Should Large ISPs Apply the Same Settlement-Free Peering Policies To Both ISPs and CDNs?

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*—Large Internet Service Providers (ISPs) often require that peers meet certain requirements to be eligible for free-settlement peering. 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. The focus of this paper is to relate the settlement-free peering requirements between two large ISPs and understand the degree to which the settlement-free peering requirements between them should apply to the peering between large ISPs and content providers.**

We analyze settlement-free peering requirements about the number and location of interconnection points (IXPs). Large ISPs often require interconnection at a minimum of 6 to 8 interconnection points. We find that the ISP's traffic-sensitive cost is decreasing and convex with the number of interconnection points. We also observe that there may be little value in requiring interconnection at more than 8 IXPs. We then analyze the interconnection between a large content provider and an ISP. 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.

***Index Terms*—Internet Interconnection, Net Neutrality, 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.

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. 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 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 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. In section III, we develop a model that will enable us to examine the effect of the number and location of IXPs. In Section IV, 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 V, we analyze settlement-free peering requirements about the number and location of IXPs. Large ISPs often require interconnection at a minimum of 6 to 8 interconnection points. We find that the ISP's traffic-sensitive cost is decreasing and convex with the number of

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

interconnection points. We also observe that there may be little value in requiring interconnection at more than 8 IXPs. In Section VI, we analyze the interconnection between a large content provider and an ISP. 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.

II. 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.

PeeringDB is a database where ISPs (and other network operators) can provide information about the interconnection of their networks [2]. 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, 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 [4], 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. [5] 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. [6] 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 [7] 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, and neither paper justifies traffic ratio requirements.

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" [8]. 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" [9].

III. Model

In this section, we develop an analytical model that is designed to foster closed-form analysis. Previously, we developed a numerical model which was designed to reflect key characteristics of the United States [10].

A. Topology

The topology of an ISP's network consists of a model of the ISP's service territory, the location of IXPs, and a model of segments of the 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 a subset of the IXPs at which they peer that 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 simplistically modeled as a rectangular abstraction US of the contiguous United States, measuring $L = 2800$ miles from west to east and $W = 1582$ miles from south to north [11]. We use a coordinate system (x, y) centered on this rectangle, i.e. $US = \left[-\frac{L}{2}, \frac{L}{2}\right] \times \left[-\frac{W}{2}, \frac{W}{2}\right]$.

We focus on the interconnection between the ISP and a single interconnecting network (e.g., another ISP or a content provider). We denote by N the number of IXPs at which the ISP and the interconnecting network agree to peer. We denote the location of IXP i ($i \in l^N = \{1, \ldots, N\}$) by $IXP(i)$. We simplistically assume that these N IXPs are located at the middle latitude $y = 0$ and at equally spaced longitudes x, i.e.

$$IXP(i) = \left(-\frac{L}{2} + \frac{L(2i-1)}{2N}, 0\right) \tag{1}$$

We denote the set of locations of the IXPs at which the ISP interconnects with this interconnecting network by $I = \{IXP(i), i \in l^N\}$. We model the ISP's network as partitioned into a single backbone network, multiple middle mile

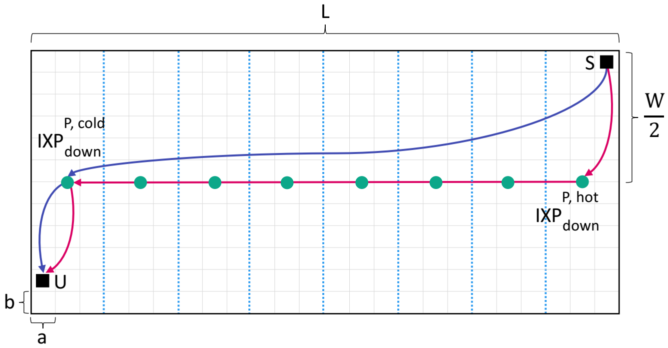

Fig. 1: Topology of an ISP's network

networks, and multiple access networks. We model each access network as a rectangle of size a miles from west to east and b miles from south to north. We index the access networks from west to east (j) and from south to north (k), so that a particular access network is referred to by the pair of indices (j,k), where $j = 1, \ldots, L/a$, and $k = 1, \ldots, W/b$. We denote by $Access(j,k)$ the geographical region of access network (j,k). We denote the location of the geographical center of access network (j,k) by

$$A(j,k) = \left(-\frac{L}{2} + \frac{(2j-1)a}{2}, -\frac{W}{2} + \frac{(2k-1)b}{2}\right) \quad (2)$$

A middle mile link is assumed to run from the geographical center of each access network (j,k) to the closest IXP. The IXPs can be used to partition the ISP's service territory into a set of regions closest to each IXP. Denote by $R(IXP(i))$ the geographical region that consists of the union of access networks for which the closest interconnection point is IXP i, namely

$$R(IXP(i)) = \bigcup_{\substack{(j,k) \mid \|A(j,k)-IXP(i)\| \le \\ \|A(j,k)-IXP(i')\| \, \forall i' \in l^N}} Access(j,k) \quad (3)$$

Figure 1(a) illustrates these regions for our model. Each region is simply a rectangle:

$$R(IXP(i)) = \left[-\frac{L}{2} + \frac{(i-1)L}{N}, -\frac{L}{2} + \frac{iL}{N}\right] \times \left[-\frac{W}{2}, \frac{W}{2}\right] \quad (4)$$

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.

We denote the probability that an end user resides within access network (j,k) by $P(j,k)$. We simply assume that end users are uniformly distributed across access networks, i.e. $P(j,k) = ab/LW$. We also simply assume that end users are uniformly distributed within each access network.

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 VI.

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

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 end user U is uniformly distributed in US, as discussed above. The IXP closest to the end user is a deterministic function of U, namely $IXP^u = (g' | \, U \in R(g'))$.

However, the IXP at which downstream traffic enters the ISP's network (IXP^p_{down}) depends on the routing policy. If the ISP and the interconnecting network use hot potato routing, then the IXP at which downstream traffic enters the ISP's network is independent of the end user, and it is the IXP closest to the source, i.e. $IXP^{p,hot}_{down} = (g | \, S \in R(g))$. Since end users are assumed to be uniformly distributed across the contiguous U.S., $IXP^{p,hot}_{down}$ is also uniformly distributed:

$$P(IXP^{p,hot}_{down} = g) = \sum_{Access(j,k) \subset R(g)} P(j,k) = \frac{1}{N} \quad (5)$$

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 is no longer independent of the end user, and it is the IXP closest to the end user, i.e. $IXP^{p,cold}_{down} = IXP^u$.

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, i.e. $IXP^{p,hot}_{up} = IXP^u$. 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 (5).

IV. Traffic-sensitive Costs

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.

The distances on each section on the ISP's network depend on the joint distribution of (IXP^p_{down}, IXP^u, U). All distances are Euclidean distances between the corresponding points on a plane. The distance 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). We denote the distance on the ISP's backbone network between these two IXPs by $D^b(IXP^p_{down}, IXP^u) = \|IXP^p_{down} - IXP^u\|$. Denote by i^p the IXP at location IXP^p_{down}, i.e. $i^p = i|(IXP^p_{down} = IXP(i))$, and denote by i^u the IXP at location IXP^u, i.e. $i^u = i|(IXP^u = IXP(i))$. The distance between two IXPs can be determined by their locations given in (1):

$$D^b(IXP^p_{down}, IXP^u) = \frac{L}{N}|i^p - i^u| \tag{6}$$

The distance 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 the end user (U) resides. We denote the distance on the ISP's middle mile network between these two locations by $D^m(IXP^u, U) = \|IXP^u - A(j,k)\|$, where $U \in R(IXP^u)$ and $(j,k)\,|\,(U \in Access(j,k))$. The distance can be determined by the locations of the IXP and the access network, given in (1)-(2):

$$D^m(IXP^u, U) = \sqrt{\left(\frac{L(2i^u-1)}{2N} - \frac{(2j-1)a}{2}\right)^2 + \left(-\frac{W}{2} + \frac{(2k-1)b}{2}\right)^2} \tag{7}$$

where $U \in R(IXP^u)$ and $(j,k)\,|\,(U \in Access(j,k))$.

The distance on the ISP's access network is a function of the location of each end user. We denote the distance on the ISP's access network by $D^a(U) = \|A(j,k) - U\|$, where $(j,k)\,|\,(U \in Access(j,k))$. The distance can be determined by the location of the end user 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 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). Thus, the average distance on the ISP's backbone network is:

$$ED^{b,hot}_{down} = \sum_{g\in I}\sum_{g'\in I} D^b(g,g')P(IXP^{p,hot}_{down} = g)P(IXP^u = g') \tag{8}$$

The distance $D^b(g,g')$ is given in closed form in (6), the probability distribution of $IXP^{p,hot}_{down}$ is given in (5), and the probability distribution of IXP^u is similarly uniformly distributed. We can use these results to give a closed-form expression:

Theorem 1:

$$ED^{b,hot}_{down} = \frac{L(N-1)(N+1)}{3N^2} \tag{9}$$

The proof can be found in the Appendix.

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, i.e. $IXP^{p,cold}_{down} = IXP^u$. Thus, the ISP does not carry traffic across its backbone, i.e.

$$ED^{b,cold}_{down} = 0 \tag{10}$$

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 is:

$$ED^m = \sum_{g'\in I}\sum_{A(j,k)\subset R(g')} D^m(g', A(j,k))P(j,k) \tag{11}$$

The distance $D^m(g', A(j,k))$ is given in closed form in (7). Also, $P(j,k) = ab/LW$. We can use these results to give a closed-form expression:

Theorem 2:

$$ED^m = \frac{abN}{LW}\sum_{k=1}^{\frac{W}{b}}\sum_{j=1}^{\frac{L}{aN}}\sqrt{\left(\frac{(2j-1)a-\frac{L}{N}}{2}\right)^2 + \left(\frac{(2k-1)b-W}{2}\right)^2} \tag{12}$$

The proof can be found in the Appendix.

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 between access networks and also within each access network, the average distance is:

$$ED^a = \frac{1}{ab}\int_{-\frac{a}{2}}^{\frac{a}{2}}\int_{-\frac{b}{2}}^{\frac{b}{2}}\sqrt{x^2+y^2}\,dy\,dx \tag{13}$$

It can be shown that:

Theorem 3:

$$ED^a = \frac{1}{12ab}\left[a^3\sinh^{-1}(\frac{b}{a}) + b^3\sinh^{-1}(\frac{a}{b}) + 2ab\sqrt{a^2+b^2}\right] \tag{14}$$

The proof can be found in the Appendix.

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.[1]

Traffic-sensitive costs are a function of both distance and traffic volume. We assume here that traffic-sensitive costs

[1] There is a small cost for each interconnection point; however, this cost is relatively small compared to transportation 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., [12]. 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 in (14) is constant. Thus, the ISP's traffic-sensitive access network cost is similarly constant. The variable portion of the ISP's traffic-sensitive cost is thus: $C = c^b\left(ED^b + \frac{c^m}{c^b} ED^m\right) V$.

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 $c^m/c^b = 1.5$. (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 and routing policies all affect ED^b and/or ED^m.

V. Number of IXPs

In this section, we examine the effect of the number of IXPs at which two ISPs agree to peer. As shown in Table I, 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. The academic literature provides little insight into why large ISPs require interconnection at a minimum specified number of IXPs.

We first assume that both ISPs use hot potato routing. The variable traffic-sensitive cost of downstream traffic is:

$$C_{down}^{hot} = c^b V_{down}\left(ED_{down}^{b,hot} + \frac{c^m}{c^b} ED^m\right) \tag{15}$$

Substituting the expressions we previously found for $ED_{down}^{b,hot}$ in (9) and for ED^m in (12), we obtain:

$$C_{down}^{hot} = c^b V_{down}\Bigg(\frac{L(N-1)(N+1)}{3N^2} + \frac{c^m}{c^b}\frac{abN}{LW} \sum_{m=1}^{\frac{W}{b}}\sum_{n=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2n-1)a - L/N}{2}\right)^2 + \left(\frac{(2m-1)b - W}{2}\right)^2}\Bigg) \tag{16}$$

Figure 2 shows the effect of the number of interconnection points (N) on the cost of downstream traffic (C_{down}^{hot}). (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 both the backbone cost and the middle mile cost.

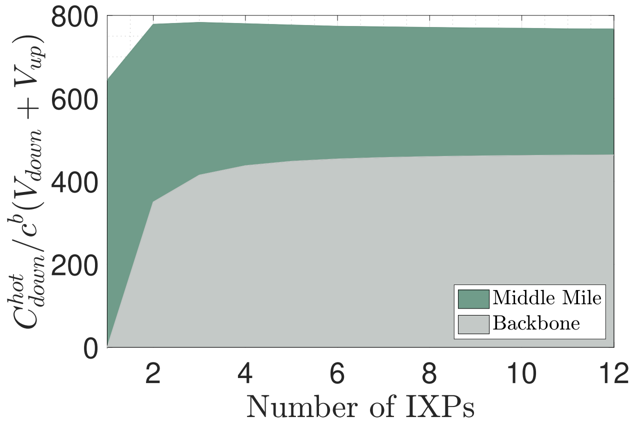

Fig. 2: Downstream costs

The average distance the ISP carries traffic across its backbone ($ED_{down}^{b,hot}$) is proportional to $1 - \frac{1}{N^2}$. Thus, the backbone cost is increasing and concave with the number of interconnection points. A larger number of interconnection points results in a larger distance between the IXP closest to the west coast and the IXP closest to the east coast. As a result, the backbone expands and the ISP carries traffic across a longer distance on this larger backbone.

The average distance the ISP carries traffic across its middle mile networks (ED^m) is a complicated function of N. However, as we see in Figure 2, the middle mile cost is decreasing and convex with the number of interconnection points. A larger number of interconnection points results in more closely spaced IXPs that are closer to the access networks. As a result, the middle mile networks shrink as the backbone expands.

Increasing the number of interconnection points increases backbone cost and decreases middle mile cost. Therefore, the cost of downstream traffic is a uni-modal function of N.

However, when two ISPs peer, there is also upstream traffic. The cost of upstream traffic using hot potato routing is:

$$C_{up}^{hot} = c^b V_{up}\left(ED_{up}^{b,hot} + \frac{c^m}{c^b} ED^m\right) \tag{17}$$

As we discussed in Section III-B, 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 traffic across its middle mile networks (ED^m) is the average distance between the center of the access network and the nearest IXP, which is the same for downstream and upstream traffic. In addition, 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. Thus, from (10), we know that:

$$ED_{up}^{b,hot} = ED_{down}^{b,cold} = 0 \tag{18}$$

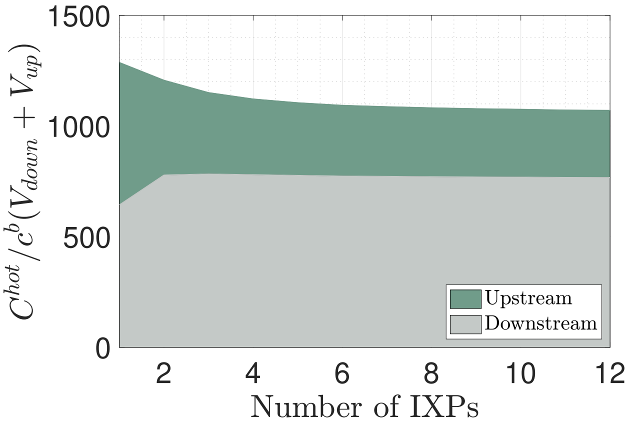

Fig. 3: Variable traffic-sensitive costs

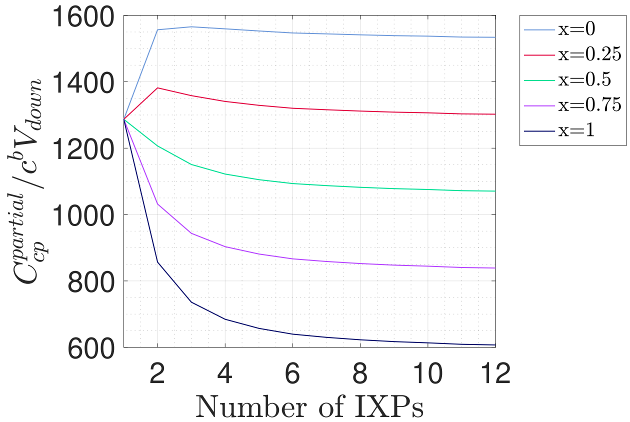

Fig. 4: Costs under partial replication

The average distance the ISP carries upstream traffic across the backbone when using hot potato routing is zero because the ISP is exchanging this upstream traffic at the nearest IXP to the access network.

Substituting the expression we previously found for ED^m in (12), we obtain:

$$C_{up}^{hot} = c^m V_{up} \frac{abN}{LW} \sum_{m=1}^{\frac{W}{b}} \sum_{n=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2n-1)a - L/N}{2}\right)^2 + \left(\frac{(2m-1)b - W}{2}\right)^2} \tag{19}$$

Figure 3 shows the effect of the number of interconnection points (N) on the cost of both downstream (C_{down}^{hot}) and upstream (C_{up}^{hot}) traffic, when there is an equal amount of downstream and upstream traffic (i.e., $V_{down} = V_{up}$). We observe that the variable traffic-sensitive cost is decreasing and convex with the number of interconnection points. We also observe that there is less than a 2% difference in the cost between $N = 8$ and $N = 12$, so this indicates there may be little value in requiring interconnection at more than 8 IXPs.

VI. Peering between a content provider and an ISP

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.

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 III-A, and the distribution of the location of end users remains the same as was presented in Section III-B. 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.

However, 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 IV, but only for downstream traffic

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

$$C_{cp}^{partial} = c^b V_{down} \left(x ED_{down}^{b,cold} + (1-x) ED_{down}^{b,hot} + \frac{c^m}{c^b} ED^m \right) \tag{20}$$

where $ED_{down}^{b,hot}$, ED^m, and $ED_{down}^{b,cold}$ are given in (9), (12), and (18) respectively.

Figure 4 shows the effect of the number of interconnection points (N) on the ISP's downstream cost, 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. The cost of the content delivered using hot potato routing dominates the ISP's downstream cost, and thus it is rational for the ISP to not agree to settlement-free peering. However, when $x > 0.3$, the cost of the locally-delivered content dominates the ISP's downstream cost, and thus the ISP benefits from increasing the number of IXPs at which they agree to peer.

Therefore, when a content provider interconnects with an ISP, 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.

VII. 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 decreasing and convex with the number of interconnection points. There may be little value in requiring interconnection at more than 8 IXPs.

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. Thus, we would expect large ISPs to have *different* settlement-free peering requirements for such content providers than for ISPs.

References

[1] Leichtman Research Group, Inc. (LRG), "About 2,950,000 Added Broadband From Top Providers In 2021," Accessed March 15, 2022.

[2] PeeringDB, "The Interconnection Database," Accessed February 6, 2022.

[3] 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.

[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] 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.

[6] 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*, pp. 1–12, 2006.

[7] 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.

[8] "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.

[9] "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.

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

[11] Maps of World, "How Wide Is The United States?," Accessed February 22, 2022.

[12] 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)*, pp. 194–205, 2011.

Appendix

A. Proof of Theorem 1

In (8), we have

$$ED^{b,hot}_{down} = \sum_{g\in I}\sum_{g'\in I} D^b(g,g')P(IXP^{p,hot}_{down}=g)P(IXP^u=g'). \quad \text{(A.21)}$$

The distance $D^b(g,g')$ is given by (6). The probability distribution of $IXP^{p,hot}_{down}$ is given in (5), and the probability distribution of IXP^u is similarly uniformly distributed. Substituting these expressions into (A.21),

$$\begin{aligned} ED^{b,hot}_{down} &= \sum_{i^g\in l^N}\sum_{i^{g'}\in l^N}\frac{L}{N}|i^g-i^{g'}|\frac{1}{N}\frac{1}{N} \\ &= \frac{1}{N^2}\sum_{i^g\in l^N}\sum_{i^{g'}\in l^N}\frac{L}{N}|i^g-i^{g'}|. \end{aligned} \quad \text{(A.22)}$$

The inner sum is the average distance from interconnection point i^g to other interconnection points. If i^g is the k^{th} interconnection point (from left to right), then the inner sum

$$\begin{aligned} &\sum_{i^{g'}\in l^N}\frac{L}{N}|i^g-i^{g'}| \\ &= \frac{L}{N}\{[|1-k|+...+|-1|]+0+[|1|+...+|N-k|]\} \\ &= \frac{L}{N}\left[\frac{k(k-1)}{2}+\frac{(N-k)(N-k+1)}{2}\right]. \end{aligned} \quad \text{(A.23)}$$

Substituting (A.23) into (A.22),

$$\begin{aligned} ED^{b,hot}_{down} &= \frac{1}{N^2}\sum_{k=1}^{N}\frac{L}{N}\left[\frac{k(k-1)}{2}+\frac{(N-k)(N-k+1)}{2}\right] \\ &= \frac{L}{N^3}\left[\frac{1}{2}(N+1)N^2+\sum_{k=1}^{N}k^2-(N+1)\sum_{k=1}^{N}k\right] \\ &= \frac{L}{N^3}\left[\frac{1}{2}(N+1)N^2+\frac{N(N+1)(2N+1)}{6}-\frac{N(N+1)^2}{2}\right] \\ &= \frac{L(N-1)(N+1)}{3N^2}. \end{aligned} \quad \text{(A.24)}$$

B. Proof of Theorem 2

In (11), we have:

$$ED^m = \sum_{g'\in I}\sum_{A(j,k)\subset R(g')} D^m(g',A(j,k))P(j,k). \quad \text{(A.25)}$$

The access networks $A(j,k)\subset R(g')$ are given by

$$\begin{aligned} \frac{(i^{g'}-1)L}{aN}+1 \leq &j \leq \frac{i^{g'}L}{aN} \\ 1 \leq &k \leq \frac{W}{b} \end{aligned} \quad \text{(A.26)}$$

The double sum in (A.25) can thus be written as

$$ED^m = \sum_{i^{g'}=1}^{N}\sum_{k=1}^{\frac{W}{b}}\sum_{j=\frac{(i^{g'}-1)L}{aN}+1}^{\frac{i^{g'}L}{aN}} D^m(g',A(j,k))P(j,k). \quad \text{(A.27)}$$

The distance $D^m(g',A(j,k))$ is given by (7):

$$D^m(g',A(j,k)) = \sqrt{\left(\frac{(2j-1)a-L/N(2i^{g'}-1)}{2}\right)^2+\left(\frac{(2k-1)b-W}{2}\right)^2} \quad \text{(A.28)}$$

The variable substitution $j' = j - \frac{(i^{g'}-1)L}{aN}$ will help simplify the equation:

$$D^m = \sqrt{\left(\frac{(2j'-1)a - L/N}{2}\right)^2 + \left(\frac{(2k-1)b - W}{2}\right)^2} \tag{A.29}$$

Changing the inner sum over j into a sum over j' and substituting $P(j,k) = ab/LW$,

$$ED^m = \frac{ab}{LW} \sum_{i^{g'}=1}^{N} \sum_{k=1}^{\frac{W}{b}} \sum_{j'=1}^{\frac{L}{aN}} D^m \tag{A.30}$$

D^m is no longer a function of $i^{g'}$, and thus we can remove the outer sum, resulting in

$$ED^m = \frac{abN}{LW} \sum_{k=1}^{\frac{W}{b}} \sum_{j'=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2j'-1)a - L/N}{2}\right)^2 + \left(\frac{(2k-1)b - W}{2}\right)^2}. \tag{A.31}$$

C. *Proof of Theorem 3*

In (13), we have

$$ED^a = \frac{1}{ab} \int_{-\frac{a}{2}}^{\frac{a}{2}} \int_{-\frac{b}{2}}^{\frac{b}{2}} \sqrt{x^2+y^2}\, dy\, dx. \tag{A.32}$$

By symmetry,

$$ED^a = \frac{4}{ab} \int_{0}^{\frac{a}{2}} \int_{0}^{\frac{b}{2}} \sqrt{x^2+y^2}\, dy\, dx. \tag{A.33}$$

We first partition the area of integration into regions below and above the 45-degree line, resulting in

$$\int_{0}^{\frac{a}{2}} \int_{0}^{\frac{b}{2}} \sqrt{x^2+y^2}\, dy\, dx = \int_{0}^{\frac{b}{2}} \int_{0}^{\frac{ay}{b}} \sqrt{x^2+y^2}\, dx\, dy + \int_{0}^{\frac{a}{2}} \int_{0}^{\frac{bx}{a}} \sqrt{x^2+y^2}\, dy\, dx \tag{A.34}$$

Converting the integral into polar coordinates by substituting $x = r\cos\theta$ and $y = r\sin\theta$, the previous expression can be written as

$$\begin{aligned}
&\int_{0}^{\tan^{-1}\frac{a}{b}} \int_{0}^{\frac{b}{2}\sec\theta} r^2\ dr\ d\theta + \int_{0}^{\tan^{-1}\frac{b}{a}} \int_{0}^{\frac{a}{2}\sec\theta} r^2\ dr\ d\theta \\
&= \frac{b^3}{24} \int_{0}^{\tan^{-1}\frac{a}{b}} \sec^3\theta\ d\theta + \frac{a^3}{24} \int_{0}^{\tan^{-1}\frac{b}{a}} \sec^3\theta\ d\theta \\
&= \frac{b^3}{48}\left[\frac{2a}{b^2}\sqrt{(\frac{a}{2})^2 + (\frac{b}{2})^2} + \sinh^{-1}(\frac{a}{b})\right] \\
&+ \frac{a^3}{48}\left[\frac{2b}{a^2}\sqrt{(\frac{a}{2})^2 + (\frac{b}{2})^2} + \sinh^{-1}(\frac{b}{a})\right] \\
&= \frac{a^3\sinh^{-1}(\frac{b}{a}) + b^3\sinh^{-1}(\frac{a}{b}) + 2ab\sqrt{a^2+b^2}}{48}.
\end{aligned} \tag{A.35}$$

Substituting this expression into (A.33),

$$ED^a = \frac{a^3\sinh^{-1}(\frac{b}{a}) + b^3\sinh^{-1}(\frac{a}{b}) + 2ab\sqrt{a^2+b^2}}{12ab}. \tag{A.36}$$