This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2023.3254367

Analysis of the Requirements of Settlement-Free Interconnection Policies

Ali Nikkhah, *Student Member, IEEE* and Scott Jordan, *Member, IEEE*

***Abstract*—Peering between two networks may be either settlement-free or paid. In order to qualify for settlement-free peering, large Internet Service Providers (ISPs) require that peers meet certain requirements. However, the academic literature has not yet shown the relationship between these settlement-free peering requirements and the value to each interconnecting network. We develop two models to analyze the value to each network from the most common and important requirements in the United States. Large ISPs in the U.S. often require potential settlement-free peers to interconnect at a minimum of 6-8 locations. We find that there is a substantial benefit from this requirement to the ISP, but little incremental benefit from a larger number of interconnection points. Large ISPs often require that the ratio of incoming traffic to outgoing traffic remain below approximately 2:1. In the case of two interconnecting ISPs, we find that this requirement ensures a roughly equal exchange of value. We also 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, but a limit on the traffic ratio is irrational.**

***Index Terms*—Internet Interconnection; Peering; Net Neutrality**

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

We focus on peering. Historically, peering was principally used by Tier 1 networks. Peering may be either paid (i.e., one interconnecting network pays the other) or settlement-free (i.e., without payment). The conventional wisdom is that two Tier 1 networks agree to settlement-free peering if and only if the two networks perceive a roughly equal exchange of value from the peering arrangement. For example, if two Tier 1 networks are both ISPs with similar numbers of customers and similar size backbones, then they may perceive a roughly equal value from the exchange of traffic with destinations in their customer cones. Large ISPs often require that peers meet certain requirements, including a specified minimum number of interconnection points, a traffic ratio less than 2:1, and symmetric routing. The conventional wisdom is that these requirements are related to the perception of roughly equal value, but the academic literature has not yet established such a relationship.

More recently, it has become common for large ISPs and large content providers or content delivery networks (CDNs) to peer. However, there have often been disagreements between them over whether the peering arrangement should be settlement-free or paid. Large ISPs advertise the same settlement-free peering requirements for content providers as for ISPs. However, large content providers do not satisfy requirements about traffic ratios, and often are more inclined to use non-symmetric routing. The academic literature has not lent much insight into when settlement-free peering between an ISP and a content provider is appropriate.

In 2013-2014, a dispute between Comcast and Netflix over terms of interconnection went unresolved for a substantial period of time. In 2014, Netflix and a few transit providers brought the issue to the attention of the United States Federal Communications Commission (FCC), which was writing updated net neutrality regulations. 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" [1]. 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" [2].

The FCC discussed the dispute in the 2015 Open Internet Order [3], and asserted oversight over such interconnection arrangements. However, in 2018, the FCC reversed itself and ended its oversight of interconnection arrangements, when it repealed most of the 2015 net neutrality regulations [4]. It is almost certain that the FCC will revisit the issue in the next few years. Recently, the FCC noted that some stakeholders are advocating that the FCC require content providers to pay a fee based on download traffic that would be used to subsidize

---

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

Scott Jordan is a Professor of Computer Science at the University of California, Irvine. He served as the Chief Technologist of the Federal Communications Commission during 2014-2016. Email: sjordan@uci.edu. Webpage: www.ics.uci.edu/ sjordan/.

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

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

broadband Internet access in rural areas and for low income consumers [5]. Such advocates are using similar arguments that large ISPs have used to argue for paid peering.

Similarly, there are ongoing debates over paid peering in South Korean and in Europe. South Korea requires paid peering between ISPs interconnecting in the country, based on the volume of traffic exchanged. These paid peering fees have often been passed on to content providers interconnecting with ISPs in South Korea. South Korea is also currently considering a proposal to require content providers to pay usage fees to ISPs based on traffic volume [6]. The European trade association representing many of Europe's ISPs has similarly recently proposed that content providers should be required to pay usage fees to ISPs based on the traffic volume [7]. European regulators, however, are worried that such usage fees could be exploited by ISPs and are skeptical of arguments that ISPs' costs are not already properly covered by ISPs' customers [8].

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

The paper is organized as follows. In Section II, we summarize the settlement-free peering requirements of the ten largest ISPs in the United States. The four largest ISPs require interconnection at a minimum of 4 to 8 interconnection points from specified lists, that incoming and outgoing traffic be roughly balanced, and that the two parties use symmetric routing. The next six largest ISPs require interconnection at a specified minimum number of interconnection points, but often less than 4, and may or may not require roughly balanced traffic. We henceforth focus on the settlement-free peering requirements of the four largest ISPs.

In Section III, we summarize the relevant research literature. Although a number of papers discuss settlement-free peering requirements, few analyze the relationship between these requirements and network costs. The academic literature thus provides little insight into why large ISPs impose settlement-free peering requirements or how these requirements are related to either the ISP's network cost or its perception of value.

In Section IV, we develop two models. The first model is an analytical model in which an ISP serves the contiguous United States with a uniformly distributed population of subscribers. The ISP's network consists of access networks, middle mile networks, and a backbone, each with a regular geometry. The second model is a numerical model in which subscribers are distributed according to census statistics. Access networks are based on counties, and interconnection points are chosen from a list of the largest exchanges in the United States. Traffic matrices are based on population, and traffic-sensitive costs are modeled as a function of both distance and traffic volume.

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

In Section VI, we analyze settlement-free peering requirements about the number and location of interconnection points between two ISPs using hot potato routing. We consider the conjecture that such requirements are related to a perception of roughly equal value. When the traffic ratio is 1:1, we show that the ISP's cost is a uni-modal function of the number of interconnection points, and that there may be little value in requiring interconnection at more than 6 Internet Exchange Points (IXPs). The ISP's cost is typically minimized by selecting interconnection points that span the country and are near population centers.

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

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

A portion of the results in this paper were previously presented in [9] and [10]. The analytical model was introduced in [9] and the numerical model was introduced in [10]. This paper includes both models and compares the results obtained from each. In addition, this paper includes several extensions, including an analysis of multiple cost ratios, an analysis of the effect of the traffic ratio on cost using the analytical model, and an analysis of full content replication and no content replication.

II. SETTLEMENT-FREE INTERCONNECTION POLICIES

We studied the settlement-free peering policies[2] of the ten largest ISPs in the United States [11]–[20]. Table I summarizes the most relevant requirements[3] of these policies.

The column labeled "subscribers" gives an estimate of the number of subscribers of each ISP in 2021 [21], as settlement-free peering policies differ with the number of subscribers.

The column labeled "peering inclination" indicates the ISP's predisposition towards or against peering, as noted by PeeringDB [22]. The four largest ISPs (Comcast, Charter, AT&T, and Verizon) are either selective (i.e., have moderate requirements) or restrictive (i.e., have strong requirements). The next four largest ISPs (Cox, CenturyLink[4], Altice, and Frontier) are selective. The ninth and tenth largest ISPs (Mediacom and TDS Telecom) are open (i.e., have low requirements).

One of the most common requirements in settlement-free peering policies is a minimum number of mutually agreeable locations at which the two parties will agree to interconnect. Tier 1 ISPs typically interconnect in at least 8 of the 10 largest Internet Exchange Point (IXPs), as shown in the first ten rows of Table II. In this paper, we use the term IXP to refer to any interconnection point where ISPs and network operators exchange traffic. While the term IXP is sometimes used to refer to neutral third-party operated exchange points, we use it more broadly to encompass all types of interconnection points. The four largest ISPs each require interconnection at a minimum of 4 to 8 IXPs (see the column labeled "minimum number of IXPs" in Table I). The next four largest ISPs are varied in their characteristics. CenturyLink's and Frontier's backbone networks likely qualify them as Tier 1 networks, but they have significantly fewer subscribers than do the four largest ISPs. Cox's and Altice's backbone networks likely qualify them as Tier 2 networks. These four moderate-size ISPs require interconnection at a minimum of 2 to 6 IXPs. The ninth and tenth largest ISPs do not specify a minimum number. Amongst ISP that do specify a minimum number of mutually agreeable locations, most also require that these IXPs must be chosen from ISP-specified lists (see the column labeled "number of IXPs in ISP-specified list" in Table I), and often require that the chosen IXPs be geographically diverse.

Another common requirement in settlement-free peering policies is the minimum amount of traffic to be exchanged between the two networks. The three largest ISPs specify a minimum of 20 or 30 Gbps in the dominant direction (see the column labeled "minimum traffic volume" in Table I). Smaller ISPs generally specify lower traffic volume thresholds, if they have such a requirement. In addition, settlement-free peering policies often require an approximate balance between incoming and outgoing traffic. Some (e.g., AT&T and Verizon) require that the ratio of incoming traffic volume to outgoing traffic volume not exceed a specified threshold. Others (e.g., Comcast) do not specify a maximum traffic ratio, but instead state the exchanged traffic should be in general balance[5].

Finally, settlement-free peering policies almost always have requirements about routing policies. Some ISPs (e.g., Charter and Verizon) require both parties to use hot potato routing[6], while other ISPs (e.g., Comcast and AT&T) only require that the two parties either both use hot potato or both use cold potato routing.

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

III. 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 [22]. Lodhi et al. [23] 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 [24], who discuss the selection of IXPs in a few networks with idealized and regular topologies.

There is some work that discusses the presence of traffic ratio requirements in settlement-free peering policies. Faratin et al. [25] noted that already by 2007 large ISPs often included traffic ratio requirements in their settlement-free peering policies, and that such a requirement impacts the relative costs of the two interconnecting networks when using hot potato routing. However, they did not analyze how traffic ratios affect the costs of each party. Dhamdhere et al. [26] constructed a model of ISP revenue and costs and used the model to compare peering policies that maximize profit to those with traffic ratio requirements. However, they did not analyze the effect of the

---

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

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

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

[5] For example, Comcast requires that "[a]pplicant must maintain a traffic scale between its network and Comcast that enables a general balance of inbound versus outbound traffic. The network cost burden for carrying traffic between networks shall be similar to justify SFI." [11].

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

TABLE I: Settlement-free peering requirements of the ten largest ISPs in the United States

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

TABLE II: The largest IXPs at which some ISPs interconnect

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

traffic ratio upon costs, and thus were not concerned with relating traffic ratio requirements to the value of the peering agreement to the ISP.

Indeed, there is some work that is skeptical that traffic ratios relate to the benefits to each interconnecting party. Zarchy et al. [27] asserted that traffic ratio requirements do not have any relevance to the economic benefits of interconnection for each party, and that an ISP's profit may be increased by examining the benefit of each potential interconnection partner and of the potential locations of interconnection.

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. [28] discussed 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. Patchala et al. [29] analyzed the economics of direct peering arrangements between a content provider and an ISP as well as peering between a CDN and an ISP. They also analyzed how sending traffic through a CDN improves quality of service for end-users compared with sending traffic through transit providers. Lee et al. [30] analyzed the effects of CDNs on content providers and ISPs in the context of Internet traffic delivery. The authors modeled the CDN as a business-to-business platform that provides caching and other services between content providers and ISPs. While this paper does not directly address the topic of peering requirements, it provides important insights into the effects of CDN-mediated delivery on content providers and ISPs and highlights the importance of considering the impact of interconnection decisions on different parties in the network. Chang et al. [31] proposed 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 [32] 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.

IV. MODEL

In this section, we develop two models that will enable our analysis of settlement-free peering policies. The analysis, presented in later sections, examines the effect of routing policies, the number of IXPs at which interconnecting networks meet, and the traffic ratio between interconnecting networks on network cost. Thus, both models focus on the characteristics that we believe are most critical to this analysis, and abstract other less critical characteristics.[7]

The analytical model is designed to foster closed-form analysis. It makes very simplistic assumptions. An Internet Service Provider (ISP) is assumed to serve the United States, which is simplistically modeled as a rectangular region. Interconnection points (IXPs) are assumed to be equally spaced throughout the region. The network is partitioned into a backbone network, middle mile networks, and access networks, each of which is similarly uniform in shape. The population is assumed to be uniformly distributed. Traffic matrices are built using these assumptions.

The numerical model is designed to reflect key characteristics of the United States. Although an ISP is still assumed to serve the United States, we now consider the actual geography of the contiguous United States. We consider the actual geographic location of the largest IXPs in the United States. Middle mile and access networks are modeled based on the U.S. counties. The density of the population is drawn from U.S. census statistics.

Although the numerical model is a more accurate representation of topologies in the United States, the analytical model also adds value to the analysis. The analytical model provides a simpler and more general understanding of the peering policies, allowing for closed-form expressions to be derived for the cost model in situations where numerical simulations are difficult due to lack of public data. Additionally, the analytical model serves as a validation tool for the numerical model by identifying potential discrepancies or errors.

Subsection IV-A introduces the topology of an ISP's U.S. network. Subsection IV-B develops the traffic matrices over this network. To help readers easily refer to the symbols used in this paper, we provide a glossary of symbols in Table III. This table includes all the symbols used in the paper and their corresponding descriptions. Throughout the paper, we use the subscript 1 to denote variables pertaining to the analytical model, the subscript 2 to denote variables pertaining to the numerical model, the subscript $down$ to denote downstream traffic, the subscript up to denote upstream traffic, the subscript cp to denote the content provider, the superscript hot to denote hot potato routing, the superscript $cold$ to denote cold potato routing, the superscript no to denote no content replication, the superscript $full$ to denote full content replication, and the superscript $partial$ to denote partial content replication.

Table IV also shows a parameterization comparison between the numerical and analytical models.

TABLE III: Glossary of Symbols

Symbol	Description
$A_1(j,k)$	Geographical center of access network j,k
$A_2(j)$	Geographical center of access network j
$Access_1(j,k)$	Geographical region of access network j,k
$Access_2(j)$	Geographical region of access network j
a	Distance from west to east of access networks
b	Distance from south to north of access networks
C	Variable portion of the ISP's traffic-sensitive cost
c^a	Cost per unit distance and volume in access network
c^b	Cost per unit distance and volume in backbone network
c^m	Cost per unit distance and volume in middle mile network
D^a	Distance on ISP's access network
D^b	Distance on ISP's backbone network
D^m	Distance on ISP's middle mile network
ED^a	Average distance on ISP's access network
ED^b	Average distance on ISP's backbone network
ED^m	Average distance on ISP's middle mile network
i^p	The IXP at location IXP^p
$IXP(i)$	Location of IXP i
IXP^u	Location of the IXP closest to the end user
IXP^p	Location of the IXP at which traffic enters/exits the ISP
I	Set of locations of the IXPs
l^N	Set of N IXPs at which the ISP agrees to peer
L	Distance from west to east of the United States
N	Number of IXPs at which the ISP agrees to peer
p	Population of the contiguous United States
$P(j,k)$	Probability that an end user resides in access network (j,k)
p_j	Population of the county associated with access network j
r	Ratio of downstream traffic to upstream traffic
$R(IXP(i))$	Geographical region of IXP i's access networks
s_j	Size of county j
S	Traffic source's location
U	End user's location
US	ISP's service region
V	The volume of traffic
W	Distance from south to north of the United States

TABLE IV: Parameterization Comparison between Analytical and Numerical Models

Parameters	Analytical	Numerical
ISP's service region	Rectangular US	Contiguous US
IXP Locations	Equally spaced	Actual locations
Access Network Topology	Uniform	Modeled by county
Population Distribution	Uniform	U.S. census
Distribution of sources	Uniform	County population
Distribution of end users	Uniform	County population
Calculation of distances	Euclidean	Great-circle

A. Topology

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

1) Service Territory: While most ISPs do not offer residential broadband Internet access service over the entire contiguous United States, we see little in their settlement-free peering policies that are specific to their service territory other than a subset of the IXPs at which they peer that are concentrated near their service territory. Thus, in both the analytical and numerical model, we focus on a single ISP whose service territory covers the contiguous United States.

In the analytical model, the ISP's service region is simplistically modeled as a rectangular abstraction US_1 of the contiguous United States, measuring $L = 2800$ miles from west to east and $W = 1582$ miles from south to north [33].

---

[7] However, outside the United States, other models may be more appropriate given differences in network topology. While the trends identified in this paper are likely to be similar in a qualitative way, the specific results may differ quantitatively. In addition, it should be noted that the terminology used in the paper (access, middle mile, and backbone) may not be applicable in other regions. For example, in Europe, a different terminology may be used, and the relative importance of different parts of the network (such as access and middle mile) may be different.

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2023.3254367

We use a coordinate system (x, y) centered on this rectangle, i.e.

$$US_1 = \left[-\frac{L}{2}, \frac{L}{2}\right] \times \left[-\frac{W}{2}, \frac{W}{2}\right] \quad (1)$$

In the numerical model, the ISP's service region is modeled as the contiguous United States, denoted US_2, using real geographical data of its boundaries. We use a coordinate system with x and y measured in degrees of longitude and latitude, respectively.

Throughout the paper, we use the subscript 1 to denote variables pertaining to the analytical model, and the subscript 2 to denote variables pertaining to the numerical model.

2) Location of IXPs: In the analytical model, 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 ($i \in l_1^N = \{1, \dots, N\}$) by $IXP_1(i)$. We simplistically assume that these N IXPs are located at the middle latitude $y = 0$ and at equally spaced longitudes x, i.e.

$$IXP_1(i) = \left(-\frac{L}{2} + \frac{L(2i-1)}{2N}, 0\right) \quad (2)$$

We denote the set of locations of the IXPs at which the ISP interconnects with this interconnecting network by $I_1 = \{IXP_1(i), i \in l_1^N\}$.

In the numerical model, we use the actual geographic locations of the $M = 12$ largest IXPs in the United States, listed in Table II [34]–[39]. The coordinates of these M IXPs are denoted by $IXP_2(i)$ i $(i = 1, ..., M)$, and the set of these IXPs are denoted by I_2^M.

We note that the largest ISPs in the United States (Comcast, Charter, AT&T, and Verizon) each interconnect at a minimum of 9 of these 12 IXPs, although a smaller ISP (Cox) interconnects at fewer IXPs; see Table II.

An ISP and an interconnecting network often agree to interconnect at a smaller number $N < M$ of IXPs. We denote the set of N IXPs at which they agree to interconnect as $l_2^N \subseteq \{1, \dots, M\}$, and we denote the set of locations of these IXPs by $I_2^N = \{IXP_2(i), i \in l_2^N\} \subseteq I_2^M$.

3) Backbone Network, Middle Mile Networks, and Access Networks: We model the ISP's network as partitioned into a single backbone network, multiple middle mile networks, and multiple access networks.

In the analytical model, we simplistically 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 $j = 1, \dots, L/a$, and $k = 1, \dots, W/b$. We denote by $Access_1(j, k)$ the geographical region of access network (j, k). We denote the location of the geographical center of access network (j, k) by

$$A_1(j, k) = \left(-\frac{L}{2} + \frac{(2j-1)a}{2}, -\frac{W}{2} + \frac{(2k-1)b}{2}\right) \quad (3)$$

A middle mile link is assumed to run from the geographical center of each access network (j, k) to the closest IXP, and points of presence connect the middle mile networks with their corresponding access networks.

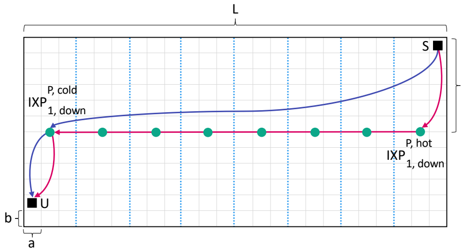

(a) Analytical Model

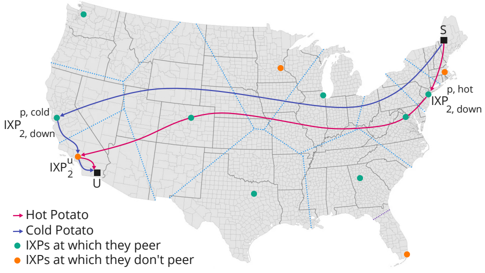

(b) Numerical Model

Fig. 1: Topology of an ISP's network

The backbone network is assumed to connect all of the IXPs at which the ISP is present. The IXPs thus serve both to offer interconnection between the ISP and other networks and to route traffic across the ISP's backbone network.

The IXPs can be used to partition the ISP's service territory into a set of regions closest to each IXP. In the analytical model, denote by $R_1(IXP_1(i))$ the geographical region that consists of the union of access networks for which the closest interconnection point is IXP i, namely

$$R_1(IXP_1(i)) = \bigcup_{\substack{(j,k) \mid \|A_1(j,k) - IXP_1(i)\| \leq \\ \|A_1(j,k) - IXP_1(i')\| \, \forall i' \in l_1^N}} Access_1(j, k) \quad (4)$$

Figure 1(a) illustrates these regions for the analytical model. Since the analytical model has a regular topology, each region is simply a rectangle:

$$R_1(IXP_1(i)) = \left[-\frac{L}{2} + \frac{(i-1)L}{N}, -\frac{L}{2} + \frac{iL}{N}\right] \times \left[-\frac{W}{2}, \frac{W}{2}\right] \quad (5)$$

In the numerical model, we model each access network as spanning a single U.S. county. While we recognize that topologies of access networks differ widely, this assumption will not significantly affect the results in this paper, since differences in network costs between various forms of peering depend more critically on the number of interconnection points than on the topologies of access networks. We index the access networks in an arbitrary order by j. We denote by $Access_2(j)$ the geographical region of access network j, and we denote by $A_2(j)$ the location of the geographical center of access network j. These locations are assigned to be the longitudes and latitudes of the center of each county in the contiguous United States [40].

In the numerical model, consider an ISP and an interconnecting network that agree to interconnect at the N IXPs

$l_2^N \subseteq \{1, \ldots, M\}$. For $i \in l_2^N$, denote by $R_2^N(IXP_2(i))$ the geographical region that consists of the union of access networks for which the closest IXP at which the ISP and the interconnecting network agree to peer is IXP i, namely

$$R_2^N(IXP_2(i)) = \bigcup_{\substack{j \,|\, \|A_2(j) - IXP_2(i)\| \leq \\ \|A_2(j) - IXP_2(i')\| \,\forall i' \in l_2^N}} Access_2(j) \quad (6)$$

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

B. Traffic Matrices

We now turn to modeling the traffic matrices over the ISP's network.

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

In the analytical model, we denote the probability that an end user resides within access network (j, k) by $P_1(j, k)$. We simply assume that end users are uniformly distributed across access networks, i.e. $P_1(j, k) = ab/LW$. We also simply assume that end users are uniformly distributed within each access network.

In the numerical model, we denote the probability that an end user resides within access network j by $P_2(j)$. We assume that end users are distributed across access networks according to the population of the county associated with the access network. We denote the population of the county associated with access network j by p_j, and we denote by $p = \sum_j p(j)$ the population of the contiguous United States. We assign these values using U.S. census data [41]. It follows that $P_2(j) = p_j/p$. We further assume that end users are uniformly distributed within each access network, and we denote the size of county j by s_j [42].

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 VIII.

In the analytical model for the ISP-ISP case, we assume that the distribution of the source S is identical to the distribution of end users, which is jointly given by $\{P_1(j, k)\}$ and the uniform distribution of end users within each access network. We assume that the source S and the end user U are independent.

In the numerical model for the ISP-ISP case, we make similar assumptions, i.e. that the distributions of S and U are both jointly given by $\{P_2(j)\}$ and the uniform distribution of end users within each access network, and that S and U are independent.

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

First, consider the analytical model. The end user U is uniformly distributed in US_1, as discussed above. The IXP closest to the end user is a deterministic function of U, namely $IXP_1^u = (g' | \, U \in R_1(g'))$.

However, the IXP at which downstream traffic enters the ISP's network ($IXP_{1,down}^p$) 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. $IXP_{1,down}^{p,hot} = (g | \, S \in R_1(g))$. Since end users are assumed to be uniformly distributed across the contiguous U.S., $IXP_{1,down}^{p,hot}$ is also uniformly distributed:

$$P(IXP_{1,down}^{p,hot} = g) = \sum_{Access_1(j,k) \subset R_1(g)} P_1(j, k) = \frac{1}{N} \quad (7)$$

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. $IXP_{1,down}^{p,cold} = IXP_1^u$.

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. $IXP_{1,up}^{p,hot} = IXP_1^u$. 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 (7).

We turn next to the numerical model. The access network of the end user is distributed according to $\{P_2(j)\}$ and the end user U is uniformly distributed within the access network, as discussed above. The IXP closest to the end user is a deterministic function of U, namely $IXP_2^u = (g' | \, U \in R_2^M(g'))$.

---

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

However, the IXP at which downstream traffic enters the ISP's network ($IXP^p_{2,down}$) depends on both the routing policy and the IXPs at which they agree to interconnect. Consider an ISP and an interconnecting network that agree to interconnect at the N IXPs in I^N_2. If the ISP and the interconnecting network use hot potato routing, then the IXP at which downstream traffic enters the ISP's network ($IXP^{p,hot}_{2,down}$) is independent of the end user, and it is the IXP closest to the source among the IXPs at which they agree to peer, i.e. $IXP^{p,hot}_{2,down} = (g|\, S \in R^N_2(g))$. For example, suppose S is in Maine and U is in Imperial county, California. Then, as illustrated in Figure 1(b), $IXP^{p,hot}_{2,down}$ might be in New York (if the two networks do not agree to peer in Boston) and IXP^u_2 is in Los Angeles. Since end users are assumed to be distributed according to U.S. county population statistics, $IXP^{p,hot}_{2,down}$ is distributed as:

$$\begin{aligned} P(IXP^{p,hot}_{2,down} = g) &= \sum_{Access_2(j) \subset R^N_2(g)} P_2(j) \\ &= \frac{1}{p} \sum_{Access_2(j) \subset R^N_2(g)} p(j) \end{aligned} \tag{8}$$

In contrast, if the ISP and the interconnecting network use cold potato routing, then the IXP at which downstream traffic enters the ISP's network ($IXP^{p,cold}_{2,down}$) is no longer independent of the end user, and it is the IXP closest to the end user at which they agree to peer, i.e. $IXP^{p,cold}_{2,down} = (g|\, U \in R^N_2(g))$. For example, suppose S is in Maine and U is in Imperial county, California. Then, as illustrated in Figure 1(b), $IXP^{p,cold}_{2,down}$ might be in San Jose (if the two networks do not agree to peer in Los Angeles) and IXP^u_2 is in Los Angeles.

For upstream traffic in the numerical model, the routes and distributions are again 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. $IXP^{p,hot}_{2,up} = (g|\, U \in R^N_2(g))$. 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 (8).

V. TRAFFIC-SENSITIVE COSTS

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

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

A. Distances

We first determine distances in the analytical model. We continue to focus on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. The distances on each section on the ISP's network depend on the joint distribution of $(IXP^p_{1,down}, IXP^u_1, U)$. All distances in the analytical model 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 ($IXP^p_{1,down}$) and the location of the IXP closest to the end user (IXP^u_1). We denote the distance on the ISP's backbone network between these two IXPs by $D^b_1(IXP^p_{1,down}, IXP^u_1) = \|IXP^p_{1,down} - IXP^u_1\|$. Denote by i^p the IXP at location $IXP^p_{1,down}$, i.e. $i^p = i|(IXP^p_{1,down} = IXP_1(i))$, and denote by i^u the IXP at location IXP^u_1, i.e. $i^u = i|(IXP^u_1 = IXP_1(i))$. The distance between two IXPs can be determined by their locations given in (2):

$$D^b_1(IXP^p_{1,down}, IXP^u_1) = \frac{L}{N}|i^p - i^u| \tag{9}$$

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_1) 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 $D^m_1(IXP^u_1, U) = \|IXP^u_1 - A_1(j,k)\|$, where $U \in R_1(IXP^u_1)$ and $(j,k)\,|\,(U \in Access_1(j,k))$. The distance can be determined by the locations of the IXP and the access network, given in (2)-(3):

$$\begin{aligned} &D^m_1(IXP^u_1, U) = \\ &\sqrt{\left(\frac{L(2i^u-1)}{2N} - \frac{(2j-1)a}{2}\right)^2 + \left(-\frac{W}{2} + \frac{(2k-1)b}{2}\right)^2} \end{aligned} \tag{10}$$

where $U \in R_1(IXP^u_1)$ and $(j,k)\,|\,(U \in Access_1(j,k))$.

The distance on the ISP's access network is a function of the location of the end user. We denote the distance on the ISP's access network by $D^a_1(U) = \|A_1(j,k) - U\|$, where $(j,k)\,|\,(U \in Access_1(j,k))$. The distance can be determined by the location of end user within the access network.

The distances in the numerical model can be similarly represented as $D^b_2(IXP^p_{2,down}, IXP^u_2) = \|IXP^p_{2,down} - IXP^u_2\|$, $D^m_2(IXP^u_2, U) = \|IXP^u_2 - A_2(j)\|$ (where $U \in R^M_2(IXP^u_2)$), and $D^a_2(U) = \|A_2(j) - U\|$, where $j\,|\,(U \in Access_2(j))$. All distances in the numerical model are great-circle distances between the corresponding points on a sphere, and are calculated using the Haversine formula. However, there are no closed form formulae for these distances.

B. Average Distances

An ISP's traffic-sensitive cost depends on the average distance of traffic on each segment of its network. As discussed above, the distances on each section of the ISP's network depend on the joint distribution of (IXP^p, IXP^u, U). This joint distribution was given in Section IV-B for both models, separately for hot potato routing and for cold potato routing.

We first determine average distances in the analytical model. We continue to focus on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. The distance on the ISP's backbone network is a function of $(IXP^{p}_{1,down}, IXP^{u}_{1})$. When hot potato routing is used, the IXP at which downstream traffic enters the ISP's network $(IXP^{p,hot}_{1,down})$ is independent of the end user and thus independent of the IXP closest to the end user (IXP^{u}_{1}). Thus, the average distance on the ISP's backbone network is:

$$ED^{b,hot}_{1,down} = \sum_{g \in I_1} \sum_{g' \in I_1} D^{b}_{1}(g, g') P(IXP^{p,hot}_{1,down} = g) P(IXP^{u}_{1} = g') \tag{11}$$

The distance $D^{b}_{1}(g, g')$ is given in closed form in (9). The probability distribution of $IXP^{p,hot}_{1,down}$ is given in (7), and the probability distribution of IXP^{u}_{1} is similarly uniformly distributed. We can use these results to give a closed-form expression:

Theorem 1:

$$ED^{b,hot}_{1,down} = \frac{L(N-1)(N+1)}{3N^2} \tag{12}$$

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 $(IXP^{p,cold}_{1,down})$ is the IXP closest to the end user, i.e. $IXP^{p,cold}_{1,down} = IXP^{u}_{1}$. Thus, the ISP does not carry traffic across its backbone, i.e.

$$ED^{b,cold}_{1,down} = 0 \tag{13}$$

The distance on the ISP's middle mile network is a function of (IXP^{u}_{1}, U). It is independent of the routing policy, and the average distance is:

$$ED^{m}_{1} = \sum_{g' \in I_1} \sum_{A_1(j,k) \subset R_1(g')} D^{m}_{1}(g', A_1(j,k)) P_1(j,k) \tag{14}$$

The distance $D^{m}_{1}(g', A_1(j,k))$ is given in closed form in (10). Also, $P_1(j,k) = ab/LW$. We can use these results to give a closed form expression:

Theorem 2:

$$ED^{m}_{1} = \frac{abN}{LW} \sum_{k=1}^{\frac{W}{b}} \sum_{j=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2j-1)a - \frac{L}{N}}{2}\right)^2 + \left(\frac{(2k-1)b - W}{2}\right)^2} \tag{15}$$

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:

$$ED^{a}_{1} = \frac{1}{ab} \int_{-\frac{a}{2}}^{\frac{a}{2}} \int_{-\frac{b}{2}}^{\frac{b}{2}} \sqrt{x^2 + y^2}\, dy\, dx \tag{16}$$

It can be shown that:

Theorem 3:

$$ED^{a}_{1} = \frac{1}{12ab} \left[a^3 \sinh^{-1}\left(\frac{b}{a}\right) + b^3 \sinh^{-1}\left(\frac{a}{b}\right) + 2ab\sqrt{a^2 + b^2} \right] \tag{17}$$

The proof can be found in the Appendix.

We next determine average distances in the numerical model. We continue to focus on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. The distance on the ISP's backbone network is a function of $(IXP^{p}_{2,down}, IXP^{u}_{2})$.

As in the analytical model, when hot potato routing is used, the IXP at which downstream traffic enters the ISP's network $(IXP^{p,hot}_{2,down})$ is independent of the end user and thus independent of the IXP closest to the end user (IXP^{u}_{2}). However, whereas in the analytical model the IXP at which downstream traffic enters the ISP's network depends only on the routing policy, in the numerical model it also depends on the IXPs at which they agree to interconnect. Consider an ISP and an interconnecting network that agree to interconnect at the N IXPs in I^{N}_{2}. The average distance on the ISP's backbone network is:

$$ED^{b,hot}_{2,down} = \sum_{g \in I^{N}_{2}} \sum_{g' \in I^{M}_{2}} D^{b}_{2}(g, g') P(IXP^{p,hot}_{2,down} = g) P(IXP^{u}_{2} = g') \tag{18}$$

The probability distribution of $IXP^{p,hot}_{2,down}$ was given in (8). The probability distribution of IXP^{u}_{2} can be similarly represented as:

$$\begin{aligned} P(IXP^{u}_{2} = g') &= \sum_{Access_2(j) \subset R^{M}_{2}(g')} P_2(j) \\ &= \frac{1}{p} \sum_{Access_2(j) \subset R^{M}_{2}(g')} p(j) \end{aligned} \tag{19}$$

When cold potato routing is used, the IXP at which downstream traffic enters the ISP's network $(IXP^{p,cold}_{2,down})$ is the IXP closest to the end user at which they agree to peer. In addition, $IXP^{p,cold}_{2,down}$ is no longer independent of IXP^{u}_{2}. Since they might not agree to peer at all IXPs, the ISP might still carry traffic across a portion of its backbone, namely from $IXP^{p,cold}_{2,down}$ to IXP^{u}_{2}, and the average such distance is:

$$ED^{b,cold}_{2,down} = \sum_{g' \in I^{M}_{2}} D^{b}_{2}(g \mid g' \in R^{N}_{2}(g), g') P(IXP^{u}_{2} = g') \tag{20}$$

The distance on the ISP's middle mile network is a function of (IXP^{u}_{2}, U). It is independent of the routing policy, and the average distance is:

$$ED^{m}_{2} = \sum_{g' \in I^{M}_{2}} \sum_{A_2(j) \subset R^{M}_{2}(g')} D^{m}_{2}(g', A_2(j)) P_2(j) \tag{21}$$

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

$$ED^{a}_{2} = \sum_{j} \frac{p_j}{ps_j} \int_{U \in Access_2(j)} D^{a}_{2}(U) \tag{22}$$

C. Cost

The ISP incurs a traffic-sensitive cost for carrying traffic over the average distances calculated in the previous subsection. We only consider here traffic-sensitive costs, because non-traffic-sensitive costs do not vary with routing policies, the number of interconnection points, or traffic ratio.[9]

Traffic-sensitive costs are a function of both distance and traffic volume. We assume here that traffic-sensitive costs are linearly proportional to the average distance over which the traffic is carried on each portion of the ISP's network, see e.g., [43]. 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 $V\left(c^b ED^b + c^m ED^m + c^a ED^a\right)$.

In the analysis below, we fix the source-destination traffic matrix, and we consider the effect of changes in the number of IXPs at which peering occurs, routing policies, and the traffic ratio.

In the analytical model, given a fixed source-destination traffic matrix, the average distance across the ISP's access networks in (17) 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:

$$C_1 = c^b\left(ED_1^b + \frac{c^m}{c^b}ED_1^m\right)V \tag{23}$$

Below we consider the effect on the variable traffic-sensitive cost (C_1) of changes in the number of IXPs at which peering occurs, routing policies, and the traffic ratio, for constant c^b and different ratios of c^m/c^b. (In the remainder of the paper, we use the term *cost* to refer to the variable traffic-sensitive cost.) We will find that changes in the number of IXPs, routing policies, and the traffic ratio all affect ED_1^b and/or ED_1^m.

In the numerical model, given a fixed source-destination traffic matrix, the average distance across the ISP's access networks in (22) is similarly constant. In addition, the average distance across the ISP's middle mile networks in (21) is constant, once we fix $M = 12$, since the IXPs at which the parties agree to peer do not affect the middle mile. The variable portion of the ISP's traffic-sensitive cost is thus only:

$$C_2 = c^b ED_2^b V \tag{24}$$

Below we consider the effect on the variable traffic-sensitive cost (C_2) of changes in the number of IXPs at which peering occurs, routing policies, and the traffic ratio, for constant c^b. We will find that changes in the number of IXPs at which peering occurs, routing policies, and the traffic ratio all affect ED_2^b.

VI. NUMBER AND LIST OF IXPS

With both the analytical and numerical models in place, we now turn to analyzing the effect on an ISP's variable traffic-sensitive costs of the number of IXPs at which peering occurs, routing policies, and the traffic ratio. We are in particular interested in explaining the settlement-free peering policies of large ISPs, why large ISPs require settlement-free peers to meet at a minimum of 4-8 IXPs, and also why these IXPs are geographically distributed across the country.

In this section, we examine the effect of the number of IXPs at which two ISPs agree to peer, and the list of IXPs at which they peer. As shown in Table 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. In addition, large ISPs often specify a list of eligible IXPs that this minimum must be chosen from. The academic literature provides little insight into why large ISPs require interconnection at a minimum specified number of IXPs, nor why they require that they be selected from a list of eligible IXPs.

In order to better understand these requirements of settlement-free peering policies, we initially focus only on downstream traffic that originates outside the ISP's network and terminates at an end user on the ISP's network. We first assume that both ISPs use hot potato routing.

Using the analytical model, the cost of downstream traffic is:

$$C_{1,down}^{hot} = c^b V_{down}\left(ED_{1,down}^{b,hot} + \frac{c^m}{c^b}ED_1^m\right) \tag{25}$$

Substituting the expressions we previously found for $ED_{1,down}^{b,hot}$ in (12) and for ED_1^m in (15), we obtain:

$$C_{1,down}^{hot} = c^b V_{down}\Bigg(\frac{L(N-1)(N+1)}{3N^2} + \frac{c^m}{c^b}\frac{abN}{LW} \sum_{m=1}^{\frac{W}{b}}\sum_{n=1}^{\frac{L}{aN}}\sqrt{\left(\frac{(2n-1)a-L/N}{2}\right)^2 + \left(\frac{(2m-1)b-W}{2}\right)^2}\Bigg) \tag{26}$$

Figure 2(a) shows the effect of number of interconnection points (N) on the cost of downstream traffic ($C_{1,down}^{hot}$), when $c^m/c^b = 1.5$. (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.

The average distance the ISP carries traffic across its backbone ($ED_{1,down}^{b,hot}$) 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_1^m) is a complicated function of N. However, as we see in Figure 2(a), the middle mile cost is decreasing and convex with the number of interconnection

---

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

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2023.3254367

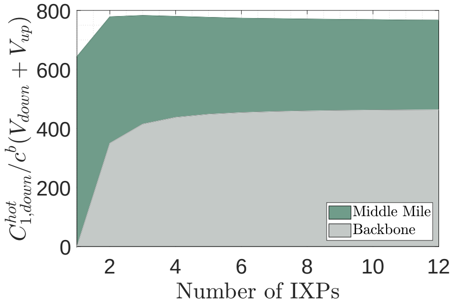

(a) $c^m/c^b = 1.5$

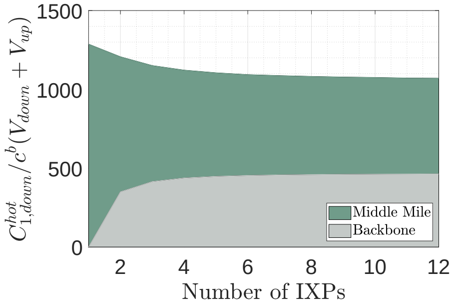

(b) $c^m/c^b = 3$

Fig. 2: Downstream costs (analytical model)

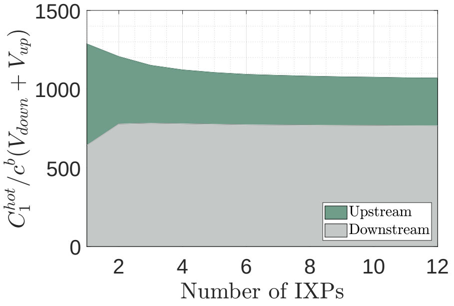

(a) $c^m/c^b = 1.5$

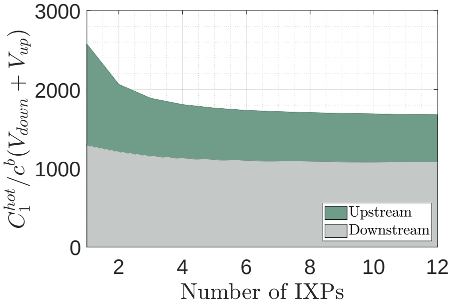

(b) $c^m/c^b = 3$

Fig. 3: Total costs (analytical model)

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.

In the analytical model, therefore, increasing the number of interconnection points increases backbone cost and decreases middle mile cost. The variation of the sum of these two costs with the number of interconnection points depends on the ratio of the per unit costs of the backbone and the middle mile. For relatively small values of c^m/c^b (e.g. 1.5), the cost of downstream traffic is a uni-modal function of N. For relatively large values of c^m/c^b, the cost of downstream traffic is decreasing with N, since the middle mile costs dominate; an example is shown in Figure 2(b) when $c^m/c^b = 3$.[10]

However, when two ISPs peer, there is also upstream traffic. Using the analytical model, the cost of upstream traffic using hot potato routing is:

$$C^{hot}_{1,up} = c^b V_{up}\left(ED^{b,hot}_{1,up} + \frac{c^m}{c^b} ED^m_1\right) \tag{27}$$

As we discussed in Section IV-B2, 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_1) 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 (13), we know that:

$$ED^{b,hot}_{1,up} = ED^{b,cold}_{1,down} = 0 \tag{28}$$

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_1 in (15), we obtain:

$$C^{hot}_{1,up} = c^m V_{up}\frac{abN}{LW} \sum_{m=1}^{\frac{W}{b}} \sum_{n=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2n-1)a - L/N}{2}\right)^2 + \left(\frac{(2m-1)b - W}{2}\right)^2} \tag{29}$$

Figure 3(a) shows the effect of the number of interconnection points (N) on the cost of both downstream ($C^{hot}_{1,down}$) and upstream ($C^{hot}_{1,up}$) traffic, when $c^m/c^b = 1.5$ and when there is an equal amount of downstream and upstream traffic (i.e., $V_{down} = V_{up}$). Figure 3(b) shows the same effect when the cost ratio $c^m/c^b = 3$.

We observe that, when the traffic ratio is 1:1, the cost is decreasing and convex with the number of interconnection points. We also observe that, over a wide range of cost ratios, 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.

---

[10] We use values of 1.5 and 3 for the parameter c^m/c^b in our analytical model to analyze the effect of interconnection points on downstream traffic cost. However, due to the lack of publicly available data on network cost structures, determining a representative value can be challenging.

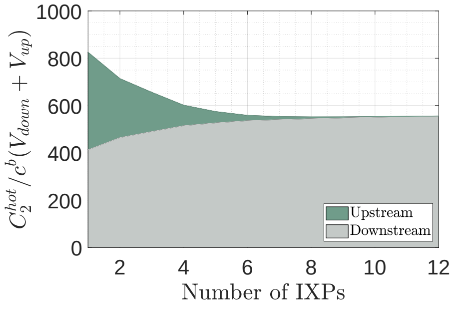

Fig. 4: Total costs (numerical model)

We turn next to our numerical model, which we expect to be more accurate, albeit without closed-form expressions. The cost of downstream traffic is:

$$C^{hot}_{2,down} = c^b V_{down} ED^{b,hot}_{2,down} \tag{30}$$

where $ED^{b,hot}_{2,down}$ is given in (18).

The cost of upstream traffic is:

$$C^{hot}_{2,up} = c^b V_{up} ED^{b,hot}_{2,up} \tag{31}$$

where $ED^{b,hot}_{2,up} = ED^{b,cold}_{2,down}$, which is given in (20).

Figure 4 shows the effect of the number of interconnection points at which they peer (N) on the cost of both downstream and upstream traffic using hot potato routing. The cost ratio c^m/c^b is no longer relevant, as only the backbone traffic is affected. However, the number of interconnection points affects the backbone cost. The cost is the function of $\|IXP^{p,hot}_{2,down} - IXP^u_2\|$ and $\|IXP^{p,hot}_{2,up} - IXP^u_2\|$ for all IXPs. As the number of IXPs at which they peer increases, the IXP closest to the end user IXP^u_2 is fixed since it is related to the location of the end user and a deterministic function of U, namely $IXP^u_2 = (g'| U \in R^M_2(g'))$. However, $IXP^{p,hot}_{2,down}$ and $IXP^{p,hot}_{2,up}$ change.

IXP^s does depend on the number of IXPs at which they peer. The IXP at which traffic enters the ISP's network is the IXP closest to the source among the IXPs at which they agree to peer, i.e. $IXP^{p,hot}_{2,down} = (g| S \in R^N_2(g))$, and thus as the number of IXPs at which they peer increases, $IXP^{p,hot}_{2,down}$ moves farther from IXP^u_2 and $\|IXP^{p,hot}_{2,down} - IXP^u_2\|$ increases. Thus, the downstream cost increases. The downstream cost is concave, because the incremental distance from $IXP^{p,hot}_{2,down}$ to IXP^u_2 associated with adding another IXP decreases, namely there are decreasing returns.

However, at the same time, the ISP exchanges upstream traffic using hot potato at the IXP closest to the end user at which the two ISPs agree to peer, i.e. $IXP^{p,hot}_{2,up} = (g| U \in R^N_2(g))$. Thus as the number of IXPs at which they peer increases, $IXP^{p,hot}_{2,up}$ moves closer to IXP^u_2 and $\|IXP^{p,hot}_{2,up} - IXP^u_2\|$ decreases. Therefore, the upstream cost decreases. The upstream cost is convex, because the incremental distance from IXP^u_2 to $IXP^{p,hot}_{2,up}$ associated with adding another IXP decreases, namely there are decreasing returns.

When the traffic ratio is 1:1, the decrease in the upstream cost exceeds the increase in the downstream cost. In the upstream route, $IXP^{p,hot}_{2,up}$ is the closest IXP to IXP^u_2, whereas, in the downstream route, $IXP^{p,hot}_{2,down}$ could be any IXP (including the closest or farthest IXP from IXP^u_2). Thus, as the number of IXPs increases, the absolute value of the slope of $\|IXP^{p,hot}_{2,up} - IXP^u_2\|$ in the upstream route is higher than the slope of $\|IXP^{p,hot}_{2,down} - IXP^u_2\|$ in the downstream route. It follows that the total cost decreases.

There are some differences between our results in our analytical and numerical models. The first difference is the middle mile cost. In our analytical model, the IXPs' locations vary with the number of interconnection points (see (2)), and thus the middle-mile cost, which varies with the distance between the access network and the closest IXP, decreases as the number of IXPs increases. However, in our numerical model, the number of IXPs at which both parties agree to interconnect does not affect the location of the IXP closest to the access network. Therefore, for fixed M, the middle mile cost is fixed. In Figure 3 for the analytical model, the upstream cost decreases but it does not decrease to zero, since the cost of the middle mile does not decrease to zero. However, in Figure 4 for the numerical model, the variable traffic-sensitive upstream cost decreases to zero, since the middle mile cost is fixed and the variable traffic-sensitive upstream backbone cost is zero at $N = M$, since there is no local delivery and thus no need to transfer traffic between IXPs.

The second difference is the traffic matrix. In our analytical model, we assume that incoming traffic is uniformly distributed over IXPs and that the population is distributed uniformly over the ISP's network. However, in our numerical model, we assume that the distribution of incoming traffic is related to the population of the region closest to an IXP and that end users are distributed across access networks according to the population of the county associated with the access network. Therefore, the slope decreases more quickly at a small number of IXPs (i.e., the second derivative is higher) in the numerical model than in the analytical model, because in the numerical model the ISP chooses an incremental IXP based on how much population would be affected by choosing that IXP.

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

We also wish to examine why ISPs require that the IXPs at which the two parties peer be selected from a specified list. To answer this question, our model selects the N IXPs at which to peer, from the list of $M = 12$ IXPs given in Table II, so as to minimize its cost (C^{hot}_2):

$$I^N_2 = \arg\min_{I^N_2} c^b \left(V_{down} ED^{b,hot}_{2,down} + V_{up} ED^{b,hot}_{2,up} \right) \tag{32}$$

The cost is typically minimized by selecting IXPs that span the country, so that the average distances the ISP carries traffic across its backbone are relatively small. Furthermore, when selecting a moderate or large number of IXPs, the cost is typically minimized by selecting more IXPs near where there are higher populations.

Comcast not only requires that potential settlement-free peering partners agree to peer at a minimum of 4 IXPs from Comcast's list of IXPs, it also requires that at least 1 of these 4 be on the west coast, that at least 1 be on the east coast, and that at least 1 be in a central region [11].[11] For $N = 4$, our numerical model chooses IXPs in Ashburn, Chicago, Los Angeles, and Atlanta, i.e. 1 on the west coast, 2 on the east coast, and 1 in the middle. All 4 of these cities are on Comcast's list.

Charter not only requires that potential settlement-free peering partners agree to peer at a minimum of 6-8 IXPs from Charter's list of IXPs[12], it also requires that at least 2 of these be in an eastern region, at least 2 be in a western region, and at least 2 be in a central region [12].[13] For $N = 8$, our numerical model chooses 4 on the east coast (Ashburn, New York, Miami, and Atlanta), 2 on the west coast (Los Angeles and Seattle), and 2 in the middle (Chicago and Dallas). All 8 of these cities are on Charter's list.

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

VII. TRAFFIC RATIO

In this section, we examine the effect of the traffic ratio on the variable traffic-sensitive cost. Large ISPs often require that the ratio of incoming traffic to outgoing traffic remain below approximately 2:1. In the case of two interconnecting ISPs, we find that this requirement ensures a roughly equal exchange of value. Two networks will agree to settlement-free peering if and only if the arrangement is superior for both parties compared to alternative arrangements including paid peering and transit. The conventional wisdom is that settlement-free peering thus occurs if and only if the two parties perceive that they are gaining an approximately equal value from the arrangement. Furthermore, the conventional wisdom, when the two parties are both Tier 1 networks, is that the perceived value is related to the size of each network, the number of customers of each party, and the ratio of traffic exchanged in each direction.

Indeed, the settlement-free peering policies of large ISPs often place limits on the ratio of downstream traffic to upstream traffic. AT&T's settlement-free peering policy requires that this traffic ratio not exceed 2:1, and Verizon's settlement-free peering policy requires that this traffic ratio not exceed 1.8:1.[14]

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

When using hot potato routing, the cost (C^{hot}) was plotted in Figures 3 and 4 for a traffic ratio of 1. For general traffic ratios, using the analytical model, we can derive the cost from (25), (27), and (28), as:

$$\begin{aligned} C_1^{hot} &= C_{1,down}^{hot} + C_{1,up}^{hot} \\ &= c^b (V_{down} + V_{up}) \frac{r ED_{1,down}^{b,hot} + \frac{c^m}{c^b}(r+1) ED_1^m}{r+1} \end{aligned} \tag{33}$$

where $ED_{1,down}^{b,hot}$ and ED_1^m are given in (12) and (15).

Figure 5(a) shows the effect of the traffic ratio on the cost, for various traffic ratios, when $c^m/c^b = 1.5$. The traffic ratio determines the trade-off between the downstream and upstream costs. For traffic ratios at or below 2:1, the variation in the upstream cost with the number of IXPs dominates, and thus interconnecting at 6 to 8 IXPs results in close to a minimum total cost. Requiring interconnection at more than 8 interconnection points is of little incremental value. However, for traffic ratios above 2:1, the variation in the downstream cost with the number of IXPs dominates, and it is rational for the ISP to not agree to settlement-free peering. The traffic ratios at which an ISP will perceive approximately equal value from peering depends on the difference in value it is willing to accept, and the alternatives it has to deliver and receive traffic.

Figure 5(b) shows the effect of the traffic ratio on the cost, for various traffic ratios, when the cost ratio is higher, at $c^m/c^b = 3$. For any traffic ratio, the variation in the middle mile cost dominates the backbone cost, and thus both downstream and upstream costs decrease as the number of IXPs increases. Therefore, the cost decreases as the number of IXPs increases for all plotted traffic ratios. As was true for a lower cost ratio, the requirement to peer at 6 to 8 interconnection points is reasonable, and requiring peering at more than 8 IXPs is of little incremental value.

Switching to the numerical model, we can derive the cost from (30) and (31), as:

$$\begin{aligned} C_2^{hot} &= C_{2,down}^{hot} + C_{2,up}^{hot} \\ &= c^b (V_{down} + V_{up}) \frac{r ED_{2,down}^{b,hot} + ED_{2,up}^{b,hot}}{r+1} \end{aligned} \tag{34}$$

Figure 6 shows the effect of the traffic ratio on the cost, for various traffic ratios using the numerical model. For traffic ratios at or below 2:1, the decrease in the upstream cost with the number of IXPs dominates the corresponding increase in the downstream cost, since the decrease in upstream cost due to *all* traffic exiting the ISP's network at a closer IXP outweighs the relatively small increase in downstream cost

---

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

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

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

[14] Some ISPs' settlement-free peering policies express an expectation of approximately equal value, but may not specify what this means, see e.g. Comcast [11].

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2023.3254367

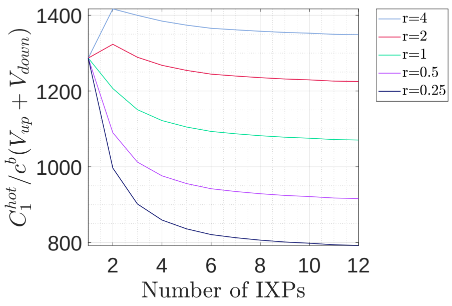

(a) $c^m/c^b = 1.5$

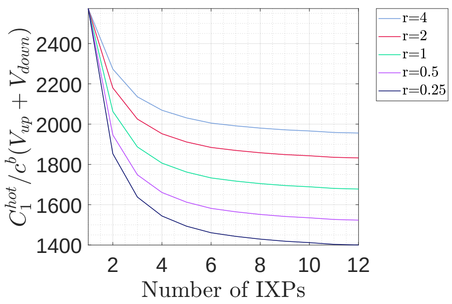

(b) $c^m/c^b = 3$

Fig. 5: Total costs for various traffic ratios (analytical model)

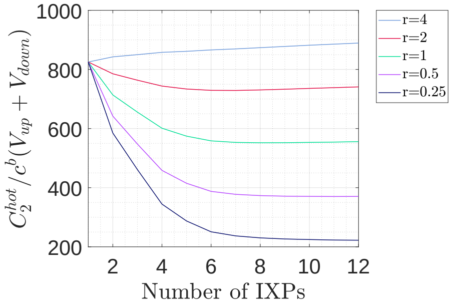

Fig. 6: Total costs for various traffic ratios (numerical model)

due to *some* traffic entering the ISP's network at a further away IXP. Recall that when the traffic ratio is 1:1, there is less than a 2% difference in the cost between $N = 6$ and the N at which cost is minimized, so there may be little value in requiring interconnection at more than 6 IXPs. We now find that when the traffic ratio is 0.5:1, there is less than a 2% difference in the cost between $N = 7$ and the N at which cost is minimized, and when the traffic ratio is 2:1, there is less than a 2% difference in the cost between $N = 4$ and the N at which cost is minimized.

In contrast, when the traffic ratio is 4:1, the increase in the downstream cost dominates the decrease in the upstream cost, since the downstream traffic volume is 4 times higher than the upstream traffic volume. As a result, the total cost increases with the number of IXPs at which they peer, and thus it is no longer rational for the ISP to agree to settlement-free peering.

In conclusion, the traffic ratios at which an ISP will perceive approximately equal value from peering depends on the difference in value it is willing to accept, and the alternatives it has to deliver and receive traffic. However, based on both models, we would expect the maximum acceptable traffic ratio to be 2:1 or less. Indeed, we observe that amongst the four largest ISPs, one specifies a maximum traffic ratio of 2:1, one specifies a maximum traffic ratio of 1.8:1, and one requires a "general balance" of traffic. In addition, we observe that for traffic ratios at or below 2:1, it remains rational to require interconnection at a minimum of 6-7 IXPs.

VIII. PEERING BETWEEN A CONTENT PROVIDER AND AN ISP

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

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

A. No Content Replication

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

The ISP network topology remains the same as was presented in Section IV-A. The distribution of the location of end users remains the same as was presented in Section IV-B1. For downstream traffic, we assume that the location of the content requested by an end user is independent of the location of the end user. We also assume that the distribution of the location of the content requested by an end user, namely the distribution of the source S, is identical to the distribution of end users. We further assume that the content provider uses hot potato routing. The routing of downstream traffic is thus identical to that considered in Section IV-B2. As a result, the distances that the ISP carries downstream traffic from a content provider across the ISP's network remain the same as was presented in Section V.

However, whereas in the ISP-ISP case there was both downstream and upstream traffic, we assume that in the CP-ISP case the volume of upstream traffic is negligible. As a

This article has been accepted for publication in IEEE Transactions on Network and Service Management. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/TNSM.2023.3254367

result, the ISP's costs are those discussed in Section V-C, but only for downstream traffic. Equivalently, we can think of this CP-ISP case as being equivalent to an ISP-ISP case with hot potato routing and a traffic ratio of infinity. It follows that the ISP's total cost in this CP-ISP case is the same as the ISP's downstream cost in the ISP-ISP case. In the analytical model, this downstream cost $C^{no}_{1,cp}$ was given by (25)-(26), and in the numerical model, this downstream cost $C^{no}_{2,cp}$ was given by (30).

The effect of the number of interconnection points (N) on the ISP's downstream cost using the analytical model is thus illustrated in figure 2. There is little, if any, decrease in the ISP's downstream cost as the number of interconnection points increases. The equivalent effect using the more accurate numerical model is illustrated as the downstream curve in figure 4. Here, the cost is minimized when $N = 1$.

We conclude an ISP has little incentive to peer at multiple IXPs with a content provider that does not replicate content and that uses hot potato routing. This is not surprising, since as we discussed at the end of Section VII, it is not rational for an ISP to agree to settlement-free peering with another ISP when the traffic ratio of downstream to upstream traffic is high.

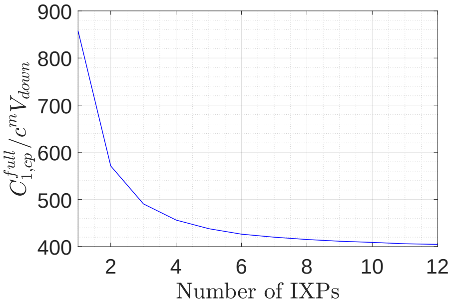

(a) Analytical Model

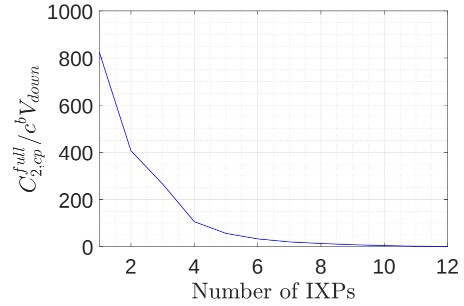

(b) Numerical Model

Fig. 7: Costs under complete replication

B. Full Content Replication

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

The ISP network topology remains the same as was presented in Section IV-A. The distribution of the location of end users remains the same as was presented in Section IV-B1. However, the location of the content is no longer the same as in previous sections. We now assume that the location of the content requested by an end user is the content server located at the IXP closest to the end user at which the content provider and the ISP agree to peer, i.e. $IXP^{p,full}_{1,cp} = (g|\, U \in R_1(g))$ in the analytical model and $IXP^{p,full}_{2,cp} = (g|\, U \in R^N_2(g))$ in the numerical model. Routing is now irrelevant, since the content is entering the ISP's network directly from the content server. The distances that the ISP carries downstream traffic from a content provider across the ISP's network are the same as the distances in the ISP-ISP case that an ISP carries downstream traffic when using cold potato routing, which were given in Section V-B. We again assume that the volume of upstream traffic is negligible. Equivalently, we can think of this CP-ISP case as being equivalent to an ISP-ISP case with cold potato routing and a traffic ratio of infinity.

As a result, the ISP's cost in this CP-ISP case using the analytical model is:

$$C^{full}_{1,cp} = c^b V_{down}\left(ED^{b,cold}_{1,down} + \frac{c^m}{c^b} ED^m_1\right) \tag{35}$$

Because the content provider provides all content to the ISP at the IXP nearest to the end user, $ED^{b,cold}_{1,down} = 0$. Substituting the expression we previously found for ED^m_1 in (15), we obtain:

$$C^{full}_{1,cp} = c^m V_{down}\frac{abN}{LW} \sum_{m=1}^{\frac{W}{b}}\sum_{n=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2n-1)a - L/N}{2}\right)^2 + \left(\frac{(2m-1)b - W}{2}\right)^2} \tag{36}$$

Similarly, using the numerical model, the ISP's cost in this CP-ISP case is:

$$C^{full}_{2,cp} = c^b V_{down} ED^{b,cold}_{2,down} \tag{37}$$

Substituting the expression we previously found for $ED^{b,cold}_{2,down}$ in (20), we obtain:

$$C^{full}_{2,cp} = c^b V_{down} \sum_{g' \in I^M_2} D^b_2(g \mid g' \in R^N_2(g), g') P(IXP^u_2 = g') \tag{38}$$

The effect of the number of interconnection points (N) on the ISP's downstream cost using both models is illustrated in figure 7. The cost decreases as the number of IXPs increases. Note that the ISP has this incentive to increase the number of IXPs at which the two parties peer *despite* the fact that the traffic ratio is infinity, unlike in the ISP-ISP case. However, there is less than a 2% difference in the cost between $N = 9$ and $N = 12$ in both models, so this indicates there may be little value in requiring interconnection at more than 9 IXPs.

Using the analytical model, we can compare the cost in the ISP-ISP case (illustrated in figure 3) to the cost in this CP-ISP case (illustrated in figure 7(a)). In the ISP-ISP case, we found that there is little value in requiring interconnection at more than 8 IXPs. In this CP-ISP case, we found that there is little value in requiring interconnection at more than 9 IXPs. The numerical model shows a similar but more pronounced pattern. Comparing the cost in the ISP-ISP case (illustrated in figure 4) to the cost in this CP-ISP case (illustrated in figure 7(b)), we find that there is little value in requiring interconnection at more than 6 IXPs in the ISP-ISP case, but there is a significant incremental value in the CP-ISP case to increasing the number of IXPs to at least 9.

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

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

C. Partial Content Replication

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

The ISP network topology remains the same as was presented in Section IV-A, and the distribution of the location of end users remains the same as was presented in Section IV-B1. 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.

For the analytical model, the ISP's cost in this CP-ISP case is:

$$C^{partial}_{1,cp} = xC^{full}_{1,cp} + (1-x)C^{no}_{1,cp} = c^bV_{down}\left(xED^{b,cold}_{1,down} + (1-x)ED^{b,hot}_{1,down} + \frac{c^m}{c^b}ED^m_1\right) \tag{39}$$

where $ED^{b,cold}_{1,down}$, $ED^{b,hot}_{1,down}$ and ED^m_1 are given in (28), (12), and (15) respectively.

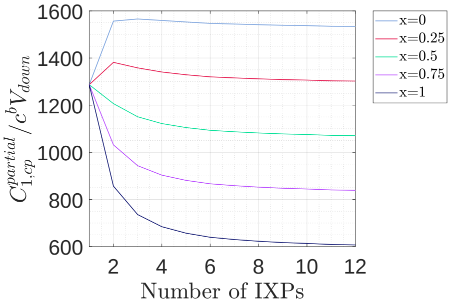

(a) Analytical ($c^m/c^b = 1.5$)

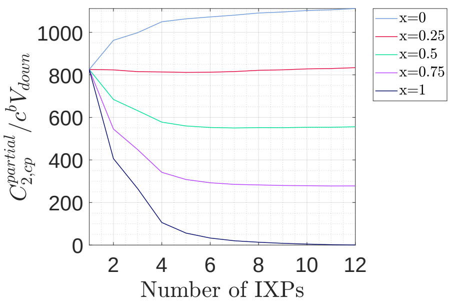

(b) Numerical Model

Fig. 8: Costs under partial replication

Figure 8(a) shows the effect of the number of interconnection points (N) on the ISP's downstream cost, for various values of the proportion x, when $c^m/c^b = 1.5$. 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 the two parties agree to peer.

We next determine the cost in the numerical model. The ISP's cost in this CP-ISP case is:

$$\begin{aligned} C^{partial}_{2,cp} &= xC^{full}_{2,cp} + (1-x)C^{no}_{2,cp} \\ &= c^bV_{down}\left(xED^{b,cold}_{2,down} + (1-x)ED^{b,hot}_{2,down}\right) \end{aligned} \tag{40}$$

where $ED^{b,hot}_{2,down}$ and $ED^{b,cold}_{2,down}$ are given in (18) and (20).

The effect of the number of interconnection points at which they agree to peer (N) on the ISP's downstream cost using the numerical model is illustrated in figure 8(b), for various values of the proportion x. The pattern is similar to the analytical model. When $x < 0.3$, too little of the downstream traffic from the content provider to the ISP is delivered locally, and as before the cost to the ISP increases as the number of IXP increases. However, when $x > 0.3$, as before the ISP benefits from increasing the number of IXPs at which the two parties agree to peer.

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

IX. CONCLUSION

In order to explain the common settlement-free peering requirements of the largest ISPs in the United States, we examined the effect of the number of interconnection points at which two networks peer, the locations of these interconnection points, and traffic ratios on an ISP's variable traffic-sensitive costs.

When two ISPs peer with a traffic ratio of 1:1, the variable traffic-sensitive cost is uni-modal, and we estimate that it is minimized with 8 IXPs. There may be little value in requiring interconnection at more than 6 IXPs. The ISP's cost is typically minimized by selecting interconnection points that span the country and are near population centers. When two ISPs using hot potato routing peer with an imbalanced traffic ratio that is below 2:1, the decrease in the upstream cost with the number of IXPs dominates the corresponding increase in the downstream cost, and thus interconnecting at 6 to 8 IXPs results in close to a minimum total cost. Requiring interconnection at more than 8 interconnection points is of little incremental value. In contrast, when the traffic ratio is above 2:1, the variation in the downstream cost with the number of IXPs dominates. As a result, the total cost increases with the number of IXPs, and thus it is no longer rational for the ISP to agree to settlement-free peering.

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

These results are not the end of the story. We considered direct interconnection between a content provider and an ISP. It would be useful to examine the decision of a content provider choosing between direct interconnection with an ISP versus indirect interconnection via a transit provider. Such an analysis could give insight into both an ISP's and a content provider's decision, as well as settlement-free peering requirements between an ISP and a transit provider. In addition, it would be worthwhile to consider other factors that can affect interconnection decisions, e.g., geographic scope, backbone capacity, and traffic volume. In particular, it would be interesting to refine the model for networks outside the United States, including differences in ISP service territories and locations of peering. It would also be interesting to examine if the use of public IXPs results in different conditions for settlement-free peering than the use of primarily private IXPs considered here.

APPENDIX

A. *Proof of Theorem 1*

In (11), we have

$$\begin{aligned} & ED_{1,down}^{b,hot} = \\ & \sum_{g\in I_1}\sum_{g'\in I_1} D_1^b(g,g')P(IXP_{1,down}^{p,hot}=g)P(IXP_1^u=g'). \end{aligned} \tag{A.41}$$

The distance $D_1^b(g,g')$ is given by (9). The probability distribution of $IXP_{1,down}^{p,hot}$ is given in (7), and the probability distribution of IXP^u is similarly uniformly distributed. Substituting these expressions into (A.41),

$$\begin{aligned} ED_{1,down}^{b,hot} &= \sum_{i^g\in l_1^N}\sum_{i^{g'}\in l_1^N}\frac{L}{N}|i^g-i^{g'}|\frac{1}{N}\frac{1}{N} \\ &= \frac{1}{N^2}\sum_{i^g\in l_1^N}\sum_{i^{g'}\in l_1^N}\frac{L}{N}|i^g-i^{g'}|. \end{aligned} \tag{A.42}$$

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

$$\begin{aligned} & \sum_{i^{g'}\in l_1^N}\frac{L}{N}|i^g-i^{g'}| \\ & = \frac{L}{N}\{[|1-k|+...+|-1|]+0+[|1|+...+|N-k|]\} \\ & = \frac{L}{N}\left[\frac{k(k-1)}{2}+\frac{(N-k)(N-k+1)}{2}\right]. \end{aligned} \tag{A.43}$$

Substituting (A.43) into (A.42),

$$\begin{aligned} & ED_{1,down}^{b,hot} = \frac{1}{N^2}\sum_{k=1}^{N}\frac{L}{N}\left[\frac{k(k-1)}{2}+\frac{(N-k)(N-k+1)}{2}\right] \\ & = \frac{L}{N^3}\left[\frac{1}{2}(N+1)N^2+\sum_{k=1}^{N}k^2-(N+1)\sum_{k=1}^{N}k\right] \\ & = \frac{L}{N^3}\left[\frac{1}{2}(N+1)N^2+\frac{N(N+1)(2N+1)}{6}-\frac{N(N+1)^2}{2}\right] \\ & = \frac{L(N-1)(N+1)}{3N^2}. \end{aligned} \tag{A.44}$$

B. Proof of Theorem 2

In (14), we have:

$$ED_1^m = \sum_{g' \in I_1} \sum_{A_1(j,k) \subset R_1(g')} D_1^m(g', A_1(j,k)) P_1(j,k). \tag{A.45}$$

The access networks $A_1(j,k) \subset R_1(g')$ are given by

$$\begin{aligned} \frac{(i^{g'}-1)L}{aN} + 1 \le & j \le \frac{i^{g'}L}{aN} \\ 1 \le & k \le \frac{W}{b} \end{aligned} \tag{A.46}$$

The double sum in (A.45) can thus be written as

$$ED_1^m = \sum_{i^{g'}=1}^{N} \sum_{k=1}^{\frac{W}{b}} \sum_{j=\frac{(i^{g'}-1)L}{aN}+1}^{\frac{i^{g'}L}{aN}} D_1^m(g', A_1(j,k)) P_1(j,k). \tag{A.47}$$

The distance $D_1^m(g', A_1(j,k))$ is given by (10):

$$D_1^m(g', A_1(j,k)) = \sqrt{\left(\frac{(2j-1)a - L/N(2i^{g'}-1)}{2}\right)^2 + \left(\frac{(2k-1)b - W}{2}\right)^2} \tag{A.48}$$

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

$$D_1^m = \sqrt{\left(\frac{(2j'-1)a - L/N}{2}\right)^2 + \left(\frac{(2k-1)b - W}{2}\right)^2} \tag{A.49}$$

Changing the inner sum over j into a sum over j' and substituting $P_1(j,k) = ab/LW$,

$$ED_1^m = \frac{ab}{LW} \sum_{i^{g'}=1}^{N} \sum_{k=1}^{\frac{W}{b}} \sum_{j'=1}^{\frac{L}{aN}} D_1^m \tag{A.50}$$

D_1^m is no longer a function of $i^{g'}$, and thus we can remove the outer sum, resulting in

$$ED_1^m = \frac{abN}{LW} \sum_{k=1}^{\frac{W}{b}} \sum_{j'=1}^{\frac{L}{aN}} \sqrt{\left(\frac{(2j'-1)a - L/N}{2}\right)^2 + \left(\frac{(2k-1)b - W}{2}\right)^2}. \tag{A.51}$$

C. Proof of Theorem 3

In (16), we have

$$ED_1^a = \frac{1}{ab} \int_{-\frac{a}{2}}^{\frac{a}{2}} \int_{-\frac{b}{2}}^{\frac{b}{2}} \sqrt{x^2+y^2}\, dy\, dx. \tag{A.52}$$

By symmetry,

$$ED_1^a = \frac{4}{ab} \int_0^{\frac{a}{2}} \int_0^{\frac{b}{2}} \sqrt{x^2+y^2}\, dy\, dx. \tag{A.53}$$

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

$$\begin{aligned} &\int_0^{\frac{a}{2}} \int_0^{\frac{b}{2}} \sqrt{x^2+y^2}\, dy\, dx = \\ &\int_0^{\frac{b}{2}} \int_0^{\frac{ay}{b}} \sqrt{x^2+y^2}\, dx\, dy + \int_0^{\frac{a}{2}} \int_0^{\frac{bx}{a}} \sqrt{x^2+y^2}\, dy\, dx \end{aligned} \tag{A.54}$$

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

$$\begin{aligned} &\int_0^{\tan^{-1}\frac{a}{b}} \int_0^{\frac{b}{2}\sec\theta} r^2\, dr\, d\theta + \int_0^{\tan^{-1}\frac{b}{a}} \int_0^{\frac{a}{2}\sec\theta} r^2\, dr\, d\theta \\ &= \frac{b^3}{24} \int_0^{\tan^{-1}\frac{a}{b}} \sec^3\theta\, d\theta + \frac{a^3}{24} \int_0^{\tan^{-1}\frac{b}{a}} \sec^3\theta\, d\theta \\ &= \frac{b^3}{48} \left[\frac{2a}{b^2}\sqrt{(\frac{a}{2})^2 + (\frac{b}{2})^2} + \sinh^{-1}(\frac{a}{b})\right] \\ &+ \frac{a^3}{48} \left[\frac{2b}{a^2}\sqrt{(\frac{a}{2})^2 + (\frac{b}{2})^2} + \sinh^{-1}(\frac{b}{a})\right] \\ &= \frac{a^3 \sinh^{-1}(\frac{b}{a}) + b^3 \sinh^{-1}(\frac{a}{b}) + 2ab\sqrt{a^2+b^2}}{48}. \end{aligned} \tag{A.55}$$

Substituting this expression into (A.53),

$$ED_1^a = \frac{a^3 \sinh^{-1}(\frac{b}{a}) + b^3 \sinh^{-1}(\frac{a}{b}) + 2ab\sqrt{a^2+b^2}}{12ab}. \tag{A.56}$$

REFERENCES

[1] "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.

[2] "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.

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

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

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

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

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

[8] Body of European Regulators for Electronic Communications, "BEREC Preliminary Assessment Of The Underlying Assumptions Of Payments From Large CAPs To ISPs," 2022.

[9] A. Nikkhah and S. Jordan, "Should Large ISPs Apply the Same Settlement-Free Peering Policies To Both ISPs and CDNs?" in *2023 IEEE 20th Annual Consumer Communications & Networking Conference (CCNC)*, 2023.

[10] ——, "Requirements of Settlement-Free Peering Policies," in *2022 IEEE Global Communications Conference (GLOBECOM)*, 2022.

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

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

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

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

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

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

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

[18] Frontier Communications, "Bilateral Peering Policy," Accessed January 26, 2022. [Online]. Available: https://ipadmin.frontier.com/bilateralpeering_policy

[19] Mediacom Communications Corporation, "Interconnection (Peering) Policy," Accessed January 26, 2022. [Online]. Available: http://mediacomcable.com/interconnection-policy

[20] TDS Telecommunications Corp, "AS4181 Peering Policies," Accessed January 26, 2022. [Online]. Available: http://as4181.net/peering.html

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

[22] PeeringDB, "The Interconnection Database," Accessed February 6, 2022. [Online]. Available: https://www.peeringdb.com/

[23] 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.

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

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

[26] A. Dhamdhere, C. Dovrolis, and P. Francois, "A value-based framework for Internet peering agreements," in *International Teletraffic Congress (ITC 22)*, 2010, pp. 1–8.

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

[28] 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.

[29] S. Patchala, S. Lee, C. Joo, and D. Manjunath, "On the economics of network interconnections and its impact on net neutrality," *IEEE Transactions on Network and Service Management*, vol. 18, no. 4, pp. 4789–4800, 2021.

[30] S. Lee, D. Manjunath, and C. Joo, "On the economics effects of cdn-mediated delivery on content providers," *IEEE Transactions on Network and Service Management*, vol. 19, no. 4, pp. 4449–4460, 2022.

[31] H. Chang, S. Jamin, and W. Willinger, "To peer or not to peer: Modeling the evolution of the internet's as-level topology," in *Proceedings IEEE INFOCOM*, 2006, pp. 1–12.

[32] 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.

[33] Maps of World, "How Wide Is The United States?" Accessed February 22, 2022. [Online]. Available: https://www.mapsofworld.com/answers/united-states/how-wide-is-united-states/#

[34] Data Center Map, "Internet Exchange Points," Accessed January 26, 2022. [Online]. Available: https://www.datacentermap.com/ixps.html

[35] TeleGeography, "Internet Exchange Maps," Accessed January 26, 2022. [Online]. Available: https://www.internetexchangemap.com/#/country/united-states

[36] Packet Clearing House, "Peering," Accessed January 26, 2022. [Online]. Available: https://www.pch.net/about/peering

[37] Network Startup Resource Center (NSRC), "Internet EXchange Points: North America Region," Accessed January 26, 2022. [Online]. Available: https://nsrc.org/ixp/NorthAmerica.html

[38] PeeringDB, "Exchanges," Accessed January 26, 2022. [Online]. Available: https://www.peeringdb.com/advanced_search?country__in=US&reftag=ix

[39] Hurricane Electric Internet Services, "Internet Exchange Report," Accessed January 26, 2022. [Online]. Available: https://bgp.he.net/report/exchanges

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

[41] ——, "Vintage county population estimates totals," Accessed January 26, 2022. [Online]. Available: https://www.census.gov/programs-surveys/popest/technical-documentation/research/evaluation-estimates/2020-evaluation-estimates/2010s-counties-total.html

[42] ——, "American Fact Finder," Accessed 2020. [Online]. Available: https://factfinder.census.gov/faces/tableservices/jsf/pages/productview.xhtml?src=bkmk

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

Ali Nikkhah is a Ph.D. candidate in the Networked Systems program at the University of California, Irvine. He obtained his B.Sc. and M.Sc. degrees in Electrical Engineering from the Sharif University of Technology in 2018 and 2016, respectively. Currently, Ali is working on an NSF grant project that focuses on internet interconnection policy. His interests lie at the intersection of telecommunication policy and data science. In addition to his research, Ali is also an active member of the IEEE and has presented his work at several conferences.

Scott Jordan (Member, IEEE) received the B.S./A.B., the M.S., and Ph.D. degrees from the University of California, Berkeley, in 1985, 1987, and 1990, respectively. From 1990 until 1999, he served as a faculty member at Northwestern University. Since 1999, he has served as a faculty member at the University of California, Irvine. In 2006, he served as an IEEE Congressional Fellow, working in the United States Senate on communications policy issues. In 2014-2016, he served as the Chief Technologist at the Federal Communications Commission, advising on technological issues across the Commission, including the 2015 Open Internet Order and the 2016 Broadband Privacy Order. His current research interests are Internet policy issues, including net neutrality, privacy, interconnection, data caps, zero rating, and device attachment.