When is Regulation of Peering or Usage Fees Warranted?

Ali Nikkhah and Scott Jordan
Department of Computer Science, University of California, Irvine, USA
Emails:{ali.nikkhah, sjordan}@uci.edu

Abstract

Debates over paid peering and usage fees have expanded from the United States to Europe and South Korea. ISPs argue that content providers should pay fees based on the amount of downstream traffic they generate. In contrast, content providers contend that customers already pay ISPs for delivering the content they request, and therefore that peering agreements should be settlement-free. The issue has arisen in debates in the United States, Europe, and South Korea over net neutrality, universal service, and infrastructure funding. Regulatory entities are considering whether to regulate peering prices and/or impose usage fees.

A key part of the debate concerns whether the market determines the socially beneficial peering price, and if not, how much of a difference there is between the socially beneficial peering price and the market-determined peering price. Our objective here is to understand the range from a cost-based peering price to a profit-maximizing peering price.

First, we determine an ISP's cost for directly peering with a content provider, by analyzing the incremental cost for transporting the content provider's traffic when it directly peers with the ISP versus when it sends its traffic through a transit provider. We find that this cost-based peering price is positive if there is little localization of content, but it is zero (i.e., settlement-free peering) if there is sufficient localization of content. We also find that the required amount of localization varies with the number of interconnection points.

Next, we determine the peering price that maximizes an ISP's profit using a two-sided market model in which a profit-maximizing ISP determines broadband prices and the peering price, and in which content providers determine their service prices based on the peering price. We find that the profit-maximizing peering price decreases with content localization.

These prices establish a range if the peering price is unregulated, from the cost-based peering price (at the low end) to the profit-maximizing peering price (at the high end). Regulatory oversight of peering prices may be warranted when there is a substantial difference between cost-based and profit-maximizing prices.

Keywords:
Usage fees, Broadband, Telecommunications Policy, Net Neutrality, Two-sided Market Model, Interconnection Agreements, Paid Peering, Peering Policies

1. Introduction

An Internet Service Provider (ISP) enables the transmission and receipt of data to and from all or almost all Internet endpoints. To offer this Internet access service, the ISP must establish connections with other networks to exchange data. An interconnection agreement is considered a *transit service* if the transit provider agrees to accept and deliver data on behalf of the ISP, regardless of the destination. On the other hand, if each network agrees to only accept and deliver data with destinations in its customer base, the interconnection agreement is referred to as *peering*. We focus on peering in this paper. Peering may be either paid (i.e., one interconnecting network pays the other) or settlement-free (i.e., without payment).

It is no longer clear who should pay whom and how much for interconnection between ISPs and content providers. Large ISPs claim that large content providers are imposing a cost on ISPs by sending large amounts of traffic to their customers. ISPs claim that it is more fair that content providers pay for this cost than consumers, because then this cost will be paid only by those consumers with high usage. In contrast, large content providers (including CDNs) claim that when they interconnect with ISPs at interconnection points (IXPs) close to consumers, they are already covering the costs of carrying traffic through the core network, and that consumers are already covering the costs of carrying traffic through the ISP's access network. These disputes between large ISPs and large content providers have recurred often during the last 10 years. When not resolved, large ISPs have often refused to increase capacity at interconnection points with large content providers and transit providers, resulting in sustained congestion which has degraded users' quality of experience because of reduced throughput, increased packet loss, increased delay, and increased jitter.

As a result, 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 the United States, between 2013 and 2014, a disagreement between Comcast and Netflix regarding interconnection terms persisted for an extended period. In 2014, Netflix and some transit providers brought the matter to the Federal Communications Commission (FCC) of the United States, which was drafting revised net neutrality regulations at the time. The debate shows contrasting views. Some large content providers and some transit 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. In response, large ISPs argue 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 large ISPs asserted that if they absorb these costs, then the ISPs would recoup these costs by increasing the prices for all subscribers, which is unfair to subscribers who do not use the services, like Netflix, that are driving the need for additional capacity.

The FCC addressed the debate surrounding interconnection arrangements in its 2015 Open Internet Order [1], and asserted oversight over interconnection arrangements. Then, in 2018, the FCC reversed its stance, as part of repealing most of the 2015 net neutrality regulations [2], ending this oversight. However, it is highly likely that the FCC will revisit this issue in the near future as part of a new net neutrality proceeding. In addition, the FCC recently reported that some stakeholders are proposing that content providers pay a fee based on their download traffic to subsidize broadband Internet access service in rural areas and for low-income consumers [3]. These advocates are using similar arguments that large ISPs have used to advocate for paid peering.

Similar debates over paid peering are also active in South Korea and in Europe. In South Korea, paid peering between ISPs is now mandatory, based on the amount of traffic exchanged. As a result, these peering fees are often passed on to content providers that interconnect with ISPs in South Korea. A proposal is currently under consideration in South Korea to also require content providers to pay usage fees to ISPs, based on traffic volume [4]. The European trade association representing numerous ISPs in Europe has recently put forward a similar proposal, suggesting that content providers should pay usage fees to ISPs, based on the volume of traffic [5]. However, European regulators are concerned that such fees could be abused by ISPs and are skeptical of the argument that ISPs' costs are not adequately covered by their customers [6].

In this paper, we address this debate over paid peering fees. We determine an ISP's cost for directly peering with a content provider. Such a cost-based peering price may be the minimum price an ISP will accept. We also identify the peering price that maximizes an ISP's profit using a two-sided market model.

Unregulated, these prices establish a range from the cost-based peering price to the profit-maximizing peering price. Regulatory oversight of peering prices may be warranted when there is a substantial difference between cost-based and profit-maximizing prices. In particular, our goal is to determine the effect of content localization and the number of interconnection points on this range of peering prices.

The paper is organized as follows. In Section 2, we develop a simplified model of backbone transportation costs. We partition an ISP's network into access networks, middle mile networks, and a backbone network. We assume that an ISP and a transit provider or content provider mutually determine a set of points at which to interconnect, chosen from a list of the largest traffic exchange locations in the United States. In order to determine the routes over which traffic flows between networks, we construct traffic matrices, using United States census statistics to determine broadband subscriber locations. We determine the distances on an ISP's backbone network over which it carries traffic to and from a subscriber, and we calculate the average distance using the traffic matrices. We then construct a simplified model of backbone transportation costs as a function of both distance and traffic volume.

In Section 3, we turn to the case of an ISP interconnecting with a transit provider. The transit provider may deliver a portion of the video traffic using cold potato routing, which localizes traffic on the ISP's network and reduces the ISP's backbone transportation costs. We model both the ISP's and the transit provider's backbone transportation costs, as a function of the number of interconnection points and the traffic ratio between the two, routing, and localization. We then consider fair cost sharing between an ISP and a transit provider. We derive the peering fee that equalizes the ISP's and the transit provider's net costs, and show how it depends not only on the traffic ratio, but also on the amount of video traffic and the amount of video content localization.

In Section 4, we determine an ISP's cost for directly peering with a directly interconnected content provider. Such a cost-based peering price may be the minimum price an ISP will accept. To do so, we analyze the incremental cost for transporting the content provider's traffic when a content provider directly peers with the ISP versus when it sends its traffic through a transit provider. We set the cost-based peering fee such that the sum of the ISP's backbone transportation costs and the peering fee is the same when the content provider is directly or indirectly connected. We show that this cost-based peering fee depends on the localization of video traffic and the number of interconnection points. If the content provider does not localize enough video traffic, the cost-based peering fee is positive. As the content provider localizes more video traffic, the cost-based peering fee should decrease.

In Section 5, we propose a model of user subscription to broadband service tiers and to subscription video on demand (SVOD) services. We consider a monopoly ISP that offers a basic tier and a premium tier differentiated by download speed, both with data caps, and an add-on that allows unlimited usage. We aggregate all SVOD providers that directly interconnect with the ISP. Consumers differ in the utilities they place on broadband service tiers and on SVOD service, and each customer chooses the service which maximizes his/her surplus. We derive the demand of each broadband service tier and of SVOD services, the associated ISP profit, and aggregated SVOD providers profit. We also show the effect of SVOD traffic localization and number of interconnection points on the monthly marginal costs of the ISP and SVOD providers.

In section 6, we determine the peering fee that maximizes an ISP's profit. The results indicate that a profit-maximizing ISP may charge SVOD providers the highest amount their willingness to pay permits. Our research reveals that as SVOD traffic localization increases, the peering price that ISPs can demand decreases due to a reduced maximum amount that SVOD providers are willing to pay for peering. Also, our results demonstrate with a low level of localization, the peering fee charged by the profit-maximizing ISP increases as the number of interconnection points rises. However, for a high level of localization, the peering fee charged by the profit-maximizing ISP declines as the number of interconnection points increases. Finally, we compare the cost-based peering fee with the profit-maximizing peering fee. We find that the ISP's profit-maximizing peering fee typically exceeds the cost-based peering fee, and this gap may diminish with increasing traffic localization.

2. Backbone Cost Model

In this section, we develop a simplified model of backbone transportation costs in the United States. This model was presented in [7], but it is summarized here for convenience.

2.1. Topology

Our model of the topology of an ISP's network consists of the ISP's service territory, the location of Internet Exchange Points (IXPs), and the segments of its network. We consider an ISP whose service territory covers the contiguous United States. We represent this service territory using the set of longitudes and latitudes of the contiguous United States.

We consider M IXPs at which an ISP may agree to peer with a transit provider or a content provider. In the numerical results below, we use the $M = 12$ largest IXPs in the United States.[1] Denote the coordinates (in longitude and latitude) of these M IXPs by $IXP(i)$, and the set of the locations of these IXPs by I^M. An ISP and a content provider often agree to interconnect at a smaller number $N < M$ of IXPs. Denote the set of N IXPs at which they agree to interconnect by $l^N \subseteq \{1, ..., M\}$, and the set of locations of these IXPs by $I^N = \{IXP(i), i \in l^N\} \subseteq I^M$.

We model an ISP's network as partitioned into a single backbone network, multiple middle mile networks, and multiple access networks. We model each access network as spanning a single U.S. county. Denote the geographical region of access network j by $Access(j)$, and the location of the geographical center of access network j by $A(j)$. We assign these locations using the longitudes and latitudes of the center of each county in the contiguous United States.

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

Figure 1 illustrates these regions when an ISP and an interconnecting network agree to interconnect at all 12 IXPs.

2.2. 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 distribution of the number of end users in each access network and, for each access network, the distribution of end users within the access network.

Denote the probability that an end user resides within access network j by $P(j)$. We assume that end users are distributed across access networks according to the population of the county associated with the access network, denoted by p_j, and we denote the population of the contiguous United States by $p = \sum_j p(j)$. These populations are taken from U.S. census data. 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.

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

[1] These are located at Ashburn, Chicago, Dallas, San Jose, Los Angeles, New York, Seattle, Miami, Atlanta, Denver, Boston, and Minneapolis.

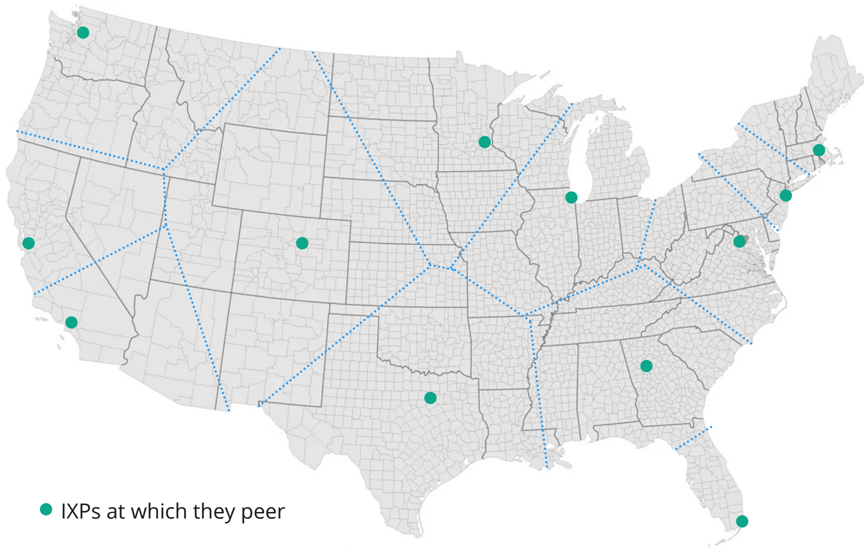

Figure 1: Topology of an ISP's network

Along the route S to U, denote the location of the IXP at which downstream traffic using hot potato routing enters the ISP's network by IXP^{down}_{hot}, the location of the IXP at which downstream traffic using cold potato routing enters the ISP's network by IXP^{down}_{cold}, and the location of the IXP closest to the end user by IXP^{u}.

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

The access network on which U resides is distributed according to $\{P(j)\}$, and U is uniformly distributed within the access network. The IXP closest to the end user is a deterministic function of U, namely $IXP^{u} = (g' | U \in R^{M}(g'))$.

The IXP at which downstream traffic using hot potato routing enters the ISP's network (IXP^{down}_{hot}) is independent of the end user, and it is the IXP closest to the source among the IXPs at which they agree to peer, i.e. $IXP^{down}_{hot} = (g | S \in R^{N}(g))$. Since end users are assumed to be distributed according to U.S. county population statistics, the distribution of IXP^{down}_{hot} 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^{down}_{cold}) is no longer independent of the end user, and it is the IXP closest to the end user at which they agree to peer, i.e., $IXP^{down}_{cold} = (g | U \in R^{N}(g))$.

For 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, i.e. $IXP^{up}_{hot} = IXP^{down}_{cold} = (g | U \in R^{N}(g))$. 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^{down}_{hot}.

2.3. Traffic-sensitive backbone costs

We now construct a simplified model of backbone transportation costs. These costs are a function of the average distance carried over a network's backbone, which in turn depends on routing and the traffic matrices.

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^{down}_{hot}) 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 $D^b(IXP^{down}_{hot}, IXP^u) = \|IXP^{down}_{hot} - IXP^u\|$. We separately consider downstream traffic using hot potato routing and downstream traffic using cold potato routing.

For downstream traffic using hot potato routing, the IXP at which the traffic enters the ISP's network (IXP^{down}_{hot}) depends on the IXPs at which they agree to interconnect. However, it is independent of the end user and thus independent of the IXP closest to the end user (IXP^u). Consider an ISP and an interconnecting network that agree to interconnect at the N IXPs in I^N.

The average distance of the traffic on the ISP's backbone network (ED^{hot}) is the average distance between the IXPs at which downstream traffic enter the ISP's network (IXP^{down}_{hot}) 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.

For downstream traffic using cold potato routing, the IXP at which the traffic enters the ISP's network is the IXP closest to the end user at which they agree to peer, i.e. IXP^{down}_{cold}. The ISP might still carry traffic across a portion of its backbone, namely from IXP^u to IXP^{down}_{cold}, and the average such distance (ED^{cold}) is the average distance between the IXPs closest to the end user at which they agree to peer and the IXPs closest to the end user.

We are concerned only with the portion of an ISP's backbone transportation costs that is sensitive to the amount of traffic, because non-traffic-sensitive costs do not vary significantly with the number of IXPs or the traffic ratio. Traffic-sensitive costs are a function of both distance and traffic volume. We model traffic-sensitive costs as linearly proportional to the average distance over which the traffic is carried on each portion of the ISP's network [8], and linearly proportional to the average volume of traffic that an ISP carries on each portion of its network.

Denote the cost per unit distance and per unit volume in the backbone network by c^b. Denote the volume of traffic by V. The ISP's traffic-sensitive backbone cost is thus Vc^bED^{hot} for downstream traffic using hot potato routing and Vc^bED^{cold} for downstream traffic using cold potato routing. The ISP's traffic-sensitive backbone cost for upstream traffic can be similarly calculated.

3. Peering Between a Transit Provider and an ISP

We now turn to peering between a transit provider and an ISP, i.e., neither is a customer of the other. In this section, we consider the traffic-sensitive backbone costs of each. In the following section, we determine the fair peering fee. An extended version of these two sections is available in [9].

Transit providers that sell transit service to content providers often promise the content provider that they will deliver the traffic to the terminating ISP using cold potato routing. The use of cold potato routing allows the transit provider to exercise greater control over the management of this traffic and thereby to potentially improve its Quality of Service (QoS). Most of this content consists of video, and in the remainder of the paper we use the term *video* to refer to it.

We partition traffic exchanged between a transit provider and an ISP into video traffic and non-video traffic. We assume that non-video traffic is transported using hot potato routing. We assume that a proportion x of video traffic transmitted to the ISP's users within each access network is delivered from the transit provider using cold potato routing. We assume that the transit provider and the ISP interconnect at all $M = 12$ IXPs, and consequently that video traffic delivered using cold potato routing is exchanged at the IXP closest to the end user. We assume that the remaining proportion $1 - x$ of video traffic is delivered using hot potato routing and that the source of this video traffic is independent of the location of the end user.

Our focus here is on traffic-sensitive costs, since they are the focus of the debate on paid peering. To analyze fair peering fees, we must define *fair* and decide if traffic-sensitive costs should be recovered from subscribers or from interconnecting networks. Economists debate how to divide costs in this two-sided market, but generally agree that subscribers should cover access and middle-mile network costs. There are several rationales for this approach. First, regulatory cost accounting often dictates that access network costs be recovered from subscribers. Second, the conditions under which two networks peer affect ISPs' backbone transportation costs. However, these same conditions do not affect ISPs' middle-mile or access network transportation costs, since an ISP must carry traffic across these portions of its network regardless.

The debate is thus over how much of the backbone network costs should be paid by subscribers versus by interconnecting networks. In the remainder of the paper, we thus focus on traffic-sensitive backbone costs.

3.1. ISP Backbone Cost

We consider downstream traffic destined for an end user located in the ISP's network. We denote the volume of non-video downstream traffic by V_d and the volume of video downstream traffic by V_v. We also consider upstream traffic originating with an end user located in the ISP's network, and denote the volume of this traffic by V_u.

We define two traffic ratios: $r = \frac{V_d}{V_u}$, the ratio of downstream non-video traffic to upstream traffic, and $r' = \frac{V_v}{V_u}$, the ratio of downstream video traffic to upstream traffic.

We denote the ISP's traffic-sensitive backbone cost by C^{ISP}, and partition it into the cost of delivering downstream non-video traffic (denoted by $C^{ISP}_{down,non-video}$), the cost of delivering downstream video traffic (denoted by $C^{ISP}_{down,video}$), and the cost of delivering upstream traffic (denoted by C^{ISP}_{up}):

$$C^{ISP} = C^{ISP}_{down,non-video} + C^{ISP}_{down,video} + C^{ISP}_{up} \quad (1)$$

The cost of delivering downstream non-video traffic using hot potato routing is:

$$C^{ISP}_{down,non-video} = c^b V_d ED^{hot}(M) \tag{2}$$

The cost of delivering downstream video traffic is:

$$C^{ISP}_{down,video} = c^b V_v \Big[xED^{cold}(M) + (1-x)ED^{hot}(M) \Big] \tag{3}$$

The first term is the ISP's backbone cost for localized video traffic, which the transit provider delivers using cold potato routing. The second term is the ISP's backbone cost for non-localized video traffic, which the transit provider delivers using hot potato routing.

The cost of delivering upstream traffic using hot potato routing is:

$$C^{ISP}_{up} = c^b V_u ED^{cold}(M), \tag{4}$$

because the average distance on the ISP's backbone network of upstream traffic with hot potato routing, when interconnecting at $M = 12$ IXPs, is the same as the average distance on the ISP's backbone network of downstream traffic with cold potato routing (i.e., $ED^{cold}(M)$).

Using the definition of the two traffic ratios r and r', and the fact that $ED^{cold}(M) = 0$, equations (1)-(4) can be simplified:

Theorem 1. *The traffic-sensitive backbone cost of the ISP when peering with the transit provider is:*

$$C^{ISP} = c^b V_u \Big[r + r'(1-x) \Big] ED^{hot}(M) \tag{5}$$

A portion of the ISP's backbone cost is to transport downstream non-video traffic over the ISP's backbone, as measured by the volume V_d of such traffic. Another portion of the ISP's backbone cost is to transport downstream non-localized video traffic over the ISP's backbone, as measured by the volume $V_v(1-x)$ of such traffic.

3.2. Transit Provider Backbone Cost

We now turn to the effect of routing policies, traffic ratios, and traffic localization on the traffic-sensitive backbone cost of the transit provider. Note that downstream traffic from the point of view of the transit provider is equal to the upstream traffic from the point of view of the ISP (V_u). Similarly, the transit provider non-video upstream traffic is equal to ISP non-video downstream traffic (V_d), and the transit provider video upstream traffic is equal to ISP video downstream traffic (V_v).

We denote the transit provider's traffic-sensitive backbone cost by C^{TP}, and partition it into the transit provider's downstream cost for delivering ISP upstream traffic (denoted by C^{TP}_{down}), its upstream cost for delivering ISP downstream non-video traffic (denoted by $C^{TP}_{up,non-video}$), and its upstream cost for delivering ISP downstream video traffic (denoted by $C^{TP}_{up,video}$):

$$C^{TP} = C^{TP}_{down} + C^{TP}_{up,non-video} + C^{TP}_{up,video} \tag{6}$$

The cost to the transit provider for delivering ISP upstream traffic using hot potato routing is given by $C^{TP}_{down} = c^b V_u ED^{hot}(M)$ and the cost to the transit provider for delivering ISP downstream non-video traffic using hot potato routing is given by $C^{TP}_{up,non-video} = c^b V_d ED^{cold}(M)$. The cost to the transit provider for delivering ISP downstream video traffic is:

$$C^{TP}_{up,video} = c^b V_v \Big[xED^{hot}(M) + (1-x)ED^{cold}(M) \Big] \tag{7}$$

The first term is the transit provider's backbone cost for localized video traffic, which the transit provider delivers using cold potato routing. The second term is the transit provider's backbone cost for non-localized video traffic, which the transit provider delivers using hot potato routing.

Using the two traffic ratios r and r', and that $ED^{cold}(M) = 0$, equations (6)-(7) can be simplified:

Theorem 2. *The traffic-sensitive backbone cost of the transit provider when peering with the ISP is:*

$$C^{TP} = c^b V_u (1 + r'x) ED^{hot}(M) \tag{8}$$

A portion of the transit provider's backbone cost is caused by the need to transport the ISP's upstream traffic over the transit provider's backbone, as measured by the volume V_u of such traffic. Another portion of the transit provider's backbone cost is caused by the need to transport ISP downstream localized video traffic over the transit provider's backbone, as measured by the volume $V_v x$ of such traffic.

3.3. A Fair Peering Fee Between a Transit Provider and an ISP

In this section, we determine the fair peering fee between a transit provider and an ISP. We define *fair* as the peering fee that equalizes the net costs of the transit provider and the ISP. Our objective is thus to derive the peering fee that ensures both parties incur the same costs for transmitting data over their backbones. To achieve this goal, we use the analysis of the cost structures of ISPs and transit providers provided in the previous section, and calculate the payment required to ensure that the costs are fairly split between the two parties. The peering fee that equalizes net costs is given by:

$$P^{TP,ISP} = \frac{1}{2}(C^{ISP} - C^{TP}) \tag{9}$$

If the ISP's traffic-sensitive backbone costs exceed those of the transit provider, then the fair peering fee is positive. If the transit provider's traffic-sensitive backbone costs exceed those of the ISP, then the fair peering fee is negative. By using Theorems 1 and 2, the fair peering fee can be expressed as:

$$P^{TP,ISP} = \frac{1}{2} c^b V_u \Big[\Big(r - 1\Big) + \Big(r'(1-x) - r'x\Big) \Big] ED^{hot}(M) \tag{10}$$

The term $\frac{1}{2} c^b V_u (r-1) ED^{hot}(M)$ is the fair peering fee resulting from any imbalance in the non-video traffic between the ISP and transit provider. The term $\frac{1}{2} c^b V_u r'(1-x) ED^{hot}(M)$

is the ISP's traffic-sensitive backbone cost incurred by non-localized video traffic. The term $-\frac{1}{2}c^bV_ur'xED^{hot}(M)$ is the transit provider's traffic-sensitive backbone cost incurred by localized video traffic. The terms can be combined to give:

Theorem 3. *The fair peering fee between the transit provider and the ISP is:*

$$P^{TP,ISP} = c^bV_u\left[\frac{1}{2}(r-1) + r'(0.5-x)\right]ED^{hot}(M) \tag{11}$$

4. Cost-Based Peering Between a Content Provider and an ISP

We now turn to peering between a content provider and an ISP. We compare such direct interconnection with the ISP with the indirect interconnection considered in the previous section, in which a content provider sends video traffic through a transit provider to the ISP. We focus on the impact of elements of peering policies, including the number of interconnection points and video traffic localization, on the cost-based peering fee. We also determine the conditions under which a content provider should be eligible for settlement-free peering.

4.1. ISP's Backbone Cost

We assume that the content provider and the ISP have agreed to peer with each other. However, peering between a content provider and an ISP differs from peering between a transit provider and an ISP for three reasons.

First, we assume that the content provider peers with the ISP at N IXPs, and that N may be less than the M IXPs at which the transit provider and the ISP peer. This could increase the ISP's backbone cost.

Second, the content provider may localize a different proportion of video traffic than does the transit provider. A higher degree of localization may reduce the backbone cost for the ISP. We assume that a proportion x^d of the video transmitted to the ISP's users within each access network is delivered from the content provider to the ISP at the IXP nearest to the end user among the IXPs at which they agree to peer. We assume that the remaining proportion $1 - x^d$ of the video transmitted to the ISP's users within each access network is delivered to the ISP at any IXP, and that the location of this content provider server is independent of the location of the end user.

Third, content providers carry only video traffic, which results in a higher ratio of downstream traffic (from the content provider to the ISP) to upstream traffic (from the ISP to the content provider) compared to transit providers. This is due to the fact that video traffic is almost entirely downstream. We continue to denote the volume of video downstream traffic by V_v.

We denote the ISP's traffic-sensitive backbone cost by $C^{ISP}_{cp,video}$.

Theorem 4. *The traffic-sensitive backbone cost of the ISP when peering with the content provider is:*

$$C^{ISP}_{cp,video} = c^bV_v\left[x^dED^{cold}(N) + (1-x^d)ED^{hot}(N)\right] \tag{12}$$

The first term, $c^b V_v x^d ED^{cold}(N)$, is the ISP's backbone cost for localized video traffic, which accounts for the ISP's transport of a proportion x^d of the video traffic from the IXP nearest to the end user among the IXPs at which they agree to peer. The second term, $c^b V_v (1 - x^d) ED^{hot}(N)$, is the ISP's backbone cost for non-localized video traffic, which accounts for the ISP's transport of a proportion $1 - x^d$ of the video traffic from any IXP where they have agreed to peer.

4.2. Cost-Based Peering Fee

We must first address the question of how to define *cost-based* in the context of direct peering between a content provider and an ISP. Should we define cost-based as the peering fee that equalizes the net costs of the content provider and the ISP, similar to our analysis above for peering between a transit provider and an ISP? Or should we define cost-based as the peering fee that results in the same ISP net costs for transporting the video traffic as in the case in which the video traffic is transported across a transit provider's network?

We believe that the appropriate definition of cost-based is the latter one. If we were to attempt to equalize the net costs of the content provider and the ISP, we would have to account for the cost to the content provider of building its CDN. However, the decision between building a CDN versus transporting video traffic across the backbone should be made on the basis of the cost of servers versus the transmission cost, not also on the peering fee. The cost-based peering fee should be determined solely by ensuring that the ISP's net costs are unaffected by the content provider's decision.

Denote the fair peering fee between the transit provider and the ISP that is related solely to video traffic (not related to upstream or non-video downstream traffic) by $P_v^{TP,ISP}$. Using Theorem 3, we can determine $P_v^{TP,ISP}$ by considering only the video traffic component of the fair peering fee between the transit provider and the ISP ($P^{TP,ISP}$). It can be expressed as:

$$P_v^{TP,ISP} = c^b V_v (0.5 - x) ED^{hot}(M) \tag{13}$$

Denote the cost-based peering fee for direct interconnection between a content provider and an ISP by $P^{CP,ISP}$. It is given by:

$$P^{CP,ISP} = P_v^{TP,ISP} + (C_{cp,video}^{ISP} - C_{down,video}^{ISP}) \tag{14}$$

where $C_{cp,video}^{ISP}$ is the traffic-sensitive backbone cost of the ISP when peering with the content provider (given in Theorem 4), and $C_{down,video}^{ISP}$ is the traffic-sensitive backbone cost of the ISP for delivering video traffic when peering with a transit provider (given in (3)).

The cost difference ($C_{cp,video}^{ISP} - C_{down,video}^{ISP}$) accounts for any changes to the ISP's cost resulting from any differences in traffic flows and localization when it peers with a content provider rather than with a transit provider. Such a cost-based peering price may be the minimum price an ISP will accept. The ISP may only consider direct peering if the content provider pays for the incremental cost that an ISP incurs by direct peering. Otherwise, the ISP has no incentive to agree to a direct peering arrangement, as the cost of peering with a transit provider would be lower.

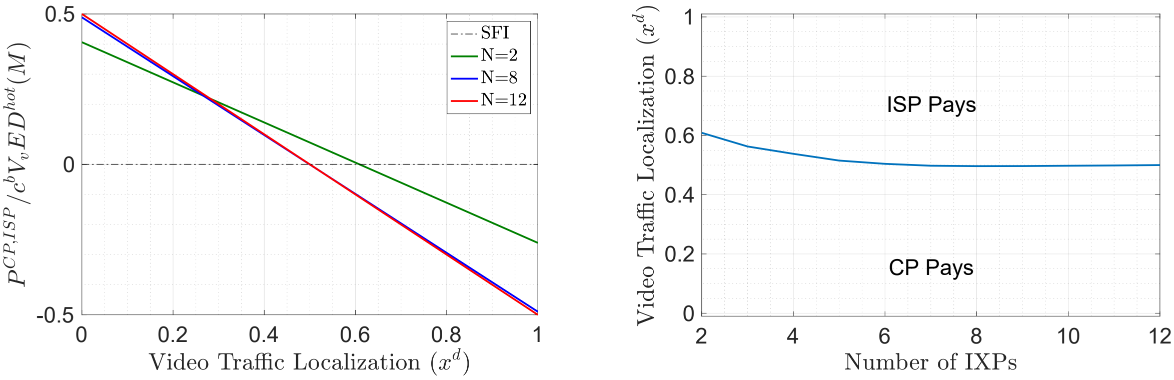

(a) Peering Fee (b) Settlement-Free Curve

Figure 2: Cost-based peering between a content provider and an ISP

Using (3), (13) and Theorem 4, we can express the cost-based peering fee in (14) as:

$$\begin{aligned} P^{CP,ISP} &= c^b V_v \Big[x^d ED^{cold}(N) \\ &\quad + (1 - x^d) ED^{hot}(N) - 0.5 ED^{hot}(M) \Big] \end{aligned} \tag{15}$$

Finally, we can rearrange the terms to separate the effects of the number of IXPs at which the content provider and the ISP peer from the effects of localization:

Theorem 5. *The cost-based peering fee between the content provider and the ISP is:*

$$\begin{aligned} P^{CP,ISP} &= c^b V_v \Big[\Big(0.5 - x^d\Big) ED^{hot}(M) \\ &\quad + \Big(1 - x^d\Big) \Big(ED^{hot}(N) - ED^{hot}(M)\Big) \\ &\quad + x^d \Big(ED^{cold}(N) - ED^{cold}(M)\Big) \Big] \end{aligned} \tag{16}$$

The first term, $c^b V_v (0.5 - x^d) ED^{hot}(M)$, represents the effect of video traffic localization on the cost-based peering fee. The second term, $c^b V_v (1 - x^d)(ED^{hot}(N) - ED^{hot}(M)) + x^d(ED^{cold}(N) - ED^{cold}(M))$, represents the effect of the number of IXPs at which they agree to peer on the cost-based peering fee.

Figure 2(a) illustrates the cost-based peering fee between the content provider and the ISP as a function of video traffic localization and the number of IXPs at which they agree to peer. At low amounts of localization, the cost-based peering fee is positive. However, as the content provider sends traffic with more localization, the cost-based peering fee decreases and at some point becomes negative (meaning that the ISP should pay the content provider).

The cost-based peering fee also varies with the number of IXPs at which they agree to peer. When localization is very low, interconnecting at more IXPs is not beneficial to the ISP, because the ISP needs to carry the video traffic over longer distances in its backbone network since the peering IXP moves farther from the IXP nearest to the end user [10]. Therefore, the cost-based peering fee increases slightly with N at very low amounts of localization. However, for moderate to high localization, interconnecting at more IXPs is beneficial to the ISP, because the ISP's backbone cost decreases since the peering IXP moves closer to the IXP nearest to the end user [10]. Therefore, the cost-based peering fee decreases with N at moderate to high amounts of localization.

4.3. Settlement-Free Peering

Finally, we wish to determine under what elements of peering policies (namely, number of interconnection points and video traffic localization), the content provider should be eligible for settlement-free peering. By setting the cost-based peering fee to zero ($P^{CP,ISP} = 0$) in Theorem 5, we can determine the number of IXPs and localization required for settlement-free peering:

Theorem 6. *The percentage of video traffic localization between a content provider and an ISP required for a cost-based peering fee of zero is:*

$$x^d = \frac{ED^{hot}(N) - 0.5ED^{hot}(M)}{ED^{hot}(N) - ED^{cold}(N)} \tag{17}$$

Figure 2(b) illustrates the settlement-free peering curve for direct interconnection between a content provider and an ISP, as a function of the amount of video traffic localization and the number of IXPs at which they peer. For peering between a transit provider and an ISP, 50% localization is sufficient for settlement-free peering when the non-video traffic ratio $r = 1$; different amounts of localization may be required for other non-video traffic ratios and it may depend on the amount of video traffic. For direct interconnection between a content provider and an ISP, however, the traffic ratio is now irrelevant because the net cost to the ISP of transporting the video traffic is solely a function of localization.[2] Indeed, if the content provider and the ISP agree to peer at all $N = 12$ locations, then 50% localization is sufficient to justify settlement-free peering.

As the number of interconnection points decreases from $N = 12$, the content provider should send an increasing proportion of video traffic locally in order to be eligible for settlement-free peering. However, when $8 \leq N \leq 12$, there is little variation in ISP's backbone cost and thus little change in amount of localization required for settlement-free peering.

4.4. Evaluation of Arguments

Large ISPs often argue that they should apply the same settlement-free peering requirements to peering ISPs, peering transit providers, and peering content providers. However, whereas when a transit provider and an ISP peer we have shown that the fair peering fee

[2]The required localization is solely dependent on the location of the interconnection points and the distribution of video traffic among the population. The other parameters of our model do not affect the settlement-free peering curve.

is a function of the traffic ratio, when a content provider and an ISP peer we have shown that the cost-based peering fee is a function of the number of interconnection points and of localization. Hence, different settlement-free peering should apply to these two situations.

Large ISPs often argue that they should be compensated by large content providers regardless of the amount of video content localization. Again we disagree. Our results show that it is rational for an ISP to agree to settlement-free peering with a content provider that delivers a sufficient amount of video traffic locally.

Large content providers sometimes argue that they should be entitled to settlement-free peering solely because the ISP's customers have already paid the ISP to transport the traffic the content providers are sending. We believe this argument is too simplistic. We have shown that the cost-based peering fee should include consideration of the ISP's backbone transportation cost and thus of the number of interconnection points and of localization. A more nuanced argument by large content providers is that they should be eligible for settlement-free peering if they bring the content close to customers. We have shown that localization should indeed play a key role in determining eligibility for settlement-free peering.

5. A Model of User Subscription to Broadband and to Video Streaming

In the previous section, we determined a cost-based peering fee between a content provider and an ISP. This cost-based peering fee also represents the minimum price an ISP may accept based on the ISP's cost for transporting the content provider's traffic using direct peering versus indirect peering.

In the remainder of the paper, we consider the peering fee that maximizes an ISP's profit. We use a two-sided market model in which a profit-maximizing ISP determines broadband prices and the peering price, and in which content providers determine their service prices based on the peering price.

The range from the cost-based peering fee to the profit-maximizing peering fee determines possible peering fees that the parties may agree to, if unregulated. The lower boundary is shaped by costs, while the upper limit is dictated by profit maximization. The upper limit of this range is the intersection of two components: the peering price that maximizes the ISP's profit, and the content provider's maximum willingness to pay.

In this section, we develop a user subscription model. In the following section, we will introduce a two-sided market model that considers how the ISP chooses charges on both sides of the market: the customers on one side and the content provider on the other.

5.1. Service offerings

Before we can analyze the paid peering prices, we need a model of user subscription to broadband service tiers and to Subscription Video on Demand (SVOD) providers. In [11], we proposed a model of user subscription to broadband and SVOD providers. However, since then, it has become common for large ISPs in the United States to levy an additional charge for unlimited monthly usage. We thus update that model to account for this new charge.

ISPs offer multiple tiers of broadband services, differentiated principally by download speed and amount of data. ISPs typically market these broadband service tiers by recommending specific tiers to consumers who engage in specific types of online activities. For example, Comcast recommends a lower service tier to consumers who principally use their Internet connection for email and web browsing, but a higher service tier to consumers who use the Internet for video streaming.

Many ISPs place a limit on the monthly usage on each service tier. The data caps placed on higher tiers have often been in the range of 1 to 1.25 TB per month, and have not increased substantially in the past few years. When these data caps were originally introduced, few users exceeded them. However, users who engage in substantial subscription video streaming are increasingly hitting these data caps. Indeed, the percentage of users whose usage exceeded 1 TB per month in the fourth quarter of 2022 was 18.7%, which is ten times the percentage observed five years before [12]. Subscribers who exceed these data caps may be charged very high per volume overage fees. In recent years, however, some large ISPs have begun to offer an add-on to each service tier that removes the usage cap. Subscribers with high SVOD usage increasingly purchase this unlimited usage add-on.

Much of the debate over paid peering concerns consumers who stream large volumes of video. Thus, we construct here a model that includes three broadband service tiers: a basic tier with a download speed intended for email and web browsing; a premium tier with a download speed intended for video streaming and gaming and with a data cap that limits the amount of video streaming, and a unlimited premium tier with the same speed as the premium tier but without any data cap. Although some ISPs offer more than three tiers, the majority of customers subscribe to a subset of three tiers, and this three-tier model is sufficient to separately evaluate the effect of paid peering prices on consumers who utilize video streaming and on consumers who don't.

Specifically, we model a single monopoly ISP that offers a basic tier at a monthly price P^b, a limited premium tier at a monthly price $P^b + P^p$, and an unlimited data option for premium tier customers for an extra monthly fee of P^h. We consider N consumers, each of whom may subscribe to the basic tier, the limited premium tier, the unlimited premium tier, or to none of these. We denote user i's utility per month from subscription to the basic tier by b_i, and user i's utility per month from subscription to the limited premium tier by $b_i + p_i$.

We focus on the aggregate of all SVOD providers that directly interconnect with the ISP and that may pay (or be paid) a fee for peering with the ISP. We model the aggregate of all plans offered by these SVOD providers, but to keep the model tractable we consider a single price of P^{svod} per month for the aggregate. We presume that a consumer who gains significant utility from SVOD subscribes to the unlimited premium tier, and that all users who subscribe to the unlimited premium tier subscribe to SVOD. As a result, SVOD users pay a total amount of $P^v = P^h + P^{svod}$ to the combination of their ISP and SVOD providers to enable SVOD. We denote user i's utility per month from the combination of subscription to SVOD providers and to the unlimited usage add-on by v_i. Consumer i's utility from all other content is included in $b_i + p_i$.

Consumers differ in the utilities they place on broadband service tiers and on SVOD.

We assume that the number of consumers N is large, and we denote the joint probability density function of their utilities by $f_{B,P,V}(b, p, v)$.

5.2. Demand functions

Each consumer thus has four choices

$$X_i \triangleq \begin{cases} n, & \text{do not subscribe} \\ b, & \text{subscribe to the basic tier} \\ p, & \text{subscribe to the limited premium tier} \\ v, & \text{subscribe to the unlimited premium tier and to SVOD.} \end{cases} \tag{18}$$

Consumer i's consumer surplus, defined as utility minus cost, under each choice is thus

$$CS_i(X_i; b_i, p_i, v_i) = \begin{cases} 0, & X_i = n \\ b_i - P^b, & X_i = b \\ b_i + p_i - P^b - P^p, & X_i = p \\ b_i + p_i + v_i - P^b - P^p - P^v, & X_i = v. \end{cases} \tag{19}$$

Each consumer is assumed to maximize consumer surplus. Thus, consumer i adopts the choice

$$X_i^*(b_i, p_i, v_i) \triangleq \arg\max_{X_i} CS_i(X_i; b_i, p_i, v_i), \tag{20}$$

and earns a corresponding consumer surplus $CS_i^* \triangleq CS_i(X_i^*)$.

Each of the N consumers makes an individual choice per (20). The consumers who choose to subscribe to the basic tier are those whose utility b_i from subscription to the basic tier exceeds its monthly price P^b, whose incremental utility p_i from subscription to the limited premium tier falls below the incremental monthly price P^p, and whose incremental utility $p_i + v_i$ from subscription to the unlimited premium tier and to SVOD falls below the corresponding incremental monthly price $P^p + P^v$. Thus, the demand[3] for the basic tier is given by:

$$N^b(P^b, P^p, P^v) = N \int_{-\infty}^{P^p} \int_{-\infty}^{P^p+P^v-p} \int_{P^b}^{\infty} f_{B,P,V}(b, p, v)\, db\, dv\, dp. \tag{21}$$

Similarly, the consumers who choose to subscribe to the limited premium tier are those whose utility $b_i + p_i$ from subscription to the limited premium tier exceeds its monthly price $P^b + P^p$, whose incremental utility p_i from subscription to the limited premium tier exceeds the incremental monthly price P^p, and whose incremental utility v_i from subscription to unlimited premium tier and to SVOD falls below the incremental monthly price P^v. Thus, the number of consumers who subscribe to the limited premium tier is given by:

$$N^p(P^b, P^p, P^v) = N \int_{-\infty}^{P^v} \int_{P^p}^{\infty} \int_{P^b+P^p-p}^{\infty} f_{B,P,V}(b, p, v)\, db\, dp\, dv. \tag{22}$$

[3]Since we model a finite number N of consumers whose utilities are given by a joint probability density function, this equation, and other similar equations below, give the *average demand*. However, for simplicity of presentation, we use the term *demand*.

Finally, the consumers who choose to subscribe to the unlimited premium tier and to SVOD are those whose utility $b_i + p_i + v_i$ from subscription to both services exceeds the combined cost $P^b + P^p + P^v$, whose incremental utility $p_i + v_i$ from subscription to only the basic tier exceeds the corresponding incremental price $P^p + P^v$, and whose incremental utility v_i from subscription to SVOD falls exceeds the incremental monthly price P^v. Thus, the demand for the unlimited premium tier customers who subscribe to SVOD is given by:

$$N^v(P^b, P^p, P^v) = N \int_{P^v}^{\infty} \int_{P^p+P^v-v}^{\infty} \int_{P^b+P^p+P^v-p-v}^{\infty} f_{B,P,V}(b, p, v)\; db\; dp\; dv. \tag{23}$$

5.3. Profits

We assume that the ISP incurs a monthly marginal cost C^b per basic tier subscriber. The ISP marginal profit per basic tier subscriber is thus $P^b - C^b$. We assume that the ISP incurs a monthly marginal cost $C^b + C^p$ per limited premium tier subscriber. The ISP marginal profit per such broadband service tier subscriber is thus $P^b + P^p - C^b - C^p$. We associate an ISP monthly marginal cost $C^b + C^p + C^v$ per unlimited premium tier customer.

We also consider a peering price of P^d per SVOD subscriber for direct interconnection between the ISP and SVOD providers. This price may be positive if the ISP charges SVOD providers for direct interconnection, negative if the SVOD providers charge the ISP for direct interconnection, or zero if the peering is settlement-free.

The ISP marginal profit per SVOD subscriber is $P^b + P^p + P^h + P^d - C^b - C^p - C^v$. The total ISP profit (excluding fixed costs)[4] is thus

$$\begin{aligned} \pi^{ISP}(P^b, P^p, P^h, P^d, P^{svod}) &= (P^b - C^b)N^b + (P^b + P^p - C^b - C^p)N^p \\ &\quad + (P^b + P^p + P^h + P^d - C^b - C^p - C^v)N^v. \end{aligned} \tag{24}$$

We assume that the SVOD providers incur a monthly marginal cost C^{svod} per subscriber. The aggregate SVOD provider marginal profit per subscriber is thus $P^{svod} - C^{svod} - P^d$, and their total profit is:

$$\pi^{SVOD}(P^b, P^p, P^h, P^d, P^{svod}) = (P^{svod} - C^{svod} - P^d)N^v. \tag{25}$$

where P^{svod} is the aggregate monthly SVOD providers price.

We assume that the market determines the aggregate SVOD price, excluding paid peering fees, when there is no regulation of prices. We denote the aggregate SVOD price, excluding paid peering fees, by P_0^v. We presume that an ISP charging peering prices would likely charge them to both directly interconnected content providers and directly interconnected transit providers. We further presume that transit providers would pass peering prices through to their customers. As a consequence, we foresee that peering prices would be paid by all large SVOD providers selling to the ISP's customers. An open question is whether the SVOD providers can pass through any peering price (P^d) to their customers by adding it to their SVOD prices. We denote the pass-through rate of the peering fee by $0 < \alpha \leq 1$:

$$P^{svod}(P^d) = P_0^v + \alpha P^d \tag{26}$$

[4]Throughout the paper, ISP profit accounts for costs that are sensitive to the number of subscribers and/or the amount of traffic, but does not account for costs that are fixed with respect to both the number of subscribers and the amount of traffic.

5.4. Costs

In order to understand the effect of number of interconnection points and SVOD traffic localization on peering agreements between the ISP and SVOD providers, we first analyze the monthly marginal costs.

We partition traffic-related costs into three categories based on the associated costs and usage:

- Subscription Video on Demand (SVOD) traffic: traffic generated by subscription-based video streaming platforms such as Netflix, Disney+, Amazon Prime, and Hulu.
- Premium traffic: video and gaming traffic that isn't generated by SVOD providers, including traffic from non-subscription video streaming, online multiplayer gaming, live-streaming platforms, video conferencing, and other high-bandwidth activities.
- Basic traffic: all Internet traffic that isn't included in the SVOD or premium traffic categories, including traffic from web browsing, email, social media, downloading or uploading files, and other common online activities that don't typically require high bandwidth or produce large amounts of data traffic.

Our assumption is that basic tier subscribers only transmit and receive basic traffic, limited premium tier subscribers transmit and receive both basic and premium traffic, and unlimited premium tier subscribers transmit and receive all types of traffic.

The ISP exchanges basic traffic through a transit provider. We assume that the ISP and the transit provider interconnect at the $M = 12$ major interconnection points, and that both use hot potato routing. We consider downstream basic traffic originating on the transit provider's network destined for an end user located in the ISP's network, and denote the volume of this traffic by V^b_{down}. We also consider upstream basic traffic originating with an end user located in the ISP's network and destined for a location on the transit provider's network, and denote the volume of this traffic by V^b_{up}. The monthly marginal cost C^b per basic tier subscriber can be defined as:

$$C^b = c^b(\overline{V}^b_{down} ED^{hot}(M) + \overline{V}^b_{up} ED^{cold}(M)) + (\overline{V}^b_{down} + \overline{V}^b_{up})(c^m ED^m + c^a ED^a) + C^f_{ISP} \tag{27}$$

where $\overline{V}^b_{down}$ is the average usage of downstream basic traffic among all ISP customers, $\overline{V}^b_{up}$ is the average usage of upstream basic traffic among all ISP customers, $ED^{hot}(M)$ is the average distance of basic traffic on the ISP's backbone network using hot potato routing, $ED^{cold}(M)$ is the average distance of basic traffic on the ISP's backbone network using cold potato routing, ED^m is the average distance of basic traffic on the ISP's middle network, ED^a is the average distance of basic traffic on the ISP's access network, and C^f_{ISP} is the non-traffic sensitive monthly marginal cost per basic tier subscriber.

The ISP also exchanges premium traffic (which excludes SVOD services that interconnect directly to the ISP) through the transit provider. Denote the volume of the downstream premium traffic by V^p_{down} and the volume of the upstream premium traffic by V^p_{up}. Recall that the monthly marginal cost per limited premium tier subscriber is $C^b + C^p$. Now C^b is given in (27), since limited premium tier subscribers transmit and receive the same volume of basic traffic as do basic tier subscribers. However, they also transmit and receive premium traffic, and this marginal traffic-sensitive cost is:

$$C^p = c^b(\overline{V}^p_{down} ED^{hot}(M) + \overline{V}^p_{up} ED^{cold}(M)) + (\overline{V}^p_{down} + \overline{V}^p_{up})(c^m ED^m + c^a ED^a) \quad (28)$$

where $\overline{V}^p_{down}$ is the average usage of downstream premium traffic among all premium tier customers and $\overline{V}^p_{up}$ is the average usage of upstream premium traffic among all premium tier customers.

The ISP receives SVOD traffic through peering with SVOD providers. Peering between SVOD providers and the ISP differs from peering between a transit provider and an ISP for three reasons. First, we assume that the SVOD provider has deployed content servers at N major interconnection points, and that N may be less than the number of interconnection points at which the transit provider and the ISP agree to peer (M). Second, the SVOD provider may localize a proportion of SVOD traffic. We assume that a proportion x^d of the SVOD traffic transmitted to the ISP's users within each access network is delivered from the SVOD provider at a server located at the IXP nearest to the end user among the IXPs at which they agree to peer. We assume that the remaining proportion $1 - x^d$ of the SVOD traffic transmitted to the ISP's users within each access network is delivered using any SVOD server, and that the location of this SVOD provider's server is independent of the location of the end user. Third, SVOD providers carry only downstream SVOD traffic. We denote the volume of downstream SVOD traffic by V^v_{down}.

Recall that the monthly marginal cost per unlimited premium tier subscriber is $C^b + C^p + C^v$. Now C^b is given in (27) and C^p is given in (28), since unlimited premium tier subscribers transmit and receive the same volume of basic and premium traffic as do limited premium tier subscribers. However, they also receive SVOD traffic, and this marginal traffic-sensitive cost is:

$$C^v = c^b \overline{V}^v_{down}\Big(x^d ED^{cold}(N) + (1 - x^d) ED^{hot}(N)\Big) + \overline{V}^v_{down}(c^m ED^m + c^a ED^a) \quad (29)$$

where $\overline{V}^v_{down}$ is the average usage of SVOD among all SVOD subscribers.

We must address whether an ISP should recover the incremental ISP cost C^v per SVOD subscriber across different parts of its network solely from its subscribers or also from interconnecting networks. We propose that customers bear middle-mile and access network costs related to SVOD traffic (denoted by C^h), and that SVOD providers bear ISP backbone costs related to SVOD traffic (denoted by C^d), where $C^v = C^h + C^d$.

The first term in (29) is the ISP backbone cost related to SVOD traffic per SVOD subscriber:

$$C^d = c^b \overline{V}^v_{down}\Big(x^d ED^{cold}(N) + (1 - x^d) ED^{hot}(N)\Big), \quad (30)$$

and the second term in (29) is the middle-mile and access network cost related to SVOD traffic per SVOD subscriber:

$$C^h = \overline{V}^v_{down}(c^m ED^m + c^a ED^a) \quad (31)$$

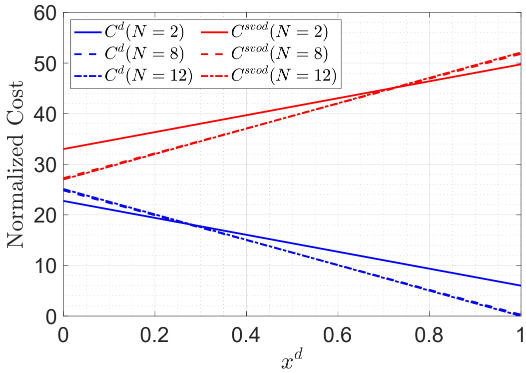

Figure 3: Effect of localization and number of IXPs on C^d and C^{svod}

We now turn to the SVOD providers' costs. Recall that we assume that a proportion x^d of the SVOD traffic transmitted to the ISP's users within each access network is delivered from the SVOD provider at a server located at the IXP nearest to the end user among the IXPs at which they agree to peer. If SVOD providers were to localize this traffic by using cold potato routing, then their the monthly marginal cost per SVOD subscriber incurred by the SVOD providers would be:

$$C^{svod} = c^b\overline{V}^v_{down}\Big(x^d ED^{hot}(N) + (1-x^d)ED^{cold}(N)\Big) + C^f_{SVOD} \tag{32}$$

where C^f_{SVOD} is the non-traffic sensitive monthly marginal cost per SVOD subscriber. The first term is the SVOD's backbone cost for localized SVOD traffic, which the SVOD providers deliver using cold potato routing. The second term is the SVOD providers' backbone cost for non-localized SVOD traffic, which the SVOD providers deliver using hot potato routing.

That said, it is common for SVOD providers to utilize CDN (Content Delivery Network) services to localize traffic, instead of using cold potato routing, if this reduces their costs. In future research, we will propose a cost model for SVOD providers using CDNs. In this paper, we simply consider (32) to be an upper bound for SVOD provider costs.

Figure 3 shows the effect of SVOD traffic localization (x^d) and number of interconnection points (N) on the backbone costs of the ISP (C^d) and of SVOD providers (C^{svod}) related to SVOD traffic. As the amount of SVOD traffic localization (x^d) increases, the ISP cost C^d is decreasing and the SVOD cost C^{svod} is increasing, because increasing localization results in the SVOD providers carrying more of the SVOD traffic on their backbone networks and handing it off to the ISP at an IXP closer to end users.

The costs of the ISP and SVOD providers also vary with the number of IXPs at which they agree to peer. When localization is very low, interconnecting at more IXPs is not beneficial to the ISP, because the ISP needs to carry the SVOD traffic over longer distances in its backbone network since the peering IXP moves farther from the IXP nearest to the end user [13]. Therefore, the ISP cost (C^d) increases slightly with N at very low amounts of localization. However, it is beneficial for the SVOD provider to peer at more IXPs when the localization is very low since the peering IXP moves closer to the IXP nearest to the source. Therefore, the SVOD cost (C^{svod}) decreases with N at low amounts of localization.

For moderate to high localization, interconnecting at more IXPs is beneficial to the ISP and the ISP's cost (C^d) decreases since the peering IXP moves closer to the IXP nearest to the end user [13]. But it is not beneficial to the SVOD provider since interconnecting at more IXPs with high amount of traffic localization increases the SVOD cost (C^{svod}) since the peering IXP moves farther from the IXP nearest to the source.

6. Peering Between a Content Provider and a Profit-Maximizing ISP

The previous section presented a model for consumer demand for broadband Internet access service and for SVOD services, resulting in the demand functions (21-23) and the corresponding ISP and SVOD provider profits (24-25). There are a number of options for modeling how the broadband service tier prices (P^b, P^p, and P^h), the aggregate SVOD price (P^{svod}), and the peering price (P^d) are determined. In this section, we explore the determination of prices when the ISP aims to maximize its profit.

6.1. Numerical parameters

This two-sided model is somewhat amenable to closed-form analysis. However, we find it useful to also examine the model under a set of realistically chosen parameters. We set out those parameters in this subsection.

The joint probability density function of user utilities for the basic tier, the limited premium tier, and the unlimited premium is represented by $f_{B,P,V}(b, p, v)$. For numerical evaluation, we assume that each utility is independent and has a Normal distribution: $B \sim \mathcal{N}(\mu_b, \sigma_b^2)$, $P \sim \mathcal{N}(\mu_p, \sigma_p^2)$, $V \sim \mathcal{N}(\mu_v, \sigma_v^2)$. We need to determine numerical values for the means and variances.

The ISP incurs a monthly marginal cost of C^b per subscriber, a monthly incremental cost of C^p per premium tier subscriber, and a monthly incremental cost C^v per SVOD subscriber. We need to determine numerical values for these three costs.

Unfortunately, direct information about user utilities and ISP costs is scarce. Instead, we choose numerical values for user utilities and ISP costs indirectly using available information about demand and prices in the United States.

There are several sets of publicly available statistics about broadband prices and subscriptions [14–17]. While the set of statistics differ, they show that roughly 85% of households in the United States subscribe to broadband service. Of subscribers, 19.5% subscribe to plans with download speeds below 100 Mbps, 61.8% subscribe to plans with download speeds above 100 Mbps but consume less than 1 TB of data, and 18.7% subscribe to plans with download speeds above 100 Mbps and consumer over 1 TB of data [12]. Hence, we wish to choose numerical values for user utilities and ISP costs so that, at the ISP profit-maximizing prices, $(N^b + N^p + N^v)/N = 0.85$, $N^b/(N^b + N^p + N^v) = 0.195$, $N^p/(N^b + N^p + N^v) = 0.618$, and $N^v/(N^b + N^p + N^v) = 0.187$.

Moreover, upon reviewing the plan and pricing web pages of major ISPs [18–20], we find that the price of the lower of the two most popular tiers is roughly \$61 per month, the price of the higher of the two most popular tiers is roughly \$83 per month, and the price of an unlimited usage add-on is roughly \$30 per month. Hence, we wish to choose numerical values for user utilities and ISP costs so that the ISP profit-maximizing prices are $P^b = \$61.00$, $P^p = \$22.00$, and $P^h = \$30.00$.

According to [21], Americans subscribe to an average of four SVOD providers. In our model, we assume that these four SVOD providers directly interconnect with the ISP. Additionally, the average monthly payment for SVOD in the United States is approximately \$54 [22]. In addition, although we consider any pass-through rate of the paid peering fee ($0 < \alpha \leq 1$), in the numerical results below we use $\alpha = 0.5$. Therefore, we assign a aggregate SVOD price, excluding paid peering fees $P_0^v = \$50$ in our model to achieve $P^{svod} = 54$ where the ISP chooses the peering price that maximizes its profit.

The average household usage is approximately 600 GB [12]. The combined traffic generated by Netflix, Disney+, Amazon Prime, and Hulu accounts for 31.6% of downstream traffic and 6.17% of upstream traffic [23]. Additionally, video and gaming traffic, excluding the aforementioned video-on-demand services, represents 50.82% of downstream traffic and 36.82% of upstream traffic [23]. The remaining 17.58% of downstream traffic and 57.01% of upstream traffic corresponds to other basic traffic [23]. Furthermore, roughly 6% of the total traffic is attributed to upstream traffic [23].

Let $\overline{V}^v$ represent the average usage of SVOD traffic (traffic generated by Netflix, Disney+, Amazon Prime, and Hulu) among subscribers who are on the unlimited premium tier plan. Similarly, let $\overline{V}^p$ denote the average usage of premium traffic (including video and gaming traffic, excluding the previously mentioned SVOD providers) among all customers on the premium tier, regardless of whether or not they have data caps. Finally, $\overline{V}^b$ denotes the average usage of basic traffic among all customers of the ISP.[5]

Let's consider the following traffic values:

$$\begin{aligned}
\overline{V}^b_{up} &= \frac{V^b_{up}}{N^b + N^p + N^v} = \frac{(600)(0.06)(57.01)}{18.7 + 61.8 + 19.5} = 21\ GB \\
\overline{V}^p_{up} &= \frac{V^p_{up}}{N^p + N^v} = \frac{(600)(0.06)(36.82)}{18.7 + 61.8} = 16\ GB \\
\overline{V}^v_{up} &= \frac{V^v_{up}}{N^v} = \frac{(600)(0.06)(6.17)}{18.7} = 12 \\
\overline{V}^b_{down} &= \frac{V^b_{down}}{N^b + N^p + N^v} = \frac{(600)(0.94)(17.58)}{18.7 + 61.8 + 19.5} = 99\ GB \\
\overline{V}^p_{down} &= \frac{V^p_{down}}{N^p + N^v} = \frac{(600)(0.94)(50.82)}{18.7 + 61.8} = 356\ GB \\
\overline{V}^v_{down} &= \frac{V^v_{down}}{N^v} = \frac{(600)(0.94)(31.6)}{18.7} = 953\ GB
\end{aligned} \tag{33}$$

We use these parameters in the remainder of the paper except as noted.

[5]Given that video traffic from large streaming providers typically has significantly more downstream traffic than upstream traffic (approximately 80 times more), the cost of upstream traffic is not taken into consideration for those subscribers in our model.

6.2. Peering Fee

Let's begin by considering the scenario in which the ISP selects all its prices in order to maximize its profit.

Once a subscriber chooses an ISP, the ISP has a monopoly on the transport of traffic within the ISP's access network that the customer resides in. In contrast, there may be a competitive market for the transport of Internet traffic across backbone networks. Correspondingly, we assume that the ISP has the market power to determine the broadband service tier prices (P^b and P^p), unlimited usage add-on price (P^h), and the peering price (P^d) to maximize its profit. However, there is a constraint on these prices given by SVOD provider's earning a positive profit. The resulting optimization problem is:

$$\begin{aligned}(P^b_{ISP}, P^p_{ISP}, P^h_{ISP}, P^d_{ISP}) = \arg \max_{(P^b, P^p, P^h, P^d)} &\ \pi^{ISP}(P^b, P^p, P^h, P^d, P^{svod}) \\ \text{s.t.} &\\ &\ \pi^{SVOD}(P^b, P^p, P^h, P^d, P^{svod}) \geq 0\end{aligned} \tag{34}$$

Equations (26) and (34) set up a two-sided model in which the ISP earns revenue from both its customers and SVOD providers (if $P^d > 0$). The combination of the two equations captures the inter-dependencies between the ISP, the SVOD providers, and the consumers. The ISP-determined peering price (P^d), along with the pass-through rate (α), leads to an aggregate SVOD price (P^{svod}). The ISP-determined broadband service tier prices (P^b and P^p), unlimited usage add-on price (P^h), and the aggregate SVOD price (P^{svod}) lead to demands for each broadband service tier (N^b and N^p) and for SVOD services (N^v). These demands in turn affect how the ISP sets each of the prices.

Since the aggregate SVOD price (P^{svod}) is solely determined by (26), we can represent the ISP's profit as a function of four variables rather than five:

$$\begin{aligned}(P^b_{ISP}, P^p_{ISP}, P^h_{ISP}, P^d_{ISP}) = \arg \max_{(P^b, P^p, P^h, P^d)} &\ \pi^{ISP}(P^b, P^p, P^h, P^d, P^v_0 + \alpha P^d) \\ \text{s.t.} &\\ &\ (P^v_0 - (1-\alpha)P^d - C^{svod})N^v \geq 0\end{aligned} \tag{35}$$

As will demonstrated by the numerical results, the profit-maximizing ISP would set the peering fee at the maximum amount that SVOD providers are willing to pay.

$$\pi^{SVOD}(P^b, P^p, P^h, P^d) = (P^v_0 - (1-\alpha)P^d - C^{svod})N^v = 0 \tag{36}$$

The peering fee that maximizes the ISP's profit, while considering the SVOD providers' willingness to pay, is determined as follows:

$$P^d_{ISP} = \frac{P^v_0 - C^{svod}}{1-\alpha} \tag{37}$$

By utilizing the upper bound on C^{svod} given in (32), we can formulate the peering price as a function of number of interconnection points and the localization of SVOD traffic:

$$P^d_{ISP} = \frac{P^v_0 - C^f_{SVOD} - c^b \overline{V}^v_{down}\Big(x^d ED^{hot}(N) + (1-x^d)ED^{cold}(N)\Big)}{1-\alpha} \tag{38}$$

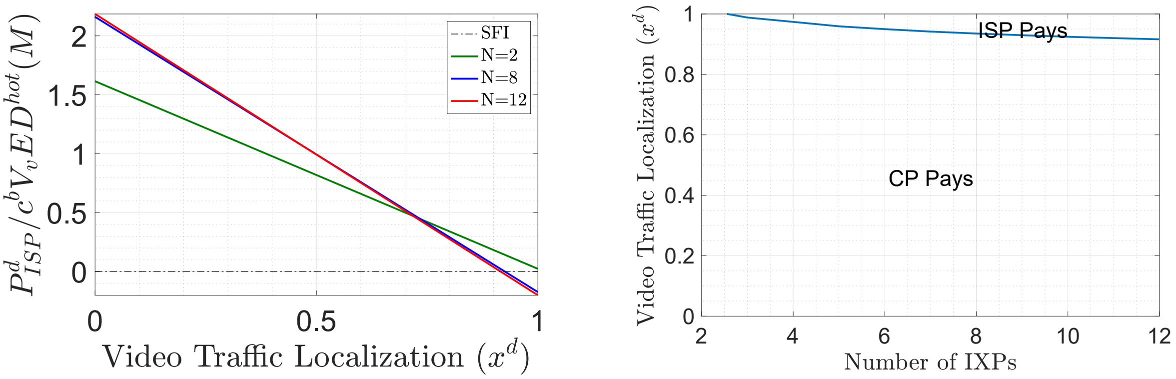

(a) Peering Fee

(b) Settlement-Free Curve

Figure 4: Peering between an SVOD provider and a profit-maximizing ISP

Note, however, that because (32) is an upper bound on SVOD costs, (38) represents a lower bound on the profit-maximizing peering price P^d.

Figure 4(a) illustrates the effect of SVOD traffic localization on the peering price (per SVOD provider) between a profit-maximizing ISP and the SVOD provider. As an SVOD provider increases the localization of its traffic, its cost (C^{svod}) increases. However, this increase in the SVOD cost results in a decrease in SVOD profit, thereby lowering the maximum amount the SVOD provider is willing to pay for peering. Consequently, with the rise in SVOD traffic localization, the peering price P^d that the ISP can demand decreases.

Figure 4(a) also illustrates the effect of the number of interconnection points on the peering price. As previously discussed in subsection 5.4, when there is a low level of localization, the cost for the SVOD provider (C^{svod}) decreases as the number of interconnection points increases. This decrease in cost raises both the profit and willingness to pay of the SVOD provider. As a result, with a low level of localization, the peering fee charged by the profit-maximizing ISP increases as the number of interconnection points rises. In contrast, with a high level of localization, the cost of the SVOD provider (C^{svod}) increases as the number of interconnection points rises. This leads to a decrease in both the profit and willingness to pay of the SVOD provider. Therefore, for a high level of localization, the peering fee charged by the profit-maximizing ISP declines as the number of interconnection points increases. Moreover, our findings suggest that there is little effect of interconnecting at more than 8 IXPs on the peering fee.

At very high levels of video traffic localization, the ISP maximizes profit by paying a negative peering fee, namely by paying SVOD providers for direct interconnection. At these levels, the ISP's backbone cost is minimal, and the SVOD provider's cost is high. As a consequence, without a small payment from the ISP to the SVOD providers, the SVOD providers can not offer this level of localization. The ISP is willing to offer this small payment in order to maintain its significant profit from its unlimited tier subscribers.

Figure 4(b) illustrates the settlement-free peering curve for direct interconnection between the SVOD provider and an ISP, as a function of the amount of SVOD traffic localization and the number of IXPs at which they peer. The curve is given by:

$$x^d = \frac{\frac{P_0^v - C_{SVOD}^f}{c^b \overline{V}_{down}^v} - ED^{cold}(N)}{ED^{hot}(N) - ED^{cold}(N)} \tag{39}$$

Finally, we examine the range from the cost-based peering fee to the profit-maximizing peering fee. The cost-based peering fee, illustrated in Figure 2, represents the minimum price an ISP may accept based on the ISP's cost for transporting the content provider's traffic using direct peering versus indirect peering. The profit-maximizing peering fee, illustrated in Figure 4, represents the maximum price a content provider may be willing to pay the ISP.

Comparing Figure 2(a) and Figure 4(a), we unsurprisingly find that the ISP profit-maximizing peering fee exceeds the cost-based peering fee for all values of video traffic localization. When the ISP and a content provider interconnect at a minimum of 8 IXPs, at low levels of localization the profit-maximizing peering fee may exceed the cost-based peering fee by over $1.50 per SVOD provider per SVOD customer. This gap may diminish with increasing localization, but a CDN cost model is required to confirm this.

Comparing Figure 2(b) and Figure 4(b), we find that when an ISP sets the peering fee to maximize profit, it is very likely to set this price to be greater than zero. In contrast, the cost-based peering fee may be zero (or negative) if the content provider localizes at least 50%-60% of its traffic.

Acknowledgment

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

7. Conclusion

Debates over paid peering and usage fees have expanded from the United States to Europe and South Korea. Our objective here is to understand an ISP's cost for directly peering with a content provider and the peering fee that maximizes an ISP's profit. The range from the cost-based peering fee to the profit-maximizing peering fee determines possible peering fees that the parties may agree to, if unregulated. The lower boundary is shaped by costs, while the upper limit is dictated by profit maximization. The upper limit of this range is the intersection of two components: the peering price that maximizes the ISP's profit, and the content provider's maximum willingness to pay.

We first considered cost-based cost sharing between an ISP and a directly interconnected content provider. The ISP's incremental cost is determined by comparing the sum of the ISP's backbone transportation costs and any peering fee when peering with another ISP or a transit provider versus peering directly with a content provider. Our results indicate that the cost-based peering fee is solely dependent on the localization of SVOD traffic and the number of interconnection points. Our results demonstrate that as the SVOD provider sends traffic with more localization, the cost-based peering fee decreases. The cost-based peering fee also varies with the number of IXPs at which they agree to peer. When localization is very low, the cost-based peering fee increases slightly with the number of interconnection points. However, for moderate to high localization, the cost-based peering fee decreases with the number of IXPs.

We then discuss the peering fee that maximizes the ISP's profit. The findings show that the profit-maximizing ISP might impose the highest fee that SVOD providers' willingness to pay can tolerate. Our research shows that with the increase in SVOD traffic localization, the peering fee that ISPs can ask for decreases, attributable to a lower maximum amount that SVOD providers are willing to pay for peering. Furthermore, our results indicate that when localization is at a low level, the peering fee charged by the ISP increases as the number of interconnection points rises. However, when localization is high, the peering fee charged by the ISP declines as the number of interconnection points increases.

Finally, we compared the cost-based peering fee with the profit-maximizing peering fee. We find that the profit-maximizing peering fee generally surpasses the cost-based fee, and this gap may diminish with increasing traffic localization.

8. Future Research

We have examined the cost-based peering price and the price that maximizes an ISP's profit. The profit-maximizing peering price is a function of an SVOD provider's willingness-to-pay, and this willingness-to-pay is in turn a function of an SVOD provider's costs. In this paper, we bounded an SVOD provider's costs using its backbone traffic costs if it delivers content using cold potato routing. However, additional research is warranted to model an SVOD provider's costs when it replaces a portion of this backbone traffic by a CDN. In addition, an analysis should determine the effect of such a cost reduction on the peering price.

We believe regulatory oversight of peering prices may be warranted when there is a substantial difference between cost-based and profit-maximizing peering prices. Research is warranted to examine how a regulator should provide such oversight. Should a regulator examine the effect of peering prices on consumer surplus or social welfare? How does the history of regulatory cost accounting and cost recovery apply here?

Our analysis of broadband subscription and streaming video subscription revealed that data caps are now playing an increasing role in consumer's streaming habits. The use of data caps have generated their own policy debates. Research is warranted to explore if and how the interconnection and data caps policy issues are connected.

Finally, another valuable angle for future research could be the exploration of international comparisons in peering arrangements and regulatory frameworks. How does the US experience compare to those of the European Union, Asia, or other global markets? What lessons might be drawn from those comparisons?

References

[1] Federal Communications Commission, "Protecting And Promoting The Open Internet, Report And Order On Remand, Declaratory Ruling, And Order, 30 FCC Rcd 5601," 2015.

[2] ——, "Restoring Internet Freedom, Declaratory Ruling, Report And Order, And Order, 33 FCC Rcd 311," 2018.

[3] ——, "Report On The Future Of The Universal Service Fund, Report, WC Docket No. 21-476," 2022.

[4] C. Gahnberg, N. de Guzman, A. Robachevsky, and A. Wan, "Internet impact brief: South korea's interconnection rules," *Internet Society*, 2022.

[5] Axon Partners Group, "Europe's Internet Ecosystem: Socio-economic Benefits Of A Fairer Balance Between Tech Giants And Telecom Operators," 2022.

[6] "BEREC preliminary assessment of the underlying assumptions of payments from large CAPs to ISPs," *Body of European Regulators for Electronic Communications*, 2022.

[7] A. Nikkhah and S. Jordan, "Analysis of the requirements of settlement-free interconnection policies," *IEEE Transactions on Network and Service Management*, 2023.

[8] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani, "How many tiers? pricing in the Internet transit market," in *Special Interest Group on Data Communication (SIGCOMM)*, 2011, pp. 194–205.

[9] A. Nikkhah and S. Jordan, "How can equitable peering be achieved between ISPs and content providers?" in *International Conference on Computer Communications and Networks (ICCCN)*, 2023.

[10] ——, "Are the settlement-free peering policy requirements for ISPs and CDNs based on network costs?" in *TPRC50: The 50th Research Conference on Communication, Information and Internet Policy*, 2022.

[11] ——, "A two-sided model of paid peering," *Telecommunications Policy*, vol. 46, no. 8, p. 102352, 2022.

[12] OpenVault, "Broadband Insights Report (OVBI)," *Retrieved from https://openvault.com. Accessed May 2, 2023.*

[13] A. Nikkhah and S. Jordan, "Requirements of settlement-free peering policies," in *IEEE Global Communications Conference (GLOBECOM)*, 2022, pp. 3617–3622.

[14] Information Technology and Innovation Foundation, "The State Of US Broadband In 2022: Reassessing The Whole Picture," *Retrieved from https://itif.org/publications/2022/12/05/state-of-us-broadband-in-2022-reassessing-the-whole-picture/. Accessed May 2, 2023.*

[15] Petroc Taylor, "Percentage Of Households With Internet Use In The United States From 1997 To 2020," *Retrieved from https://www.statista.com/statistics/189349/us-households-home-internet-connection-subscription/. Accessed May 2, 2023.*

[16] Pew Research Center, "Internet/Broadband Fact Sheet," *Retrieved from https://www.pewresearch.org/internet/fact-sheet/internet-broadband/. Accessed May 2, 2023.*

[17] The Internet & Television Association (NCTA), "Internet/Broadband Fact Sheet," *Retrieved from https://www.ncta.com/broadband-facts. Accessed May 2, 2023.*

[18] Comcast Cable Communications, "Explore Speeds And Prices," *Retrieved from https://www.xfinity.com/learn/internet-service. Accessed June 21, 2023.*

[19] Charter Communications, "Reliable Internet Plans And Pricing," *Retrieved from https://www.spectrum.com/internet/plans. Accessed Accessed June 21, 2023.*

[20] Cox Communications, "Cox internet plans and pricing," *Retrieved from https://www.cox.com/residential/internet. Accessed June 21, 2023.*

[21] Zippia, "33 trending streaming statistics," *Retrieved from https://www.zippia.com/advice/streaming-statistics. Accessed June 21, 2023.*

[22] Statista, "Cost of video streaming services in the U.S." *Retrieved from https://www.statista.com/statistics/831235/video-streaming-service-monthly-cost. Accessed June 21, 2023.*

[23] Sandvine, "Global Internet Phenomena," *Retrieved from https://www.sandvine.com/phenomena. Accessed May 2, 2023.*