

**The Economics of Internet Interconnection: Insights from the Comcast-
Level3 Peering Dispute**

Gregory Rose
Econometric Research and Analysis

March 28, 2011

Abstract: The dispute between Level3 and Comcast with regard to termination of Netflix's traffic is only the latest in a series of disputes over interconnection and who pays whom for the transport of internet traffic. Because the circumstances of the dispute have been unusually public and protracted, and because of the unusually large number of residential end users potentially impacted, the dispute provides an opportunity to assess the validity of various models for pricing internet traffic, and whether leaving such services unregulated represents a greater risk to the public welfare than efforts to regulate such traffic.

An examination of the circumstances of the Level3/Comcast dispute and comparison with the previous Level3/Cogent dispute casts doubt on the validity of the "bill-and-keep" model for internet traffic, and suggests that the "two-sided market" or "n-sided market" analysis more closely models the existing market structure. Using two-sided market analysis, it appears that while there is the potential for the market to resolve issues without regulatory intervention, there is also the potential for substantial anti-consumer outcomes in the absence of regulation. The Commission should therefore undertake to gather sufficient data to conduct a proper analysis so that the question of what regulations may be required to maximize pro-consumer outcomes can be resolved.

Introduction

On November 29, 2010 Level3 Communications issued a release reporting that Comcast had informed Level3 that it was demanding "a recurring fee from Level3 to transmit Internet online movies and other content to Comcast's customers who request such content."¹ Characterizing Comcast's payment demand as a "take it or leave it" matter, unilaterally revising their previous peering arrangement, Level3 announced that it had acceded to Comcast's demand on November 22.² There followed an exchange of releases and ripostes between Comcast and Level3 which portrayed the matter from Comcast's viewpoint as a reasonable demand for compensation for a move from a 2:1 Level3-to-Comcast traffic asymmetry to a 5:1 traffic asymmetry favoring Level3 arising

¹ "Level 3 Communications Issues Statement Concerning Comcast's Actions," Level3 Communications, November 29, 2010 (<http://www.level3.com/en/About-Us/Newsroom/Press-Release-Archive/2010/2010-11-29-level3-statement-comcast.aspx>).

² *Idem*.

from an agreement between Level3 and Netflix,³ while Level3 presented Comcast's demand as a use of market power to extract monopoly rents from Level3.

Ultimately Level3 addressed an appeal to Federal Communications Commission Chairman Julius Genachowski, enquiring whether the F.C.C.'s Open Internet Order applies to internet peering arrangements. Eventually Comcast, as well as such incumbents as Verizon, AT&T, and the National Cable and Telecommunications Association which had also weighed in on Comcast's side, were signaled by Genachowski in a response to queries in front of the House Energy and Commerce's Subcommittee on Communications and Technology that the F.C.C. regarded internet peering arrangements as private business matters to which the F.C.C. Open Internet Order did not apply.⁴

While Comcast's action is prima facie quite close to what then-SBC CEO Edward Whitacre proposed doing to Google which ignited the controversy over net neutrality,⁵ it is important to note that this controversy does not appear to be so much about the predation of Comcast on Level3 as about the emergence of new business models at which Comcast has simply arrived earlier than the rest of the industry (although, as we shall see, Level3 tried something very similar in its peering dispute with Cogent Communications in 2005). A more detailed examination of the facts underlying the dispute between

³ Joe Waz, Senior Vice President for External Affairs and Public Policy Counsel, Comcast Communications, "20 Q's – with Accurate A's – about Level 3's Peering Dispute," Comcast Voices, December 7, 2010 (<http://blog.comcast.com/2010/12/20-qs---with-accurate-as---about-level-3s-peering-dispute.html>).

⁴ Amy Schatz, "FCC Chairman: Net Neutrality Rules Don't Cover Comcast-Level 3 Dispute," *Wall Street Journal*, February 16, 2011 (<http://online.wsj.com/article/BT-CO-20110216-718576.html>).

⁵ "At SBC, It's All About 'Scale and Scope,'" Bloomberg Businessweek, November 7, 2005.

Comcast and Level3 and assessment of the content delivery network (CDN) and broadband access network (BAN) relationship as a variety of two-sided market suggests that the Comcast-Level3 peering arrangement dispute is merely an adumbration of a fundamental change in business models which has radical implications for the regulation of intercarrier compensation and, more importantly, for the openness of the internet. These harbingers of change suggest that regulators urgently need to gather pertinent data on the current and emerging state of peering and other transit arrangements so that the dynamic interaction of market power and pricing strategy in this industry can be evaluated

Setting the Stage: Peering Arrangements and Comcast Peering Policy

A few definitions are necessary before embarking on discussion of the Comcast-Level3 peering dispute as the harbinger of major changes. First, a **peering arrangement** is an agreement to maintain the interconnection of administratively and economically separate networks for the purpose of exchanging traffic between the customers of each network. Such arrangements are a necessary, but not sufficient condition, for the existence of the internet. Originally such arrangements were **settlement-free**, i.e., neither carrier paid the other for the exchanged traffic.⁶

Note that there is a difference between CDN transit covered by peering arrangements and paid off-net transit over an internet backbone. While on-net and off-net traffic are carried over the same infrastructure, on-net CDN traffic has been historically settlement-free and off-net traffic costed at market rates. In other words, customers of the CDN, like Netflix, paid CDNs like Level3 for caching and on-net transit

⁶ An excellent explanation of peering and transit, as well as a history of their development, can be found in P. Faratin, et al. (2008)

of content, while customers/subscribers of a BAN, like Comcast, paid for the ability to receive on-net traffic, while, theoretically, the traffic passed settlement-free between the CDN and the BAN. Off-net traffic was covered by market-determined agreements.

Among broadband access networks in the U.S., only Comcast deviates from the settlement-free model of peering. Comcast has always seen peering as a potential profit center. In recent F.C.C. filings Comcast has reported that both Akamai and Limelight, as well as other CDNs, have been paying settlement fees for access to Comcast subscribers, despite the near-universal practice of settlement-free peering. However, it is difficult to verify Comcast's claim. Virtually all parties to these arrangements have nondisclosure agreements which prevent public scrutiny. There is no way to tell whether other CDNs are paying Comcast just for access to its subscribers, for routing traffic to its ultimate destination, for paying ISPs to haul this traffic themselves, or as a cost-recovery measure. This is an area where only regulatory scrutiny can pierce the veil of nondisclosure agreements to disclose the truth of the matter.

Comcast also purchases IP transit services from Tata Communications in order to reach networks with which it does not have peering arrangements. Adam Rothschild reports that "...Comcast runs its ports to Tata at capacity, *deliberately*, as a means of degrading connectivity to networks which won't peer with them or pay them money."⁷ nLayer Chief Technical Officer Richard Steenbergen makes a cogent point relating Comcast's relationships with Tata and Level3:

The true power of Comcast isn't in the size or scope of its network, it's in the captivity of its customer base.

⁷ Adam Rothschild, "Peering Disputes: Comcast, Level3, and You," voxel, December 2, 2010, (<http://www.voxel.net/blog/2010/12/peering-disputes-comcast-level-3-and-you>).

If Level 3 turned off Comcast for refusing to pay their contractually obligated transit bills, the traffic would be forced through massively congested Tata transit ports, and a huge number of Level 3's customers would take their business elsewhere as a result. If Comcast intentionally congests its transit providers and provides terrible service to its end users, which it has been doing for several months now, most of those users have no real alternatives to switch to.

In other words, content is mobile, eyeballs are not. Comcast realizes that they can (ab)use their captive eyeballs to force content to pay them for access, without having to create a "100% down" partition like Cogent has done in the past. For a Netflix customer, 20% packet loss is effectively just as down as a hard partition.

[...] while it should be every network's right to choose who they do and don't peer with, or buy transit from, things start to get murky when one network is abusing their franchise agreements and near monopoly or duopoly status in many markets. If users had an actual choice, and could get comparable broadband access elsewhere, then Comcast would be free to congest their network however they see fit. But that isn't the case, and this is where government involvement and Net Neutrality start to have legitimate grievances with Comcast's actions.⁸

This is not, however, entirely a case of Comcast's throwing around its weight.

Level3 has recently shifted its business model from an infrastructure and wholesale IP transit provider to a CDN. Furthermore, its existing relationships with a number of large DSL and cable providers have been conditioned on standard provider-customer off-net terms rather than CDN peering agreements and it is clearly Level3's desire to generalize this model to one in which CDNs do not interconnect with BANS via peering arrangements, but as standard incoming off-net traffic for which BANS must pay. This is paralleled by Level3's attempt to exercise market power in terminating its peering arrangement with Cogent Communications in 2005. While Level3 ironically raised some of the same arguments as Comcast about traffic asymmetry, it became quickly apparent that Level3's principal objective was to discipline a rival backbone operator for

⁸ Idem.

undercutting Level3's pricing policy by suspending the peering arrangement. The adverse reaction of capital markets to Level3's sudden, unilateral suspension of the peering arrangement with Cogent led to Level3's quick backtracking, acknowledged by Level3 CEO O'Hara in a conference call with analysts: "I apologize to both sets of customers. We recognize that we have an obligation to customers of the Internet, and in this instance, we contributed to letting them down."⁹ Sharper elbows and hardball measures against competitors were becoming more common amongst both internet backbone providers and CDNs.

It is in this context that Comcast's demand on Level3 arose, occasioned by the agreement of Netflix to switch a significant portion of its streaming video content back from Akamai to Level3. Comcast has claimed that this has occasioned a radical increase in the traffic asymmetry between Level3 and Comcast. This claim is questionable on the face of it for two reasons. First, the contract between Netflix and Level3 shifts traffic which was going through Akamai to BANS, including Comcast's, to Level3. While there is evidence that demand for streaming video from Netflix is increasing, in the short term the shift from Akamai to Level3 scarcely affects the net amount of traffic coming onto Comcast's network. Second, the increasing demand for Netflix content is from Comcast subscribers – if Comcast incurs costs in meeting increasing demands from Comcast's subscribers, then those costs should be recouped from those subscribers, not from the CDN which passes that content to Comcast.

Furthermore, there is no necessary connection between traffic rations and costs:

⁹ Stacy Cowley, "Level 3, Cogent resolve peering dispute, renew deal," Computerworld, Oct. 28, 2005 (http://www.computerworld.com/s/article/105790/Level_3_Cogent_resolve_peering_dispute_renew_deal).

It's a common misconception that the benefit an ISP derives from peering depends upon the direction of the flow of traffic. According to this way of thinking, if YouTube peers with an ISP, this benefits YouTube more than it does the ISP (since YouTube sends so much data but receives comparatively little). But in practice, the flow of traffic is not an issue for an interconnect. Whether it goes to or from the network, companies still need the same Cisco equipment.

In practice, it is actually quite likely that the ISP side of an ISP-YouTube relationship would see the greatest savings both in absolute costs and as a percentage of total traffic costs. Most ISPs have less traffic (and buy less transit) than YouTube and its parent Google have. Their buying power therefore is less than that of YouTube/Google, so their price per Mbps/month for transit is likely to be higher. Given that the amount of traffic saved from transit is by definition equal for both YouTube and the ISP, it follows that the ISP is saving more money.¹⁰

When is a Monopoly a Monopoly

Definitionally Comcast, like all BANs,¹¹ is a terminating access monopoly, i.e., there is no other way for a CDN's customers, content providers, to reach Comcast subscribers except via Comcast's last-mile connection. While there are competitive options for content providers in the selection of CDN and BANs have competitive options in the selection of transit providers, content providers have no alternative but Comcast in delivering their content of Comcast end-users.

However, there is a sense in which Comcast verges on a monopoly for its end-users which decidedly affects consumer welfare. End-users are often limited to a duopoly: seeking service from the local cable company or the local telephone company.

As Level3 pithily sums up the situation:

¹⁰ Rudolph van der Berg, "How the 'Net Works: an introduction to peering and transit," *ars technica*, September 2, 2008.

¹¹ Technically Comcast is more than a BAN: it is also an internet backbone, a provider of wholesale VOIP and IP transit services, a major provider of video for other cable systems, a content creator and distributor, and the owner of a major television network and film archive.

While Comcast and others talk about theoretical competition from broadband cellular or broadband over power line, for almost all Americans, broadband to the home means service from either the cable or phone company. Anyone who believes that broadband wireless is a substitute for the broadband access services provided by cable and phone companies should try connecting their cell phone to their TV to try to watch an online movie or TV show.¹²

Technically, there is a way for Comcast subscribers to access content which is not mediated by an agreement between Comcast and a CDN: transit links. However, attempting to access content via a transit link puts the Comcast subscriber at the mercy of hugely congested Tata Communications links. Trying watching a streaming video over a highly congested off-net transit network if one wants to understand the phrase “exercise in futility.” Even where the Comcast subscriber has a potential BAN competitor to which to turn, such subscribers are effectively deterred by switching costs; an ample literature exists on the role of switching costs in limiting competitive choice for consumers (for literature reviews and particularly salient contributions see J. Farrell and C. Shapiro, 1988; P. Klemperer, 1995; R. Agarwal and M. Gort, 1996; J.-J Laffont and J. Tirole, 2000; S. Klepper, 2002; O. Shy, 2002; J. Krafft and E. Salies, 2006; and J. Farrell and P. Klemperer, 2006).

All this would give Comcast an effective means of extracting monopoly rents simply because Level3 was acquiring additional revenue from Netflix, which, in turn, depends on access to Comcast subscribers who have chosen to order Netflix services. But, as indicated above, treating Comcast simply as a BAN may be useful to analysis of its terminating access monopoly. When Comcast is seen in a larger context, the potential

¹² “Level 3 Releases Statement to Clarify in the Comcast/Level3 Interconnection Dispute,” Level Communications, December 3, 2010 (<http://www.level3.com/About-Us/Newsroom/Press-Release-Archive/2010/2010-12-03-statement-interconnection-dispute.aspx>).

for abuse in control over access to its subscribers by content providers through CDNs. Comcast is itself a major, competing source for content and already enjoys substantial advantages over other content providers by virtue of the way it chooses to configure its cable and broadband services to access Comcast-provided content: it is simply easier for a Comcast subscriber to find and access Comcast content. Further, by virtue of being such a large content provider Comcast can reduce the costs to itself implicit in the threat to discontinue the Level3 CDN arrangement if monopoly rents are not paid: end-users are less likely to try switching when a plethora of Comcast content is available. Finally, Comcast's role as a major content provider potentially provides incentives to erect additional barriers to the entry of other content providers into competition for Comcast subscribers' "eyeballs" above and beyond those which exist even if Comcast is not acting anti-competitively.

The Capacity Paradox

E. Jahn and J. Prüfer (2006) predicted the emergence of a crisis affecting internet backbone providers: the capacity paradox. They argue:

The main problem of IP-Transit sellers is that they offer a homogeneous product being characterized by excess capacity and the absence of bottlenecks, which could involve market power. This makes IP-Transit a commodity. Standard economic theory would predict that prices approach marginal costs, which are zero. Again, this is exactly what is happening in practice. Nevertheless, price erosion has come in line with demand expansion. It is interesting to note that *TeleGeography*... points to the fact that during 2003 and 2004 the demand expansion effect exceeded the price erosion effect in many countries, including the United States, Germany, China, India, Brazil and also for Trans-Atlantic and Trans-Pacific traffic. Therefore, revenues have increased despite declining prices. However, in the long-term revenues will erode more and more. This results from the fact that demand expansion has an upper boundary (set by the world's online population) while the prices' lower boundary is zero — and from the definition of revenues as price times quantity sold. As a consequence, allowing for the fact that marginal costs are zero and networks only have fixed costs for operations and maintenance, long-term profits of Tier 1 networks will vanish.

According to Jahn and Prüfer (2006), all but the very small networks could make use of selling Paid Peering agreements to even smaller players, thereby tapping another source of revenue. Unfortunately, income from Paid Peering is closely (and positively) correlated with prices for IP-Transit as Peering is, in part, a substitute for IP-Transit. On the basis of the above facts one might conclude that network operators focusing on sustainable revenues from Paid Peering may well be disappointed.

To summarize, IBPs have four main sources of revenue. Of those, selling access to end-users, selling IP-Transit (by Tier 1 networks) and selling Paid Peering arrangements are homogeneous goods with virtually no capacity restrictions. Hence, in these segments profit potential is low to negative. In the fourth segment, bundling Internet access for large corporate customers with value-added services, the current profit situation due to larger possibilities for product differentiation looks somewhat better. Nevertheless, a reasonable share of those profits is taken by suppliers of Internet infrastructure and equipment. Furthermore, as the de facto revenue source is not selling access to the Internet but selling complementary services (and perhaps consulting) which are less scalable, barriers to market entry are not very high. (J. Prüfer and E. Jahn, 2007, 147-48)

The declining profitability of internet backbone providers in an environment of increasing capacity presents them with an even more dire scenario in the context of the decline of excess capacity and the emergence of capacity constraints:

...capacity constraints per se would have a negative impact on the quality of high-end broadband applications, since those require high-quality transmission levels. A *capacity paradox* exists: advanced IP-services depend on excess capacity that, in turn, erodes networks' profit potential. This is why innovative broadband applications may fail due to capacity constraints ensuring profits (and survival) of networks on the one hand but impeding innovation on the other. To put it more succinctly: the fundamental trade-off in this scenario is reaping profits from Cournot competition (including less market exits) versus increasing consumer surplus by higher network quality resulting from Bertrand competition with excess capacity (leading to more market exits). (J. Prüfer and E. Jahn, 2007, 150)

Prüfer and Jahn (2007, 151) identify a relatively limited number of market strategies for a backbone provider like Level3 to escape the capacity paradox: horizontal differentiation (which they reject as unlikely), becoming a seller of services to end-users

(essentially becoming a CDN), and price discrimination. Prüfer and Jahn further argue that becoming a seller of services to end-users is complicated by the fact that it depends on the lack of close substitutes provided by competitors to maximize profits, an argument which suggests that strategies erecting barriers to entry for potential competitors is crucial to success as well as price discrimination for customers who agree to collude to discriminate against other competitors. It is possible to see in Prüfer and Jahn's analysis the forces which have pushed Level3 from internet backbone provider to CDN and which explain Level3's increasingly sharp elbows with competitors to guarantee the firm's price floor, as well as its attempt to move from on-net to off-net pricing in the distribution of content to BANS.

Bill-and-Keep

The game theoretic literature on "bill-and-keep" and the decision making of firms regarding internet interconnection has emerged from application of a more classical economic analysis (M. Jackson and A. Wolinsky 1996; J. Cremer P. Rey, and J. Tirole, 2000; W. Norton, 2002; N. Badasyan and S. Chakrabarti, 2003, 2004, 2008; F. Bloch and M. Jackson, 2004; G. Shrimali and S. Kumar, 2008). Much of this research has focused on problems arising from the facts that

...[T]here is a general convention that the providers peer if they perceive equal benefits from peering, and have transit arrangements otherwise. There has been a growing concern with regard to the increasing number of transit agreements replacing previously peering agreements. There is a debate whether the large providers are unwilling to peer with small providers, or that transit arrangements are related to the actual differences in the relative costs incurred by the providers. (N. Badasyan and S. Chakrabarti, 2008)

It uses game theoretic techniques to plumb the factors which determine whether interconnection will occur through peering or transit agreements. An example of this sort of analysis is the simple game developed by N. Badasyan and S. Chakrabarti (2008) which finds three possible equilibrium outcomes:

- When both providers demand payment for the exchanged traffic, the providers will arrive at a peering arrangement;
- If provider A offers to pay and provider B demands payment, the providers will arrive at a transit agreement; and
- If provider B offers to pay and provider A demands payment, the providers will arrive at a transit agreement.

Badasyan and Chakrabarti derive in conclusion that

The larger provider may demand a transit arrangement, if it believes that the smaller provider might free ride on its infrastructure investments, because the larger provider might be the one transporting most of the traffic between them. The results suggest that the providers do not necessarily exploit market power when refusing to peer. Moreover, the joint profits are maximized under the transit arrangement. Peering partners do not get any compensations for carrying each other's traffic, while incurring costs of transporting each other's traffic. Under the transit agreement, on the other hand, the downstream provider pays for the traffic carried by the upstream provider, while getting benefits of dumping most of the traffic on the upstream provider. These two effects cancel out when maximizing the joint profits of the transit partners. The paper argues that the market forces determine the decisions of peering and transit, and, given the current system of peering and transit arrangements, there is no need for a regulation on interconnection policies.

For them the crucial determinant of whether a peering arrangement or a transit agreement will arise is traffic ratios and willingness of providers to demand payment or pay. There are, however, serious problems with the assumption that traffic ratios measure actual costs. If, as Prüfer and Jahn argue, the marginal cost of traffic approaches zero, then decision makers should be indifferent to traffic ratios unless capacity constraints are present, i.e., serious network congestion arises. The failure of Badasyan and Chakrabarti, like most analysts of “bill-and-keep” agreements, to take measures of network congestion and capacity constraint into consideration – and that failure is probably as much a tractability issue as substantive – calls into question whether real market forces are at action here in such a straightforward way as the simplified model suggests. What is important here is this sort of game theoretic analysis appears to give credence to Comcast’s argument about traffic ratio asymmetry costs underlying its pricing and peering policies. While this paper does not concur with the general approach taken by

game theorists analyzing “bill-and-keep” in the internet interconnection market, it is a literature not infrequently cited by market actors in an appeal to economic theory to justify business practices, particularly to regulators.

A more interesting finding in this literature is to be found in G. Shrimali and S. Kumar (2008), who model internet interconnection as both single-shot and iterated games. They found that “providers that are better off in the one-shot game can cooperate using threat strategies in an infinitely repeated game and can each be even better off.” While the games examined by Shrimali and Kumar, strictly speaking, modeled only smaller ISPs, this finding may be generalizable and, thus, it may explain, in part, the increasingly hardball strategies adopted by BANs and CDNs in negotiating peering and transit relationships.

J.-J Laffont, S. Marcus, P. Rey, P. and J. Tirole, (2003) evaluate bill-and-keep using a two-sided model, finding that bill-and-keep models are consistently optimal only under perfect competition assumptions:

Under perfect competition, firms make zero profit; they are thus indifferent as to the level of the access charge and should not resist a regulation of the access charge that implements the second-best optimum. Under perfect competition, firms make zero profit; they are thus indifferent as to the level of the access charge and should not resist a regulation of the access charge that implements the second-best optimum. In practice, backbones have historically opted for “bill and keep” ($a = 0$), which minimizes transaction costs. Bill and keep is favorable to websites, which might have been a good idea to promote the development of Internet-based services. Now that many web services are available, and the emphasis is more on encouraging consumers to connect and use these services, absent significant transaction costs, bill and keep is unlikely to be close to optimal (2003, 379).

Even under a perfect competition assumption,

Perfect competition implies that backbones’ budgets are always balanced, whatever the access charge. But through its allocation of costs between end users, the access charge plays a central role in achieving economic efficiency. We show below that the Ramsey access charges, i.e., the access charges that maximize social welfare, must take into account not only the demand elasticities of the two segments, but also the externality that each side exerts on the other (2003, 377).

This line of argument calls into question the value of models which do not acknowledge the internet interconnectedness market as two-sided, as well as provides a rationale for some CDNs seeking off-net pricing rates.

Two-Sided Markets

A relatively new approach to analysis of internet interconnection pricing is the modeling of such interconnection as a two-sided market. Excellent reviews of this literature are provided by J.-C. Rochet and J. Tirole (2004 and 2005), M. Armstrong (2005), B. Jullien (2005), and R. Roson (2005).

A two-sided market is one in which a) two sets of actors interact through platform, and b) the decisions of each set of actors affects the outcomes of the other set of actors through externalities. Both Rochet and Tirole (2004 and 2005) and Roson (2005) provide slightly different formal definitions of a two-sided market, e.g., Rochet and Tirole find that the inapplicability of the Coase theorem¹³ as a necessary, but insufficient, condition for two-sidedness, which Roson rejects while tightening the definition to exclude some cases which Rochet and Tirole accept. Rochet and Tirole's formulation has found wider acceptance in the discipline thus far.

It is easy to conceptualize the interaction of CDNs and BANs like Level3 and Comcast, as a two-sided market: content providers, the sellers, have their access to end-users, the buyers, mediated by a platform, the relationship between the CDN and the BAN, and, indeed, study of two-sided markets has emerged largely from examination of credit card arrangements, computer operating systems, and internet interconnection.

¹³ The Coase theorem states that if trade in an externality is possible in the absence of transaction costs bargaining will lead to an efficient outcome irrespective of the initial allocation of property rights. Given the fact that many empirical exceptions to the Coase theorem arise, this is not a particularly restrictive assumption.

As Roson (2005, 144) observes, “Network externalities are said to exist when consumer utility in a certain market depends (usually, in a positive way) on consumption of the same good or service by other agents.” In general the sellers and buyers are subject to network effects. In two-sided markets the platform’s value to any given user largely depends on the number of users on the network’s other side.

Two-sided markets typically present unique complexities which often seem to confute traditional industrial organization and antitrust theory. Market power is difficult to assess in two-sided markets. Evans (2003) makes the point that

Market share as a proxy for market power is problematic in many circumstances but is especially so for businesses that compete in multi-sided platform markets. Economists have shown that Cournot-competition or differentiated-market Bertrand competition among firms in single-sided markets implies that the equilibrium prices will depend on some function of market shares. Those models do not apply when looking at just one side of multi-sided platform businesses. Pricing power on each side depends on the degree of competition on both sides.

This has immediate implications for the Comcast-Level3 dispute: if the degree of competition on each side of the market is asymmetrical, the side with less competition may enjoy the ability to exercise market power. The extent to which Comcast end-users are locked into Comcast as their BAN by switching costs is not matched by a similar lock-in of content providers into Level3 (as evidenced by Netflix’s seamless shifts of traffic from Level3 to Akamai and back again). The exercise of market power in a two-sided market is not necessary under such conditions, but theory clearly indicates a greater likelihood. Evans also suggests a strategy for evaluating market power in a two-sided markets:

More sophisticated analyses do not rely on market share as a proxy but instead seek to determine directly whether the firm under consideration prices above marginal cost by a significant amount. As seen earlier, however, there is no necessary relationship between price and marginal cost on any side of multi-sided

platform markets. In fact, the price on one side of the market could be well above marginal cost, while the price on the other side of the market could be below marginal cost. To analyze market power from this perspective, one has to examine whether the total price is significantly above total marginal costs.

This, however, is an empirically difficult task due to the inaccessibility of relevant data shielded from examination by nondisclosure agreements among the parties. Here, again, regulatory intervention to secure the necessary data for public examination seems the only method likely to make evaluation of market power in a two-sided market possible. Evans' analysis is supported by the findings of E. Farhi and A. Hagiu (2008) who conclude that "the strategy space and the logic of competitive advantage are fundamentally different in two-sided markets relative to one-sided markets."¹⁴

T. Tangerås (2009) arrives at interesting results with implications for two-sided internet interconnection markets by a critique of the price neutrality result reported by J.-J. Laffont, P. Rey, and J. Tirole (1998a) in their study of access charges in interconnected telephony markets. Tangerås compellingly argues that the price neutrality result in J.-J. Laffont, P. Rey, and J. Tirole (1998a) is knifed-edged, and hence remarkably non-robust as well as equally non-robust in its policy implications:

The networks by their choice of access charge affect the social surplus, as well as the intensity of competition for subscribers through the effect on call tariffs. The social surplus is divided between the industry and the consumers in proportion to the price sensitivity of subscription demand - the subscription elasticity. When subscription demand is elastic, competition for subscribers is intense. With intense competition the equilibrium subscription fees are low and consequently most of the surplus goes to the consumers. When demand is inelastic, competition is weak and most of the surplus goes to the industry. Network profit is independent of the access charge if and only if the subscription elasticity is proportional to consumer net surplus at symmetric prices.

This implies that access charges for participation in a two-sided market can be

¹⁴ See also J. Wright (2003) for additional insights into this problem.

manipulated to the price advantage of one side of the market. Interestingly, this is confirmed in a special case examined by J.-J. Laffont, P. Rey, and J. Tirole (1998b) showing that the price neutrality result does not hold if networks price discriminate between on-net and off-net traffic. There exists substantial likelihood that such price discrimination in part underlies the Comcast-Level3 peering dispute and, again, calls out for regulatory intervention to pierce the veil of nondisclosure agreements so that empirical data for evaluation can be obtained. Weyl (2006) provides a general analysis of the price structure of two-sided markets which finds that “competition, price controls and subsidies always reduce the *price level*, defined as the sum of prices on the two sides of the market. However, price controls and competition that are “unbalanced” may raise prices on one side of the market.” Given that Roson (2005), following B. Jullien (2004), confirms that Ramsey-Boiteux pricing¹⁵ is commonplace in two-sided markets, any ability of a party to the two-sided market to manipulate price structure has social welfare implications.

It is clear from even a cursory examination of the two-sided market literature that it is not possible to adequately model internet interconnection arrangements, like the Comcast-Level3 dispute, nor to assess the antitrust implications of such arrangements without relying two-sided market theory.

The Evolution of Complexity in Interconnection Models

P. Faratin, et al. (2008) points out how changing market conditions have increased empirical diversity and complexity in internet interconnection modalities which have, to

¹⁵ Under Ramsey-Boiteux pricing, in a monopoly, the price markup is inverse to the price elasticity of demand: the less the elastic demand for the product, the larger the price markup.

some extent, moved beyond the peering-transit antimony. They focus on the evolution of increasingly market-based incentives and strategies for networks in approaching interconnection arrangements rather than direct externality effects (e.g. H. Nuechterlein and P. Weiser 2005) or agency models (e.g., P. Milgrom, B. Mitchell, and P. Srinagesh 2000). In particular, they observe that

Over time and with the growth of Internet traffic, the idea that ISPs of a certain size were more or less the same has eroded. We have seen the emergence of "eyeball" heavy broadband access networks such as Verizon, AT&T or Comcast, on the one hand, and "content" heavy networks such as Abovenet or Cogent (that host a lot of content servers) and large content-providers such as Google and Yahoo. These networks have asymmetric traffic flows. Users on eyeball networks send small requests to content servers on content networks, while the servers on a content network send large replies. The "eyeball" customers want the content since that is part of the reason they pay for broadband service; the content networks need the eyeballs because that is what they sell to advertisers and the "eye balls" are the end-users who may subscribe directly to pay-to-view content. Thus, there are demand complementarities across distinct end-host markets (content providers and consumers) who are customers of ASes. Such markets also exhibit strong indirect externalities, where consumption by one side of the market increases as the consumption of the other market grows. The question of who should pay whom to recover the costs of supporting that interconnection is ambiguous in this asymmetric world. (P. Faratin, et al., 2008, 38)

This transition from relatively symmetrical networks to content-heavy CDNs and eyeball-heavy BANs precisely conditions the emergence of disputes like Comcast-Level3 over peering. Faratin, et al. (2008, 59) is prescient in projecting the emergence of such disputes and speculating on underlying causes for them.:

We observe in practice that most content-heavy networks are more open in their peering policies than are most eyeball-heavy networks. We can speculate on a number of reasons for this difference:

- As opposed to early access networks where switching costs for consumers were insignificant (because they could call any local modem bank ISP), modern broadband consumers may feel that switching costs are relatively higher, assuming they even have a choice of providers. Therefore eyeball networks may perceive that they have some increased bargaining power because they "own" the eyeballs.
- Eyeball networks believe that the "natural" direction of value flow is toward them, rather than away from them. The growth of Internet advertising suggests that content-providers place high value on reaching end-users on eyeball networks.

- The last-mile networks of the broadband eyeball networks are more capital intensive, often involving "lumpy" investments, than are the long-haul and backbone networks of content-providers. Consequently, the cost recovery challenge of the last-mile networks is greater (although as noted earlier, it is not clear that their incremental costs for delivery are higher).

Furthermore,

- If the eyeball network has a very strong bargaining position, it might try to force the content provider to purchase transit service from the eyeball network. However, apart from cost, the content provider may not need access to all of the Internet, but only to the cone of prefixes that belong to the eyeball network.
- Both the eyeball network and the content network can purchase transit from third parties, in which case both are worse off, and the third parties are the only beneficiaries.
- If the eyeball network already peers with some third network, the content network can negotiate a transit arrangement with that third network. Again, in this case, only the third network is better off. (P.Faratin, et al., 2008, 60)

New Business Models at the Root of the Conflict

At the heart of the Comcast-Level3 peering dispute is the emergence of new business models which directly impact the way traffic is moved from source to end-user. From the beginning of the internet the compensation mechanism for interconnection has remained relatively stable, primarily because network structure and traffic appeared to be relatively symmetric and stable between interconnecting networks and, frankly, the details of such arrangements were relative non-salient to firms which did not focus on transporting traffic from place to place. However, over time there has been a significant increase in end-user demand for traffic as users shifted from traditional wireline service to VOIP and downloaded and streaming video has become the largest source of traffic volume. This has become patent at exactly the same time as the rapid expansion of broadband subscribership has begun to taper off. While the contention that broadband

prices are declining as the broadband market has approached saturation (there are real constraints to expanding, for example, rural subscribership in the U.S. which arises from pressure by capital markets on major providers) may be controversial, there is no question that the rate of subscriber recruitment has begun to seriously drop. A consequence of reduced recruitment rates is pressure broadband providers to cut costs and find new revenue streams to maintain their rate of profit. New profit centers and new models of doing business have resulted from these pressures. It is precisely this pressure which motivates Comcast to resist any attempt by Level3 to reduce the revenue which Akamai formerly provided for access to Comcast's "eyeballs" by shifting from traditional CDN peering to a transit model for Netflix content. In this new world BANS like Comcast are likely to be able to use market power derived from the relative lock-in of their subscribers to extract rents from CDNs/transit providers like Level3. Likewise, the declining profitability of backbone provision has forced Level3, like other providers, to focus on content distribution as a profit center and to seek to move from traditional, settlement-free peering to a transit/paid peering model for the exchange of traffic with BANS.

The review, undertaken above, of the economic literature pertinent to internet interconnection suggests the importance of regulatory scrutiny to ensure that anticompetitive or predatory practices do not emerge in the two-sided market structure which has emerged in internet interconnection. This is particularly the case where the literature indicates that there can be behaviors which are prima facie anticompetitive in traditional market, but which are not suitable matters for antitrust action due to the specific dynamics encountered in two-sided markets. This means that the transparency in

peering and transit arrangements which the industry has shunned through nondisclosure agreements becomes fundamental to determining whether any parties to interconnection are rent-seeking. Such transparency can only come into being by regulatory intervention. Only when regulatory agencies have forced industry to disclose the necessary data for genuine antitrust evaluation can the public interest be satisfied that new business models in internet interconnection are not facilitating reduction of social welfare. This, in the long term, may well be the most important lesson derived from the Comcast-Level3 dispute.

REFERENCES

- R. Agarwal and M. Gort (1996). "The evolution of market and entry, exit, and survival of firms," *Review of Economics and Statistics* 78:3, 489–98.
- M. Armstrong (2005). "Competition in Two-Sided Markets," *RAND Journal of Economics*, 37:3, 668-691.
- N. Badasyan and S. Chakrabarti (2003). "Private Peering Among Internet Backbone Providers." Working paper.
- _____ (2004) "Inter and Intra Backbone Peering Decisions Among Internet Service Providers." Working paper.
- _____ (2008). "A Simple Game Theoretic Analysis of Peering and Transit Contracting among Internet Access Providers," *Telecommunications Policy*, 32:1, 4-18.
- F. Bloch and M. Jackson (2004). "The Formation of Networks with Transfers among Players." Working paper.
- J. Cremer P. Rey, and J. Tirole. (2000). "Connectivity in the Commercial Internet," *The Journal of Industrial Economics*, 48:4, 433-472.
- D. Evans (2003). "The Antitrust Economics of Multi-Sided Platform Markets," *Yale Journal on Regulation*, 20:2, 325–82.
- P. Faratin, et al. (2008). "The Growing Complexity of Internet Interconnection," *Communication and Strategies* 72, 51-71.
- E. Farhi and A. Hagiu (2008). "Strategic Interactions in Two-Sided Market Oligopolies" Working paper.
- J. Farrell and C. Shapiro (1988). "Dynamic competition with switching costs," *RAND Journal of Economics*, 19:1, 123-137.
- J. Farrell and P. Klemperer (2006). "Coordination and Lock-In: Competition with Switching Costs and Network Effects." Working paper.
- M. Jackson and A. Wolinsky (1996). "A Strategic Model of Social and Economic Networks"; *Journal of Economic Theory*, 71:1, 44-74.
- E. Jahn and J. Prüfer (2006). "Interconnection and competition among asymmetric networks in the internet backbone market." Working paper.
- B. Jullien (2004). "Two-Sided Markets and Electronic Intermediaries." CESifo Working Paper 1345.
- B. Jullien (2005). "Two-Sided Markets and Electronic Intermediaries," *CESifo Economic Studies*, 51, 235–262.

P. Klemperer (1995). "Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade," *The Review of Economic Studies*, 62:4, 515–539.

S. Klepper (2002). "Firm survival and the evolution of oligopoly," *RAND Journal of Economics* 33(1), 37–61.

J. Krafft and E. Salies (2006). "The cost of switching Internet providers in the French broadband industry, or why ADSL has diffused faster than other innovative technologies," OFCE Working paper.

J.-J. Laffont, P. Rey, and J. Tirole (1998a). "Network Competition: Overview and Nondiscriminatory Pricing," *RAND Journal of Economics*, 29, 1-37.

_____ (1998b). "Network competition: Price Discrimination," *RAND Journal of Economics*, 29:1, 38-56.

J.-J. Laffont and J. Tirole (2000), *Competition in Telecommunications*. Cambridge, MA: MIT Press.

J.-J. Laffont, S. Marcus, P. Rey, P. and J. Tirole, (2003). "Internet Interconnection and the Off-Net-Cost Pricing Principle," *RAND Journal of Economics*, 34:2, 370-390.

P. Milgrom, B. Metchell, and P. Srinagesh (2000). "Competitive Effects of Internet Peering Policies", in I. Vogelsang and B. Compaine, eds. *The Internet Upheaval*. Cambridge, MA: MIT Press, 175-195.

W. Norton (2002). "A Business Case for ISP Peering." Equinix White Papers.

J. Prüfer and E. Jahn (2007). "Dark clouds over the internet?" *Telecommunications Policy*, 31, 144-154.

J.-C. Rochet and J. Tirole (2004). "Two-Sided Markets: An Overview." Working paper.

J.-C. Rochet and J. Tirole (2005). "Two-Sided Markets: A Progress Report." Working paper.

R. Roson (2005). "Two-Sided Markets: A Tentative Survey," *Review of Network Economics*, 4:2, 142-160.

G. Shrimali and S. Kumar (2008). "Bill-and-Keep Peering," *Telecommunications Policy*, 32:1, 19-32.

O. Shy (2002). *The Economics of Network Industries*. Cambridge, MA: Cambridge University Press.

T. Tangerås (2009). "On the profit neutrality of access charges in network competition." Working paper.

E. Weyl (2006). "The Price Theory of Two-Sided Markets." Working paper.

J. Wright (2003). "One-sided Logic in Two-sided Markets." Working paper, National University of Singapore.