Peering Among Content-Dominated Vertical ISPs

Prasun Kanti Dey, *Student Member, IEEE*, and Murat Yuksel, *Senior Member, IEEE*

***Abstract*—Content providers (CPs) typically control the digital content consumption services and are getting the most revenue by implementing an "all-you-can-eat" model via subscription or hyper-targeted advertisements. Revamping the existing Internet architecture and design, a *vertical* integration where a CP and access Internet service provider (ISP) will act as unibody in a *sugarcane* form seems to be the recent trend. In this letter, using the U.S. as a case study, we show the overlaps between access ISPs and CPs to explore the viability of a future in terms of peering between these new emerging content-dominated *sugarcane* ISPs and the healthiness of Internet economics.**

***Index Terms*—Content providers, network economics, network topology, peering evolution, vertical integration.**

I. Introduction

THE INTERNET has been designed to be open and neutral for everyone from the very beginning and as the complexity grows due to the increasing number of users utilizing diversified applications, key stakeholders like Internet eXchange Points (IXPs) and Internet Service Providers (ISPs), be that content, transit or access ISPs, have been introduced to the current architecture. While content providers/ ISPs generate various contents to be consumed by the end customers, transit and access ISPs are responsible for delivering the data to end users smoothly by setting up new fibers and maintaining the existing infrastructure. IXPs house multiple carrier ISPs[1] to exchange their traffic in numerous strategically located facilities.

In the traditional horizontally organized system, carrier ISPs charge end-user a fee for connectivity, and content providers (CPs) for ensuring high-speed data delivery or improved stream quality. As most of the money is in the content business [7], the trend is to invest more in creating own content and serve it or acquire an existing CP to gain its control.

Recent ruling on Net Neutrality favors carrier ISPs and allows them to *legally* prioritize data before delivering to end users and it is possible that a user may experience delay while accessing certain CP. Breaking the status quo means carrier ISPs now have an unfair advantage over CPs if there raises any conflict of interest between a CP vs. another provider affiliated to a particular ISP. Merger of AT&T and Time Warner ignites the following question, *what if* AT&T starts favoring own content over its competitors? On the opposite side, some CPs have also started provisioning access or making "paid peering" deals.

It appears that the existing Internet structure, which separates the providers *horizontally*, will no longer be applicable, rather a *vertical* integration [11] of multiple players from different layers seems to emerge. This new architecture will eventually eliminate the *typical* access and transit ISPs. In this new Internet design, CPs will likely to dominate, sitting on the top, and the means of content delivery (i.e., transit and access) will be vertically integrated to the CPs all through the carrier ISPs to the end users. We name such *vertically* integrated ISPs as *sugarcane* ISPs. This can be in two ways: A carrier ISP can acquire a CP (Verizon acquiring Yahoo!) or vice versa.

Akin to how existing ISPs benefit from peering, the new *sugarcane* ISPs also need to peer between themselves. We call such agreement as "*sugarcane* peering". Without peering, end users could be secluded and forced to see only specific contents from selected group of providers, which is unacceptable and violates the ground rule of the *Open Internet*.

This letter explores the possibility of *vertical* integration of ISPs by inspecting the geographic coverage of access ISPs vs. CPs within the United States (U.S.), and tries to understand the plausibility of such a content-dominated vertical ISP market by exploring *sugarcane* peering. Globally, beyond U.S., all network infrastructure sectors struggle to maintain the profit margins while facing vertical market forces [3]. Thus, our observations and findings apply to the countries and regions where vertical integration is prominent.

A. Motivation

Figure 1 presents an *approximate* timeline of the evolution of peering. In the legacy model, CPs and access ISPs were horizontally separated and had to purchase transit service from transit ISPs. Content Delivery Networks (CDNs) gained momentum [13] when companies started relying on CDN services after *Akamai* received significant market exposure. During 2009 [2], having large enough fiber footprint of their own, CPs started to by-pass transit and CDN providers, and established direct peering with access ISPs following a "Donut Peering" model. Since the traffic ratio between CPs and access ISPs are not even and, for CPs, putting caches directly at the access ISPs' end [5] proved to be more beneficial, generic peering policy was not enough, which triggered access ISPs to introduce "Paid Peering" to charge CPs. As CPs continue to earn larger share of the revenue, they will either expand their own fiber network or buy services from access ISPs or datacenters, and will gradually minimize their dependency on transit ISPs.

CPs like Google, Facebook, and Netflix want to ensure that the consumers– without being worried to exceed data

Manuscript received March 15, 2019; revised May 11, 2019; accepted June 5, 2019. Date of publication June 12, 2019; date of current version August 22, 2019. This work was supported in part by NSF under Award 1814086. The associate editor coordinating the review of this paper and approving it for publication was B. Kantarci. *(Corresponding author: Prasun Kanti Dey.)*

The authors are with the Department of Electrical and Computer Engineering, University of Central Florida, Orlando, FL 32816 USA (e-mail: prasun@knights.ucf.edu; murat.yuksel@ucf.edu).

Digital Object Identifier 10.1109/LNET.2019.2922208

[1]A carrier ISP can be either an access/eyeball and/or transit ISP.

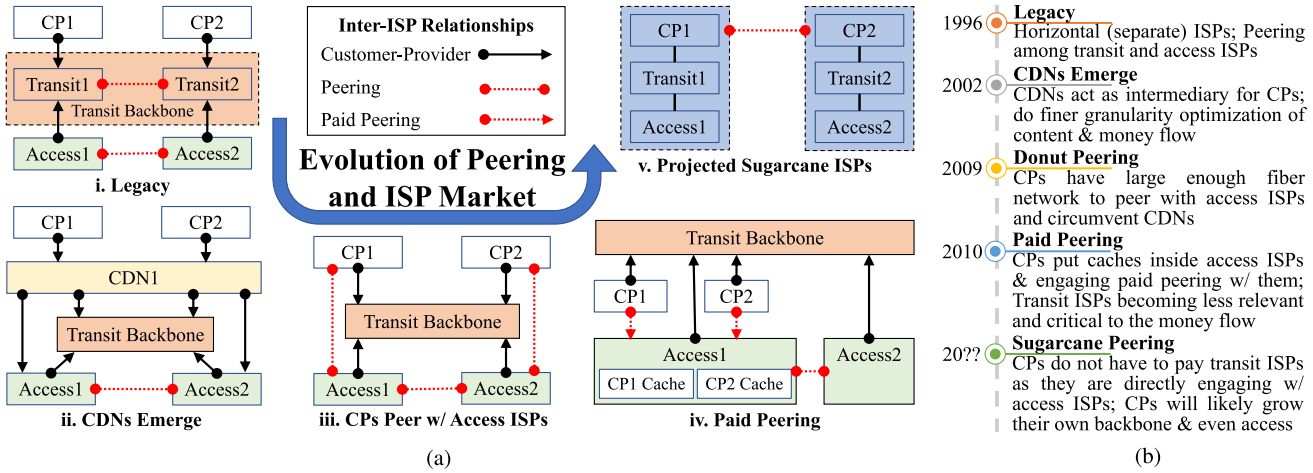

Fig. 1. Peering Evolution [2], [5], [10], [13]. (a) Architectural Overview of Peering Evolution in the Internet. (b) Timeline.

caps– get the best experience while streaming videos, and thus sometimes intentionally downgrade the video quality to reduce the data consumption. For example, Netflix checks end user device and throttle-down the streaming quality [9]. Throttling traffic quality rejuvenates the net neutrality issue, as the Federal Communications Commission (FCC) imposed *non-discriminating treatment towards data* on carrier ISPs with the hope that it would ensure the creation and unrestricted distribution of content or services. Although, the absence of the net neutrality do not impede the *vertical* integration between CPs and carrier ISPs, not much work has been done on detailing the architecture, economic perspective or peering settlement among those vertical ISPs. Earlier study [8] explored the *vertical* integration and argued how it may motivate access providers not to block or slow down competitors' content.

Such *vertical* integration will certainly improve the end-user experience as the CP already possesses the detailed information about the traffic volume and can prepare its downstream carrier network to handle any sudden burst of traffic. Having said that, *vertical* integration brings new challenges:

a) How will CPs establish the end-to-end network?
b) How big their footprints (geographically) will be? How will they inter-operate?
c) Once vertical merging is complete, how these new sugarcane ISPs will peer with each other since they will grow bigger in size and their business strategy may shift?

Regarding challenge-*a*, the new management will most likely rely on the access ISPs' already established infrastructure and will fine-tune their specific requirements to offer more curated services as a bundle. Challenge-*b* can be visualized as the *union set* of previously separated access and/ or transit ISPs and CPs existing coverage footprint. Primary concern will be whether there will be a healthy peering policy among these ISPs when *vertical* integration becomes the new norm. We aim to answer challenge-c in this letter.

II. Sugarcanes: To Peer or Not to Peer

In peering, ISPs carry their own traffic to another ISPs' points of presence (PoPs) and agree to exchange traffic without paying any fees [10] to gain the reciprocal access to each other's customers. Otherwise, ISPs have to purchase transit service for global reachability. Two ISPs will likely peer if:

a) they are similar in (customer cone) size and market value;
b) they cover multiple locations and their coverage areas are mostly non-overlapping; and
c) they generate similar traffic volumes.

We consider conditions *a* and *b* only, as these information are publicly available and condition *c* is mostly proprietary. We show how much overlapping of coverage area exist between ISPs of different types and their market values. ISPs operating in different locations will be motivated to peer with each other to expand their networks' reach. It is also expected that if the market values of ISPs are close to each other, there is a higher possibility of peering. This is simply because if an ISP is bigger in size and value, it will charge smaller ISPs.

A. ISP Peering Locations and Coverage Area

Typically, ISPs are restrictive about disclosing their internal topology, rather, they share their PoPs so that other ISPs may consider them for potential peering. Geographical scope is a key parameter as overlapping coverage among two ISPs will reduce the likelihood of peering.

We consider the following 37 major U.S.-based ISPs:

Access ISPs: AT&T, CenturyLink, Charter, Comcast, Cox, Google Fiber, GTT, Hotwire, Liberty, Mediacom, PenTeleData, Sonic.net, TDS Telecom, Time Warner, WOW.

Content Providers: Amazon, Facebook, Google, IBM, Microsoft, Netflix, Spotify, Verizon, Yahoo, Yelp!

Transit ISPs: Cogent, Coresite, Frontier, General, Hurricane Electric, IIJ, Level3, NTT, PCCW Global, Qwest (CenturyLink), Sprint, Verizon, WOW, Zayo.

We categorized ISPs according to PeeringDB (https://peeringdb.com). We listed all the peering locations (latitude, longitude), and PoP facilities count in a city for each ISP, and then calculated the *geometric median* of those PoP locations. We consider this metric to represent the focal point of an ISP's coverage area because, sum of distances from this point to all other PoPs is minimum. The fact that we can

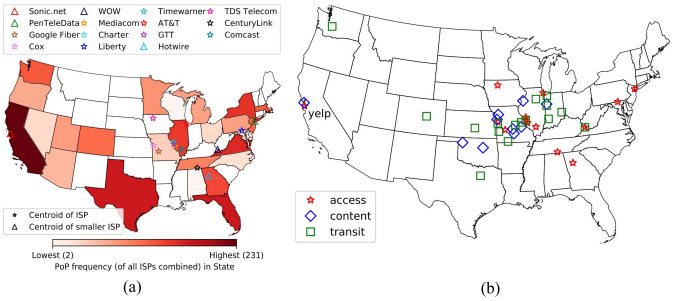

Fig. 2. ISP coverage areas as centroids. (a) For Access ISPs only. (b) For all ISPs.

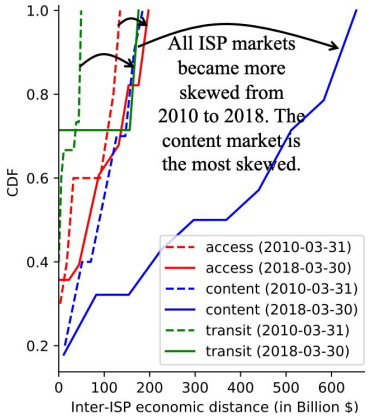

Fig. 3. Inter-ISP *economic* distance CDF.

generalize *geometric median* to include weighted distances and convert it to a 'Weber problem' [6] also motivated us in selecting the metric. We shall call it as 'centroid' onward.

An alternative could be calculating the *geographic center* to represent the centroid. But, in this case, two providers with PoPs in completely different locations may end up having their *geographic center* nearby. In contrast, *geometric median* always tries to be closer to where most of the PoPs are present. Based on our data, access ISPs are mostly East Coast oriented (*Sonic* operates dedicatedly in West Coast) as shown in Figure 2(a). We categorized access ISPs (specially) into small and big groups according to Wikipedia [1] since it marked an ISP with more than 1 million residential customers as big ISP and ISPs with fewer numbers were marked as small. Higher population attracts ISPs to expand and this reflects in more access PoPs in East Coast. Figure 2(b) presents an overview of all ISPs' focus areas and we can see that except *Yelp!*, most of the content ISPs' centroids are condensed in the center of the U.S., which means they are operating coast-to-coast with a coverage area spanning the whole country.

B. Market Capital

Using Intrinio (https://intrinio.com) and Macrotrends (https://www.macrotrends.net/), we collected (quarterly) market capital of 26 ISPs from either New York Stock Exchange or Nasdaq Stock Market for the period of March 31, 2005 to June 30, 2018. We used the last business day of each month. However, Level3 merged with Centurylink in 2017; as a result, we have market capital for Level3 till 2017. On the other hand, Facebook announced their initial public offering (IPO) on February 2012, and so, we have its data since then. On average, we have around 80% of the data points for each ISP.

III. Results

A. Inter-ISP Economic Distance

A key measure to determine the peering likelihood of two ISPs is comparing their market value. To quantify how similar two ISPs are to each other in terms of value, we look at *inter-ISP economic distance*, the absolute difference of market capitals of a pair of ISPs, and present the CDFs in Figure 3. Our observations from this measurement are multi-fold.

First, it is uncanny that economic growth patterns for access and transit ISPs remain almost identical for the entire period except that the inter-ISP economic distance is higher among access than transit ISPs. This means there is a chance for small access ISPs to survive even if they serve only to a small number of customers. But, it will be exceptionally challenging to run a transit business with little market capital as it requires a bigger infrastructure to maintain in comparison to access.

Second, inter-ISP economic distance is consistently increasing for CPs. From economic point of view, this can be either good or bad. For instance, it represents a flourishing market with new contestants coming forward with innovations and value creation, which is good. Or, it may indicate an unhealthy competition where only few players are dominating.

Figure 3 also compares CDFs of market capital differences at two different timestamps, 8 years apart from each other. During this time, neither access nor transit market has found any clear dominance, while in the content market, some ISPs have taken a profound lead over others. This observation supports our claim for the future of a *vertically* integrated ISPs, where a CP can acquire an access and/or transit ISP.

The trends in the ISP markets show that they are becoming more skewed in terms of ISPs' market value (see Figure 3). Of them, the content market has grown to be a much more skewed one. A potential drawback is that peering may not be extensive in a highly skewed market since the market values of ISPs will not be similar. The market skew increases the incentive of ISPs to not collaborate. This is of particular concern for a content-dominated ISP market where highly skewed CPs will likely drive the contracts and peering agreements. Further, the skew will also increase the incentive of CPs to acquire others, which may drive towards an unhealthy oligopoly market.

In a content-dominated *sugarcane* ISP architecture, access and transit infrastructures will be used to carry traffic according to CPs' peering policy. Since, access providers are the downstream retailer, peering among *sugarcane* ISPs will be mandatory; unless ISPs want to build the underlying physical infrastructure for their own. Annual spending on broadband infrastructure in U.S., an ambivalent indicator whether ISPs' coverage area are expanding, has just recovered from its consecutive two years of downfall and hit $76.3B mark in 2017, still less than 2014's expenditure of $78B [4]. Net neutrality repeal, in this regard, can be treated as an initiative to persuade ISPs for rising their network expansion investments and to attract more market competition. Yet, it is questionable if the *sugarcane* ISPs will be incentivized to peer since the existing CP market is highly skewed. If the skewness stays even after vertical integration, *sugarcane* ISPs will be less eager to peer which may degrade the overall end-to-end Internet experience.

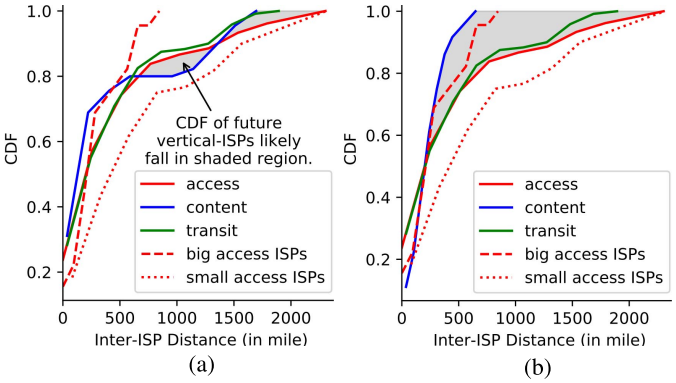

Fig. 4. Inter-ISP distance. (a) With Yelp! (b) without Yelp!

B. Inter-ISP Overlap

To estimate how much an ISP overlaps with another one, we measured the distance between their centroids. We calculated 'inter-ISP distance' between pairs of same type ISPs and plotted the CDFs in Figure 4. The outcome is pretty revealing. Content providers are located very close to each other (except *Yelp!*) with the least inter-ISP distance, and access ISPs are the farthest from each other, while transit ISPs are in between.

It is relatively easy to expand the business coverage area for CPs. With already created contents, only requirement is to peer or contract with another carrier provider. CPs get an upper hand here as they continue to penetrate into different locations. As their coverage area expands, their centroids concentrate at the geographical center of the country. Transit ISPs have centroids gathering at the center of the countrty, though they are a bit more dispersed than CPs. They have strong backbone and usually lay their network in major cities where access or CPs purchase transit support. They do not want to spread their coverage as wide as CPs, but their footprints are *complete enough* [12] to cover the whole country.

Contrary to content and transit ISPs, access ISPs have significantly more dispersed centroids. It requires large investments to provide Internet access to a location. Unless an access provider has significant enough capital, it is bound to serve only regional consumers. To get further insight, we plot CDFs for small and big access ISPs separately. Bigger ones are more sparse, while the smaller ones are more regional and oriented towards a specific area. This presents a possible merger or peering scenario for small access ISPs with CPs, and may explain the recent trend of paid peering agreements among content and regional access ISPs.

Although it is hard to predict precisely, it is likely that a content-dominated *sugarcane* ISP in future will have centroids dispersed more than the current CPs but less than the current access ISPs. This prediction assumes that existing content and access ISPs will merge, which translates into a inter-ISP CDF in between the current content and access ISPs' CDF plots as illustrated by *gray* color in Figure 4.

IV. Summary and Future Work

Using three datasets, we compared 37 ISPs from the U.S. market to find the possibility of a *vertically*-integrated ISP market where content will be dominating and inaugurated the idea of *sugarcane* ISP. To forecast on a future market of content-dominated *sugarcane* ISPs, we utilized geographic distance between existing ISPs and introduced an inter-ISP *economic* distance to quantify the thriving progress of CPs over carrier ISPs. Based on our analysis, CPs are clearly dominating in market value and their coverage areas overlap the most. As CPs are vertically integrating with (or merely purchasing) access ISPs, it seems likely that the future *sugarcane* ISPs will have centroids closer to each other than the existing access ISPs, which may reduce the incentive for peering and hence reduction in the overall end-to-end performance of the Internet. As for inter-ISP economic distance, the CP market is most skewed, implying less incentive for peering and collaboration.

ISPs controlling huge swathes of Internet market may yield higher prices for end customers and feel less compelled to be innovative, but, as long as technology evolves and reliance on Internet continues to strengthen, ISPs will consistently keep investing in their infrastructure. More research is needed to understand the trend of vertical ISP integration, particularly in terms of peering quality and inter-ISP overlap quantity using distances among PoPs of potentially peering ISPs, and regulating such a vertical market as new incentives and dynamics may emerge between new *sugarcane* ISPs. Lack of any regulatory body in the world of *sugarcane* ISPs may lead to an unfortunate scenario where an ISP may deliberately degrade its competitors content to increase their own demand. For those ISP markets owned/ dominated by state or implementing regulations that impose horizontal competition, new studies are needed to understand the increasing dominance of content.

References

[1] *List of Broadband Providers in the United States*. Accessed: Feb. 2019. [Online]. Available: https://en.wikipedia.org/wiki/List_of_broadband_providers_in_the_United_States

[2] The 'Donut Peering' Model: Optimizing IP Transit for Online Video. *MZIMA*. Accessed: Sep. 2009. [Online]. Available: https://web.archive.org/web/20091116164326/http://www.mzima.net/pdf/donut_peering.pdf

[3] What Future for Internet Service Providers? *British Computer Society*. Accessed: Jan. 2009. [Online]. Available: https://www.bcs.org/content/ConWebDoc/24111

[4] P. Brogan, *U.S. Broadband Investment Rebounded in 2017*, U.S. Telecommun. Res. Brief, Oct. 2018. [Online]. Available: https://tinyurl.com/y3krg8dz

[5] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the expansion of Google's serving infrastructure," in *Proc. ACM IMC*, 2013, pp. 313–326.

[6] R. Chandrasekaran and A. Tamir, "Algebraic optimization: The Fermat–Weber location problem," *Math. Progr.*, vol. 46, nos. 1–3, pp. 219–224, 1990.

[7] P. K. Dey, M. A. Canbaz, M. Yuksel, and M. H. Gunes, "On correlating ISP topologies to their businesses," in *Proc. IEEE ICC*, 2018, pp. 1–7.

[8] J. Farrell and P. J. Weiser, "Modularity, vertical integration, and open access policies: Towards a convergence of antitrust and regulation in the Internet age," *Harvard J. Law Technol.*, vol. 17, p. 85, Oct. 2003.

[9] S. Musil. *Netflix: We're the Ones Throttling Video Speeds on AT&T and Verizon. GeoMidpoint*. Accessed: Nov. 2018. [Online]. Available: https://www.cnet.com/news/netflix-admits-throttling-video-speeds-on-at-t-verizon/

[10] W. B. Norton. *The History of Internet Peering. The Internet Peering Playbook*. Accessed: Nov. 2018. [Online]. Available: http://drpeering.net/white-papers/History-Of-Internet-Peering/History-Of-Internet-Peering.html

[11] J. P. Choi and B.-C. Kim, "Net neutrality and investment incentives," *RAND J. Econ.*, vol. 41, no. 3, pp. 446–471, 2010.

[12] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with rocketfuel," *IEEE/ACM Trans. Netw.*, vol. 12, no. 1, pp. 2–16, Feb. 2004.

[13] G. Andrli, *The History of Content Delivery Networks (CDN)*, GlobalDots, London, U.K., Dec. 2012. [Online]. Available: https://tinyurl.com/y6oe56rv