Interconnection in the Internet: the policy challenge

David Clark[1]

William Lehr[2]

Steven Bauer[3]

Massachusetts Institute of Technology[4]

August 9, 2011

Prepared for

The 39th Research Conference on Communication, Information and Internet Policy (www.tprcweb.com), George Mason University, Arlington, VA, September 23-25, 2011.

Abstract

In days past, Internet Service Providers (ISPs) relied on two basic types of contracts for exchanging traffic (peering and transit) and ISP interconnection was not regulated. As we explained in (Faratin, Clark et al. 2007), the world of Internet interconnection is no longer so simple. The increased complexity poses significant challenges for policymakers who might contemplate regulating Internet interconnection, so it is perhaps lucky that calls for Internet interconnection regulation have been muted to date. That quietude was threatened in late 2010 by two events: the issuance of the FCC's Network Neutrality order[5] and the dispute between Level 3 and Comcast over their interconnection agreements.[6]

Regardless of how one views the Level 3/Comcast dispute or its relationship to on-going discussions about broadband access regulation,[7] we believe that there remains an

---

[1] Corresponding author: ddc@csail.mit.edu, 617-253-6003

[2] wlehr@mit.edu.

[3] bauer@mit.edu.

[4] The authors would like to acknowledge support from NSF Awards 1040020, 1040023, the MIT Communications Futures Program (http://cfp.mit.edu), and the MITAS Project (http://mitas.csail.mit.edu). All opinions expressed herein are those of the authors alone.

[5] See *Report and Order*, In the Matter of Preserving the Open Internet and Broadband Industry Practices, Before the Federal Communications Commission, GN Docket No. 09-191 and WC Docket No. 07-52, Adopted December 21, 2010 (available at: http://www.fcc.gov/Daily_Releases/Daily_Business/2010/db1223/FCC-10-201A1.pdf).

[6] See, Ante, Spencer and Amy Shatz (2010), "Web-traffic spat over Netflix highlights new tensions," Wall Street Journal, November 30, 2010.

[7] As of February 2011, the FCC Commissioner was claiming that the FCC's Net Neutrality order did not apply to disputes like that between Comcast-Level 3 (see, "WSJ Update: FCC Chairman –

enduring public interest in ensuring a healthy Internet interconnection market and that a better understanding of the underlying economics impacting those markets is important to frame appropriate policies. In this paper, we examine the changing dynamics of Internet interconnection economics, focusing on the challenge of interconnecting CDNs and access ISPs.

Interconnection agreements do not just route traffic in the Internet, they also route money. Allowing money to flow from the end-users (the ultimate source of all funding other than that provided by public subsidies) to the providers of infrastructure services to allow them to recover their capacity-related costs is necessary in order to sustain infrastructure investment and a healthy Internet ecology; however, concerns about abuses of potential market power raise valid policy concerns. Our analysis provides a basis for understanding why revenue neutral peering with traffic-balance requirements may be yielding to new models of paid peering between CDNs and access ISPs. We conclude that it may be efficient for payments to flow between CDNs and access ISPs that may be justified as contributing to covering the increased costs incurred in delivering high volumes of content traffic downstream. While payments might be warranted in either direction, we would expect it to be more common to see those payments flowing from CDNs to access ISPs. Although this might be perceived to pose a risk to competition, we believe that competitive pricing for Internet transit prices provides an effective bound against abuse of such payment mechanisms by any access ISP who may be deemed to have market power. We also expect that usage-based retail pricing, probably in the form of tiered pricing, is likely to become more prevalent. While we believe that that is a justifiable outcome, we believe it is appropriate and not surprising that such pricing will attract the watchful eye of regulators. On the whole, we are cautiously optimistic about the competitiveness of interconnection markets, but believe efforts to enhance transparency into how these markets work would be beneficial. Transparency would be enhanced with better information about traffic patterns, the incremental costs of supporting increased usage, and about the terms, conditions, and norms that are emerging as interconnection markets continue to evolve.

1. Introduction

In the early days of the commercialization of the Internet, two sorts of interconnection agreements among ISPs emerged. One was a traditional customer-supplier arrangement, in which one ISP purchased *transit* service from another, perhaps larger ISP. An ISP that offers traditional transit service agrees to provide access to the entire Internet for its customers. The other arrangement was *peering*, in which two ISPs that each had traffic for the other agreed to interconnect to exchange that traffic directly. A peering arrangement between two ISPs does not give either ISP access to the entire Internet via the other; normally each ISP exchanges with the other ISP only traffic that is local to the region of that ISP and its customers. In other words, a peering agreement implies a *routing restriction* with respect to the traffic exchanged: the only traffic exchanged

Net Neutrality Rules Don't Cover Comcast-Level 3 Dispute," February 16, 2011). See Clark et al. (2009) for our comments to the FCC on its original order in this matter.

originates from the source ISP and its customers and terminates in the destination ISP and its customers.

In contrast to a transit arrangement, which is a commercial arrangement between a buyer and a seller in which monetary payments flow from the buyer to the seller as compensation for services rendered, peering (as the name might suggest) has been viewed as an interconnection among approximate equals, with value to both and no *a priori* obvious direction for monetary payments to flow. In early negotiations among potential peering partners, it became clear that it would be very difficult to determine if the balance of values favored one or the other ISP, and the convention emerged that peering was "settlement-free," or "revenue-neutral." Given the perception that negotiation about relative value would be costly (i.e. incur potentially high transaction or bargaining costs) and under the assumption of approximately equal value accruing from the relationship on each side, the approximation of settlement-free could be seen as economically efficient.

For a number of years, these two options—transit and revenue-neutral peering—captured the set of expectations among parties that negotiated about interconnection; they represented an informal norm or bargaining regime. But like many such informal norms, revenue-neutral peering has been breaking down slowly as the various parties to potential peering relationships no longer see the simplicity of the balanced value approximation as serving their needs. It is being replaced (as has happened in other venues like the bargaining over international trade agreements) with a period of more unconstrained bilateral negotiation among the parties. There are a number of reasons why this seems to be happening.

- In the past, many ISPs were more or less similar. There were small ISPs, and larger ISPs with larger footprints, of course, but many ISPs had the same mix of customers. When similar ISPs established peering interconnections, the similarity made the assumption of balanced value plausible. Today, ISPs are more specialized, with some serving broadband residential customers (sometimes called access or "eyeball" networks), some serving enterprise customers, perhaps with a highly distributed footprint, some serving high-volume content providers and so on.
- Networks that serve different sorts of customers may have very different internal cost structures. A residential broadband provider, with an extensive "outside plant" of fiber, HFC or copper pairs, may have a much higher cost, measured as a function of peak rate mb/s carried, than an ISP that only serves large customers with high speed connections running 1 gb/s or 10 gb/s. Under-sea providers again have a different internal cost structure.

At the same time, it is clear to many parties that they can save money by finding a way to negotiate a peering arrangement. If the alternative is for two ISPs to exchange traffic with each other by each purchasing transit service from a third intermediate, then both ISPs pay out to have that traffic transported. The idea of cutting a middleman out of the path and saving the cost of transit is obviously appealing, which leads to the objective of finding and negotiating some basis for direct connection, even if revenue-neutral peering is not an acceptable outcome to one of the parties. And thus other approaches to interconnection emerge, including the option of *paid peering.*

It might seem that one outcome of this negotiation could be that one ISP would purchase transit service from the other to achieve direct connection, but the objective of peering is different from transit, as we observed above—in peering two ISPs each want to exchange traffic that is restricted to their own network and that of their customers. Transit and paid peering are two distinct arrangements, and some ISPs may not be in the business of selling transit service (e.g. access to all the rest of the Internet) even if the other ISP were prepared to pay for transit. However, although transit and peering are distinct arrangements, we will argue that *the outcome of bargaining with respect to paid peering is constrained by the option of transit.* One of the conclusions of this paper will be to relate the pricing for transit to some hypotheses about the outcome of peering negotiations.

It may be that over time, this period of bilateral negotiation over peering will lead to the emergence of new norms—new regimes—for how peering should be handled. One possible benefit of such norms would be to reduce the transaction costs of negotiation by providing a commonly accepted framework and set of assumptions about the outcome of negotiation, just as revenue-peering did in its time.[8] Periods of unconstrained bilateral agreements often suggest new norms and patterns that better serve the evolved market. Another objective of this paper is to speculate about what this new generation of norms might be.

1.1. High volume content and its providers

A particular focus of attention today and for this paper is the sub-case of interconnection between residential broadband ISPs (access networks) and networks that deliver high volume commercial content (content delivery networks or CDNs). In the first quarter of 2011, 49.2% of peak-time traffic coming into residential access networks was real-time entertainment, with a single content provider (Netflix) accounting for 60% (Sandvine 2011). YouTube, from Google, represented another 23% of the real time entertainment content. The Netflix content is delivered by three commercial CDN operators: Akamai, Limelight and Level 3. So these three providers plus Google originate and control over 80% of all real-time entertainment content, or over 40% of all peak-time residential download traffic.[9] Collectively, we will call these *content delivery* networks, in contrast to *access* networks. A cost-effective connection between a content delivery network and an access network will often be of advantage to both, because of the high volumes of data being transferred.

[8] Common agreement on norms may reduce transaction costs by reducing information asymmetries or learning costs, and provide a substitute for complete contracts that might otherwise be needed to mitigate concerns of opportunistic behavior (see for example, den Butter 2010). On the other hand, if the equilibria for efficient agreements shift, the constraints imposed by current norms could be seen as an impediment to efficient negotiation. New norms may be needed to accommodate exogenous marketplace developments (see for example, Bertrand 2009). Thus, dynamic optimization of transaction costs may require the evolution of norms.

[9] Data from Sandvine is aggregated across all of their North American clients. Of course, ISPs' traffic experiences differ, and even within a single ISPs network, there is significant variation across subscribers, time, and geography.

Negotiation between a content delivery net and an access net about direct connection is at heart a peering negotiation, because it would normally imply the same *routing restriction* as we described above. Specifically, the CDN is normally only trying to get access to the customers of the access network.[10] But because of the high volumes of traffic potentially involved and the fact that the traffic is highly asymmetric (from content to access), negotiation about payment may be commonplace. The carriage of this high-volume traffic will generate costs for access networks, which will naturally attempt to recover those costs; while at the same time one might argue that access to the content creates value for the access network's customers. The outcome of these negotiations seems to be, in practice, that the content delivery networks are making payments to access networks.[11] We are interested in understanding both the direction and magnitude of these sorts of payments.

Interconnection between content delivery and access networks has attracted much attention recently. From a business point of view, these interconnection arrangements are of considerable interest because of the high volumes of data exchanged, and the implications of this high volume for internal costs. Industry observers and regulators are interested because the high content-related value of this data raises questions as to whether one or another actor will have sufficient market power to benefit by extracting rents that derive from the value of the content, not its delivery.

The desire to separate content and delivery or "conduit" concerns has a long regulatory history. Traditional telecommunications regulation focused on "conduit" regulation and the need to ensure common-carriage access to basic telecommunications services, for which public utility regulation was long justified by the concern that telephone networks were a natural monopoly.[12] On the other hand, regulation of content was manifest in broadcast and programming regulation and media cross-ownership rules.[13] The regulation

[10] In practice, there may be exceptions to this exact form of the restriction; we will return to this issue later in the paper.

[11] Our conclusion is tentative because there is limited publicly verifiable data on what is going on in interconnection markets. Our conclusions rely on the mosaic of anecdotal reports that appear periodically in the trade press, on email lists and blogs, and on our communications with industry contracts.

[12] Basic telecommunication services are regulated as common carrier services under Title II of the United States Communications Act of 1934 (as amended). Over-time, the FCC has grappled with defining the boundary between what services should or should not be regulated as basic telecommunication services (see, for example, Brock, 1994). In a succession of decisions since 2002, the FCC sought to reclassify broadband services as "information services," rather than basic telecommunication services. By so doing, they opened the possibility of shifting to a more light-handed form of regulation that could be more technology neutral, and evolve beyond the silo-based legacy regulatory models that characterized traditional telecommunications and cable television regulation. For a discussion of evolving broadband policy, see *Notice of Proposed Rulemaking*, In the Matter of Preserving the Open Internet and Broadband Industry Practices, Federal Communications Commission, GN Docket No. 09-191 and WC Docket No. 07-52, Released October 22, 2009 (hereafter, NPRM2009).

[13] First Amendment "Free speech" concerns induce a special level of concern for content-based regulation in the U.S., providing additional motivation for separating content and conduit

of cable television services, where investor interests in content and conduit services were conjoined, led to its own silo of regulation in the United States. With the rise of the Internet as the "new PSTN," policymakers are faced with the question of how much of the traditional regulatory model for telecommunications ought to be mapped over to the new world of the broadband Internet.[14] In their efforts to craft a framework that is more technology agnostic (i.e., is less focused on whether the underlying infrastructure has evolved from a telephone or cable television network or supports fixed or mobile services), policymakers have opted for increased reliance on market forces and more light-handed regulatory approaches.[15] This debate continues in its current incarnation as the debate over Network Neutrality.[16]

Policymakers are concerned that ISPs might engage in discriminatory network or traffic management practices that may interfere with competition, or worse, limit end-users access to content or applications of their choice. The implicit presumption is that access ISPs may have market power[17] and may seek to use such market power to earn monopoly profits from over-charging end-users or other participants in the value chain, such as application service providers or content providers. The focus of our concern here is on the potential threat to content providers and content delivery networks. To the extent that any discrimination might distort competition in content markets, it is especially worrisome to policymakers for the reasons noted above.

---

regulation as much as possible. Traditional content regulation includes such things as program access rules (e.g., to ensure that incumbents make programming available to other distribution channels), public programming obligations (e.g., children and news programming), and censorship rules (e.g., pornography restrictions).

[14] Although we expect most of our readers are familiar with the acronym, PSTN is short for the Public Switched Telecommunications Network.

[15] For example, a significant goal of the Telecommunications Act of 1996, which embodied significant reforms to the Communications Act of 1934, was to significantly expand the scope of competition in all telecommunication services and provide a roadmap for further deregulation. In the years that followed, talk turned to notions of layered regulation which would be more appropriate for the emerging world of facilities-based competition between multi-service, platform networks offered by telephone, cable television, and potential new entrants (see, for example, Sicker et al. 1999, Werbach 2002, or Friedan 2004.

[16] See for example, Jordan (2007) who relates the layered-regulation and network neutrality debates; or Lehr, Peha, et al . (2007).

[17] For example, some have argued that access ISPs have a terminating monopoly because subscribers face switching costs in switching to another broadband provider. This source of market power might exist even if the market for selecting a broadband service provider is competitive (which is, itself, something that folks have disputed). See NPRM2009, note 12 *supra* at paragraph 73. For a critique of this view, see *Testimony of Jeffrey Eisenach, PhD,* Before the Subcommittee on Communications, Technology and the Internet, Committee on Energy and Commerce, United States House of Representatives, April 21, 2010 (available at: http://www.naviganteconomics.com/docs/Eisenach%20Broadband%20Testimony%20042110%20With%20Attachments.pdf).

In this paper, we will try to maintain a clean distinction between two sorts of payments that might flow among parties. *Transport payments*, which we will indicate with the notation P_t, are payments among parties that cover the internal costs related to the delivery of flows of data. *Content payments*, P_c, are payments that relate to the value of the commercial content itself, not its transport.[18] While this distinction may not always be precise, we believe that it is an acceptable simplification for the purpose of this paper.[19]

Traditionally, there has been no P_c component in the payment that a residential customer pays for broadband access. If the consumer wants to receive fee-based commercial content, such as Netflix or the Wall Street Journal, the relevant content payments flow directly from the consumer to the content provider. However, this is starting to break down in small but suggestive ways. ESPN3, an online source of streaming sports content, has negotiated an arrangement with certain broadband providers in which they purchase bulk access to ESPN3 content for all their broadband customers, which implies that the cost of this is recovered by a part of the monthly bill that each customer pays. Right now, the per-customer P_c related to ESPN3 is about \$0.10 a month[20] but other content arrangements may follow.

Similarly, there has been no P_c component in the historical pricing of ISP interconnection. Transit is a service that provides access to any end-point on the network, not specific high-value content; to date and for the most part, it seems to be a competitive, commodity business.[21] However, the special case of interconnection between the content and access network raises the possibility that there is a P_c component to the negotiated payment. One of the goals of this paper is to propose an approach to determining whether the payment between the parties is likely to include a P_c component.

---

[18] We have introduced the notation of P_t and P_c to indicate two general classes of payment: transport and content, without relating them to any particular point of payment. When we want to discuss a particular payment flow, as from content delivery to access network, we will write P(CDN->A), which will imply a direction.

[19] The separation is not precise because how content is delivered may impact its quality, and hence the value to the end-user. For example, high-definition content may be down-coded for delivery to the small screen of a mobile handset without any deterioration in the user-experience as a way to economize on delivery (congestion) costs; alternatively, increased latency or congestion may severely adversely impact the user-experience, potentially to the point where the viewer decides to forego watching the content altogether.

[20] See http://finance.yahoo.com/news/ESPN-Charges-Broadband-Firms-ibd-345214765.html?x=0&.v=1

[21] We see no evidence from their published annual reports that CDNs are earning significant supranormal profits. Moreover, anecdotal evidence suggests that the prices for CDN delivery have followed the downward trend in transit pricing, and are quite modest. See, for example, http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php and http://blog.streamingmedia.com/the_business_of_online_vi/2010/06/data-from-q1-shows-video-cdn-pricing-stabilizing-should-be-down-25-for-the-year.html. This suggests that the ability of CDNs to capture more than the cost of content delivery from content-owners is limited. Of course, as with other types of interconnection, the markets for transit are also changing, with ISPs offering partial transit (i.e., transit with routing restrictions), variations on traffic commitments, and other innovations.

But we stress that payment from a content delivery network to an access network (e.g. a form of paid peering) does not automatically imply that there is a P_c component in the payment. The payment may be only a P_t in which the access network has successfully negotiated to recover some of its internal cost to carry the content.[22]

2. A simple model of interconnection and cost

Let us consider a single access network A, and a single content delivery network CD. Typical examples of A would be Comcast, Verizon, Time Warner or AT&T. Typical examples of CD would be Google (YouTube), Akamai or Limelight.

The emergence of high-volume commercial content has added cost to A, which we will denote C_A. CD also has costs to deliver this traffic, which include the costs of servers and communication links, and we will denote these as C_{CDN}.

There is a slight asymmetry to the situation of CD and A, which is worth keeping in mind. A has a substantial base cost which existed prior to the growth of high-volume commercial content. C_A is an incremental, usage-based component of the total cost of A. CD, on the other hand, exists only to deliver this content, and will tend to view their total costs as content delivery related. Of course, both parties have the goal of reducing their respective costs.

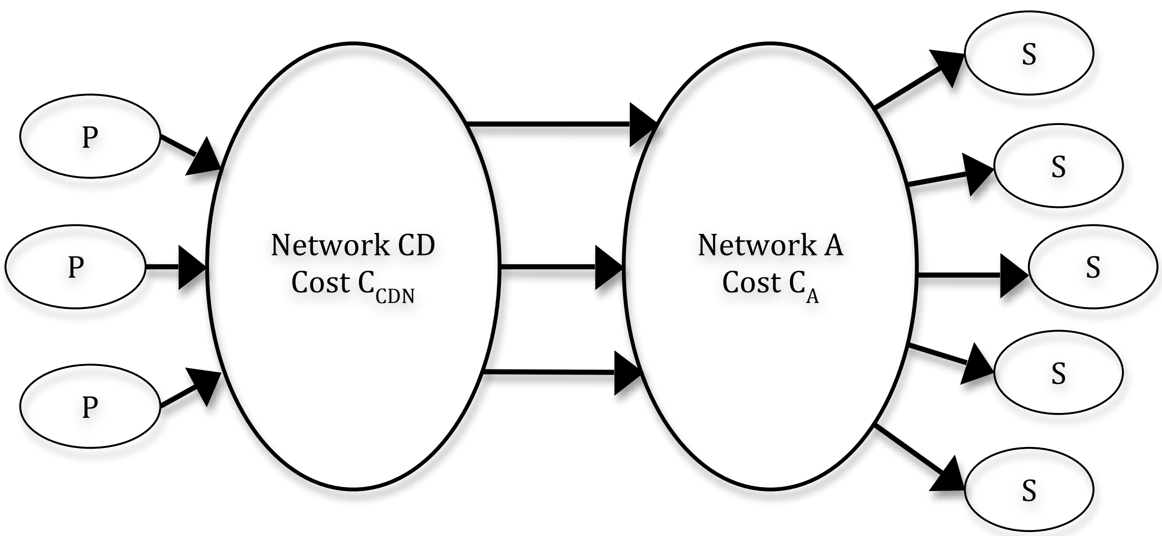

Figure 1: CDN interconnects with Access ISP A

Figure 1 illustrates this simplified model: Content traffic flows from the content producers or owners ("P") to the CDN ("CD") and then onto the Access ISP ("A"),

[22] Moreover, we assume that content production/creation are competitive, so that the ability of content-owners to extract excessive value from CDNs or ISPs is also limited. (That implies that P(CP->CDN) cannot be negative—CDNs do not pay the content providers for the privilege of carrying their content.)

eventually terminating at A's subscribers ("S"). C_{CDN} and C_A denote the incremental costs incurred by CD and A in delivering this traffic.[23]

2.1. A reality-check: the size of C_A

It would be possible to pose a model without discussing the actual values of the parameters. However, some understanding of real magnitudes proves helpful in understanding the current context for negotiations. The lack of publicly verifiable information about the terms and conditions for interconnection agreements makes it difficult to estimate C_A, but we believe it is possible to rely on anecdotal evidence from various papers and blog postings to provide some reasonable bounds on what these costs may be. In an earlier paper, we estimated the cost at $0.10/GB (Clark 2008),[24] *not including costs related to the access network*. Other writers have suggested lower numbers, in the range of $0.07 to $0.10.[25] Depending on the extent to which *access* network costs are allocated as fixed or usage-related, the per-GB costs may be considerably higher, perhaps $0.20 to $0.30/GB. These estimates are applicable for large urban/suburban broadband wireline access networks with low costs for their transit.[26] At one extreme, Netflix estimated that the costs were in the range of $0.01/GB.[27] We believe such an estimate is unreasonably low. At the other extreme, some have pointed to the $1.00/GB overage fees charged by some ISPs for subscriber usage that exceeds their usage-tier monthly allotment. These overage prices are not intended to reflect the incremental cost of additional usage, but to provide a strong inducement for subscribers to self-select into the appropriate usage tier. While these estimates cover a broad range, it is not so broad that we cannot reach some useful conclusions.

For example, Sandvine (2011) reported that average (mean) monthly download usage was currently 18.6 GB per customer (household).[28] If usage costs $0.20/GB, that would mean that the per-customer usage-related monthly costs are close to $4. To watch a 90-minute movie in HD from Netflix (at about 5 mb/s) would cost about $0.65. Although not precise, these numbers imply that *the emergence of high-volume content (video) has*

---

[23] Aguapong and Sirbu (2011) also examine the relationship between CDNs and access ISPs, presenting a model of how routing/interconnection choices might influence total costs and potential payment flows.

[24] This measurement in terms of cost per GB is perhaps confusing. As written, it is not a *rate* but a volume. That is, the implication is that it costs the ISP (say) $0.10 to deliver a GB of data, independent of rate. This characterization is obviously a simplification, but it implies that the cost to deliver a GB in one unit of time at one rate, or the cost to deliver that same GB at half the rate over twice the time is more or less the same. Another way of saying this is that $0.10/GB is a contraction of "$0.10/month for each GB/month".

[25] An estimate of $0.08/ GB (Canadian) is at http://www.michaelgeist.ca/content/view/5727/125/ and http://www.michaelgeist.ca/content/view/5952/125/.

[26] In contrast, these numbers will not apply to rural ISPs that may be far from peering and transit interconnection points, or smaller networks, wherever they are located

[27] See https://prodnet.www.neca.org/publicationsdocs/wwpdf/051211netflix.pdf.

[28] In contrast, the median download usage was 6 GB. The mean/median ratio over 3 is an indication of the heavy-tailed distribution of usage: some heavy users are very heavy indeed.

generated substantial new costs C_A. These costs are not so great as to destroy the viability of the service, but they are large enough that we can expect access networks to take explicit steps to recover these costs.[29]

2.2. Managing access costs C_A

There are roughly four ways that an access network A can manage usage-related costs C_A.

- By careful design of their network, specifically with attention to where the content delivery networks interconnect with them, they may be able to reduce the actual C_A incurred.
- They can negotiate to have the content network CD compensate them for some of these costs.
- They can increase the price of service to their retail customers:
  - They can allocate an equal share of C_A to all customers.
  - They can create mechanisms that discriminate among users based on some proxy of usage, and increase the price of service for these users.
- They can become less profitable (but if they are competitive, then that is not a sustainable option).

These seem the only general options; we will assert that all responses by A will be some mix of these.

2.3. The flow of payments

While Figure 1 may be helpful in understanding the physical flow of traffic, we are really interested in the flow of payments. Figure 2 redraws Figure 1 with arrows to suggest the possible direction of payments. The producers P are paying CD to redistribute their content, so payments will flow from P to CD. Payments will also flow from S, subscribers for broadband Internet access to the access ISP A. The remaining question is whether and how money flows between CD and A.

Figure 2 illustrates several possible cases:[30]

- P = 0: Traditional revenue neutral peering, where each network covers all its internal costs from its own customers.
- Payment from CD to A, where the payment from CD to A helps to cover some of the internal costs of A. Presumably, these costs are then passed through by CD to the content producers, who pay more to CD and thus indirectly cover the cost of

---

[29] This analysis contains a potentially dangerous simplifying assumption, which is that usage-related costs scale linearly with usage. In some cases this is true (each line card added to a router costs the same) and in other cases it is not true. For example, prices for transit are non-linear, with substantial discounts for larger volumes. Prices for usage also will come down over time. So these numbers should be thought of as a rough "tangent to the curve" approximation valid at a given time for a given overall level of usage.

[30] Figure 2 still represents an over simplification since it excludes potential direct payment flows from CuS to Pr (e.g., for subscriptions to premium content like Netflix of the WSJ) or potential payment flows from advertisers.

transporting their content across A. This outcome is fairly common in today's CDN market.

- Payment from A to CD. This outcome seems uncommon, but makes sense in certain circumstances. Consider the case where A is a small, rural ISP.[31] If there is no direct connection between CD and A, all of the content from the producers will come into A over a potentially very expensive transit link. Having CD make a direct connection to A may greatly reduce A's costs. However, if A is small, it may not be cost-effective for CD to connect to A; the connection might actually increase C_{CDN}, not reduce it. In this case, it might make sense for A to pay CD.

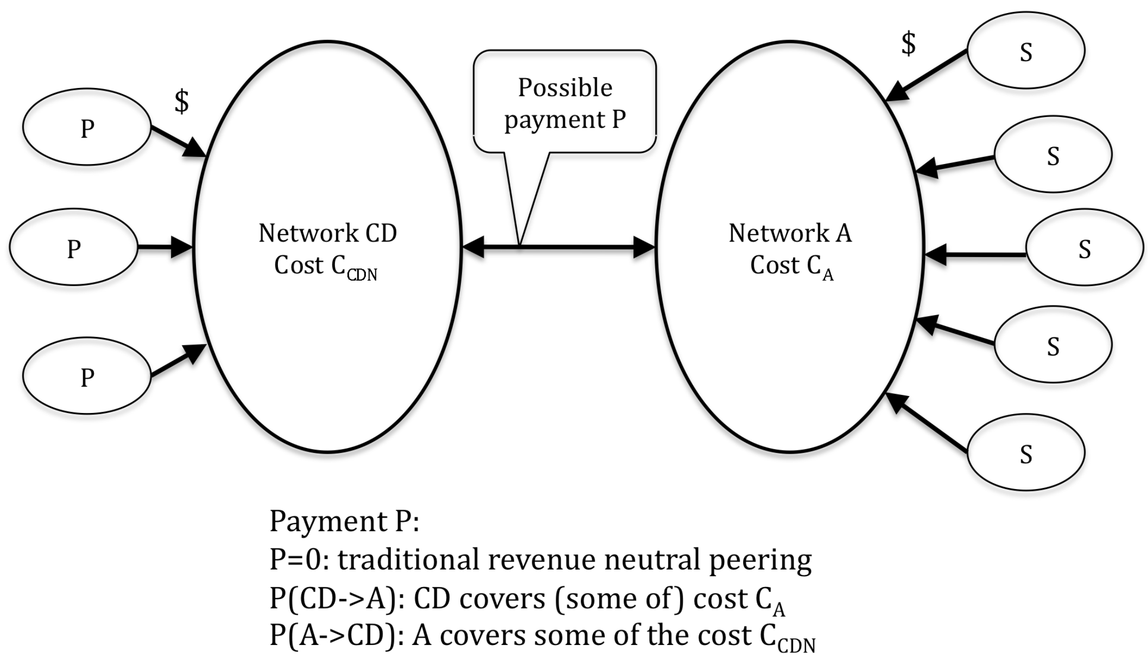

Figure 2: CDN to ISP A Money Flows

2.4. Why does CD pay A?

The figure above illustrates the payment patterns that we observe in practice today. When CDNs connect to large access networks, payment seems to most commonly flow from CD to A. Since both have costs, it is worth asking why payment should flow from CD to A instead of the other way.

One possible answer is that A has a better bargaining position, because it holds a terminating monopoly with respect to its customers. However, the pattern of payment may be as much a result of a persistent norm or common practice that is older than the emergence of high-value commercial content: a presumption that money flows and packet flows go in the same direction. That is, if X is delivering packets to Y, then (if there is payment) X pays Y and not the other way round. Of course, the difficulty of

[31] This may be a more common occurrence internationally, especially in light of the fact that so much of the content on the Internet still originates from the US.

justifying this assumption is what led to revenue-neutral peering in the first place. One rule that ISPs use to determine if they will agree to revenue-neutral peering is that the traffic in the two directions is roughly in balance. But if they go out of balance, one or the other party may use it as a useful context for renegotiating the agreement.[32] However, if there is complaint about the imbalance in costs associated with the asymmetric traffic, anecdotally it is often the party receiving the excess traffic that complains. But increased traffic (say from X to Y) adds to the costs both for X and Y. So why would X pay Y?

There seems to be an unstated assumption that a transfer is of more benefit to the originator than the receiver, so the sender should be expected to cover more of the delivery costs. This assumption is not always true: for example when a user downloads a large open-source software package (e.g. a Linux release) the benefit is essentially all to the receiver. However, these cases seem to be ignored as part of the current regime of bargaining.

So long as payment between the parties is an acceptable outcome, the parties will find it profitable to interconnect so long as the net benefit to all parties (including the payments) is positive. Paid peering can lead to more direct connections, which presumably reduces overall system cost and increases total surplus.[33]

This discussion concerns how costs of delivery are covered: they concern *transport payments,* not *content payments*. However, the possibility of non-zero payments also raises the possibility that one actor (for example the access network A) might have enough market power (e.g, because it is a terminating monopoly with respect to its customers) to demand a payment from CD that exceeds its internal costs C_A. In this case, we should assume that the payment is not just a *transport payment* P_t, but includes as well a *content payment*, P_c. Regulators and industry observers have worried that access networks might have enough power to demand content payments, and this would signal the potential for unacceptable discrimination and manipulation in the business of content production. So the obvious question follows: if we allow non-zero values for P, how can we distinguish content payments from transport payments?

3. Bounding the outcome of peering negotiations

One way to try to understand the context of negotiation between CD and A is to speculate about their relative market power. A has a terminating monopoly with respect to its customers, but CD may be hosting valuable content that the customers of A demand. So can CD hold up A, or can A hold up CD? Earlier we noted our assumption that content-producers and CDNs are unlikely to possess significant market power, but even if one were to relax those assumptions (and we expect that some will argue for just such a case), we do not expect such speculation to be productive: one can likely find specific

[32] For example, the original settlement free peering arrangement may have been perceived as marginal originally or have become so over time (e.g., because one ISP grew much larger than the other).

[33] See Dhamdhere, Dovrolis and Francois (2010) for an interesting alternative model for interconnection.

circumstances in which one or the other outcome seems to hold.[34] Ultimately, we believe that the determination of whether market power exists and on which side will depend on empirical facts that may vary case-by-case, and we do not wish to engage in those debates herein.

3.1. Finding limits on payment P

Rather than make assumptions about market power, we seek constraints or bounds on the outcome of negotiation between CD and A that might allow policymakers to infer whether the result of the interconnection negotiation were about a reasonable allocation of delivery-related costs or about an unreasonable allocation of the surplus associated with end-users' willingness-to-pay for content, above whatever it costs to efficiently deliver that content to the end-users. The topology of figure 1 (or the "money map" of figure 2) is actually deceptive, in that it suggests that the outcome of bargaining depends only on the actors illustrated in the figures. In fact, CD and A are typically embedded in a rich complex of interconnection agreements, and these agreements will limit the bargaining power of the two networks CD and A. We need a more complex picture to discuss the more realistic constraints.

3.1.1. Transit as a limit

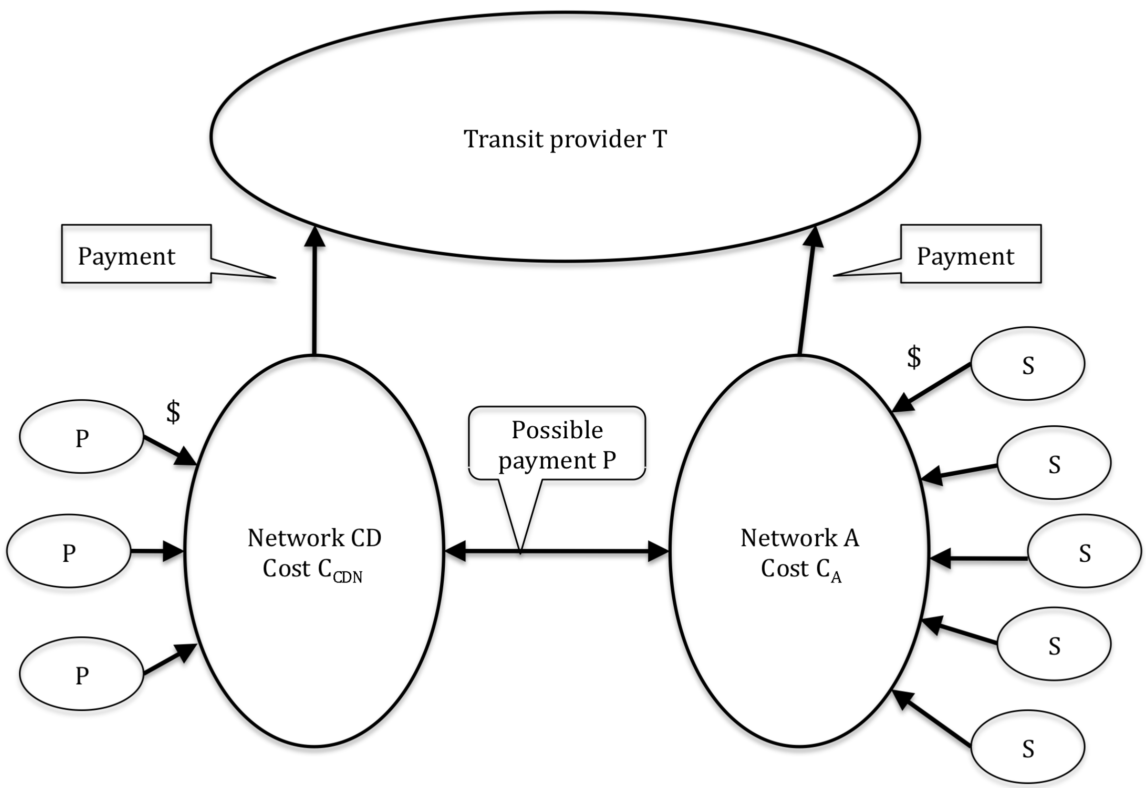

Figure 3: CDN, ISP A and Transit provider S

[34] We noted the case above of ESPN3, which had content of sufficient popularity that it bargained with access networks such as Comcast to pay ESPN3 a per-customer fee, which presumably is then passed on to the customers.

Figure 3 illustrates a common case, in which CD and A, in addition to bargaining about direct connection, purchase transit service from a third provider T. Figure 3 shows a single provider, but the logic of what follows does not change if CD and A purchase transit from different providers, who in turn peer.

In this case, imagine that A attempts to extract a large payment from CD. CD has the option of sending the traffic to A via T. This will increase total cost to CD, but will also increase cost to A. A, instead of seeing payments flowing in, now see payments flowing out to T. We have heard that this sort of thing happens in interconnection negotiations. This suggests that independent of market power, the amount that A can extract from CD is related to $P_{transit}$ from CD to T. CD will pick the lower cost option, if performance is equal: it will pay $P_{transit}$ if it is lower than the proposed direct payment from CD to A, and "punish" A.[35] The price $P_{transit}$ may not be the exact cap on the outcome of the negotiation. On the one hand, CD might choose to "punish" A unless it can negotiate a payment that is a discount off $P_{transit}$. On the other hand, if A offers CD valuable performance enhancements or other quality assurances, CD might be willing to pay a premium above $P_{transit}$. However, we argue that *P (CD->A) will be capped in practice by some function based on the customary cost of transit*.[36]

While (as we noted above) it is difficult to get internal cost numbers for C_A, since most ISPs consider these proprietary, we can speculate that C_A is larger than current values of $P_{transit}$, and in many cases substantially larger, measured in dollars/GB transferred. Access networks have large outside plant or access networks, and to the extent that these have a usage-sensitive cost component, these are likely to be much larger than transit costs. We speculated above that C_A might be between $0.10 and $0.30/GB. Typical $P_{transit}$ for large volume agreements (e.g. with negotiated discounts) are currently as low as $1/mb/s per month (peak rate). Assuming average link loading of 70%, 1 mb/s is about 162 GB/month. Put otherwise, $1/mb/s is the same as $.0062/GB—less than one cent. So as a practical matter, we speculate that even if CD is persuaded by A to pay a premium over $P_{transit}$, they will by no means be able to recover all of their internal costs, which makes it unlikely that there will be a content payment P_{CDN} that is part of the negotiated payment from C.

There is anecdotal evidence that the relationship we predict here is true in practice. Bill Norton, who tracks peering and transit issues closely, reported in January 2011 that: "The metered rate [of Comcast paid peering] is rumored to be in the $2-$4/Mbps price range, in the same ballpark as the market price of transit."[37] Our intention in this paper is to explain why this relationship might hold.

---

[35] There are exceptions to this analysis; the most obvious being the case where the access network A is also a tier-1 transit provider, in which case they do not purchase transit from any other provider. This option will weaken the bargaining position of CD.

[36] Thus, one might observe paid peering payments that exceed transit, but we would not expect any such excess to be large.

[37] http://drpeering.net/AskDrPeering/blog/articles/Ask_DrPeering/Entries/2011/1/14_Internet_Peering,_Paid_Peering_and_Internet_Transit.html.

3.1.2. Single-hop access

Figure 4 illustrates another pattern of interconnection and interconnection agreement, which we have called "single-hop access".

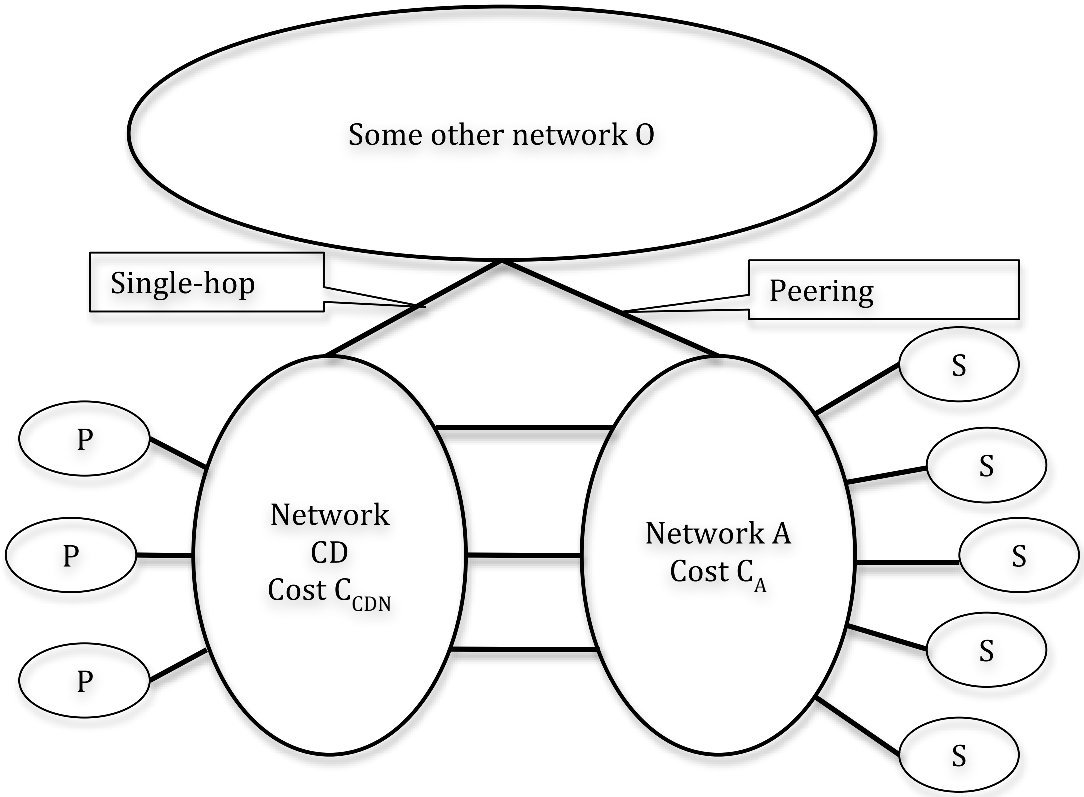

Figure 4: Configuration of connections for single-hop access.

Single-hop access is an interconnection service that can be offered by some network, such as O in the figure, which has a peering agreement with A. This concept is particularly appealing to O if the peering agreement is revenue-free. Once the peering agreement is in place between O and A, O then solicits CD to connect to its network at the same physical location as the link from O to A, ideally to a port on the same router. The internal cost C_O of this service to is very low: only the load on the router backplane passing the traffic from one port to another. (Hence the name "single-hop; this configuration has also been called "backplane access".) So O can set a very low price for the single-hop access that it sells to CD, presumably much lower even than the typical cost of transit.

If A feels that the resulting flow of payments is inequitable, its only option is to demand of O that the peering agreement be renegotiated, and made into a paid peering agreement. Such a demand, of course, is complex and implies the potential cost and difficulty of negotiation and a loss of goodwill.[38]

[38] Negotiation (and renegotiation) is costly because there are direct costs associated with establishing or re-arranging high-capacity physical interconnects, as well as the potential that

3.1.3. Changing the routing restriction

We described two sorts of traditional interconnection: *transit*, which gives one party access to all of the Internet via the other, and *peering*, which implies a routing restriction on each party that the traffic exchanged between them is local to them and their customers. Normally, a network would not agree to route traffic coming in from one peering partner out to another peering partner: it would be forwarding traffic without being paid by either partner. However, once paid peering is an option, more variants open up for different sorts of routing restrictions.[39]

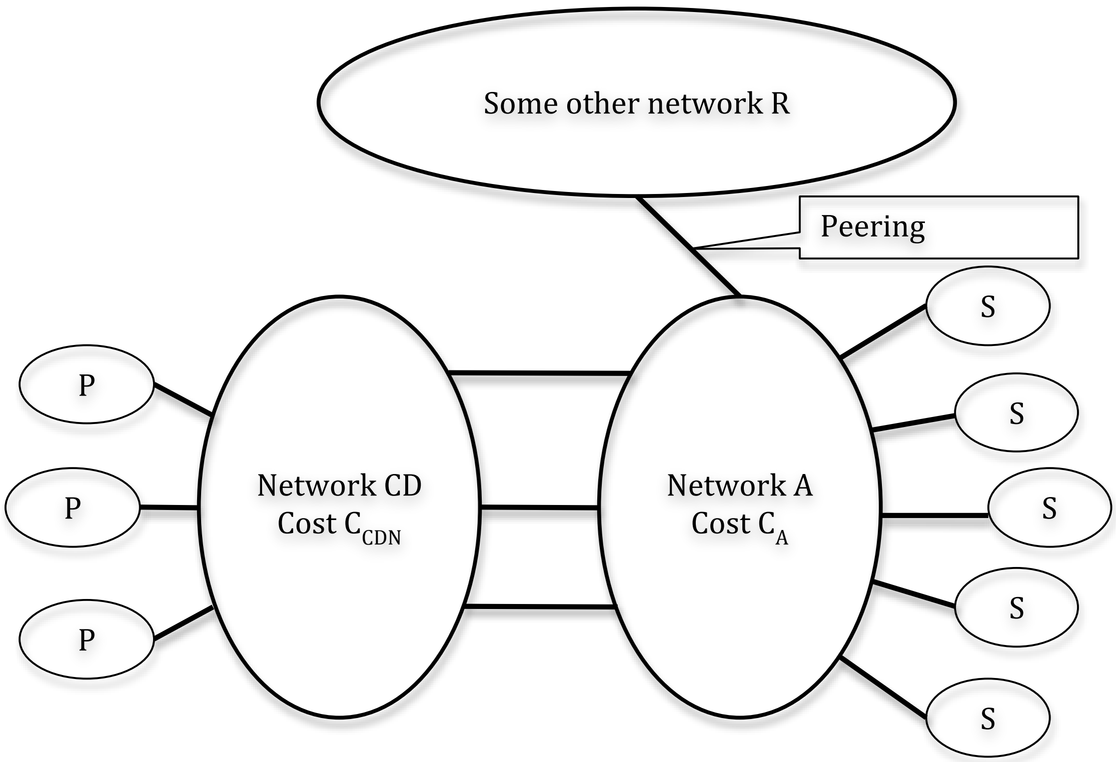

Figure 5: Network CD negotiates to gain access to S via the peering arrangement between S and A.

Figure 5 is almost the same as figure 4, but in this case CD has no direct access to R. It may not be cost-effective for CD and R to peer, but CD may still want a cost-effective path to S. CD might negotiate with A, as part of their interconnection agreement, to allow CD to reach S via A. This would make business sense only if CD pays A for access rights, otherwise A is incurring uncompensated costs to carry traffic from CD to S. So the option of payment may facilitate more direct connections, and also may lead to greater diversity in the negotiated routing restrictions.

---

negotiations may break down, resulting in traffic disruptions and the potential for damage to an ISP's brand image with customers.

[39] Another reason why a provider might agree to route "one-hop" traffic might be to offset traffic imbalances associated with other interconnection agreements, and thereby allow the provider to stay within traffic bounds for its peering agreements.

3.2. The real picture

In figures 3, 4 and 5 we have added a third network to the mix of CD and A, but in the real world a large access network might have several tens of peers, and might purchase transit from several providers.[40] We drew network CD as one oval, but a real content delivery network CD might have thousands of servers, each able to serve the same content. So CD might have thousands of choices as to how to source content flowing into A, and can use this in very nuanced fashions to control overall flows. If, for example, CD has an agreement with some other network S that it can deliver only so much traffic across a link, it can control which one of its servers generates that traffic to exactly load the link to the agreed capacity.

We noted above that traffic from networks of type CD into networks of type A now represents at least 40% of all traffic coming into A (Sandvine 2011). Because CD controls the routing of this traffic, what this means is that control of routing has moved away from the low-level routing protocols, and into a space controlled by higher-level business agreements and by subtle control over the dynamics of how the large CDNs choose one source rather than another for content. Given the rich complexity of the real picture, networks of type CD have a great deal of control over flows, and thus a great deal of bargaining power, independent of whether A is a terminating monopoly with respect to its customers. The CDN's bargaining power inheres in its ability to influence the costs realized by the access networks to which it delivers content. In the past, this control may have been limited to choosing between hot/cold potato routing,[41] whereas today, CDNs may have much finer-grained control.[42]

3.3. Norms of negotiation

As we noted in the beginning of the paper, as the old regime of revenue-neutral peering started to break down, to be replaced by less constrained negotiation, we could expect over time to see the emergence of new norms and regimes of interconnection. The previous discussion hints at two sorts of norms. First are criteria by which one ISP would consider agreeing to revenue-neutral peering. One source[43] examines a number of existing peering policies, and identifies 25 criteria, of which perhaps 10 are commonly used. One such criteria is balance of flows, in which the data rates between the two

---

[40] ISPs might purchase transit from more than one provider for several reasons. A simple reason is redundancy. But just as we are seeing more diversity in the routing restrictions on "peering", we are also seeing more diversity in transit agreements. For example, a transit provider might offer "low cost paths to Asia". This so-called *partial transit* further blurs the boundary between transit and peering.

[41] In hot potato routing the source ISP passes traffic to the destination ISP as soon as possible, thereby minimizing the resources used by the traffic on the source ISP and maximizing the resources used on the destination ISP. Cold potato routing reverses that allocation. By choosing between these options or some mix in between, the source ISP may be able to directly influence the usage-based costs incurred by the destination ISP.

[42] An Alcatel-Lucent white paper argues that the rise of CDNs poses a threat to ISP profits, by sucking transit revenues out of the Internet ecosystem (see, Alcatel-Lucent, 2011).

[43] http://drpeering.net/white-papers/Peering-Policies/A-Study-of-28-Peering-Policies.html.

parties are roughly in balance (perhaps no more than 2 to 1 in the peak direction).[44] Balance of flows is a rather rough approximation for balance of value, as we discussed above, but it can be used to impose limits on behavior such as single-hop access. If all parties understand up front the maximum amount of imbalance that A will tolerate, this can avoid the pain of after-the-fact attempts to renegotiate a peering agreement.

When revenue-neutral peering is not agreeable to both parties, we have speculated that a new norm might emerge to bound the price that might be charged for paid peering, which is that the rate for paid peering would be related to the price of transit by some function—perhaps a discount, perhaps a premium, but not wildly divergent. A proposal for a paid peering fee that greatly exceeds the customary price of bulk transit would be seen as evidence that the network proposing that fee does indeed have market power that allows it to distort the market. But a non-zero peering fee is not in itself a signal of such power.

Other norms might emerge, such as other proxies for cost (e.g. average route miles internal to an ISP), or industry average costs for outside plant. Average route miles could be used to bargain over the relative benefit of hot-potato vs. cold-potato routing.

4. Charging the consumer

As we noted above, the only options for access network A to recover costs C_A are to impose fees on the interconnected content networks CD or to impose additional fees on their own subscribers—the residential broadband consumers. (We ignore the final option of becoming less profitable in this analysis.) There are two basic ways that fees charged to customers could be structured. In one, the total C_A to be recovered from the customers is divided equally among them (flat rate pricing), in the other the fees are imposed on specific customers in proportion to their usage.[45] This option would suggest per/GB pricing, or more probably price tiers, which provide a lumpy form of per/GB pricing that seems to be more tolerable in the market since the consumers can self-select the right tier after which they see a fixed price.

Price tiers seem like a reasonable way to allocate fees in rough proportion to costs, but there is a very important consequence of usage-based pricing, which is that it can completely change the balance of power in negotiation about interconnection fees. To see how this can happen, one must look at the larger picture surrounding usage-based customer pricing. In countries where residential broadband access is sold with rather low monthly usage caps, too small to permit substantial downloads of high-volume commercial content, some ISPs are offering a "premium service" to their content network partners. With this premium service, the content delivery network CD pays a per-GB fee

[44]Nine of the 28 peering agreements included a requirement for traffic ratios. ISPs with traffic ratio requirements included AboveNet, Comcast, Verizon, ATT, CableVision and Quest. Several of these are what we classify as access networks, which supports the hypothesis that these networks are especially concerned with how peering agreements with CDNs are negotiated.

[45] Costs could be allocated on some other basis, of course, but we want to focus on usage-based contexts.

to the access network A, in exchange for which the consumer can download the content without having it count against their monthly quota[46].

With usage-based fees, and especially with a scheme like "unmetering", the whole landscape of negotiation is changed. Whereas before the access network A faced a range of interrelated interconnection agreements that the content networks CD could "play off" against each other, in this case, A can impose a consistent bargain on CD, no matter how the content arrives: either pay the set fee for unmetering, which A can control, or have the user face such a low monthly quota (and high usage charges) that the consumer forgoes the download. By setting a consumer-facing price for usage, A has in effect set a per-GB price for interconnection. Further, A may have limited the interconnection options open to CD, since unless CD delivers the traffic over a link where the traffic subject to the premium service tariff can be segregated and metered, the option of paying the unmetering fee may be precluded.

For these reasons, access ISPs that move toward usage tiers and similar forms of usage-based pricing can expect to receive increased attention from consumer advocates and regulators, not because of price tiers by themselves, but because of the resulting potential for influence over interconnection agreements. Moreover, in the United States, where consumers have long been accustomed to flat-rate (non-usage-sensitive) pricing, any movement to usage-based pricing will attract significant attention.[47]

One could ask whether usage-based pricing makes good sense, either from a business or "fairness" point of view. In flat rate pricing, the smaller users subsidize the larger users. It is the magnitude of the subsidy that matters. If the average subsidy were a fraction of a dollar, we would expect few to argue for usage tiers. The cost of implementing them might swamp the benefit to the smaller users, and there is considerable evidence that flat rate pricing has encouraged experimentation by users, and thus driven innovation and the creation of fresh value. On the other hand, if the cross subsidy were (say) $10, it would be hard to imagine that flat rate pricing could survive. It is not "fairness" that would drive toward usage tiers, but the opportunity for competitive advantage. Most users are smaller users, or in other words, the distribution of usage is heavy-tailed. The many small users are subsidizing a smaller number of large users. If the cross-subsidy were substantial, there would be a strong temptation in a competitive market for one ISP to offer a cheaper price tier targeted to the large number of small users, leaving the other ISP only with the

[46] In Australia, this service is called "unmetered content". This Wikipedia page briefly describes the service: http://en.wikipedia.org/wiki/Internet_television_in_Australia. See http://apextelecom.com.au/support/unmetered.php for an example of how unmetered content is marketed to the consumer. In Europe, it is sometimes called "zero data charge". See http://www.mblox.com/products/sender-pays-data/ for a product that is described as being trialed in the UK.

[47] For example, flat rate pricing of local telephone calls was cited as one reason for the more rapid takeoff of dial-up Internet access in the U.S. than in Europe and other markets where local measured pricing tariffs were more common. Also, with the great market success of block calling (non-distance sensitive) calling plans for mobile telephones and "all-you-can-eat" flat rate tariffs for broadband, it has been difficult for service providers to succeed with usage-based offerings.

expensive larger users. That ISP would have to increase prices accordingly, and price tiers would emerge.

Again, it is hard to get exact values for cost, and somewhat challenging to get information about the distribution of usage by different broadband customers, but the Sandvine data quoted above suggests that the mean usage is around 20 GB/m, but the medial usage is around 5 GB/m. If usage costs $.20/GB, that means the medial user is paying around $3/month to subsidize the heavy users. Reasonable people could disagree about what conclusion to draw from this $3/month number, but it seems large enough to pay attention to, but not so big that usage tiers will necessarily emerge. As evidence of this ambivalence, we see differences in approach, both within U.S. ISPs and more clearly across markets in different parts of the world. We mentioned Australia above, and one factor that is notable about Australia is that per-GB charges appear to be much higher there, presumably because a significant portion of the content comes over expensive inter-continental undersea cables. So perhaps it is not surprising that we see usage tiers in countries like Australia.

4.1. What is reasonable?

We have argued that negotiation to recover operating costs related to usage is reasonable. We called these *transport payments,* or P_t. We observed that negotiation that led to fees that seem unrelated to transport costs would raise the concern that transport providers were using market power to extract a part of the content payments P_c. One way to mitigate this concern would be for industry to provide additional information to make these negotiations more transparent. A variety of information might contribute to such transparency. First, information to allow some consensus to emerge about an aggregated cost model that would allow advocates and regulators to see if the proposed fees are consistent with costs would help. We could also debate whether negotiations that result in the content networks C paying *all* the C_A associated with their content is reasonable. But with cost models, we would at least have some confidence we knew what we were debating.

Second, information about traffic trends and distributions (e.g., about how heavy and light users differ) would help verify cost claims. As part of a project at MIT, we are engaged in examining such data, which has been provided voluntarily by an international group of ISPs.[48]

Third, information about actual interconnection agreements, which may include public commitments to interconnection or peering policies may help in better understanding how interconnection markets are changing and what the emerging "norms" are. Because there are significant shared and non-traffic sensitive costs that need to be recovered[49] and because the details of interconnection agreements may convey sensitive strategic

[48] See http://mitas.csail.mit.edu.

[49] Economics does not provide clear guidance on how to apportion the burden for recovering shared and non-traffic-sensitive costs that comprise a significant share of total costs, and so bargaining may result in a variety of allocations.

information,[50] we are not surprised that ISPs are reticent to share publicly the terms that are negotiated. If needed, the proprietary information might be shared with regulators with adequate disclosure restrictions to protect confidentiality.

In each case, we hope that industry and independent analysts will take the lead in expanding the information available to all stakeholders, since ultimately, we think that the market may do a better job of ensuring transparency than regulatory interventions. However, we recognize that more activist interconnection policies may be needed in the future if those markets provide clear evidence of market failures. Our policy recommendations focus on enhancing transparency because 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 importantly, even if we were to see a need for such regulation, we are concerned any such regulation might cause more harm than good. 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.

5. Summary conclusions

The key observations and conclusions of the paper are as follows: The norm of revenue neutral peering has been breaking down for some time, and being replaced by less constrained negotiation about paid peering. Paid peering may increase the number of direct peering connections, reduce transit costs, and thus reduce ISP operating cost, but it also may increase transaction costs around peering agreements. To the extent that related costs can be contained (perhaps by the emergence of new norms and conventions to negotiate new forms of peering), total surplus in the system may increase.

The emergence of networks that deliver high-volume commercial content (our networks of type CD) raise potential concerns about market power by the various actors, given that there are payments both for content P_c and transport P_t within the ecosystem. Interconnection between CD and A are a type of peering agreement,[51] but a very asymmetric interconnection influenced by the peculiarities of high-volume commercial content flowing from producer to consumer. Even if market power were not an issue, we believe there may be obvious economic efficiency rationales for expecting to see CDN-> Access ISP payments occurring.

Even if networks of type A have some market power because of their terminating monopoly, the complex mesh of interconnections, with diverse pricing models, constrains the range of negotiating positions that can be sustained by A. In particular, we assert that

[50] For example, ISPs might regard as sensitive information about how they manage their traffic and provision their networks.

[51] That is, they are like "peering" because traffic only flows from CD to A that is destined for customers of A, which is in contrast to what happens in a "transit" relationship. However, since effectively all of the traffic flows in the direction from CD to A, there is no reciprocal exchange of traffic (and the usage-cost causing implications implied by such traffic).

the limit on the payment that A can extract from CD will be related in some way to the current customary price for transit, which is a commodity product with no P_c component.

Analysis of negotiation options between CD and A based on a simple model that represents only those two networks will not be realistic. The presence of additional forms of interconnection must be included in the model.

Retail pricing based on usage tiers is a reasonable way to allocate the cost of usage to the relevant customers, but at the same time, especially if combined with schemes that let the content networks CD pay the price (in some form) on behalf of the user, change the negotiating options for the parties, improves the bargaining position of networks of type A, and will attract attention from regulators and consumer advocates. In the extreme, it may lead to regulation of retail broadband pricing.

Interconnection policy is going to become the battleground for the new telecom regulatory debates -- re-imagined as Internet or broadband regulation in the context of the Network Neutrality wars. To the extent that regulators are concerned about the power of the access networks—our networks of type A—it will not be possible to disentangle how these providers treat their two sorts of interconnections: their retail customers and their connections to other networks, in particular CDNs.

Before more interventionist regulatory approaches are applied, we believe any policy focus should be on improving transparency into the workings of the Internet ecosystem in general and interconnection markets, more specifically. We identified three categories of information that would contribute to improved transparency: (a) information about industry-wide cost models; (b) information about traffic trends and distributions, and (c) information about interconnection agreement terms and conditions. Better data on these will contribute to the public debate, even if some categories of information are deemed too sensitive to be shared except under some restrictive confidentiality protections. It will help provide a foundation on which a better assessment might be made of need for further regulations; and should such a need be identified, in helping to craft suitable rules.

References:

Agyapong, P. and M Sirbu (2011) "Economic Incentives in Content-Centric Networking: A Network Operator's Perspective," *39th Research Conference on Communications, Information, and Internet Policy (TPRC)*, George Mason University, Arlington, VA, September 2011.

Alcatel-Lucent (2011), "Analysis: Content Peering and the Internet Economy," an Alcatel-Lucent White Paper, April 19, 2011, available at: http://www2.alcatel-lucent.com/blogs/techzine/2011/analysis-content-peering-and-the-internet-economy/.

Bertrand, E. (2011) "What Do Cattle and Bees Tell Us About the Coase Theorem?" *European Journal of Law and Economics*, vol 31 (2011) 39-62.

Brock, G. (1994), *Telecommunications Policy for the Information Age: from Monopoly to Competition*, Harvard University Press: Cambridge, 1994.

Clark, David (2008), "A simple cost model for broadband access: What will video cost", 36th Research Conference on Communication, Information, and Internet Policy (www.tprcweb.com), George Mason University, Arlington, VA.

Clark, D., W. Lehr, and S. Bauer (2009), "Comments of David Clark, William Lehr, and Steve Bauer," In the Matter of Preserving the Open Internet and Broadband Industry Practices, Before the Federal Communications Commission, GN Docket No. 09-191 and WC Docket No. 07-52. Washington, DC.

den Butter, F. (2010), "Transaction Management: Value Creation by Reducing Transaction Costs," *Tinbergen Institute Discussion Paper, T1 2010-051/3, VU University of Amsterdam.*

Dhamdhere, A., C. Dovrolis and P. Francois (2010) "A Value-Based Framework for Internet Peering Agreements," *Proceedings of 22nd International Teletraffic Congress (ITC 22), September 7-9, 2010.* Amsterdam, Netherlands: IEEE, 1-8.

Faratin, P., D. Clark, P. Gilmore, S. Bauer, A. Berger, and W. Lehr (2007), "Complexity of internet interconnections: Technology, incentives and implications for policy," 35th Research Conference on Communication, Information, and Internet Policy (www.tprcweb.com), George Mason University, Arlington, VA.

Frieden, R. (2004), "Adjusting the horizontal and vertical in telecommunications regulation: a comparison of the traditional and a new layered approach", *Federal Communication Law Journal*, Vol. 55 No.207, pp.209.

Jordan, Scott (2007), "A Layered Network Approach to Net Neutrality," *International Journal of Communication* 1 (2007), 427-460.

Lehr, W., J. Peha, and S. Wilkie (eds) (2007), Special Issue on Network Neutrality, International Journal of Communication, Volume 1 (August 2007), available at: http://ijoc.org/ojs/index.php/ijoc/issue/view/1.

Sandvine (2011), Global Internet Phenomena Report, Spring 2011. Available through: http://www.sandvine.com/news/global_broadband_trends.asp

Sicker, D., J. Mindel, and C. Cooper (1999), "The Internet Interconnection Conundrum," unpublished FCC working paper (1999)

Werbach, Kevin (2002), "A layered model for Internet policy," *Journal of Telecommunications and High Technology Law*, vol 1 (2002) pages 58-64.