Information Economics and Policy 25 (2013) 235–245

Contents lists available at ScienceDirect

Information Economics and Policy

journal homepage: www.elsevier.com/locate/iep

The evolution of Internet interconnection from hierarchy to "Mesh": Implications for government regulation

Stanley M. Besen [a], Mark A. Israel [b,*]

[a] Charles River Associates, United States
[b] Compass Lexecon, United States

ARTICLE INFO

Article history:
Received 12 July 2012
Received in revised form 9 July 2013
Accepted 10 July 2013
Available online 24 July 2013

Keywords:
Internet transit
Internet peering
Secondary peering
Settlement-free peering
Regulation
Regulatory costs
ISP
CDN

ABSTRACT

The Internet has evolved from a "hierarchy"—in which interconnection was achieved by having Internet Service Providers (ISPs) purchase transit services from top-level backbones and top-level backbone providers engage in direct settlement-free peering—to a "mesh" in which peering occurs among a much larger number of participants and some peering arrangements involve payments from one peer to another. In this new environment, backbone providers, ISPs, and suppliers of content have a far wider array of interconnection alternatives, both technical and financial, than they did only a short time ago. As is often the case, the introduction of new alternatives and contractual arrangements has led to calls to regulate which alternatives and arrangements are acceptable. In this paper, we explain why such regulation would be harmful, as it would (i) reduce the incentives of industry participants to minimize total costs; (ii) lead to higher access prices to end users; (iii) result in prices that do not adequately reflect costs; and (iv) create regulatory inefficiencies. We also explain why the alternative interconnection arrangements to which Content Delivery Networks (CDNs) (and their content provider clients) and ISPs generally have access already impose limits on the exercise of market power, thus obviating any need for regulation.

1. Introduction

A decade or so ago, the Internet could be fairly described as a system in which interconnection was achieved by having Internet Service Providers purchase transit services from top-level backbones and having top-level backbone providers engage in direct settlement-free peering with one another. In the intervening period, however, this hierarchical system has been displaced by one with a far richer array of arrangements – what might be thought of as a "mesh" – including peering among a much larger number of participants, secondary peering, and peering arrangements that involve payments from one peer to another. Moreover, this system has continued to evolve as industry participants discover new and creative ways to interconnect and financial arrangements take a wider variety of forms. Significantly, the system has adapted to the rapid changes in both the amount and nature of Internet traffic with few, if any, government restrictions on the nature and form of the interconnection arrangements that are permitted.

This paper describes this evolution in some detail and explains why an attempt to impose detailed regulation on Internet interconnection would be an extraordinarily difficult undertaking. Regulation of interconnection would reduce the incentives of industry participants to minimize total costs, lead to higher access prices to end users, result in prices that do not adequately reflect costs, and create regulatory inefficiencies. There is no compelling reason to bear such regulatory costs: The alternative interconnection arrangements to which Content Delivery Networks (CDNs) (and their content provider clients) and Internet Service Providers (ISPs) generally have access already impose limits on the exercise of market power. Hence, detailed government regulation of Internet interconnection would be likely to do more harm than good.

* Corresponding author. Address: Compass Lexecon, 1101 K Street NW, 8th Floor, Washington, DC 20005, United States. Tel.: +1 202 589 3484; fax: +1 202 589 3480.
E-mail address: misrael@compasslexecon.com (M.A. Israel).

0167-6245/$ - see front matter © 2013 Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.infoecopol.2013.07.003

2. Transit and peering primer

2.1. Traditional transit and peering arrangements

The Internet consists of approximately 40,000 networks ("autonomous systems" or, "ASs") around the world that have agreed to use a common protocol and addressing system known as the Internet Protocol ("IP").[1] In the first decade of the commercial Internet, the relationships among these IP networks were largely "hierarchical": business and residential users typically connected to the Internet via a local Internet service provider ("ISP"), and each local ISP typically relied on one or more national backbone providers for connections to other ISPs.[2] For example, if a residential subscriber wished to download content from a website, his request would travel from his home to his ISP, which would hand it "up" (sometimes through a separate regional access provider) to an Internet backbone provider. Unless that backbone provider also served the (typically different) ISP that served the content provider, it would hand the request to another backbone provider, which, in turn, would hand the request "down" to the content provider via the latter's ISP. The content provider, in turn, would reverse the process when transmitting the requested content back through the latter ISP and backbone operator to the residential end user. The following diagram illustrates the basic hierarchical structure of these inter-network relationships:[3]

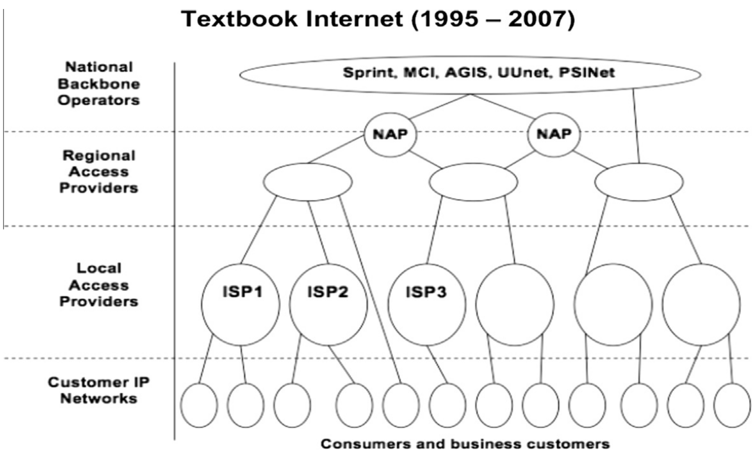

Throughout this period, networks typically exchanged traffic through two types of commercial arrangements adopted through bilateral negotiations. The first was *transit*, where one IP network (such as a local ISP) pays another IP network (such as a national backbone provider) to arrange for the transmission of traffic between its end users and *all* other sites on the Internet. The second was *peering*, where two IP networks exchange traffic from and to each other's end users (or the end users of their respective transit customers). "Settlement-free" peering, in which no cash payments are made between the parties exchanging traffic, has been the norm for network providers that exchange roughly comparable volumes of traffic.[4]

Internet backbones have always been divided into several classes, determined primarily by their reach, extent of connectivity with other networks, and settlement relationships that, for the most part, are privately negotiated and not publicly disclosed. According to Peyman Faratin and David Clark, "[t]he so-called Tier 1 ISPs are the set of ASes that do not purchase transit from any other AS, and thus must peer with every other Tier 1 ISP. The Tier 1 ISPs collectively form a complete mesh of peering arrangements. Tier 1 ISPs are large, with global scope."[5] Today, ten networks are generally viewed as Tier 1 backbones: Deutsche Telekom, Level 3, AT&T, Verizon, CenturyLink, Inteliquent, Sprint, NTT, TeliaSonera, and Tata.[6] Tier 2 networks typically peer with some networks and purchase transit from others, while Tier 3 networks generally only purchase transit.

[1] *See* CIDR Report for April 9, 2012, http://www.cidr-report.org/as2.0/ (viewed April 9, 2012).

[2] *See*, e.g., Christopher S. Yoo, "Network Neutrality or Internet Innovation?," *Regulation*, at 22-24 (Spring 2010) (*hereinafter*, "Yoo, *Regulation*"); Stanley M. Besen, et al., *Evaluating the Competitive Effects of Mergers of Internet Backbone Providers*, 2 *ACM Transactions on Internet Technology* 187, at 189-90 (2002) (*hereinafter*, "Besen et. al."). Some commentators use the term "ISP" to encompass both local access providers and backbone providers. Ultimately, it would be futile to attach precise definitions to terms describing various types of IP networks because major IP network operators are increasingly likely to perform multiple roles in the Internet, including local access, long-haul transport, and content distribution.

[3] Craig Labovitz et al., *ATLAS Internet Observatory Annual Report*, at 9 (2009), <http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf> (*hereinafter* Labovitz I) (visited March 24, 2011).

[4] *See* Bruce M. Owen, *The Internet Challenge to Television* at 202 (1999) ("The big ISPs…connect with one another and exchange roughly equal volumes of traffic. It is customary for such 'peer' networks not to pay one another for accepting the other's traffic, presumably because transactions costs can be avoided in that way with no significant revenue loss on either side").

[5] Peyman Faratin, David Clark et al., *The Growing Complexity of Internet Interconnection*, 72 Communications & Strategies 51, at 55 (4Q 2008) (*hereinafter*, "Faratin and Clark").

[6] For a list of Tier 1 networks see http://en.wikipedia.org/wiki/Tier_1_network (visited April 9, 2012).

2.2. Recent developments in peering and transit relationships

2.2.1. Description of recent developments

In recent years, several new types of commercial arrangements have emerged to supplement—and in some cases supplant—traditional transit and peering relationships. We focus here on *secondary peering* (a broad concept that includes paid peering) and *partial transit* (in which one IP network pays another IP network to arrange for the transmission of traffic between its end users and a subset of the other sites on the Internet).[7] These arrangements "are not new in the Internet, but have been in use for over a decade. What is new is the trend toward increased reliance on such contracts."[8]

Secondary peering. Secondary peering occurs when smaller IP networks interconnect directly with each other for various reasons, such as limiting transit costs or increasing the quality and predictability of the connection between the two networks.[9] When secondary peering arrangements are established, some traffic that would previously have flowed across transit connections is delivered instead over secondary peering connections.

Paid peering differs from settlement-free peering in that compensation flows from one party to another, but it is like traditional peering in that the peers only exchange traffic that is bound to and from their respective customers. Paid peering arrangements have become more common for two main reasons. First, ISPs have built up and expanded their networks, so they now have significantly more "backbone" facilities, which makes it easier for them to connect directly with a greater number of other networks. Second, the aggregation of content has enabled the "disintermediation" of conventional transit providers.[10] Exemplifying this phenomenon is the rise of Content Delivery Networks (CDNs) such as Akamai and Limelight.[11] Many Internet content providers contract with CDNs to cache content in servers distributed broadly across the Internet at points that are relatively close to the residential subscribers served by ISPs, or at least close to regional aggregation points serving those end users. Although CDNs can—and often do—pay one or more third-party transit providers to bridge the gap between their servers and some ISP networks, they may also pay ISPs to peer with them directly.[12] The result has been a significant evolution in the basic physical architecture of the Internet, as illustrated in the following diagram:[13]

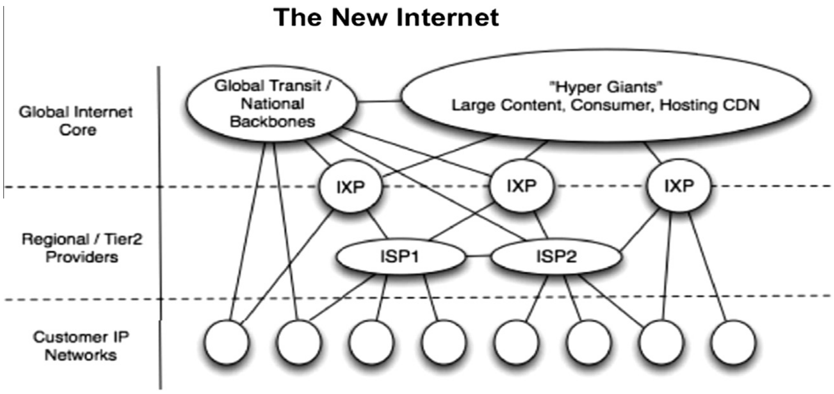

Partial transit. The proliferation of peering arrangements has coincided with the rise of a related commercial innovation known as "partial transit".[14] Under this

[7] Another type of arrangement is *mutual-transit*, which "allows a pair of administrative domains to provide connectivity to the rest of the Internet for each other. This…is typically between two administrative domains such as small ISPs who are located close to each other and who cannot afford additional Internet services for better connectivity." (L. Gao, "On Inferring Autonomous System Relationships in the Internet," http://www.renesys.com/blog/2009/02/gao.pdf, viewed September 7, 2012).

[8] Faratin and Clark at 53, n.2.

[9] Besen et al. at 197. Phillipa Gill, Martin Arlitt, Zongpeng Li and Anirban Mahanti, "The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse," May 21, 2008, http://www.hpl.hp.com/techreports/2008/HPL-2008-47.pdf, viewed September 7, 2012, explain that these developments "give content providers more control over their end-to-end applications performance."

[10] *See* Rudolph van der Berg, *How the 'Net works: an introduction to peering and transit*, Ars Technica (September 2, 2008) (*hereinafter*, "van der Berg"), http://arstechnica.com/old/content/2008/09/peering-and-transit.ars (visited March 24, 2011), ("Being a pure transit provider with only Autonomous Systems as customers puts a network in a weird spot. Such a network's business case is built on being the intermediary in the flow of traffic, so it tries to charge all of the other autonomous systems for their traffic. The problem for a pure transit provider is that its customers are always looking at ways to lower their transit fees, and lower transit fees can be had by switching to a competitor or by not using the transit provider at all. So disintermediating the transit provider is standard behavior for the transit provider's customers").

[11] For a relatively comprehensive and current list of CDNs see www.cdn-list.com (visited April 9, 2012).

[12] According to one observer, "Time Warner Cable and other providers … have paid peering products. Several Tier 1 ISPs and AOL have had paid peering products for years …" Ask Dr. Peering, Editor's Note for *Paid Peering and Net Neutrality* (November 5, 2009), http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2009/11/5_Paid_Peering_and_Net_Neutrality.html (visited January 7, 2011).

[13] Labovitz I at 17.

[14] *See* Christopher S. Yoo, "Innovations in the Internet's Architecture that Challenge the Status Quo," 8 Journal *of Telecomm & High Tech. Law 79, at 95 (2010) (hereinafter, "Yoo, Architecture").*

arrangement, a CDN that, for example, wishes to deliver content to subscribers of seven different ISPs can negotiate peering relationships with three of them and purchase partial transit from a backbone provider to reach the other four.

Multihoming. In an arrangement known as *multihoming*, a CDN —or any IP network—can purchase partial transit from several competing backbone providers and choose among them opportunistically and often in real-time. "[T]he ability of ISPs to establish multihomed connections, and the introduction of content delivery mechanisms, have enabled lower-tier providers, including ISPs and web-farms, to build redundancy into their networks, insulating them to an extent from the low quality of Internet traffic exchanges" and helping them "reduce their reliance on the top-level [backbone providers]."[15] Content providers, in turn, maximize their flexibility by establishing multihomed connections with competing CDNs. For example, Netflix reportedly has CDN contracts with Limelight, Akamai, and Level 3.[16]

Finally, Internet traffic exchanges exhibit an additional dimension of complexity in that any given peering arrangement (paid or settlement-free) typically includes traffic for which the peering partners have assumed responsibility under the transit services (partial or comprehensive) that they have sold to their individual customers (including content providers and other ISPs). Moreover, each of those underlying transit relationships is potentially "recursive": "[a] small AS A could purchase transit service from a medium sized AS M, which in turn could purchase transit service from X."[17] In the aggregate, "[t]he recursive combination of these standardized bilateral peering and transit contracts [has] created the complex web of interconnections wherein networks become resellers of transport and peer with one another in what has been a mutually beneficial and hence largely stable manner."[18]

2.2.2. The quantitative significance of these developments

It is difficult to quantify the effects of these developments, largely because many peering links are invisible to the monitors that collect the data used to measure their importance.[19] Haddadi et al. have noted that "For years, researchers have modeled the Internet AS topology using topologies collected from BGP and traceroute data. These data [miss] many peering links and [underestimate] the effects of…peering links on the AS topology structure."[20] Oliveira et al. report that publicly available routing data, which are obtained from remote monitors, tend to miss more than 85% of the links for large content networks,[21] and Dhamdhere et al. report this figure at 75%.[22] Edwards et al. report that, between 2002 and 2010, the observed Average Path Length in the Internet increased and the clustering coefficient—a measure of the extent to which neighbors of a node in a network are neighbors of each other—decreased. They conjecture that "the shift is likely caused by changes in peering policies that affect the hidden Internet and cannot be measured with public [Border Gateway Protocol] dumps…content providers are routing more traffic over hidden peer-to-peer links, and relying less on the more publicly visible Internet infrastructure."[23]

A number of studies have attempted to overcome the limitations of publicly available data by pursuing alternative approaches to measure the extent of Internet peering. For example, Dhamdhere and Dovrolis report that the median peering degree of Content/Access/Hosting Providers increased from 2 to 10 between 2003 and 2007.[24] Dhamdhere et al. obtained routing information for the subset of ASes that provided routing feeds to Routeview/RIPE collectors, which they call *usable monitors*. When they use these data to analyze the period 2006–2010, they find that, when Content Providers first appear in the data, the largest percentage does not use transit. They explain that "these networks tend to create a large number of non-transit (mostly settlement-free or paid-peering) links."[25] Using a different approach to overcome the limitations of publicly available data, Gill et al. reported in 2008 that "Microsoft connects to at least 24 different ASes, Google to at least 23, and Yahoo! to at least 18."[26]

In commenting on these developments, Labovitz et al. note that "In the emerging new Internet, the majority of traffic by volume flows directly between large content providers, datacenter/CDNs and consumer networks. In many cases, CDNs…are directly connected with both consumer networks and tier-1/tier-2 providers."[27] Similarly, Gill

[15] Besen et al. at 190.

[16] Jefferies & Company, Inc., *Limelight Networks: A Great Way to Play the Rise of Online Video, But Profits Are Still Elusive*, at 2 (December 2, 2010).

[17] Faratin and Clark at 54.

[18] Faratin and Clark at 52.

[19] Ricardo Oliveira, Beichuan Zhang, and Lixia Zhang, "Observing the Evolution of Internet AS Topology," SIGCOMM '07, August 27–31, 2007 note that "Besides the invisible links, which the monitors are unable to capture, there are many hidden links that can be captured, but are missing in the routing tables on the particular sampling day."

[20] Harmed Haddadi, Steve Uhlig, Andrew Moore, Richard Mortier, and Miguel Rio, "Modeling Internet Topology Dynamics," ACM SIGCOMM Computer Communications Review, Volume 38, April 2008, p. 66.

[21] Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, and Lixia Zhang, "In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure," ACM SIGMETRICS '08, June 2–6, 2008, http://www.cs.arizona.edu/~bzhang/paper/08-sigmetrics-completeness.pdf (viewed September 21, 2012).

[22] Amogh Dhamdhere, Himalatha Cherukuru, Constantine Dovrolis, and Ke Claffy, "Measuring the Evolution of Internet Peering Agreements," (hereafter Dhamdhere et al.), http://www.caida.org/publications/papers/2012/measuring_evolution_internet_peering/measuring_evolution_internet_peering.pdf (viewed September 21, 2012).

[23] Benjamin Edwards, Steven Hofmeyr, George Stelle, and Stephanie Forrest, "Internet Topology over Time," http://arxiv.org/pdf/1202.3993v1.pdf, viewed September 7, 2012, pp. 5–6. Although Amogh Dhamdhere, and Constantine Dovrolis, "Twelve Years in the Evolution of the Internet Ecosystem", *IEEE/ACM Transactions in Networking*, 2011, (hereafter Dhamdhere and Dovrolis I), report that the average path length has remained almost constant between 2000 and 2011, they attribute this to "the increasing multihoming degree of transit and content/access providers." (p. 1420).

[24] Amogh Dhamdhere and Constantine Dovrolis, "Ten Years in the Evolution of the Internet Ecosystem," IMC '08, October 20–22, 2008 (hereafter Dhamdhere and Dovrolis II).

[25] Dhamdhere et al., op. cit.

[26] Gill et al., op. cit.

[27] Craig Labovitz, Scott Ikel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian, "Internet Inter-Domain Traffic," (hereafter Labovitz II) http://www.sigcomm.org/sites/default/files/ccr/papers/2010/October/1851275-1851194.pdf, viewed September 7, 2012.

et al. observe that "the Internet topology is becoming flatter, as large content providers are relying less on Tier-1 ISPs, and peering with larger numbers of lower tier ISPs."[28]

2.3. Factors that affect parties' choices among alternative peering and transit arrangements

Historically, *settlement free* peering has been the arrangement of choice only when the costs and benefits of a peering arrangement are roughly equivalent for both peers. In such situations, an exchange of payments is unnecessary, and indeed imposes needless transaction costs. As Faratin and Clark explain, "settlement free peering arose as a defensible approximation only in the context of assumed symmetry in value flow. If the difference in actual value were small, bargaining costs would swamp the benefits of negotiating a price. However, once the assumption of symmetric value starts to break down, the binary world of transit and settlement free peering will break down."[29]

Network operators rely on several key criteria that, in practice, help determine each network's perception of the value of peering with another network. These include, among others, the traffic ratio between the two networks, the geographic diversity of the other network, the traffic volume exchanged between the two networks, and minimum backbone capacity and number of points of interconnection.[30] It is significant that "[i]n settlement-free peering relationships with very large networks, there is frequently a requirement to keep traffic 'in ratio.' The traffic going from A to Z is measured, and the traffic going from Z to A is measured. If the two numbers are not close enough, peering will be denied. For very large networks, the traffic ratio requirement is usually 2:1, and sometimes 1.5:1."[31]

Traffic ratio requirements help indicate whether the costs and benefits of a peering arrangement are approximately the same for both parties. Where those costs and benefits are comparable, the parties may find it more economical simply to "barter" traffic than to incur the transaction costs involved in making bilateral payments.

Large traffic imbalances suggest that the costs and benefits of peering may not be comparable for the two peers.[32] In such cases, although peering may still be efficient, it may no longer be appropriate for the parties to exchange traffic on a settlement-free basis. Instead, they may enter into a paid peering relationship. To be clear, the payments may not exactly match the distribution of traffic between two networks, and settlement-free peering *may* still be the appropriate solution even for two asymmetric Internet providers. For example, the amounts that ultimately would have to change hands may not be large enough to justify the transaction costs of monitoring and accounting for the traffic. Moreover, additional factors beyond traffic ratios may also be relevant to each provider's assessment of the value of direct peering. However, any *presumption* in favor of settlement-free peering is no longer correct because the relationship no longer exhibits symmetric traffic and equal costs.

In late 2005, Level 3 briefly depeered Cogent Communications after concluding that the ratio of traffic between the two networks had fallen out of balance.[33] This had the effect of "denying Level 3's customers access to Cogent's customers and denying Cogent's customers access to Level 3 customers," although only if those customers were single-homed—*i.e.*, relied exclusively on a single Tier 1 backbone to route their traffic.[34] Eventually, Cogent agreed to pay Level 3 for peering if, among other criteria, its traffic ratios exceeded a prescribed threshold.[35]

In commenting on this and other similar developments, Faratin and Clark concluded that, "[t]there is little evidence, aside from a few highly visible events such as depeering actions, that the range of negotiated contracts, whether discriminatory or not, has harmed the overall

[28] Gill et al., op. cit.

[29] Faratin and Clark at 60.

[30] Faratin and Clark at 56. Of course, in a range of circumstances, it may be uneconomic for two networks to have *any* peering relationship, settlement-free or paid. That may be the case, for example, if one of the parties is very small, in which case contracting and technical costs may swamp the benefits of direct interconnection.

[31] Faratin and Clark. For examples of peering policies *see*, AT&T's Global IP Network Settlement-Free Peering Policy, http://www.corp.att.com/peering/ (visited December 26, 2010); Suddenlink Communications' Settlement-Free Interconnection (Peering) Policy, http://www.suddenlink.com/terms-policy/peering.php (visited January 5, 2011); Verizon Business Policy for Settlement-Free Interconnection with Internet Networks, http://www.verizonbusiness.com/terms/peering (visited January 5, 2011); Qwest's North America IP Network Peering Policy, http://www.qwest.com/legal/peering_na.html (visited January 5, 2011); and Deutsche Telekom's Settlement Free Peering Policy, Version 3.0 (March 10, 2009) ("The ratio of aggregate traffic exchanged must be roughly balanced and shall not exceed 1.8 to 1. Ratio will be calculated using average monthly traffic sent from Peering partner to Deutsche Telekom's network divided by average monthly traffic sent from Deutsche Telekom to Peer."). Amogh Dhamdhere, Constantine Dovrolis, and Pierre Francois, "A Value-based Framework for Internet Peering Agreements," October 2010, [http://www.caida.org/~amogh/depeering_itc10.pdf] describe the traffic ratio rule as "widely used." (visited April 11, 2012).

[32] A sufficient condition for traffic ratios to be balanced is that (a) the ratio of subscribers to content providers is the same for both parties and (b) that subscribers to each party download the same amount of content from each of the content providers that are served by the other peer.

[33] PR Newswire, *Level 3 Issues Statement Concerning Internet Peering and Cogent Communications* (October 7, 2005), http://www.prnewswire.com/news-releases/level-3-issues-statement-concerning-internet-peering-and-cogent-communications-55014572.html (visited March 25, 2011). Another example is PSINnet's depeering of Exodus Communications [See P. Fusco, "PSINet, Exodus Terminate Peering Agreement," *Internet News.com*, April 5, 2000, http://www.internetnews.com/isp-news/article.php/334471/PSINet+Exodus+Terminate+Peering+Agreement.htm].

[34] The Register, *Level 3 depeers Cogent* (October. 6, 2005), http://www.the-register.co.uk/2005/10/06/level3_cogent (visited March 25, 2011).

[35] PR Newswire, *Level 3 and Cogent Reach Agreement on Equitable Peering Terms* (October 28, 2005), http://www.prnewswire.com/news-releases/level-3-and-cogent-reach-agreement-on-equitable-peering-terms-55637437.html (visited March 25, 2011). ("The modified peering arrangement allows for the continued exchange of traffic between the two companies' networks, and includes commitments from each party with respect to the characteristics and volume of traffic to be exchanged. Under the terms of the agreement, the companies have agreed to the settlement-free exchange of traffic subject to specific payments if certain obligations are not met."). One industry participant notes that "[m]any sender-keep-all [*i.e.*, settlement-free] peering arrangements are designed to convert to paid peering if either party exceeds the specified threshold." Declaration of Lyman Chapin at n.16, attached to Comments of Verizon, *Developing a Unified Intercarrier Compensation Regime*, Federal Communications Commission CC Docket No. 01-92 (May 23, 2005).

connectivity of the Internet."[36] To the contrary, the history of peering and transit relationships, including the emergence over the past decade of "[p]artial transit and paid peering [,] may be seen as efficiency-enhancing responses to changing market conditions."[37]

3. The costs of regulation

Since the beginning of the commercial Internet in the early 1990s, governments have not regulated Internet peering and transit agreements.[38] Nonetheless, despite the acknowledged success of the Internet in adapting flexibly to significant changes in both the amount and nature of the traffic, some have argued that the system of interconnection should be subject to detailed regulation, both of the forms of interconnection that are permitted and of the prices that can be charged for it.[39]

Detailed regulation—particularly regulation that would limit the ability for private parties to negotiate contracts with flexible payment and service terms for Internet transport, including terms that involve payment for peering services that have traditionally been settlement free— would be a mistake. The logic is simple. First, imposing restrictions on the ability for private firms freely to enter contracts that involve payments among them risks raising total costs of Internet transport, as well as prices to end consumers. Second, more generally, regulation is likely to generate substantial inefficiencies, yielding prices that do not adequately reflect costs and thus create improper incentives, as well as substantial risk of investment-retarding uncertainty or regulatory paralysis. Finally, concerns about the exercise of market power by particular ISPs should be assuaged by the existence of a large number of competitive alternatives to any peering arrangement. Each of these points is explained in the following sections.

3.1. Increased transport costs and prices to end consumers

3.1.1. Lack of incentives to minimize total costs

Regulation that limits the ability of private parties to pay one another for Internet transport may prevent the adoption of cost-minimizing arrangements for the exchange of traffic. Negotiations between ISPs and CDNs illustrate this point. When a CDN transports traffic from its customer (a content provider) to an ISP's network en route to the ISP's subscribers, both the CDN and the ISP bear costs of facilitating the transaction. Suppose that a CDN bears a cost of $3 million for transporting the content to an ISP, which then bears a cost of $8 million to transport this content to end users. Now suppose that the ISP becomes aware of a method for delivering the traffic that would increase its own costs by $1 million, to $9 million, but would reduce the CDN's costs by $2 million, to $1 million.[40] It would clearly be economically efficient for the ISP to adopt the new method, because doing so would increase its own costs by less than the reduction in the costs incurred by the CDN. In this example, if the CDN were to pay the ISP an amount between $1 million and $2 million to adopt the new method, both the CDN and the ISP (and ultimately their customers) would be better off.[41]

However, if regulation were to prohibit paid peering (or restrict substantially the terms that payments could take), it is unlikely that this efficient outcome would be achieved. The ISP would have no incentive to increase its own costs because it could not capture a portion of the reduction in the CDN's costs.

This example illustrates a general economic point: Price regulation—such as a requirement that certain types of interconnection be settlement-free—may be highly inefficient, particularly when different costs are subject to the control of different parties. Only if the parties can freely set payment terms as part of their negotiations will they have incentives to adopt distribution methods that minimize their combined costs.

3.1.2. Higher access prices to end users

The lack of incentive to minimize total costs is likely to raise prices to end users. However, even if costs do not increase, prices to end users are still likely to increase because many networks that make up the Internet can be conceptualized as platforms that link two "sides" of a

---

[36] Faratin and Clark at 67.

[37] *Id.* at 68.

[38] As a key Administration official recently remarked: "Interconnection is, of course, the single greatest imperative for a network of networks. And here the absence of governmental or intergovernmental controls is particularly striking. The physical and economic arrangements necessary for—that in a real sense constitute—interconnection have been worked out through normal adherence to international technical standards and through commercial negotiation. National governments, let alone international institutions, have not intervened to direct the creation of the controlling technical standards, have not mandated that the standards be observed, and have not prescribed the economic transfers that take place between and among the participating networks." Ambassador Philip L. Verveer, "Internet Must Stay Free of Government Control," The European Institute (January 2011) (*hereinafter,* "Verveer"), http://www.europeaninstitute.org/EA-January-2011/internet-must-stay-free-of-intergovernmental-control.html (visited March 24, 2011).

[39] Perhaps the most prominent example of recent proposals to regulate interconnection is Level 3's proposal to mandate settlement-free peering. *See,* Letter from John M. Ryan, Assistant Chief Legal Officer, Level 3 Communications, Inc. to Marlene H. Dortch, Secretary, Federal Communications Commission, December 7, 2010; Letter from John M. Ryan, Executive Vice President and Chief Legal Officer, Level 3 Communications, Inc. to Marlene H. Dortch, Secretary, Federal Communications Commission, January 14, 2011; and Letter from John M. Ryan, Chief Legal Officer, Level 3 Communications, Inc. to Marlene H. Dortch, Secretary, Federal Communications Commission, February 22, 2011. In its February 22, 2011 letter (at 3), Level 3 refers to its proposal as the "open, regional no-charges, interexchange model". See also Letter from Paul Kouroupas, Global Crossing, to Marlene Dortch, FCC, GN Docket No. 09-191 (February 4, 2011).

[40] These cost savings may, for example, reflect more efficient capacity utilization, which changes the rate at which each party needs to build out its own network—with associated capital costs—to accommodate increasing Internet traffic.

[41] Alternatively, suppose that a CDN can adopt a new method that would reduce its own costs by $3 million but which would raise the costs of the ISP by $5 million, thus raising total costs by $2 million. The CDN will have an incentive to adopt this *inefficient* method unless it is forced to pay at least a portion of the increased cost incurred by the ISP. Analysys Mason, *Overview of recent changes in the IP interconnection ecosystem,* May 2011, at 29, describes how a sending network can impose costs on a receiving network by handing off traffic at a location that imposes additional costs on the recipient while relieving some of its own costs.

market. For example, when a provider acts in its role as an ISP for university campuses, it connects end-user students with the full range of Internet traffic delivered by each interconnecting IP network. Similarly, when a provider acts as a residential access ISP, it connects its subscribers with the full range of Internet traffic provided by each interconnecting IP network. Accordingly, a useful framework for evaluating the effect of regulating the price of peering is the economics of "two-sided markets."

In two-sided markets, the platform that links multiple "sides" can, in principle, collect payments for its services in a number of different ways. For example, to the extent that a network acts as a residential access ISP, it could receive payments (1) only directly from its subscribers, (2) only directly from interconnecting IP networks supplying content and applications to those subscribers, or (3) from both. Indeed, charges in two-sided markets can take a variety of forms. For example, developers of computer operating systems traditionally charge only end users while not charging (perhaps even subsidizing) applications developers, whereas newspapers typically receive payments both from their readers and their advertisers.

In two-sided markets, revenue collected from one side of the market reduces the cost that the other side of the market must bear, lowering the prices charged to that other side. Put differently, in two-sided markets in which users on each side of the market benefit from transactions with users on the other side, if prices are higher on side A, the platform has a greater incentive to attract users to side B so that those users can attract and transact with the users on side A. As Rochet and Tirole explain:

> The linkage. .. shows up in the form of a simple "seesaw principle": a factor that is conducive to a high price on one side, to the extent that it raises the platform's margin on that side, tends also to call for a low price on the other side as attracting members on the other side becomes more profitable.[42]

In the present context, the seesaw principle says that if regulation were to mandate lower revenues from interconnecting CDNs, this would be associated with higher prices charged to subscribers.

Although the seesaw principle presupposes that both sides of the market benefit from transactions with each other, it does not require any assumptions about the degree of competition on each side of the market. Moreover, it is not unique to the Internet. For example, if a newspaper were unable to collect payments from advertisers, it would be expected to charge more to subscribers. In the present context, as one industry observer explains, "preventing network providers from exercising pricing flexibility... would simply increase the proportion of the network costs that providers must recover directly from end users. This simultaneously raises the prices paid by consumers and decreases the likelihood that the capital improvements [necessary to accommodate ever-greater traffic volumes] will ever be built."[43]

3.2. Other regulatory inefficiencies

Beyond the higher costs and higher prices already described, regulation is likely to generate substantial inefficiencies. In particular, given the expected increases in Internet traffic, the need for sophisticated congestion-based pricing systems seems likely to grow in the future, but would be particularly difficult to implement if Internet transport arrangements are highly regulated. The remainder of this section develops the economic logic behind this conclusion in more detail.

3.2.1. Prices that do not adequately reflect costs

The complexities in matching Internet usage to congestion costs make it likely that charges between networks will play an important part in an efficient system. To understand the need for sophisticated congestion-based pricing systems (likely including charges between networks), note that the cost of serving incremental demand on the network is not constant and is not always predictable in advance. Instead, it varies with the congestion on the relevant interconnected networks at any given time. As Christopher Yoo has explained:

> The classic economic solution to congestion is to set the price of incremental network usage equal to the congestion costs imposed on the network by that usage. However, determining the congestion cost imposed by any particular user at any particular time can be quite complex.... The contribution of any particular usage cannot be determined simply by counting the number of bits being transmitted. The overall impact of any particular increase in network usage can only be determined in light of other subscribers' Internet usage. Thus it may make sense to charge different amounts to users who are using the Internet to access the same content or application if a sufficient number of other users sharing the same bandwidth are using the network at the same time.[44]

In the absence of regulation, we would expect interconnecting IP networks to adopt pricing algorithms that produce different interconnection rates to CDNs at different times.

At the very least, market participants should be allowed to experiment with alternative pricing mechanisms that serve a variety of interrelated objectives, including recovery of network investments, avoidance of excessive complexity in consumer pricing, and increased Internet adoption and use. As Yoo has observed:

> A determination of the most efficient institutional form would require detailed analysis of the relevant cost data and network flows and would likely vary from network to network. Indeed, one might well expect different networks to pursue different pricing strategies. In addition, the data would need to be updated constantly in response to technological changes.[45]

[42] Jean-Charles Rochet and Jean Tirole, "Two-sided markets: a progress report," 37 *Rand Journal of Economics* 645, at 659 (2006).

[43] Yoo, *Architecture*, at 98.

[44] Yoo, *Architecture*, at 93.

[45] *See* Christopher S, Yoo, *Network Neutrality and the Economics of Congestion*, 94 Georgetown Law Journal 1847, at *1876* (2006) (*hereinafter*, "Yoo, *Network Neutrality*").

Yoo concludes that these considerations "provide one of the most powerful arguments against mandating or foreclosing any particular institutional arrangement."[46]

It is impossible to be definitive about how precisely CDNs and content providers will adapt to the imposition of congestion-based pricing, but we can suggest a number of possibilities. If they have to pay higher prices for using congested routes, they may make greater efforts to shift traffic to less congested ones. Content providers may also adjust to congestion-based prices by adapting their offerings to consumers to reflect these prices. Among the possible adaptations are: (1) offering unlimited consumer usage only during off-peak hours, much like the "free" evening and weekend service that is offered by some cellular operators; (2) downloading content during off-peak hours for storage and later viewing by consumers (rather than on-demand); (3) adopting more sophisticated compression schemes or reducing the amount of traffic that they generate in other ways;[47] and (4) imposing premium charges on consumers who demand service during congested periods.

Although the responses to congestion-based pricing are difficult to predict, and are likely to vary among CDNs and content providers, it is important not to restrict through regulation the options for recovery of the costs of interconnection in order to encourage both efficient investment in, and efficient usage of, the Internet infrastructure.[48] Clark, Lehr, and Bauer conclude that "...it may be efficient for payments to flow between [Content Delivery Networks] and access ISPs that may be justified as contributing to covering the increased costs incurred in delivering high volumes of content traffic downstream."[49]

3.2.2. Other regulatory inefficiencies

Attempts to regulate prices in two-sided markets require "a detailed knowledge of demand that may not be available to a regulator" and can potentially produce large inefficiencies.[50]

Regulation of peering rates in particular would require many complicated technical questions to be resolved. It would require, as a threshold matter, the determination of *which* IP networks would be required to peer with which others and on what terms (i.e. paid peering, settlement free peering, or transit). The fact that IP networks often (and increasingly) play multiple roles would make this an extremely difficult task since many simultaneously serve many *originators* and many *recipients* of Internet traffic. Moreover, as the FCC has explained, any given end user can be simultaneously both a "content provider" *and* a "content recipient,"[51] so it is artificial to place end users into one category or another. This rapidly evolving multiplicity of roles played by each participant in today's Internet would subvert any regulatory effort to divide the Internet neatly into, for example, "backbone" and "eyeball" networks for purposes of regulatory intervention.

Price regulation would also be likely to draw regulators into disputes about where, and on what technical terms, networks would have to interconnect. As discussed above, the relative costs two networks bear for transporting traffic depend on their physical points of interconnection and the distance each network must carry traffic, among other factors. If one network hands off traffic to another as soon as possible (so-called "hot potato routing"), it imposes on the other network the very substantial additional costs of carrying the traffic a longer distance. Regulators could not simply entitle every originating network to drop off its traffic to each recipient ISP at any point of the originating network's choosing, no matter how costly and inefficient that point might be from the recipient's perspective. Instead, if they were to attempt to mandate settlement-free interconnection, regulators would likely be asked, in a variety of contexts, to prescribe the precise physical points of interconnection between the two networks, which would likely differ among pairs of interconnecting networks depending on factors including their particular architectures, the locations of their respective customers, and other variables. Such prescriptions would be particularly difficult, and thus particularly likely to generate inefficiencies, in an environment that is changing as rapidly as Internet interconnection.

Regulators also would have to resolve other complex questions. As one example, because content providers, CDNs, and other interconnecting networks can redirect large streams of traffic on a real-time basis, regulators would have to decide what obligations an ISP had to maintain capacity in anticipation of such changes in traffic patterns, or traffic growth more generally. If a content provider decided to shift large amounts of traffic from one CDN to another, would an ISP be required to have in place "extra" capacity in its interconnection facilities with the second CDN because of the possibility of that shift? Would it have to leave the existing capacity in place with the first CDN in case the content provider later changed its mind? Who would pay for this "excess" capacity?

The complexities involved in a scheme to regulate peering leave two possibilities. First, a regulator could leave key aspects of the rules unspecified and defer them

[46] *Id.*

[47] An example is the decision by Netflix to reduce the amount of data that it generates in Canada with only a "minimal" effect on motion picture quality. See Alastair Sharp, *Netflix cuts data use on Canada online service*, Reuters, March 29, 2011, http://www.reuters.com/article/2011/03/29/us-netflix-canada-idUSTRE72S3BT20110329 (visited April 11, 2011).

[48] Although congestion-based pricing may also be applied to consumers, that would not eliminate the value from adoption of efficient pricing for transactions between networks.

[49] David Clark, William Lehr, and Steven Bauer, "Interconnection in the Internet: the policy challenge," August 9, 2011, presented at the 39th Research Conference on Communication, Information and Internet Policy, September 23–25, 2011, at 2 http://people.csail.mit.edu/wlehr/Lehr-Papers_files/Clark%20Lehr%20Bauer%20TPRC2011%20Interconnection.pdf).

[50] Weyl at 1666; Faratin and Clark at 68 ("We also have a cautionary conclusion: if one should be motivated (for whatever reason) to contemplate some regulatory rule to manage interconnection, the design of such a rule will be both complex and informationally demanding."). More recently, Clark, Lehr, and Bauer, op. cit. (at 21) argued that: "...we are not convinced by the evidence we have seen to date that more activist policies (e.g., direct regulation of Internet interconnection) is warranted; and equally important, even if we were to see a need for such regulation, we are concerned any such regulation might cause more harm than good."

[51] *See, e.g.*, Federal Communications Commission, *Report and Order In the Matter of Preserving the Open Internet, Broadband Industry Practices*, GN Docket No. 09-191, WC Docket No. 07-52, Adopted: December 21, 2010, Released: December 23, 2010, §20.

to case-by-case adjudication. That approach would leave the operators of Internet networks uncertain about both (1) the costs associated with mandated interconnection and (2) the circumstances in which they would or would not be able to charge other networks to recoup some of those costs. Such uncertainty would deter investments in Internet networks at a time when such investment is particularly critical, given rapidly expanding traffic demands.[52] Indeed, the economics literature confirms that price regulation would be a "particularly unattractive" solution in the type of two-sided market at issue here, where each side of the market benefits from greater participation on the other side.[53] In such markets, the observation that price regulation tends to reduce quality. for reasons including the investment-deterring effects of regulation, is "even stronger" than usual.[54]

Alternatively, the regulator could (at least in theory) impose precise rules specifying each of the relevant terms. That approach, however, would prevent firms from experimenting with new approaches to interconnection peering and pricing. That inflexibility, moreover, would also come a particularly inauspicious time, when the Internet in general, and traffic flows in particular, are changing rapidly, and when regulators and providers alike have only limited information concerning other relevant factors, including consumer receptivity to usage-based pricing.[55]

Price regulation also risks creating regulatory arbitrage opportunities, as parties design business models and practices aimed at taking advantage of "free" interconnection. For example, if some connections to an ISP's network were free or priced below cost, traffic would migrate to those connections even if they were not the most efficient ones to use. The result of such arbitrage would be to shift substantial costs from content providers, CDNs, and other interconnecting networks to ISPs and their customers.

3.3. Constraints on the exercise of market power

Given the inefficiencies associated with price regulation discussed above, one would need to demonstrate that ISPs have such strong market power over CDNs (and their content provider clients) that ameliorating its effects through government regulation would more than offset these inefficiencies. However, CDNs (and their content provider clients) and ISPs have alternatives to direct peering, and those alternatives limit whatever negotiating leverage an ISP would otherwise have. As discussed, any CDN or other IP network normally has a choice of several alternative paths into an ISP's network, and it is capable of rerouting traffic among these paths in real time. For example, CDNs can—and we understand that Akamai, Limelight, and others do—deliver traffic bound for an ISP's end users by purchasing transit services from one or more of that ISP's peering partners, which in turn exchange traffic with the ISP *on settlement-free terms*.

In some cases, a CDN can even send traffic over the ISP's own paid transit connections, in which case the ISP *pays for the traffic*. As Faratin and Clark explain, if one network denies settlement-free peering privileges to others, those other networks, "if they can control the routing of their traffic," can "cause their traffic to/from the prospective peer to route over the peer's transit connection to raise the peer's transit costs in order to induce it to peer."[56]

These transit alternatives play an important role in negotiations between access ISPs and CDNs that serve large content providers. As discussed above, large content providers are generally multihomed and thus divide their traffic among multiple CDNs or other IP networks. A CDN, in turn, may divide traffic that is bound for a particular access ISP among many different transit providers and shift among them in real time in response to congestion delays and other factors. Each of the CDN's transit providers may itself be either a settlement-free peer of the ISP, in which case the ISP receives no compensation for receiving the traffic,[57] or a transit provider to that ISP, in which case the ISP generally *pays* the transit provider for the greater traffic volume. In sum, in negotiations with an ISP about the terms of paid peering, a CDN can threaten to exploit transit alternatives that would leave the ISP worse off than if it had entered into a reasonably priced paid peering

---

[52] The "real options" model—a standard model in economics and finance—implies that increasing the uncertainty faced by a firm reduces its incentives to make investments (or, more precisely, creates incentives to delay investments until more of the uncertainty is resolved). See Avinash K. Dixit and Robert S. Pindyck, *Investment Under Uncertainty* (1994). Substantial empirical work has confirmed that conclusion. See, *e.g.*, Laarni T. Bulan, "Real Options, Irreversible Investment and Firm Uncertainty: New Evidence from U.S. firms," 14 *Review of Financial Economics* 255 (2005) and Timothy Dunne and Xiaoyi Mu, Investment Spikes and "Uncertainty in the Petroleum Refining Industry," 58 *The Journal of Industrial Economics* 190 (2010).

[53] Weyl at 1667. The observations in this paragraph are valid only for this type of two-sided market. For example, they would not necessarily hold for the two-sided market mediated by newspapers, where readers may prefer less advertising.

[54] *Id.* at 1666.

[55] Verveer, op. cit., notes: "A very important aspect of the absence of closely controlling laws is the ability of the Internet to evolve quickly and freely in response to changes in technology, commercial practice, and consumer preference. This dynamism is critically important. Its protection is one of the central concerns of the Internet policies of the United States. These policies are rooted in the view that nothing should be done that would impair the ability of the Internet to continue to grow and evolve based on the needs and desires of its users. A different way to say the same thing is that, in the view of U.S. policy, the world's people should continue to benefit from the efficiencies that advancing technology and business practice make available on an ongoing basis." Clark, Lehr, and Bauer, op. cit., at 21, make a similar point: "The Internet ecosystem is evolving rapidly and the expansion in interconnection agreements seems consistent with the need to accommodate new types of business relationships and service requirements."

[56] Faratin and Clark at 63; *see also* van der Berg at 2 ("Allegedly, a big American software company was refused peering by one of the incumbent telco networks in the north of Europe. The American firm reacted by finding the most expensive transit route for that telco and then routing its own traffic to Europe over that link. Within a couple of months, the European CFO was asking why the company was paying out so much for transit. Soon afterward, there was a peering arrangement between the two networks."). According to the chief executive officer of Exodus Communications, when PSINnet depeered Exodus Communications, "Thirty minutes after the lines went down, Exodus traffic was transferred through Sprint," [See P. Fusco, op. cit].

[57] By shifting traffic among transit providers, a CDN (or its content-provider customers) can help prevent traffic imbalances that would threaten any given transit provider's settlement-free peering relationship with a given "eyeball" ISP.

relationship with the CDN. Clark, Lehr, and Bauer have observed that:

> ...the complex mesh of interconnections, with diverse pricing models, constrains the range of negotiating positions that can be sustained by [an access network]. In particular, we assert that the limit on the payment that [an access network] can extract from [a content delivery network] will be related in some way to the customary price for transit, which is a commodity product... [58]

It is also important to observe that higher peering prices charged to a CDN affect content providers only if the CDN passes along those prices to its content-provider customers. Given this pass-through, however, content providers would be expected to react. Either directly or via the CDNs that serve it, a content provider can inform an ISP that it will offer particular packages of content to the ISP's subscribers only if the ISP limits the price it charges for peering. Alternatively, those content providers that charge end users directly can charge higher rates to the particular ISP's end users but not to the end users of rival ISPs. Either would have the effect of harming the ISP's broadband offering and potentially causing it to lose subscribers to other providers.[59]

These options are not hypothetical; indeed, content providers currently employ them. For example, ESPN has successfully induced most of the nation's leading ISPs to *pay* fees as a precondition to giving those ISPs' subscribers access to premium content, lest the ISPs lose those subscribers to competitors that do pay ESPN's fees.[60] ESPN imposes additional pressure by displaying a message to end users who seek to access to its website but subscribe to a non-"participating" ISP, listing ISPs that carry the website and encouraging the users to switch.[61]

4. Conclusion

The interconnection arrangements that currently characterize the Internet are far different from those that existed only a decade or so ago. In this new environment, backbone providers, ISPs, and suppliers of content have a far wider array of interconnection alternatives, both technical and financial, than they did only a short time ago. This new regime creates incentives for industry participants to minimize interconnection costs, permits costs to be recovered from both sides of "two sided" markets, and encourages the efficient use of prices to ration scarce network capacity. Given the Internet's success in adapting its interconnection arrangements to significant changes in both the amount and nature of traffic, it would be unwise to introduce what would almost certainly be costly and inefficient government regulation in an attempt to improve upon it.

Acknowledgements

The authors would like to acknowledge the valuable comments of Samir Jain, Jon Nuechterlein, and an anonymous referee on earlier versions of this paper. Although portions of the paper were developed as part of a project funded by Comcast, the final paper represents the independent views of the authors.

References

47 U.S.C. § 254 (g).

Ask Dr. Peering, 2009. Editor's Note for Paid Peering and Net Neutrality. November 5, 2009. <http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2009/11/5_Paid_Peering_and_Net_Neutrality.html> (retrieved 07.01.11).

AT&T's Global IP Network Settlement-Free Peering Policy, 2010. <http://www.corp.att.com/peering/> (retrieved 26.12.10).

Besen, S.M. et al., 2002. Evaluating the competitive effects of mergers of Internet backbone providers. ACM Transactions on Internet Technology 2 (3), 187–204.

Bulan, L.T., 2005. Real options, irreversible investment and firm uncertainty: new evidence from U.S. firms. Review of Financial Economics 14, 255.

CIDR Report, April 9, 2012. <http://www.cidr-report.org/as2.0/> (retrieved 09.04.12).

Clark, D., Lehr, W., Bauer, S., 2011. Interconnection in the Internet: the policy challenge. In: Presentation at the 39th Research Conference on Communication, Information and Internet Policy, September 23–25, 2011, p. 2. <http://people.csail.mit.edu/wlehr/Lehr-Papers_files/Clark%20Lehr%20Bauer%20TPRC2011%20Interconnection.pdf>.

Deutsche Telekom's Settlement Free Peering Policy, Version 3.0, March 10, 2009.

Developing a Unified Intercarrier Compensation Regime, 2005. Federal Communications Commission, May 23, 2005. CC Docket No. 01-92.

Dhamdhere, A., Dovrolis, C., 2008. Ten Years in the Evolution of the Internet Ecosystem. IMC '08, October 20–22, 2008.

[58] Op. cit., at 21–22. Although a recent paper by Steffen Lippert and Giancarlo Spagnolo (2008, "Internet peering as a network of relations," *Telecommunications Policy*, 32:33–49) argues (p. 34) that "...*bilateral* monetary payments—bilateral paid peering—may actually end up harming the ecosystem of the peering networks...", it is comparing bilateral paid peering not to *unpaid* peering but rather to *paid multilateral* peering. Indeed, the paper is clear that allowing paid peering can only expand the set of feasible contractual arrangements, whether bilateral or multilateral, and thus can only improve the efficiency of the outcome (p. 43). Although, the paper makes the straightforward point that multilateral contractual relationships may create more opportunities for beneficial trade than do bilateral ones, regulation is likely to be particularly ineffective in promoting multilateral peering relationships, as these are precisely the types of complex relationships that may arise in markets but which regulation is likely to be ill-suited to anticipate and accommodate. Moreover, introducing additional regulations, particularly those that limit the use of paid peering generally, could prevent both efficient bilateral and multilateral arrangements that might otherwise occur.

[59] This is one of several respects in which the Internet differs markedly from the Public Switched Telephone Network context. If a local telephone carrier charges supracompetitive "terminating access charges" to the calling party's long-distance carrier, its customers—the called parties—are insulated from the effects of higher terminating access charges because the long-distance carrier paying those access charges has no means of passing them through to those called parties. *See, e.g.*, 47 U.S.C. § 254(g) (requiring long distance carriers to charge "averaged" rates to all customers within a service class). As a result, called parties in the PSTN environment have no incentive to switch to a local telephone carrier that pays efficiently low terminating access charges. By contrast, an ISP's end users are *not* necessarily insulated from any fees the ISP imposes on interconnecting networks.

[60] *See* Analysys Mason, op. cit. at 13, for a brief description of this arrangement.

[61] *Help/Faq*, ESPN3, http://espn.go.com/espn3/faq#1 (visited March 25, 2011) ("How do I get access to ESPN3? ESPN3 is available nationwide, but you must subscribe to a participating high speed internet service provider. Click here for a list of participating providers. Click here to find out more on how you can request access to ESPN3 or *switch your service to a participating high speed internet service provider*.") (emphasis added).

Dhamdhere, A., Dovrolis, C., 2011. Twelve years in the evolution of the Internet ecosystem. IEEE/ACM Transactions in Networking 19 (5), 1420–1433.

Dhamdhere, A., Dovrolis, C., Francois, P., October 2010. A Value-based Framework for Internet Peering Agreements. <http://www.caida.org/~amogh/depeering_itc10.pdf> (retrieved 11.04.12).

Dhamdhere, A., Cherukuru, H., Dovrolis, C., Claffy, K., 2012. Measuring the Evolution of Internet Peering Agreements. <http://www.caida.org/publications/papers/2012/measuring_evolution_internet_peering/measuring_evolution_internet_peering.pdf> (retrieved 21.12.12).

Dixit, A.K., Pindyck, R.S., 1994. Investment Under Uncertainty. Princeton University Press.

Dunne, T., Mu, X., 2010. Investment spikes and uncertainty in the petroleum refining industry. The Journal of Industrial Economics 58, 190.

Edwards, B., Hofmeyr, S., Stelle, G., Forrest, S., 2012. Internet Topology over Time, pp. 5–6. <http://arxiv.org/pdf/1202.3993v1.pdf> (retrieved 07.09.12).

Faratin, P., Clark, D., et al., 2008. The Growing Complexity of Internet Interconnection. Communications & Strategies. No. 72, 4th quarter 2008, pp. 51–71.

Fusco, P., April 5, 2000. PSINet, Exodus Terminate Peering Agreement. Internet News.com. <http://www.internetnews.com/isp-news/article.php/334471/PSINet+Exodus+Terminate+Peering+Agreement.htm>.

Gao, L., 2001. On inferring autonomous system relationships in the Internet. IEEE/ACM Transactions on Networking 9 (6), 733–745.

Gill, P., Arlitt, M., Li, Z., Mahanti, A., May 21, 2008. The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse. <http://www.hpl.hp.com/techreports/2008/HPL-2008-47.pdf> (retrieved 07.09.12).

Haddadi, H., Uhlig, S., Moore, A., Mortier, R., Rio, M., 2008. Modeling Internet topology dynamics. ACM SIGCOMM Computer Communications Review 38 (2), 65–68.

Jefferies & Company, Inc.,2010. Limelight Networks: A Great Way to Play the Rise of Online Video, But Profits Are Still Elusive, at 2 (December 2, 2010).

Kende, M., May 2011. Overview of Recent Changes in the IP Interconnection Ecosystem. Analysys Mason.

Kouroupas, P., February 4, 2011. Global Crossing. Letter to Marlene H. Dortch, Secretary, Federal Communications Commission. GN Docket No. 09-191.

Labovitz, C., et al., 2009. ATLAS Internet Observatory Annual Report. <http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf> (retrieved 24.03.11).

Labovitz, C., Ikel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F., 2010. Internet Inter-Domain Traffic. <http://www.sigcomm.org/sites/default/files/ccr/papers/2010/October/1851275-1851194.pdf> (retrieved 07.09.12).

Level 3 and Cogent Reach Agreement on Equitable Peering Terms, 2005. PR Newswire, October 28, 2005. <http://www.prnewswire.com/news-releases/level-3-and-cogent-reach-agreement-on-equitable-peering-terms-55637437.html> (retrieved 25.03.11).

Level 3 depeers Cogent, 2005. The Register, October 6, 2005. <http://www.theregister.co.uk/2005/10/06/level3_cogent> (retrieved 25.03.11).

Level 3 Issues Statement Concerning Internet Peering and Cogent Communications, 2005. PR Newswire, October 7, 2005. <http://www.prnewswire.com/news-releases/level-3-issues-statement-concerning-internet-peering-and-cogent-communications-55014572.html> (retrieved 25.03.11).

Lippert, S., Spagnolo, G., 2008. Internet peering as a network of relations. Telecommunications Policy 32, 33–49.

Oliveira, R., Zhang, B., Zhang, L., 2007. Observing the Evolution of Internet AS Topology. SIGCOMM '07, August 27–31, 2007.

Oliveira, R., Pei, D., Willinger, W., Zhang, B., Zhang, L., 2008. In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure. ACM SIGMETRICS '08, June 2–6, 2008. <http://www.cs.arizona.edu/~bzhang/paper/08-sigmetrics-completeness.pdf> (retrieved 21.09.12).

Owen, B.M., 1999. The Internet Challenge to Television. Harvard University Press, p. 202.

Qwest's North America IP Network Peering Policy, 2011. <http://www.qwest.com/legal/peering_na.html> (Retrieved 05.01.11).

Report and Order In the Matter of Preserving the Open Internet, Broadband Industry Practices, 2010. Federal Communications Commission. GN Docket No. 09-191, WC Docket No. 07-52, Adopted: December 21, 2010, Released: December 23, 2010, §20.

Rochet, C.J., Tirole, J., 2006. Two-sided markets: a progress report. Rand Journal of Economics 37, 645–667.

Ryan, J.M., December 7, 2010. Assistant Chief Legal Officer, Level 3 Communications, Inc. Letter to Marlene H. Dortch, Secretary, Federal Communications Commission.

Ryan, J.M., January 14, 2011. Executive Vice President and Chief Legal Officer, Level 3 Communications, Inc. Letter to Marlene H. Dortch, Secretary, Federal Communications Commission.

Ryan, J.M., February 22, 2011. Chief Legal Officer, Level 3 Communications, Inc. Letter to Marlene H. Dortch, Secretary, Federal Communications Commission.

Sharp, A., March 29, 2011. Netflix Cuts Data use on Canada Online Service. Reuters. <http://www.reuters.com/article/2011/03/29/us-netflix-canada-idUSTRE72S3BT20110329> (retrieved 11.04.11).

Suddenlink Communications' Settlement-Free Interconnection (Peering) Policy, 2011. <http://www.suddenlink.com/terms-policy/peering.php> (retrieved 05.01.11).

Van der Berg, R., September 2, 2008. How the 'Net Works: an Introduction to Peering and Transit. Ars Technica. <http://arstechnica.com/old/content/2008/09/peering-and-transit.ars> (retrieved 24.03.11).

Verizon Business Policy for Settlement-Free Interconnection with Internet Networks, 2011. <http://www.verizonbusiness.com/terms/peering> (retrieved 05.01.11).

Verveer, P.L., January 2011. Internet Must Stay Free of Government Control. The European Institute. <http://www.europeaninstitute.org/EA-January-2011/internet-must-stay-free-of-intergovernmental-control.html> (retrieved 24.03.11).

Yoo, C.S., 2006. Network neutrality and the economics of congestion. Georgetown Law Journal 94, 1847–1876.

Yoo, C.S., 2010. Network neutrality or Internet innovation? Regulation, 22–29.

Yoo, C.S., 2010. Innovations in the Internet's architecture that challenge the Status Quo. Journal of Telecommunications and High Technology Law 8, 79–100.