ISPs, but these papers do not examine settlement-free peering requirements of a minimum number of interconnection points. Some papers examine peering between a content provider and an ISP, but none analyze the number or location of interconnection points, routing, or traffic ratios. The academic literature thus provides little insight into why large ISPs impose settlement-free peering requirements or how these requirements are related to either the ISP's network cost or its perception of value.

In Section 4, we develop a model in which subscribers are distributed according to census statistics. The ISP's network consists of access networks, middle mile networks, and a backbone. Access networks are based on counties, and interconnection points are chosen from a list of the largest exchanges in the United States. Traffic matrices are based on population, and traffic-sensitive costs are modeled as a function of both distance and traffic volume.

In Section 5, we determine the distances on each portion of an ISP's network over which it carries traffic to and from an end user. We calculate the average distance using traffic matrices. We model the average traffic-sensitive cost associated with carrying the traffic over these average distances.

In Section 6, we analyze settlement-free peering requirements about the number and location of interconnection points between two ISPs using hot potato routing. In hot-potato routing, an ISP hands off traffic to a downstream ISP as quickly as it can. Cold-potato routing is the opposite of hot-potato routing, where an ISP carries traffic as far as possible on its own network before handing it off to a downstream ISP. We consider the conjecture that such requirements are related to a perception of roughly equal value. When the traffic ratio is 1:1, we show that the ISP's cost is a uni-modal function of the number of interconnection points, and that there may be little value in requiring interconnection at more than 6 IXPs. The ISP's cost is typically minimized by selecting interconnection points that span the country and are near population centers.

In Section 7, we analyze interconnection between a large content provider and an ISP. Large ISPs have often asserted that content providers should meet the same settlement-free peering requirements on the number of interconnection points and the traffic ratio as do ISPs in order to qualify for settlement-free peering. However, it is not clear the degree to which the settlement-free peering requirements between two ISPs should apply to interconnection between a large content provider and an ISP. We first consider a content provider that does not replicate its content and delivers traffic using hot potato routing. We show that the ISP has little incentive to engage in settlement-free peering. We next consider a content provider that replicates all of its content at peering points and delivers 100% of traffic to the ISP locally. We show that it is rational for an ISP to agree to settlement-free peering, if the content provider agrees to interconnect at a minimum of 9 IXPs. Finally, we consider a content provider that hosts a content server at peering points, but that replicates only a portion of this content on each of these servers. We show that it is rational for an ISP to agree to settlement-free peering, if the content provider agrees to interconnect at a specified minimum number of interconnection points and to deliver a specified minimum proportion of traffic locally. However, we show that a limit on the traffic ratio is not rational.

2. Settlement-Free Interconnection Policies

We studied the settlement-free peering policies[4] of the ten largest ISPs in the United States [2–11]. Table 1 summarizes the most relevant requirements[5] of these policies.

[4]Often settlement-free peering policies are called settlement-free interconnection policies.

[5]However, meeting the requirements of a settlement-free peering policy is not a guarantee that the ISP will agree to peer on a settlement-free basis; instead, it is typically viewed as the minimum requirements to start a more detailed discussion with the ISP.

Table 1

ISP	Subscribers	Peering Inclination	Minimum Number of IXPs	Number of IXPs in ISP-Specified List	Minimum Traffic Volume	Traffic Ratio	Routing
Comcast	31,901,000	Selective	4	12	20 Gbps	Balanced	Common Policy
Charter	30,089,000	Selective	6-8	15	30 Gbps	-	Hot Potato
AT&T	15,504,000	Selective	6	12	30 Gbps	2:1	Common Policy
Verizon	7,365,000	Restrictive	8	-	12 Gbps	1.8:1	Hot Potato
Cox	5,530,000	Selective	2	15	1 Gbps	Balanced	Common Policy
CenturyLink	4,519,000	Selective	6	10	10 Gbps	1.5:1	Common Policy
Altice	4,386,200	Selective	2	-	0.1 Gbps	1.8:1	Hot Potato
Frontier	2,799,000	Selective	3	6	-	-	Hot Potato
Mediacom	1,463,000	Open	-	5	5 Gbps	-	Common Policy
TDS Telecom	526,000	Open	-	9	-	-	-

The column labeled "subscribers" gives an estimate of the number of subscribers of each ISP in 2021 [12], as settlement-free peering policies differ with the number of subscribers.

The column labeled "peering inclination" indicates the ISP's predisposition towards or against peering, as noted by PeeringDB [13]. The four largest ISPs (Comcast, Charter, AT&T, and Verizon) are either selective (i.e., have moderate requirements) or restrictive (i.e., have strong requirements). The next four largest ISPs (Cox, CenturyLink[6], Altice, and Frontier) are selective. The ninth and tenth largest ISPs (Mediacom and TDS Telecom) are open (i.e., have low requirements).

One of the most common requirements in settlement-free peering policies is a minimum number of mutually agreeable locations at which the two parties will agree to interconnect. Tier 1 ISPs typically interconnect in at least 8 of the 10 largest interconnection points (IXPs), as shown in the first ten rows of Table 2. The four largest ISPs each require interconnection at a minimum of 4 to 8 IXPs (see the column labeled "minimum number of IXPs" in Table 1). The next four largest ISPs are varied in their characteristics. CenturyLink's and Frontier's backbone networks likely qualify them as Tier 1 networks, but they have significantly fewer subscribers than do the four largest ISPs. Cox's and Altice's backbone networks likely qualify them as Tier-2 networks. These four moderate-size ISPs require interconnection at a minimum of 2 to 6 IXPs. The ninth and tenth largest ISPs do not specify a minimum number. Amongst ISPs that do specify a minimum number of mutually agreeable locations, most also require that these IXPs be chosen from ISP-specified lists (see the column labeled "number of IXPs in ISP-specified list" in Table 1), and often require that the chosen IXPs be geographically diverse.

Another common requirement in settlement-free peering policies is the minimum amount

[6]Throughout this paper, we consider CenturyLink as it was circa 2019. In 2020, it changed its name to Lumen, and it is currently attempting to sell a substantial portion of its business.

Table 1

List of Major U.S. IXPs	Comcast	Charter	AT&T	Verizon	Cox	Latitude	Longitude
Ashburn	√	√	√	√	√	39.0438° N	77.4874° W
Chicago	√	√	√	√	√	41.8781° N	87.6298° W
Dallas	√	√	√	√	√	32.7767° N	96.7970° W
San Jose	√	√	√	√	√	37.3382° N	121.8863° W
Los Angeles	√	√	√	√	√	34.0522° N	118.2437° W
New York	√		√	√	√	40.7128° N	74.0060° W
Seattle	√	√	√	√		47.6062° N	122.3321° W
Miami	√		√	√		25.7617° N	80.1918° W
Atlanta	√	√	√		√	33.7490° N	84.3880° W
Denver	√	√		√		39.7392° N	104.9903° W
Boston	√				√	42.3601° N	71.0589° W
Minneapolis		√				44.9778° N	93.2650° W

of traffic to be exchanged between the two networks. The three largest ISPs specify a minimum of 20 or 30 Gbps in the dominant direction (see the column labeled "minimum traffic volume" in Table 1). Smaller ISPs generally specify lower traffic volume thresholds, if they have such a requirement. In addition, settlement-free peering policies often require an approximate balance between incoming and outgoing traffic. Some (e.g., AT&T and Verizon) require that the ratio of incoming traffic volume to outgoing traffic volume not exceed a specified threshold. Others (e.g., Comcast) do not specify a maximum traffic ratio, but instead state the exchanged traffic should be in general balance[7].

Finally, settlement-free peering policies almost always have requirements about routing policies. Some ISPs (e.g., Charter and Verizon) require both parties to use hot potato routing[8], while other ISPs (e.g., Comcast and AT&T) only require that the two parties either both use hot potato or both use cold potato routing.

In the remainder of this paper, we focus on the settlement-free peering requirements of the four largest ISPs.

3. Research Literature

Although there are many papers in the academic literature that consider various aspects of peering, there are few that analyze the common requirements of settlement-free peering policies, and fewer yet that attempt to relate these requirements to the value of the peering agreement to each interconnecting network.

[7]For example, Comcast requires that "[a]pplicant must maintain a traffic scale between its network and Comcast that enables a general balance of inbound versus outbound traffic. The network cost burden for carrying traffic between networks shall be similar to justify SFI." [2].

[8]The settlement-free peering policies commonly use the term "shortest-exit routing".

PeeringDB is a database where ISPs (and other network operators) can provide information about the interconnection of their networks [13]. Lodhi et al. [14] 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, i.e., large ISPs interconnect at many IXPs, and that ISPs with large traffic volumes and a large number of subscribers are more likely to be classified by PeeringDB as having a selective or restrictive peering inclination. However, they did not analyze the particular requirements in settlement-free peering policies (e.g., the minimum number of IXPs or the traffic ratio), instead relying on PeeringDB's more coarse classification of peering inclination (i.e., restrictive, selective, or open). We have not found any academic papers that do. The closest may be Johari and Tsitsiklis [15], who discuss the selection of IXPs in a few networks with idealized and regular topologies.

In addition, there is some work that points out that traffic ratio requirements are not directly relevant to the case in which an ISP interconnects with a content provider or CDN. Clark et al. [16] discuss how interconnection between a content provider and an ISP differs from the interconnection between two ISPs. They suggest a simple model of interconnection between a content provider and an ISP, and use this model to consider settlement-free peering and paid peering. In the case of paid peering, they suggest that payment may be based either on bargaining power or on traffic ratio, but point out that traffic ratio may not be an accurate representation of benefit. However, they do not analyze the effect of the number of interconnection points nor the effect of routing upon an ISP's costs.

There are also some papers that model the benefits and costs of peering between a CDN and an ISP. Chang et al. [17] propose benefit-based and cost-based frameworks for interconnection decisions by ISPs. They suggest that large ISPs choose peers based on their geographic scope and number of customers, and the traffic ratio. Agyapong and Sirbu [18] examined the relationship between ISPs and CDNs and proposed a model of how routing or interconnection choices might influence total costs and potential payment flows. However, neither paper considers the number or location of interconnection points, nor routing.

As a result, the academic literature provides limited insight into how to judge disputes between ISPs and content providers over interconnection. In 2014, as part of the Federal Communications Commission's net neutrality proceeding, some large content providers and some large ISPs disagreed over the appropriate requirements for settlement-free peering between content providers and ISPs. For example, Verizon asserted that "[i]f parties exchange roughly equal amounts of traffic ..., then the parties may exchange traffic on a settlement-free basis", but that "when the traffic exchange is not roughly balanced, then the net sending party typically makes a payment in order to help compensate the net receiving party for its greater relative costs to handle the other party's traffic" [19]. In contrast, Netflix asserted that "[traffic] [r]atio-based charges no longer make economic sense since traffic ratios do not accurately reflect the value that networks derive from the exchange of traffic" [20].

4. Model

In this section, we develop a model that will enable our analysis of settlement-free peering policies. The analysis, presented in later sections, examines the effect on network cost of routing policies and the number of IXPs at which interconnecting networks meet. Thus, our model focuses on the characteristics that we believe are most critical to this analysis, and abstracts other less critical characteristics.

Our model is designed to reflect key characteristics of the United States. An Internet Service Provider (ISP) is assumed to serve the contiguous United States. We consider the geographic locations of the largest interconnection points (IXPs) in the United States. The network is partitioned into a backbone network, middle mile networks, and access networks. Middle mile and access networks are modeled based on U.S. counties. The density of the population is drawn from U.S. census statistics. Traffic matrices are built using these assumptions.

Subsection 4.1 introduces the topology of an ISP's U.S. network. Subsection 4.2 develops the traffic matrices over this network.

4.1. Topology

The topology of an ISP's U.S. network consists of a model of the ISP's service territory, the location of IXPs, and a model of segments of the network.

4.1.1. Service Territory

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 a single ISP whose service territory covers the contiguous United States.

The ISP's service region is modeled as the contiguous United States using geographical data of its boundaries. We use a coordinate system with x and y measured in degrees of longitude and latitude, respectively.

4.1.2. Location of IXPs

We focus on the interconnection between the ISP and a single interconnecting network (e.g., another ISP or a content provider). We use the geographic locations of the $M = 12$ largest IXPs in the United States, listed in Table 2 [21–26]. The coordinates of these M IXPs are denoted by $IXP(i)(i = 1, ..., M)$. We note that the largest ISPs in the United States (Comcast, Charter, AT&T, and Verizon) each interconnect at a minimum of 9 of these 12 IXPs, although a smaller ISP (Cox) interconnects at fewer IXPs; see Table 2.

An ISP and an interconnecting network often agree to interconnect at a smaller number $N < M$ of IXPs.

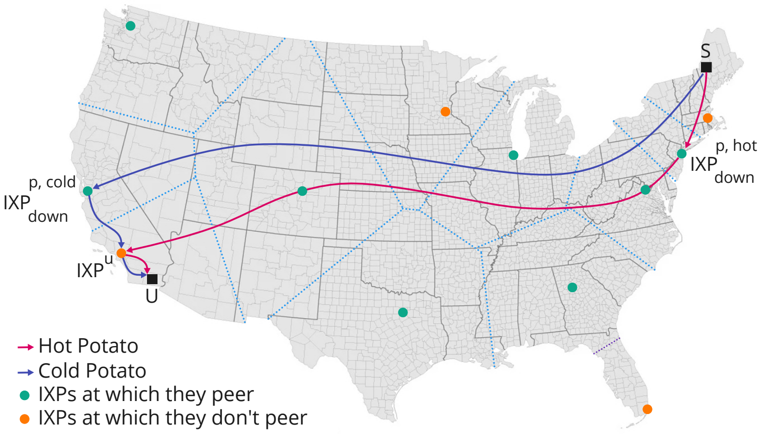

Figure 1: Topology of an ISP's network

4.1.3. Middle Mile Networks and Access Networks

We model the 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 differences in network costs between various forms of peering depend more critically on the number of interconnection points than on the topologies of access networks.

The locations of the geographical center of access networks are assigned to be the longitudes and latitudes of the center of each county in the contiguous United States [27]. A middle mile link is assumed to run from the geographical center of each access network to the closest IXP.

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

Figure 1 roughly illustrates these regions when the ISP and an interconnecting network agree to interconnect all 12 IXPs.[9] (We will discuss the case when $N = 8$ below.)

4.2. Traffic Matrices

We now turn to modeling the traffic matrices over the ISP's network. A traffic matrix represents the volume of traffic between all pairs of sources and destinations.

4.2.1. Distribution of Sources and End Users

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

[9] In the figure, the partition of the regions is only roughly illustrated. More precisely, they should follow county boundaries.

number of end users in each access network and (b) for each access network, the distribution of end users within the access network.

We 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. We denote the population of the county associated with access network j by p_j, and we denote by $p = \sum_j p(j)$ the population of the contiguous United States. We assign these values using U.S. census data [28]. 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 [29].

We focus here 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 by U. We consider two cases. When we consider interconnection between the ISP and another ISP (which we call the ISP-ISP case), the source S is on the other ISP's network. When we consider the interconnection between the ISP and a content provider (which we call the CP-ISP case), the source S may be at an IXP at which the content provider has a server. We consider the ISP-ISP case here, and we consider the CP-ISP case in Section 7.

We assume that the distribution of the source S is identical to the distribution of end users U, which is jointly given by $\{P(j)\}$ and the uniform distribution of end users within each access network. We assume that the source S and the end user U are independent.

4.2.2. Routes

In the analysis below, we distinguish between several points along traffic routes. We continue to focus on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. Along the route from the source S to the end user U, denote the location of the IXP at which downstream traffic enters the ISP's network with hot potato routing by $IXP^{p,hot}_{down}$, the location of the IXP at which downstream traffic enters the ISP's network with cold potato routing by $IXP^{p,cold}_{down}$, and the location of the IXP closest to the end user by IXP^u. These points are illustrated in Figure 1.

The ISP offers a portion of the route from a source S to an end user U. It carries traffic on its backbone from the IXP at which traffic enters the ISP's network (IXP^p_{down}) to the IXP closest to the end user (IXP^u), and it carries traffic on a middle mile network and access network from the IXP closest to the end user (IXP^u) to the end user (U). The portion of the route on the ISP's network thus depends on the joint distribution of (IXP^p_{down}, IXP^u, U).

The access network of the end user is distributed according to $\{P(j)\}$ and the end user U is uniformly distributed within the access network, as discussed above. The IXP closest to the end user is a deterministic function of U.

However, the IXP at which downstream traffic enters the ISP's network (IXP^p_{down}) depends on both the routing policy and the IXPs at which they agree to interconnect. Consider an ISP and an interconnecting network that agree to interconnect at N IXPs. If the ISP and the interconnecting network use hot potato routing, then the IXP at which downstream traffic enters the ISP's network ($IXP^{p,hot}_{down}$) is independent of the end user, and it is the IXP closest to the source among the IXPs at which they agree to peer. For example, suppose S is in Maine and U is in Imperial county, California. Then, as illustrated in Figure 1,

$IXP^{p,hot}_{down}$ might be in New York (if the two networks do not agree to peer in Boston) and IXP^u is in Los Angeles. Since end users are assumed to be distributed according to U.S. county population statistics, the distribution of $IXP^{p,hot}_{down}$ is given by the population of the corresponding region.

In contrast, if the ISP and the interconnecting network use cold potato routing, then the IXP at which downstream traffic enters the ISP's network ($IXP^{p,cold}_{down}$) is no longer independent of the end user, and it is the IXP closest to the end user at which they agree to peer. For example, suppose S is in Maine and U is in Imperial county, California. Then, as illustrated in Figure 1, $IXP^{p,cold}_{down}$ might be in San Jose (if the two networks do not agree to peer in Los Angeles) and IXP^u is in Los Angeles.

In the ISP-ISP case, there is also upstream traffic. The routes and distributions are similar, but inverted. If the ISP and the interconnecting network use hot potato routing, then the IXP at which upstream traffic enters the interconnecting network is the IXP closest to the end user at which they agree to peer. If the ISP and the interconnecting network use cold potato routing, then the IXP at which upstream traffic enters the interconnecting network is independent of the end user and follows a distribution similar to $IXP^{p,hot}_{down}$.

5. Traffic-sensitive Costs

Although we know that an ISP's traffic-sensitive cost is a complicated function of the topology of the network, to make the analysis tractable, we abstract the network geographically into three non-overlapping sections: backbone, middle mile, and access. We define the backbone network as the set of links between IXPs. We define the middle mile networks as the set of links between the geographical center of each access network and the closest IXP. We define the access networks as the set of links that connect the middle mile networks to end users.

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

5.1. 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 $D^b(IXP^p_{down}, IXP^u)$, the distance between IXP^p_{down} and 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 $D^m(IXP^u, U)$, the distance between the IXP closest to the end user (IXP^u) and the location of the geographical center of access network on which U resides.

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 $D^a(U)$, the distance between the location of the end user (U) and the geographical center of access network on which U resides. The distance can be determined by the location of U within the access network.

All distances in the our model are great-circle distances between the corresponding points on a sphere, and are calculated using the Haversine formula.

5.2. Average Distances

An ISP's traffic-sensitive cost depends on the average distance of traffic on each segment of its network. As discussed above, the distances on each section of the ISP's network depend on the joint distribution of (IXP^p, IXP^u, U).

We continue to focus on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. The distance on the ISP's backbone network is a function of (IXP^p_{down}, IXP^u).

When hot potato routing is used, the IXP at which downstream traffic enters the ISP's network ($IXP^{p,hot}_{down}$) 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 routing policy and the IXPs at which they agree to interconnect. Consider an ISP and an interconnecting network that agree to interconnect at N IXPs. The average distance on the ISP's backbone network ($ED^{b,hot}_{down}$) is the average distance between the IXPs at which downstream traffic enter the ISP's network ($IXP^{p,hot}_{down}$) and the IXPs closest to the end user (IXP^u). The probability distribution of IXP^u is determined by the population of each corresponding region.

When cold potato routing is used, the IXP at which downstream traffic enters the ISP's network ($IXP^{p,cold}_{down}$) is the IXP closest to the end user at which they agree to peer. In addition, $IXP^{p,cold}_{down}$ is no longer independent of IXP^u. Since they might not agree to peer at all IXPs, the ISP might still carry traffic across a portion of its backbone, namely from $IXP^{p,cold}_{down}$ to IXP^u, and the average such distance ($ED^{b,cold}_{down}$) is the average distance between the IXPs closest to the end user at which they agree to peer ($IXP^{p,cold}_{down}$) and the IXPs closest to the end user (IXP^u).

The distance on the ISP's middle mile network is a function of (IXP^u, U). It is independent of the routing policy, and the average distance (ED^m) is the average distance between the IXPs closest to the end user and the center of the access network on which the end user resides.

The distance on the ISP's access network is a function of U. It is also independent of the routing policy. Since end users are uniformly distributed within each access network, but not between access networks, the average distance (ED^a) is the weighted average of average distances within each access network.

5.3. Cost

The ISP incurs a traffic-sensitive cost for carrying traffic over the average distances calculated in the previous subsection. We consider here only traffic-sensitive costs, because non-traffic-sensitive costs do not vary with routing policies or the number of interconnection points.[10]

Traffic-sensitive costs are a function of both distance and traffic volume. We assume here that traffic-sensitive costs are linearly proportional to the average distance over which the traffic is carried on each portion of the ISP's network, see e.g., [30]. We also assume that traffic-sensitive costs are linearly proportional to the average volume of traffic that an ISP carries on each portion of its network. Although the cost might be an increasing concave function of traffic volume (or a piecewise constant function), the linear model will suffice for our analysis.

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, the cost per unit distance and per unit volume in the middle mile networks by c^m, and the cost per unit distance and per unit volume 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 fixed source-destination traffic matrix, the average distance across the ISP's access networks is constant. In addition, the average distance across the ISP's middle mile networks is constant, once we fix $M = 12$, since the IXPs at which the parties agree to peer do not affect the middle mile. The variable portion of the ISP's traffic-sensitive cost is thus only $C = c^b ED^b V$

Below we consider the effect on the variable traffic-sensitive cost (C) of changes in the number of IXPs at which peering occurs and routing policies, for constant c^b. (In the remainder of the paper, we use the term cost to refer to the variable traffic-sensitive cost.) We will find that changes in the number of IXPs at which peering occurs and routing policies affect ED^b.

6. Number and List of IXPs

We now turn to analyzing the effect on an ISP's variable traffic-sensitive costs of the number of IXPs at which peering occurs, and routing policies. We are in particular interested in explaining the settlement-free peering policies of large ISPs.

In this section, we examine the effect of the number of IXPs at which two ISPs agree to peer, and the list of IXPs at which they peer. As shown in Table 1, large ISPs require other ISPs 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.

[10] There is a small cost for each interconnection point; however, this cost is relatively small compared to transportation costs.

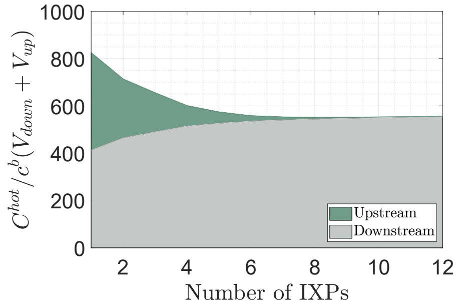

Figure 2: Total costs

In order to better understand these requirements of settlement-free peering policies, we initially focus only on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. We first assume that both ISPs use hot potato routing. The cost of downstream traffic is $C^{hot}_{down} = c^b V_{down} ED^{b,hot}_{down}$, where $ED^{b,hot}_{down}$ is the average distance on the ISP's backbone network.

As we discussed in Section 4.2.2, the route that upstream traffic takes when using hot potato routing is the same route (but in the opposite direction) that downstream traffic takes when using cold potato routing. The average distance the ISP carries upstream traffic across the backbone when using hot potato routing is the same as the average distance the ISP carries downstream traffic across the backbone when using cold potato routing. The cost of upstream traffic is $C^{hot}_{up} = c^b V_{up} ED^{b,hot}_{up}$, where the average upstream distance on the ISP's backbone network using hot potato is equal to the average downstream distance on the ISP's backbone network using cold potato routing.

Figure 2 shows the effect of the number of interconnection points at which they peer (N) on the cost of both downstream and upstream traffic using hot potato routing. (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 number of interconnection points affects the backbone cost. The cost is a function of the average distances between $IXP^{p,hot}_{down}$ and IXP^u, and of the average distances between $IXP^{p,hot}_{up}$ and IXP^u for all IXPs. As the number of IXPs at which they peer increases, the IXP closest to the end user IXP^u is fixed since it is related to the location of the end user and a deterministic function of U. However, $IXP^{p,hot}_{down}$ and $IXP^{p,hot}_{up}$ change.

IXP^s depends on the number of IXPs at which they peer. The IXP at which traffic enters the ISP's network ($IXP^{p,hot}_{down}$) is the IXP closest to the source among the IXPs at which they agree to peer, and thus as the number of IXPs at which they peer increases, $IXP^{p,hot}_{down}$ moves farther from IXP^u and the distance between $IXP^{p,hot}_{down}$ and IXP^u increases. Thus, the downstream cost increases. The downstream cost is concave, because the incremental distance from $IXP^{p,hot}_{down}$ to IXP^u associated with adding another IXP decreases, namely there are decreasing returns.

However, at the same time, the ISP exchanges upstream traffic using hot potato routing at the IXP closest to the end user at which the two ISPs agree to peer ($IXP^{p,hot}_{up}$). Thus as the number of IXPs at which they peer increases, $IXP_{up}^{p,hot}$ moves closer to IXP^u and the distance between $IXP_{up}^{p,hot}$ and IXP^u decreases. Therefore, the upstream cost decreases. The upstream cost is convex, because the incremental distance from IXP^u to $IXP_{up}^{p,hot}$ associated with adding another IXP decreases, namely there are decreasing returns.

When the traffic ratio is 1:1, the decrease in the upstream cost exceeds the increase in the downstream cost. In the upstream route, $IXP_{up}^{p,hot}$ is the closest IXP to IXP^u, whereas, in the downstream route, $IXP_{down}^{p,hot}$ 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 the distance between $IXP_{up}^{p,hot}$ and IXP^u in the upstream route is higher than the slope of the distance between $IXP_{down}^{p,hot}$ and IXP^u in the downstream route. It follows that the total cost decreases.

We observe that, when the traffic ratio is 1:1, the 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. Based on Table 1, 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 given in Table 2, so as to minimize the ISP's cost (C^{hot}).

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.

Comcast not only requires that potential settlement-free peering partners 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[2].[11] For $N = 4$, our model chooses IXPs in Ashburn, Chicago, Los Angeles, and Atlanta, i.e. 1 on the west coast, 2 on the east coast, and 1 in the middle. All 4 of these cities are on Comcast's list.

Charter not only requires that potential settlement-free peering partners agree to peer at a minimum of 6-8 IXPs from Charter's list of IXPs[12], 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[3].[13] 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

[11]Comcast does not specify which IXPs are considered to be on the west coast, on the east coast, or in the central region. It also requires that the IXPs be "mutually agreeable".

[12]It requires a minimum of 6 IXPs when the 95th percentile of traffic exchanged is less 500 Gbps in the dominant direction, and it requires a minimum of 8 IXPs when it exceeds this threshold.

[13]Charter also specifies the IXPs in each of these regions.

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.

7. Peering between a content provider and an ISP

There have been an increasing number of disputes over interconnection between large ISPs, on one side, and large content providers and transit providers, on the other side. In 2013-2014, a dispute between Comcast and Netflix over terms of interconnection went unresolved for a substantial period of time, resulting in interconnection capacity that was unable to accommodate the increasing Netflix video traffic. In 2014, Netflix and a few transit providers brought the issue to the attention of the United States Federal Communications Commission (FCC), which was writing updated net neutrality regulations. The FCC discussed the dispute in the 2015 Open Internet Order [31].

The FCC first summarized the arguments of large content providers and transit providers. It noted that "[content] providers argue that they are covering the costs of carrying [their] traffic through the network, bringing it to the gateway of the Internet access service". Large content providers and transit providers argued that they should be entitled to settlement-free peering if the interconnection point is sufficiently close to consumers. The lack of willingness of large ISPs to offer settlement-free peering with large content providers, and to augment the capacity of existing interconnection points with transit providers with which they had settlement-free peering agreements, had led to the impasse. The FCC noted that "[s]ome [content] and transit providers assert that large [ISPs] are creating artificial congestion by refusing to upgrade interconnection capacity ... for settlement-free peers or CDNs, thus forcing [content] providers and CDNs to agree to paid peering arrangements."

The FCC then summarized the arguments of large ISPs. It noted that "large broadband Internet access service providers assert that [content] providers such as Netflix are imposing a cost on broadband Internet access service providers who must constantly upgrade infrastructure to keep up with the demand". The large ISPs explained that the network upgrades include adding capacity in the middle mile and access networks. The FCC noted that the large ISPs asserted that if they absorb these costs, then the ISPs would recoup these costs by increasing the prices for all subscribers, and that the large ISPs argued that "this is unfair to subscribers who do not use the services, like Netflix, that are driving the need for additional capacity".

Both large ISPs and large content providers agree that settlement-free peering is appropriate when both sides perceive equal value to the relationship. However, whereas large content providers assert that carrying their traffic to an interconnection point close to consumers is of value, large ISPs assert that "if the other party is only sending traffic, it is not contributing something of value to the broadband Internet access service provider".

In 2015, the FCC was concerned about the duration of unresolved interconnection disputes and about the impact of these disputes upon consumers. However, it concluded that in 2015 it was "premature to draw policy conclusions concerning new paid Internet traffic exchange arrangements between broadband Internet access service providers and [content] providers, CDNs, or backbone services." Thus, in 2015 the FCC adopted a case-by-case approach in which it would monitor interconnection arrangements, hear disputes, and ensure that ISPs are not engaging in unjust or unreasonable practices. However, in 2018, the FCC reversed itself and ended its oversight of interconnection arrangements, when it repealed most of the 2015 net neutrality regulations [32]. It is almost certain that the FCC will revisit the issue in the next few years.

In the paper so far, our focus has been on the ISP-ISP case, in which two ISPs interconnect with each other. Settlement-free peering policies were originally constructed for peering between two Tier 1 ISPs. However, it has become common for large content providers to peer with ISPs. We call this the CP-ISP case. It is not clear the degree to which the settlement-free peering requirements discussed above should apply to the CP-ISP case.

Large ISPs do not generally have different settlement-free peering policies for content providers than for ISPs and transit providers. In addition, they have often asserted that content providers should meet the same requirements on the number of interconnection points and traffic ratio to qualify for settlement-free peering. We will show that if a content provider delivers traffic to the ISP locally, then a requirement to interconnect at a minimum number of interconnection points is rational, but a limit on the traffic ratio is not rational. We will also show that if a content provider does not deliver traffic locally, the ISP is unlikely to perceive sufficient value to offer settlement-free peering. In Section 7.1, we consider a content provider that does not replicate its content and delivers traffic using hot potato routing. In Section 7.2, we consider a content provider that replicates all of its content and delivers 100% of traffic to the ISP locally. Finally, in Section 7.3, we consider a content provider that replicates only a portion of its content and delivers only that portion to the ISP locally.

7.1. No Content Replication

We first consider a content provider that does not replicate its content and delivers traffic using hot potato routing.

The ISP network topology remains the same as was presented in Section 4.1. The distribution of the location of end users remains the same as was presented in Section 4.2.1. For downstream traffic, we assume that the location of the content requested by an end user is independent of the location of the end user. We also assume that the distribution of the location of the content requested by an end user, namely the distribution of the source S, is identical to the distribution of end users. We further assume that the content provider uses hot potato routing. The routing of downstream traffic is thus identical to that considered in Section 4.2.2. As a result, the distances that the ISP carries downstream traffic from a content provider across the ISP's network remain the same as was presented in Section 5.

However, whereas in the ISP-ISP case there was both downstream and upstream traffic, we assume that in the CP-ISP case the volume of upstream traffic is negligible. As a result, the ISP's costs are those discussed in Section 5.3, but only for downstream traffic. Equivalently, we can think of this CP-ISP case as being equivalent to an ISP-ISP case with hot potato routing and a traffic ratio of infinity. It follows that the ISP's total cost in this

CP-ISP case is the same as the ISP's downstream cost in the ISP-ISP case. This downstream cost C_{cp}^{no} is equal to C_{down}^{hot}.

The effect of the number of interconnection points (N) on the ISP's downstream cost is illustrated in figure 2. The cost is minimized when $N = 1$. We conclude an ISP has little incentive to peer at multiple IXPs with a content provider that does not replicate content and that uses hot potato routing. This is not surprising, since as we discussed in [1], it is not rational for an ISP to agree to settlement-free peering with another ISP when the traffic ratio of downstream to upstream traffic is high.

7.2. Full Content Replication

We next consider a content provider that hosts a content server at each IXP at which it agrees to peer with an ISP, that replicates all of its content on each of these servers, and that transmits all of its traffic locally.

The ISP network topology remains the same as was presented in Section 4.1. The distribution of the location of end users remains the same as was presented in Section 4.2.1. However, the location of the content is no longer the same as in previous sections. We now assume that the location of the content requested by an end user is the content server located at the IXP closest to the end user at which the content provider and the ISP agree to peer. Routing is now irrelevant, since the content is entering the ISP's network directly from the content server. The distances that the ISP carries downstream traffic from a content provider across the ISP's network are the same as the distances in the ISP-ISP case that an ISP carries downstream traffic when using cold potato routing, which were given in Section 5.2. We again assume that the volume of upstream traffic is negligible. Equivalently, we can think of this CP-ISP case as being equivalent to an ISP-ISP case with cold potato routing and a traffic ratio of infinity. As a result, the ISP's cost in this CP-ISP case is $C_{cp}^{full} = c^b V_{down} ED_{down}^{b,cold}$.

The effect of the number of interconnection points (N) on the ISP's downstream cost is illustrated in figure 3(a). The cost decreases as the number of IXPs increases. Note that the ISP has an incentive to increase the number of IXPs at which the two parties peer *despite* the fact that the traffic ratio is infinity, unlike in the ISP-ISP case. However, there is less than a 2% difference in the cost between $N = 9$ and $N = 12$, so this indicates there may be little value in requiring interconnection at more than 9 IXPs.

Comparing the cost in the ISP-ISP case (illustrated in figure 2) to the cost in this CP-ISP case (illustrated in figure 3(a)), we find that there is little value in requiring interconnection at more than 6 IXPs in the ISP-ISP case, but there is a significant incremental value in the CP-ISP case to increasing the number of IXPs to at least 9.

The number of interconnection points at which the cost curve flattens is higher in the CP-ISP case than in the ISP-ISP case. In the CP-ISP case, the cost is entirely incurred by carrying downstream traffic, which is localized. In the ISP-ISP case, there are cost components for both downstream and upstream traffic. The upstream cost using hot potato routing is the same as the downstream cost in the CP-ISP case, which is similarly decreasing in N. However, the downstream cost using hot potato routing in the ISP-ISP case is increasing with N, which causes the total cost to flatten out at lower values of N.

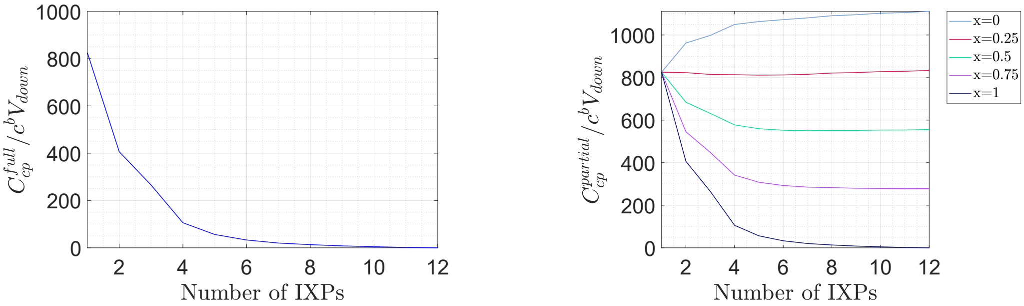

(a) Complete Replication

(b) Partial Replication

Figure 3: Costs

This comparison indicates that it is likely rational for an ISP to agree to settlement-free peering with a content provider that replicates its content at all agreed peering points and delivers all traffic locally, as long as it agrees to interconnect at a minimum of 9 IXPs. We would thus expect large ISPs to have *different* settlement-free peering requirements for such content providers than for ISPs. First, we would expect the minimum number of interconnection points to be higher for content providers than ISPs. Second, we would certainly expect there to be *no* traffic ratio requirements for content providers. Third, we expect there to be some type of traffic localization requirement. We turn to this last requirement in the next subsection.

7.3. Partial Content Replication

Finally, we consider a content provider that hosts a content server at each IXP at which it agrees to peer with an ISP, but that replicates only a portion of this content.

The ISP network topology remains the same as was presented in Section 4.1, and the distribution of the location of end users remains the same as was presented in Section 4.2.1. However, the location of the content is no longer the same as in previous sections. We assume that, within each access network, a proportion x of requests is served by the content server located at the IXP closest to the end user at which the content provider and the ISP agree to peer. We also assume that, within each access network, the remaining proportion $1-x$ of requests is served by a content server that is independent of the location of the end user, and that the distribution of the location of this content server is identical to the distribution of end users. We further assume that the content provider uses hot potato routing for non-locally delivered content.

The ISP's cost in this CP-ISP case is:

$$C_{cp}^{partial} = xC_{cp}^{full} + (1-x)C_{cp}^{no} = c^b V_{down}\left(xED_{down}^{b,cold} + (1-x)ED_{down}^{b,hot}\right) \qquad (1)$$

The effect of the number of interconnection points at which they agree to peer (N) on the ISP's downstream cost is illustrated in figure 3(b), for various values of the proportion x. When $x < 0.3$, too little of the downstream traffic from the content provider to the ISP is delivered locally, and as the cost to the ISP increases as the number of IXP increases.

However, when $x > 0.3$, the ISP benefits from increasing the number of IXPs at which the two parties agree to peer.

We conclude that it is likely rational for an ISP to agree to settlement-free peering with a content provider that provides partial replication and delivers that portion locally. We expect that the ISP may require a specified minimum amount of traffic to be delivered locally. We expect the ISP to require interconnection at a specified minimum number of interconnection points, although the number may depend on the amount of traffic delivered locally. However, we certainly expect there to be *no* traffic ratio requirements.

8. Conclusion

In order to explain the common settlement-free peering requirements of large ISPs, we examined the effect of the number of interconnection points at which two networks peer and the locations of these interconnection points on an ISP's variable traffic-sensitive costs. When two ISPs peer with a traffic ratio of 1:1, the variable traffic-sensitive cost is uni-modal, and we estimate that it is minimized with 8 IXPs. There may be little value in requiring interconnection at more than 6 IXPs. The ISP's cost is typically minimized by selecting interconnection points that span the country and are near population centers.

However, when a content provider interconnects with an ISP, it is not clear the degree to which the settlement-free peering requirements between two ISPs discussed above should apply. Large ISPs often argue that large content providers should meet the same requirements as other ISPs to qualify for settlement-free peering. If a content provider does not replicate its content and uses hot potato routing, an ISP is unlikely to perceive sufficient value to offer settlement-free peering. However, if a content provider does replicate its content, it is rational for an ISP to agree to settlement-free peering if the content provider agrees to interconnect at a specified minimum of IXPs and to deliver a specified minimum proportion of traffic locally. However, we show that a limit on the traffic ratio is not rational.

References

[1] A. Nikkhah and S. Jordan, "Requirements of settlement-free peering policies," in *2022 IEEE Global Communications Conference (GLOBECOM)*, 2022.

[2] Comcast, "Settlement-Free Interconnection (SFI) Policy," Accessed January 26, 2022. [Online]. Available: https://www.xfinity.com/peering

[3] Charter, "Settlement-Free Interconnection (SFI) Policy," Accessed January 26, 2022. [Online]. Available: https://www.spectrum.com/policies/ip-interconnection

[4] AT&T, "Global IP Network Peering Policy," Accessed January, 2020. [Online]. Available: http://www.att.com/peering/

[5] Verizon, "Interconnection Policy For Internet Networks," Accessed January 26, 2022. [Online]. Available: https://enterprise.verizon.com/terms/peering/

[6] Cox, "Internet Peering / Interconnection," Accessed January 26, 2022. [Online]. Available: http://www.cox.com/peering

[7] CenturyLink, "North America IP Network Peering Policy," Accessed April 21, 2022. [Online]. Available: https://web.archive.org/web/20190722152324/http://www.centurylink.com/legal/peering_na.html

[8] Suddenlink, "Communications' Settlement-Free Interconnection (Peering) Policy," Accessed January 26, 2022. [Online]. Available: https://www.suddenlink.com/terms-policy/peering

[9] Frontier Communications, “Bilateral Peering Policy,” Accessed January 26, 2022. [Online]. Available: https://ipadmin.frontier.com/bilateralpeering_policy
[10] Mediacom Communications Corporation, “Interconnection (Peering) Policy,” Accessed January 26, 2022. [Online]. Available: http://mediacomcable.com/interconnection-policy
[11] TDS Telecommunications Corp, “AS4181 Peering Policies,” Accessed January 26, 2022. [Online]. Available: http://as4181.net/peering.html
[12] Leichtman Research Group, Inc. (LRG), “About 2,950,000 Added Broadband From Top Providers In 2021,” Accessed March 15, 2022. [Online]. Available: https://www.leichtmanresearch.com/about-2950000-added-broadband-from-top-providers-in-2021/
[13] PeeringDB, “The Interconnection Database,” Accessed February 6, 2022. [Online]. Available: https://www.peeringdb.com/
[14] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, and K. Claffy, “Using PeeringDB to understand the peering ecosystem,” *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 2, pp. 20–27, 2014.
[15] 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, 2004, pp. 1556–1561.
[16] D. D. Clark, W. Lehr, and S. Bauer, “Interconnection in the Internet: the policy challenge.” The Research Conference on Communications, Information, and Internet Policy (TPRC), 2011.
[17] H. Chang, S. Jamin, and W. Willinger, “To peer or not to peer: Modeling the evolution of the internet’s as-level topology,” in *Proceedings IEEE INFOCOM*, 2006, pp. 1–12.
[18] M. A. Sirbu and P. Agyapong, “Economic incentives in content-centric networking: Implications for protocol design and public policy.” The Research Conference on Communications, Information, and Internet Policy (TPRC), 2011.
[19] “Reply Comments of Verizon and Verizon Wireless, In the Matter of Framework for Broadband Internet Service (GN Docket No. 10-127) and Open Internet Rulemaking (GN Docket No. 14-28),” September 15, 2014.
[20] “Comments of Netflix, In the Matter of Framework for Broadband Internet Service (GN Docket No. 10-127) and Open Internet Rulemaking (GN Docket No. 14-28),” July 15, 2014.
[21] Data Center Map, “Internet Exchange Points,” Accessed January 26, 2022. [Online]. Available: https://www.datacentermap.com/ixps.html
[22] TeleGeography, “Internet Exchange Maps,” Accessed January 26, 2022. [Online]. Available: https://www.internetexchangemap.com/#/country/united-states
[23] Packet Clearing House, “Peering,” Accessed January 26, 2022. [Online]. Available: https://www.pch.net/about/peering
[24] Network Startup Resource Center (NSRC), “Internet EXchange Points: North America Region,” Accessed January 26, 2022. [Online]. Available: https://nsrc.org/ixp/NorthAmerica.html
[25] PeeringDB, “Exchanges,” Accessed January 26, 2022. [Online]. Available: https://www.peeringdb.com/advanced_search?country__in=US&reftag=ix
[26] Hurricane Electric Internet Services, “Internet Exchange Report,” Accessed January 26, 2022. [Online]. Available: https://bgp.he.net/report/exchanges
[27] United States Census Bureau, “The U.S. Gazetteer Files,” Accessed January 26, 2022. [Online]. Available: https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.2021.html
[28] ——, “Vintage county population estimates totals,” Accessed January 26, 2022. [Online]. Available: https://www.census.gov/programs-surveys/popest/technical-documentation/research/evaluation-estimates/2020-evaluation-estimates/2010s-counties-total.html
[29] ——, “American Fact Finder,” Accessed 2020. [Online]. Available: https://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?src=bkmk
[30] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani, “How many tiers? pricing in the Internet transit market,” in *Proceedings of the ACM Special Interest Group on Data Communication Conference (SIGCOMM)*, 2011, pp. 194–205.

[31] Federal Communications Commission, “Protecting And Promoting The Open Internet, Report And Order On Remand, Declaratory Ruling, And Order, 30 FCC Rcd 5601,” 2015.
[32] ——, “Restoring Internet Freedom, Declaratory Ruling, Report And Order, And Order, 33 FCC Rcd 311,” 2018.