


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



ELSEVIER

Available online at www.sciencedirect.com

Telecommunications Policy 32 (2008) 33–49

www.elsevierbusinessandmanagement.com/locate/telpolTELECOMMUNICATIONS
POLICY

Internet peering as a network of relations

Steffen Lippert^{a,*}, Giancarlo Spagnolo^b

^a*Department of Commerce, Massey University Auckland, PB 102 904, NSMC, Auckland, New Zealand*

^b*Faculty of Economics, Department of Economics and Institutions, University of Rome Tor Vergata, Via Columbia 2, I-00133 Rome, Italy*

Abstract

This paper applies results from recent theoretical work on networks of relations to analyze optimal peering strategies for asymmetric Internet Service Providers (ISPs). From a network of relations perspective, ISPs' asymmetry in bilateral peering agreements need not be a problem, since when these form a closed network, asymmetries are pooled and information transmission is faster. Both these effects reduce the incentives for opportunism in general, and interconnection quality degradation in particular. The paper also explains why bilateral monetary transfers between asymmetric ISPs (Bilateral Paid Peering), though potentially good for bilateral peering, may have negative effects on the sustainability of the overall peering network.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Internet Service Provider; Internet peering agreements; Transit; Networks of relations; Quality degradation; Implicit contracts

1. Introduction

Consumer welfare from Internet services, such as exchanging e-mails, accessing web sites, or concluding transactions with users located elsewhere, depends crucially on the ability of Internet Service Providers (ISPs) to connect their users with users of other ISPs anywhere in the world. There are two ways in which ISPs can be interconnected. The most common one is *transit*, whereby ISPs buy connectivity from large Internet Backbone Providers (IBPs) that have the ability to carry traffic all over the Internet. The second way to interconnect is *peering agreements*, that is, reciprocal relationships where two ISPs establish and maintain a direct interconnection through which they exclusively exchange their own end-customers' traffic with each other (Giovannetti, Neuhoff, & Spagnolo, 2005). Unless two ISPs exchange very little traffic, peering is the most efficient form of interconnection as—being a direct connection—it does not route traffic through other intermediate networks. This avoids potential congestion on other paths and increases overall traffic speed.

A frequently voiced problem with reciprocal peering is that traffic flows between two ISPs are often highly asymmetric. This means that ISPs have asymmetric incentives to engage in peering and to maintain a high quality of interconnection, up to the point where situations with a high level of asymmetry have led to the most extreme form of interconnection quality degradation: complete depeering.¹ Monetary compensation to

*Corresponding author. Tel.: +64 9414 0800; fax: +64 9414 8177.

E-mail addresses: s.lippert@massey.ac.nz (S. Lippert), giancaspagnolo@yahoo.com (G. Spagnolo).

¹Historical attempts at depeering include BBNPlanet vs Exodus, PSI vs Cable and Wireless, AOL/ADTN vs Cogent, Teleglobe vs Cogent, France Telecom vs Cogent, Level (3) vs XO, and Level (3) vs Cogent.

balance asymmetric allocations of the costs and benefits of peering, so-called *paid peering*, is seen by many as a natural, though partial, solution to this problem (see, for example, Jahn & Prüfer, 2004; Miller, 2002; Norton, 2002).

Typically, these studies as well as the peering practice view the peering decision as an inherently *bilateral* one. Two ISPs have to decide *bilaterally* whether or not it is worthwhile establishing, maintaining, and ensuring a certain quality of service on a direct interconnection. The embeddedness of these decisions into a network both of traffic exchange and of decisions to peer is ignored.

This paper follows a different path of investigation. It applies recent theoretical results on networks of long-term cooperative relations to the analysis of the optimal interconnection strategies of asymmetric ISPs when quality maintenance is not easily monitored, so that there is scope for *moral hazard* in the form of interconnection quality degradation. It is shown that if ISPs take an explicit network perspective—as appropriate, though opposed to the bilateral perspective typical of the discussion and literature on peering—and if the network is *closed*, asymmetric traffic flows need not lead to any asymmetric incentives to peer. A closed network of peering relations pools asymmetries across peering partners and facilitates information sharing on the behavior of each ISP, thereby fostering cooperation in terms of high interconnection quality maintenance.

The paper proceeds by clarifying why monetary transfers may indeed facilitate bilateral peering relations. It goes on to show, however, that, when an appropriate network perspective is taken, *bilateral* monetary payments—bilateral paid peering—may actually end up harming the *ecosystem* of the peering networks by making a greater number of bilateral peering relations sustainable on their own, independent of other peering relations. This reduced interdependence of the individual elements of the network may prevent it from being able to sustain a high quality of service, as it may lack the ability to deter opportunistic, quality-reducing strategies via the threat of a *domino effect* on interconnection quality.

This study concludes that large ISPs should adopt a global network approach when defining their interconnection strategies, and value more asymmetric peering relationships with smaller ISPs when these close the network and speed up information transmission. It also concludes that, even when ISPs take a network perspective, they will still be likely to undervalue (at least from one side) asymmetric peering relations by not fully internalizing the *social* value of a closed network, but that public policies (for example, subsidies) aimed at correcting this market failure might have the same negative consequence of monetary payments in terms of reduced interdependence and cooperation on quality maintenance in the network. The remainder of the paper is organized as follows: Section 2 discusses the determinants of ISPs' peering strategies, Section 3 proposes a simple model of cooperation in peering, Section 4 introduces the effect of peering networks' *closure* on the sustainability of asymmetric peering relationships, and Section 5 studies the information-sharing effects of peering networks. In addition, both Sections 4 and 5 highlight positive and negative effects of bilateral paid peering. Finally, Section 6 summarizes and proposes policy and strategy implications.

2. The bilateral peering decision

This section examines the peering decision of two ISPs in a purely bilateral setting.

2.1. Peering vs transit

Global connectivity is largely provided to the ISPs by IBPs in exchange for a payment to carry the traffic. IBPs invest in, and maintain, backbone capacity, that is, typically large long-distant fiber optic cables with a huge capacity of data throughput. An ISP connects to the IBP, who charges a fee to the ISP for this connection. Such a fee may be flat for access, related to the *capacity* of the connection link, or may be usage dependent (whereby the traffic is metered and billed) or may be a combination of all three. This contractual relationship between the IBP and the ISP is called *transit*. It is the primary way in which ISPs provide their customers with access to web pages (Miller, 2002). The backbone services market is characterized by huge fixed investments, leading to a small number of IBPs, which, in turn, are able to exercise market power (see Giovannetti & Ristuccia, 2005). This market power often takes the form of discrimination in terms of connection quality (see Crémer, Rey, & Tirole, 2000), and depends on the relative size of the networks




Fig. 1. Two ways of interconnecting: transit and peering.

exchanging data. Another way of providing access to web pages hosted by an ISP different from that of the end user is by means of *peering* agreements. Peering is a bilateral, reciprocal relationship in which two ISPs exclusively exchange their own end-customers' traffic directly with each other (Giovannetti et al., 2005). Typically, these agreements are settlement free, that is, traffic and access are not billed, even though Miller (2002) reports that some networks have recently begun "...charging for peering, [...] because the value proposition is *unbalanced* in some way" (Fig. 1).

To engage in peering, two ISPs have to physically connect their networks. There are two commonly used options for doing so. One is *circuit-based*. In this system, two ISPs create a physical connection (a direct exchange point) between their networks, that is, they run a cable directly from one ISP's network to that of the other, the cost of which is typically shared equally among the two networks. The second option is *exchange-based*. In this system, both ISPs connect a cable to a switch at the same Internet Exchange Point (IXP). An IXP is a physical infrastructure where several different ISPs interconnect with each other. The costs for operating an IXP are shared by its member ISPs. Engaging in peering is, therefore, *economically viable* only from a certain traffic size.

2.2. The advantages of peering

Baake and Wichmann (1999), amongst others, argued that transit implies that a package has to take "a detour... [through] ...several 'routers' that determine which way it has to take... [which] ...typically implies a lower quality of service: transferring the same amount of data takes longer and response time (the so-called 'latency') increases." Peering avoids this detour and, thus, often the time from the start of packet transmission to the start of packet reception (latency) is reduced. Secondly, as only the two ISPs, the one originating the request and the one hosting the web page, are involved in the exchange, they determine the speed of the transmission (the time from the start of packet reception to the end of packet reception) with their own bandwidth and the bandwidth of the connection between themselves. No congestion due to a third network has an impact on the connection quality. Finally, no upstream provider (Internet backbone) needs to be paid for the traffic, which reduces the variable cost of exchanging data between the two ISPs. Besides these general advantages, peering at an IXP often involves other forms of cooperation, such as information sharing or free mutual technical help forums. Thus, the advantages of peering are a reduction in *latency* and an increase in *speed*, and possibly a reduction in the variable *cost* of exchanging data.

2.3. Cooperation in peering

Despite these advantages, the bilateral exchange of data within a peering agreement between ISPs may not always come into existence, even though it may be technically feasible and economically viable. There are three important technological features of the Internet, which determine the *necessity and the sustainability of cooperation* within peering agreements (see Giovannetti et al., 2005): Firstly, "...only the amount of traffic exchanged between two networks can be measured, and not the paths followed by each packet." This means that the quality of connections is not contractible. Secondly, given that transmission within one's own network is costly,² and given that ISPs often have multiple points of interconnection, it is common practice to pass

²This cost includes, for example, the maintenance of routers and long-distance cables as well as the avoidance of congestion (see Buccirossi, Ferrari Bravo, & Siciliani, 2005).




Fig. 2. Hot potato routing (adapted from Kende, 2000).

traffic off to another network as quickly as possible, thus using the other network for wide-area transit. This practice is called *Hot Potato Routing* (see also Laffont, Marcus, Rey, & Tirole, 2001, 2003 for such *hot potato* behavior among Internet backbones). This phenomenon is described in Fig. 2. ISP 1 routes the request of its end user, C1, out of its own network as quickly as possible. ISP 2 does the same with the response of its customer (a web site), C2.

Thirdly, the speed of the connection between two end users is determined by the congestion in the networks on its path. As the hot potato routing feature leads traffic out of an ISP's own network as soon as possible, the most congested network on the path contributes most to any delay. Each *single* ISP's efforts that are aimed at avoiding congestion, that is, maintenance of bandwidth capacity or of redundancy, benefit *both* peering ISPs, leading to free riding and, thus, too low a level of maintenance being carried out. As these efforts are *not contractible*, overcoming this free-riding problem requires *cooperation*, which may be enforced in the dynamic course of interaction between peering ISPs³ through the threat of retaliation. Constantiou and Courcoubetis (2002) describe the informational problems between ISPs and the resulting individually rational *moral hazard* behavior. The inability to perfectly monitor the effort of a peering partner after concluding the peering agreement may induce partners to alter their efforts opportunistically. As a result:

- An ISP may not keep upgrading its network capacity after an interconnection agreement. This will result in poorer servicing of the partner's traffic. As peering agreements are currently based on best effort services, such behavior cannot be easily verified.
- An ISP may actively discriminate against IP packets that enter into its network from the interconnected partner when its network has large amounts of local traffic.
- An ISP may overbook its network in order to maximize economies of scale. To avoid congestion the ISP may delay or not admit interconnected traffic. This is not the predictable outcome under 'naturally' arising congestion, but is the result of intentional unilateral overbooking.

2.4. Asymmetries and bilateral cooperation in peering

In particular, *asymmetries in the traffic flows* present obstacles to the incentives to engage in cooperative peering agreements, the refusal of which is an extreme form of the quality degradation described by Crémer et al. (2000). The main reasons for a refusal to enter peering agreements are *backbone free-riding* and *business stealing* effects. Various studies (e.g. Norton, 2003) highlight the point that most relationships between ISPs are asymmetric and—as emphasized by Giovannetti et al. (2005)—unbalanced situations have, in some cases, led to the discontinuation of the peering arrangement. Large ISPs may not want to engage in peering with smaller, often regional, providers as they fear *backbone free riding* (see Baake & Wichmann, 1999; Kende, 2000 for the IBP market). Fig. 3 describes this phenomenon: The smaller, regional ISP 2 benefits from being connected to a larger, national ISP 1, which maintains a national backbone capacity by being able to connect its customer C2 to the national ISP's geographically distant customer, C1, without having to invest in its own backbone capacity.

³Indeed, given the necessity to cooperate, Giovannetti et al. (2005) find in an empirical study that geographic proximity, which helps the enforcement of implicit contracts, has an impact on the decision of ISPs of whether to peer or exchange traffic via transit.




Fig. 3. Backbone free riding (adapted from Kende, 2000).

		ISP 2	
		e_2^h	e_2^l
ISP 1	e_1^h	$B_1^h - C_1^h, B_2^h - C_2^h$	$B_1^l - C_1^h, B_2^h - C_2^l$
	e_1^l	$B_1^h - C_1^l, B_2^l - C_2^h$	$B_1^l - C_1^l, B_2^l - C_2^l$

Fig. 4. Normal form representation of the effort decision.

Peering agreements induce a reduction in latency, as compared to transit. Without peering, only the customers of the ISP hosting a web site would enjoy this low latency. With a peering agreement, the end users of the competing ISP also enjoy a low latency, thereby making customers more willing to move over to the other service provider. This effect has been described as a *business stealing effect* (Baake & Wichmann, 1999), as it induces large ISPs to reduce the interconnection quality in peering agreements with smaller ISPs in order to steal their customers (see Crémer et al., 2000). Both the backbone free-riding and the business stealing effects imply that the larger network has a greater incentive to renege on interconnection quality than does the smaller one. As expressed by Huston (1999): “The larger provider often provides more traffic to a smaller attached provider than it receives from that provider [...] traffic-receiving volumes typically coincide with the relative interconnection benefit to the two providers”. As peering is technically efficient, solutions have been sought to overcome the asymmetric incentive problems. One natural solution that has emerged is *paid peering* (Miller, 2002; Norton, 2002).

3. A simple model of cooperation in peering

Consider two ISPs, ISP 1 and ISP 2, that must decide whether to start and maintain a direct peering link between each other, or instead indirectly exchange traffic through transit purchased from a large backbone operator. If they purchase transit, let their payoff be t_1 and t_2 , respectively. If they peer, they simultaneously decide what their effort will be in providing interconnection quality. For simplicity, assume that each ISP can only choose either a high effort, e_i^h , or a low effort, e_i^l . If an ISP provides a high (low) effort, his peering partner will receive a high (low) quality interconnection and a benefit of B_i^h (B_i^l), where $B_i^l < B_i^h$. Providing a high (low) effort leads to a cost of C_i^h (C_i^l), where $C_i^h > C_i^l$. The normal form of the game, shown in Fig. 4, represents the strategic situation of the two ISPs.

Assuming $B_i^h - C_i^h > B_i^l - C_i^l$, this payoff matrix fulfills the properties of a prisoner’s dilemma. Both players benefit from a unilateral deviation to providing the lower effort: They save fully in terms of cost, bear none of the loss in terms of benefits from customer satisfaction, and may even win over some of the other ISP’s customers since the connections are of high quality within their own network. This is true, notwithstanding what level of effort the other ISP is providing. If both are providing a low level of effort, however, both are

worse off than if both were providing a high level of effort. The theory of repeated games⁴ postulates that if players maximize the present value of their utility stream discounted with the relevant factor, δ , and there is complete information on each party's past behavior, the outcome (e_1^h, e_2^h) can be sustained in equilibrium if and only if both

$$\frac{(B_1^h - C_1^h)}{1 - \delta} \geq (B_1^h - C_1^l) + \frac{\delta(B_1^l - C_1^l)}{1 - \delta},$$

and

$$\frac{(B_2^h - C_2^h)}{1 - \delta} \geq (B_2^h - C_2^l) + \frac{\delta(B_2^l - C_2^l)}{1 - \delta}.$$

The payoffs in the normal form are ISP specific, indicating potential asymmetries. Let, without loss of generality, ISP 1 be larger than ISP 2. Then, typically, ISP 2 benefits more from the peering relation than ISP 1. That is,

$$B_1^h - B_1^l < B_2^h - B_2^l.$$

This may lead to a situation where

$$\frac{(B_1^h - C_1^h)}{1 - \delta} < (B_1^h - C_1^l) + \frac{\delta(B_1^l - C_1^l)}{1 - \delta},$$

while

$$\frac{(B_2^h - C_2^h)}{1 - \delta} \geq (B_2^h - C_2^l) + \frac{\delta(B_2^l - C_2^l)}{1 - \delta}.$$

In this case, in a peering relation, a low level of effort would be provided and a peering agreement would only be concluded and sustained if the payoff from a peering agreement with these low levels of effort, $B_1^l - C_1^l$ and $B_2^l - C_2^l$, were larger than that from transit, t_1 and t_2 , for each of the two operators. Paid peering has emerged and is advocated as a possible, natural way of addressing the traffic asymmetry problem in bilateral peering relationships (e.g. Miller, 2002; Norton, 2002). Indeed, standard game theoretic arguments suggest the possibility that operating monetary transfers across parties may facilitate cooperation in asymmetric bilateral relationships. If ISP 2 transfers an amount T to ISP 1 in each period, maintaining a high quality peering link will be in equilibrium if

$$\frac{(B_1^h - C_1^h) + T}{1 - \delta} \geq (B_1^h - C_1^l) + T + \frac{\delta(B_1^l - C_1^l)}{1 - \delta}$$

and

$$\frac{(B_2^h - C_2^h) - T}{1 - \delta} \geq (B_2^h - C_2^l) + \frac{\delta(B_2^l - C_2^l)}{1 - \delta}.$$

These two conditions imply the following proposition, the proof of which is relegated to the Appendix A.

Proposition 1. *A monetary transfer, $T \geq 0$ (bilateral paid peering), facilitates high quality peering connections among asymmetric ISPs if, and only if,⁵*

$$\delta \left((B_2^h - B_2^l) - \frac{C_2^h - C_2^l}{\delta} \right) - \left((B_1^h - B_1^l) - \frac{C_1^h - C_1^l}{\delta} \right) \geq 0.$$

Proposition 1 states that, given the prevailing discount factor δ at which the larger ISP has a negative net gain from maintaining the high quality peering connection, a positive *bilateral* monetary transfer (i.e. *bilateral*

⁴See Abreu (1988).

⁵In the stylized prisoner's dilemma formulation, improved cooperation on a peering connection's quality takes the form of a reduction in the minimum level of the discount factor necessary to sustain high quality as an equilibrium outcome path. As usual, it would be straightforward to build a richer (but more cumbersome) model where parties have more than two actions and where improved cooperation takes the dual form of higher equilibrium connection quality given the discount factor.

		Player 2	
		C	D
Player 1	C	c_1, c_2	l_1, b_2
	D	b_1, l_2	d_1, d_2

Fig. 5. Normal form representation of a prisoners' dilemma.

paid peering) may indeed help in sustaining a cooperative, high quality peering relationship. This will only occur, however, if the small ISP has enough to gain, as compared to the larger ISP. Sections 4.3 and 5.3 will show that when taking a network perspective, other considerations may run against this simple bilateral argument.

4. Pooling asymmetries with complete information

This section will first briefly discuss some theoretical results of Lippert and Spagnolo (2005),⁶ from here on LS (2005), and then apply them to shed some light on when and how the *network* structure of peering relationships among ISPs may substantially reduce the problem of asymmetry in bilateral relations. It will finally show that introducing *bilateral* paid peering is only an imperfect substitute for making use of the network structure of the peering relationships. In the discussion that follows, it is important to distinguish between the *data exchange network* and the *relational network*, which is spanned by the ISPs' peering agreements. Defining the latter as the network of peering relationships, which are part of a multilateral agreement, it is useful to keep in mind that not every pair of ISPs who exchange traffic within a peering relationship need to be part of the multilateral relational network agreement.

4.1. Pooling asymmetries in the theory of relational networks

LS (2005) showed that, if players repeatedly interact in bilateral prisoner's dilemmas with generic, hence asymmetric net gains from cooperation, they may cooperate for a larger range of discount factors if they are able to pool payoff asymmetries in a multilateral punishment mechanism.⁷

Consider Fig. 5, which represents the normal form of a generic prisoner's dilemma, in which c_i stands for the cooperation payoff, d_i for the payoff if both defects, b_i stands for the payoff player i gets if he betrays while the other player cooperates, and l_i stands for the loss payoff if player i is betrayed while cooperating.

In the infinitely repeated version of this game with players maximizing their discounted present value, which is discounted with the given, common discount factor δ , let

$$\frac{1}{1-\delta}c_1 - b_1 - \frac{\delta}{1-\delta}d_1$$

be player 1's net gain from cooperating with player 2, and let the respective expression with payoffs of player 2 be player 2's net gain from cooperating with player 1. In general, this net gain can be positive or negative.

⁶There is a large and very interesting body of literature on network formation (for a survey of the literature, see Jackson, 2004) which deals with the question of predicting an equilibrium network structure and its social properties, given the value of a network, the cost of establishing a link, and the distribution of the network's value among the network members. However, this paper does not deal with network *formation* but with the enforceability of implicit cooperative agreements in dynamic settings within an existing network. To the authors' best knowledge, for this purpose, the closest reference is Lippert and Spagnolo (2005).

⁷Note that symmetric gains from cooperation are a non-generic, empirically irrelevant special case of the model presented here.




Fig. 6. Non-circular networks.




Fig. 7. Circular network.

In order to sustain C_1 , C_2 forever as an equilibrium of the bilateral repeated prisoner's dilemma from Fig. 5, it is necessary and sufficient that *both* players have a positive net gain from cooperating.⁸

Assume that player 1 has a negative net gain, whereas player 2 has a positive net gain. Let these two players be aware of the fact that they are interacting not only with each other, but also with other players and assume for the sake of simplicity that each player observes the history of all other players in the network.

Denoting players as nodes of a network and the cooperative relations between these players as arcs, where an incoming arrow signifies a non-negative net gain from cooperation it is possible to depict a collection of such cooperative relations in a network. Take Figs. 6 and 7 as an illustration. In these figures, the incoming arrow for player 2 means that, in his relation with player 1, $(1/(1 - \delta))c_2 - b_2 - (1/(1 - \delta))d_2 \geq 0$. In other words, 2 has a non-negative net gain from cooperation with 1. Player 1 on the other hand does not have an incoming arrow in his relation with 2, i.e. in his relation with player 2, $(c_1/(1 - \delta)) - b_1 - (\delta d_1/(1 - \delta)) < 0$: player 1 has a negative net gain from cooperation with player 2.

In the networks in Fig. 6, the cooperative outcome is not an equilibrium, as player 1 only has relation(s) with a negative net gain from cooperating. In the network shown in Fig. 7, however, cooperation may be an equilibrium as each player has a relation with a positive net gain from cooperation. Define Strategy Profile 1 as follows.

Strategy Profile 1: Every player, who is part of the multilateral cooperative agreement:

1. starts playing C with every neighbor in the agreement,
2. continues playing C with every neighbor in the agreement as long as he observes C from every player in the agreement, and
3. reverts to D with every neighbor in the agreement forever otherwise.

Strategy Profile 1 is an equilibrium as long as the sum of the net gains of each player for all of its relations is positive. As shown formally in LS (2005), this implies that, with complete information on the history of play, if each of the bilateral relations is such that—due to an asymmetry in the payoffs—one player has positive net gains from cooperation and the other has negative ones,⁹ then, *absent transfers*:

1. Cooperation in non-circular networks, i.e. in lines, stars, or other forms of trees, is not sustainable because of an *end-network effect* (somewhat analogous to well-known *end game effects* of finitely repeated games). Players at the end nodes of such networks always have incentives to defect from cooperation.

⁸Note that the terms in the net gains from cooperating correspond to the payoffs a player receives from playing Friedman's (1971) grim trigger strategies.

⁹Again note that this situation, which may appear uncommon *prima facie*, in fact proxies the real world situation in which many quality levels can be chosen and the discount factor binds (limits quality) with unequal intensity on the two sides of each node. In other words, the most common situation one can think of, perfect bilateral symmetry, is a non-generic, statistically irrelevant special case.




Fig. 8. Three bilateral peering decisions and the backbone free-riding problem.




Fig. 9. Network representation of the bilateral (left) and the multilateral (right) peering decision.

2. Cooperation in circular networks is instead sustainable if, and only if, the sum of the net gains from cooperating is positive for each player.

4.2. Pooling asymmetries in peering

The results just examined have implications for peering agreements among ISPs. Consider an example where the asymmetries of ISPs come from having different geographical home markets. Suppose there to be three ISPs, 1, 2, and 3, with home markets A, B, and C, respectively, in which they are the market leader. Let ISPs 1, 2, and 3 have a subsidiary in the foreign markets C, A, and B, respectively, each having a small market share. Further, assume that a large part of the traffic originates and terminates in the same region.¹⁰ The situation is depicted in Fig. 8.

Fig. 9 depicts the situation described in a network representation. The three ISPs are the nodes of the network and their peering agreements are represented by the edges of the graph. Let the arrows once more signify the net gains of cooperation, i.e. an ISP i has an incoming arrow if, and only if, $((B_i^h - C_i^h)/(1 - \delta)) \geq (B_i^l - C_i^l) + \delta(B_i^l - C_i^l)/(1 - \delta)$. In the left-hand panel, the shadowed representation of the peering relations between ISPs 1 and 3 and ISPs 2 and 3 signifies that these are not taken into account when ISPs 1 and 2 make the decision about peering in market A. In that market, ISP 1 has a negative net gain

¹⁰This is increasingly the case, as noted by Giovannetti and Ristuccia (2005), since there is a growing cultural and linguistic differentiation of web content especially in Europe, and a simultaneous proliferation of regional European IXPs.

from peering with ISP 2 and, bilaterally, peering would not be sustainable. Nevertheless, taking into account that there is a region B and a region C and that the decision whether to, and how to, peer with ISP 3 are to be taken, the ISPs might implement the following multilateral peering agreement:

- The three pairs of ISPs conclude peering agreements.
- Each ISP provides a high level of effort in maintaining interconnection quality in its peering relations as long as it has not observed too low a level of quality provided by any other ISP which is also part of the multilateral agreement.
- Each ISP reverts to the low level of effort in *both* peering relations (or simply depeers) as soon as it observes too low a level of quality provided by any ISP which is part of the multilateral agreement.

This multilateral agreement sustains the high quality connection if, and only if, the sum of the net gains from cooperation with both peering partners is non-negative for each ISP. If any of the three providers starts providing a low level of effort, there will be depeering in both markets, and that ISP will also lose the beneficial high quality peering relationship in the other market. This multilateral mechanism is depicted in the right-hand panel of Fig. 9. This mechanism can be refined in the spirit of Green and Porter (1984) to account for the uncertainties in the market, such that ISPs carry out a limited time punishment whenever the connection quality is lower than a certain threshold. It can also be adapted to an environment where each ISP only observes the quality of its own traffic exchanges. The main insight stays the same and is summarized by Proposition 2.

Proposition 2. *Appropriate pooling of asymmetries through multilateral network agreements enable asymmetric ISPs to sustain high quality peering without transfers as long as the sum of each ISP's net gain from cooperating in peering is greater than zero. This condition may be met by ISPs that are not able to sustain non-paid peering bilaterally because of asymmetries.*

4.3. Why pooling asymmetries in a network agreement may be a better idea than bilateral paid peering

Introducing bilateral paid peering may make a high effort in asymmetric bilateral peering relations possible. As this subsection shows, however, bilateral paid peering is an imperfect substitute for a multilateral agreement that pools payoff asymmetries. Consider the situation depicted in Figs. 8 and 9, and assume that the stage game between each pair of players is that in Fig. 10, where L and R can take the values of 1, 2, and 3, and ISP L is to the left of ISP R (that is, the markets being examined, are ISP L's home and ISP R's foreign markets).

If, in each period, ISP R transfers an amount T to ISP L, maintaining a high quality peering link will be an equilibrium if, and only if,

$$\frac{(B_L^h - C_L^h) + T}{1 - \delta} \geq (B_L^h - C_L^h) + T + \frac{\delta(B_L^l - C_L^l)}{1 - \delta}$$

and

$$\frac{(B_R^h - C_R^h) - T}{1 - \delta} \geq (B_R^h - C_R^h) + \frac{\delta(B_R^l - C_R^l)}{1 - \delta}.$$

		ISP R	
		e_R^h	e_R^l
ISP L	e_L^h	$B_L^h - C_L^h, B_R^h - C_R^h$	$B_L^l - C_L^h, B_R^h - C_R^l$
	e_L^l	$B_L^h - C_L^l, B_R^l - C_R^h$	$B_L^l - C_L^l, B_R^l - C_R^l$

Fig. 10. Normal form representation of the effort decision.

There are, however, situations where these two incentive constraints are violated, whereas those in a multilateral mechanism are not. The multilateral mechanism is an equilibrium if, and only if,

$$\frac{(B_L^h - C_L^h)}{1 - \delta} - (B_L^h - C_L^l) - \frac{\delta(B_L^l - C_L^l)}{1 - \delta} + \frac{(B_R^h - C_R^h)}{1 - \delta} - (B_R^h - C_R^l) - \frac{\delta(B_R^l - C_R^l)}{1 - \delta} \geq 0.$$

Assume that this condition holds with equality. Then, adding up the constraints for enforcing high efforts in a bilateral paid peering relation and rearranging the equations requires that

$$\frac{(B_L^h - C_L^h) + T}{1 - \delta} - (B_L^h - C_L^l) - T - \frac{\delta(B_L^l - C_L^l)}{1 - \delta} + \frac{(B_R^h - C_R^h) - T}{1 - \delta} - (B_R^h - C_R^l) - \frac{\delta(B_R^l - C_R^l)}{1 - \delta} \geq 0$$

or

$$-T \geq 0.$$

As this cannot be satisfied, the multilateral mechanism sustains high efforts for discount factors, for which bilateral paid peering does not. This argument holds as long as it is impossible for ISP R to withhold the transfer to ISP L in the same period where ISP L provided a low level of effort. Proposition 3 summarizes this result.

Proposition 3. *Bilateral paid peering is only an imperfect substitute for a multilateral punishment mechanism.*

Nevertheless, if the pooling of payoff asymmetries across peering relations is not possible (for example, because one of the ISPs is an end node) then, as Proposition 1 has shown, having a bilateral paid peering agreement may enable the ISPs involved to peer.

4.4. Multilateral paid peering

So far, this paper has concentrated on comparing bilateral paid peering, the focus of the Internet policy debate, with multilateral network mechanisms for informal cooperation without transfers, the focus of LS (2005). Having made the step toward a network view on peering agreements, it is only logical to also conceive of a multilateral system of transfer payments \mathbf{T} among peering partners.¹¹ Any agreement implemented by a network of relations without transfer payments is equivalent to such an agreement with zero transfer payments $\mathbf{T} = (0, \dots, 0)$. Thus, any feasible multilateral agreement involving positive transfers, if it is chosen over the always available $\mathbf{T} = (0, \dots, 0)$ will Pareto-improve on the outcome without transfers. Following the logic of Proposition 1, such an agreement is feasible as long as the small ISPs have enough to gain, as compared to the larger ISPs.

Proposition 4. *A feasible multilateral network agreement involving positive transfers may further improve on the interconnection quality sustained by the peering network relative to a multilateral network agreement without transfers (and, a fortiori, by bilateral paid peering).*

A possible problem of such a mechanism could be its relatively high complexity relative to the alternatives upon which this essay mainly focuses. This complexity would probably require some form of centralized coordination and monitoring. Such a task could perhaps be taken up in the future by IXPs. Further work on this particular issue would be welcome, as it cannot be examined in this essay due to reasons of space and focus.

5. Imperfect information transmission

A feature of direct peering connections is that the investments in quality maintenance undertaken by the each of the two ISPs sharing the connection are not directly observed by the other ISPs within the network. The directly connected peering parties, however, will more easily perceive the effects of quality degradation on the other side of the connection, and can exchange/transmit this kind of information to other ISPs. Indeed,

¹¹The authors are grateful to an anonymous referee for pointing out this argument.




Fig. 11. Circular network with four players.

informal information circulation is a frequently mentioned benefit from peering, and (as will be shown) it may help overcome the problem of interconnection quality maintenance in peering. This section will show that, once a network perspective is taken, this feature has implications for the peering decisions of large ISPs with smaller ISPs as well as for bilateral paid peering.

As in the previous section, this section will also first briefly discuss the theoretical results of LS (2005), and then apply these results to shed light on the role information exchange may play within peering relationships among ISPs if they make use of the network structure of the peering relationships. It will finally apply these insights to provide a reason why *bilateral* paid peering may be bad for the *peering ecosystem*.

Once more, in the discussion that follows, it is useful to keep in mind that not every pair of ISPs who exchange traffic within a peering relationship, need be part of the *relational network* of peering relations formed by a multilateral agreement. Bilateral and multilateral agreements may co-exist.

5.1. Imperfect information transmission in the theory of relational networks

LS (2005) showed that the insight that pooling payoff asymmetries in a multilateral punishment mechanism may help to sustain more cooperation than bilateral agreements would (see review in Section 4.1), generalize to environments in which players do not observe the histories of actions undertaken by all other players which are part of the multilateral agreement.

No information transmission and contagion. Consider the situation from Fig. 11, where each pair of players interacts in an infinitely repeated prisoner's dilemma game like the one in Fig. 5, and where again an incoming arrow signifies that a player has a non-negative net gain from cooperation in that particular relation. Assume that every player only observes the history of his own interactions (for example, 4 does not observe the history of 1 and 2) and that communication about the history of play among the players is not possible. LS (2005) showed that a simple modification of Strategy Profile 1 may sustain cooperation in this case:

Strategy Profile 2: Every player, who is part of the multilateral agreement:

1. starts playing C with every neighbor in the agreement,
2. continues playing C with every neighbor in the agreement as long as he observes C from each neighbor in the agreement, and
3. reverts to D with every partner in the agreement forever otherwise.

With the *contagious* Strategy Profile 2, if player 1 defects against player 2, he is punished by player 4 with permanent defection two periods later,¹² and LS (2005) showed that, given the common discount factor, cooperation in a multilateral agreement spanning a circular network where each player has a relation with a positive net gain, and one with a negative net gain from cooperating, is sustainable in equilibrium by this strategy profile if and only if, for every player in the agreement, the sum of the net gains from cooperation, discounted appropriately, is non-negative.

Word-of-mouth communication. Consider now the case where each player can pass on to neighbors information on the history of his play as well as information received from other neighbors. Assume that in

¹²Given this and that he has a positive net gain from the relation with 4, his optimal defection is to first defect from the relation with 2 and two periods later from the relation with 4.




Fig. 12. Mixed network.

every period of interaction, players can meet with more than one neighbor to exchange information. In this environment LS (2005) found that private information about defections would *still not* be transmitted to neighbors. Given that there is no return to cooperation in the future after a defection is observed and the contagious punishment phase starts, every player is better off by keeping that private information for himself and reaping the benefits of defecting, while his other neighbor cooperates one last time. LS (2005) go on to define alternative multilateral *repentance* strategies which overcome this problem as follows.

Strategy Profile 3: Every player, who is part of the multilateral agreement:

1. starts with C with every neighbor in the agreement, and
2. continues playing C with every neighbor in the agreement as long as no deviation of a neighbor is observed or reported,
3. if a player observes D from neighbor j in the agreement, he
 - sends a message about the deviation and goes on playing C with his other neighbors in the agreement,
 - plays D with the cheater until the cheater has played C with him for T_j periods,
 - sends a message about the end of punishment and goes back to 2 thereafter,
4. if a neighbor of the cheater j receives the message of j 's cheating, he
 - plays D with the cheater until both, the cheater has played C with him for T_j periods, and he receives the message from the original sender of the message about the carrying out of punishment,
 - goes back to 2 thereafter,
5. deviants from 3 to 4 are subject to the same punishment.

LS (2005) showed that, with these alternative strategies, information will be exchanged in the network as: (a) there is the prospect of a fast return to cooperation after the punishment phase, the earlier and the faster information circulates in the network and (b) imposing the punishment benefits the punisher. They also showed that players will be able to sustain more cooperation than with Strategy Profile 2, and that (contrary to standard results for the repeated prisoner's dilemma) in the case of networks of relations only repentance strategies (Strategy Profile 3) form an optimal penal code because: (a) they make use of fast information transmission and (b) as the duration of the punishment phase T_j can be individualized, they give an expected continuation payoff of a cheater in a punishment phase equivalent to his maximin payoff. In Fig. 11, player 4 could get to know about 1's betrayal against 2 already during the period of the actual betrayal, leading to his punishment against 1 occurring after one period, as compared to after two periods with the contagious strategies. LS (2005) also studied the impact of making bilaterally sustainable relations part of a multilateral agreement when there is imperfect information transmission. In the network illustrated in Fig. 11, increase the cooperation payoff of player 1 in his relation with player 2, such that he has a positive net gain from cooperating. This results in the network in Fig. 12.

For this case, LS (2005) showed that if: (1) the beliefs¹³ of player 1, in the case where he observes player 4 deviate and player 2 cooperate, is such that he does not put sufficient weight on player 2 having been the first to have deviated in the network *and* (2) players adhere to Strategy Profile 2 then cooperation in the network may break down due to the fact that 1 does not have a sufficiently high incentive to enter into a punishment

¹³One needs to amend Strategy Profile 2 by beliefs, as now there is not anymore an automatic incentive to punish. See LS (2005) for a more detailed treatment of that matter.




Fig. 13. Circular peering network.

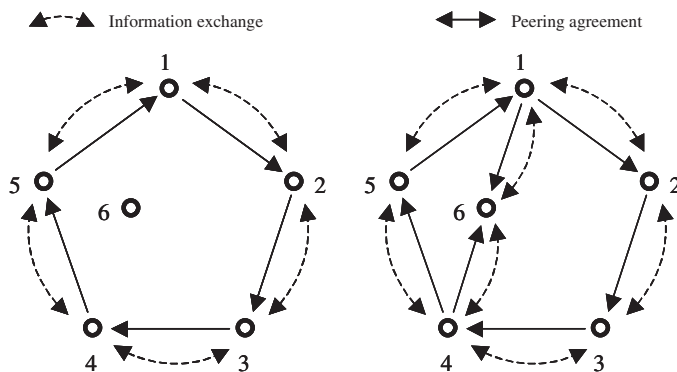


Fig. 14. ISPs 1 and 4 peer with small ISP 6 for faster information transmission.

phase with player 2. Strategy Profile 3, on the other hand, does not suffer from this problem, as entering a punishment benefits the punisher.

5.2. Word-of-mouth communication and peering networks

A frequently mentioned benefit from peering is the exchange of information. In a network of peering relations, ISPs would also benefit from transmitting information about the *quality of service* within the bilateral peering relationships at the IXP, which proxies for the level of effort chosen by the peering partner. This increases the speed of targeted multilateral quality degradation with respect to ISPs that exert a low level of effort and, thus, provides higher incentives to exert a high level effort.

Consider Fig. 13, where four ISPs are active in a total of four markets, with each of them being active in two of these markets. If the quality of service between each pair of ISPs is their private information, ISP 4, for example, does not know whether ISP 1 has provided a high level of effort in its peering relation with ISP 2. ISP 4 observes only the quality of service provided by ISP 3 and ISP 1 in their peering agreement with ISP 4. If ISP 2 transmits information about ISP 1’s quality to ISP 3 and, if ISP 3 passes on this information to ISP 4, then ISPs 2 and 4 can quickly degrade the quality provided to ISP 1. The faster this information travels within the community, the stronger are the incentives to provide high efforts.

Benefits from fast information transmission may in fact be large enough to justify peering with smaller ISPs and should be taken into account in the peering decisions. Suppose that two ISPs with a large installed base consider peering with an ISP with a small installed base within a network of peering agreements. If this peering agreement with the small ISP speeds up information transmission it may help sustain high effort peering agreements in the rest of the network. Fig. 14 describes the situation.

If, in the left-hand panel of Fig. 14, ISP 5 provides low efforts in its peering relation with ISP 1, information about this must travel through three links before it arrives at ISP 4 and, thus, before ISP 4 can also enter targeted quality degradation in its peering relation with ISP 5. If ISP 1 and ISP 4 peer with the smaller ISP 6, then the information must travel through only two links before it arrives at ISP 4. This reduces the delay of the punishment.

Proposition 5 summarizes the arguments.

Proposition 5. 1. *Perfect information about interconnection quality is not necessary for an appropriate pooling of asymmetric peering relationships in a multilateral network agreement to enable ISPs to peer without transfers.*

2. *If ISPs can exchange information on connection quality, they can implement high quality connections at lower discount factors with multilateral strategies having targeted, time-limited, harsh punishment of cheaters,¹⁴ than with multilateral strategies which punish cheating by depeering.¹⁵*

5.3. Paid peering in peering networks

Section 4.3, has shown that introducing bilateral paid peering is only an *imperfect* substitute for a multilateral peering agreement which pools payoff asymmetries. Building on Section 5.1, this section will argue that it may even be harmful to other ISPs' abilities to peer in multilateral peering agreements and, thus, be ecologically bad.

Bilateral paid peering may make the choice of applying a high level of effort enforceable on a bilateral basis. Thus, if two ISPs agree on bilateral paid peering, their relationship becomes sustainable without the rest of the network. The peering network will, therefore, be a mixed network, with the result of a reduced interdependence of the individual elements of the network. If the reaction (punishment) to a low level of effort of a neighboring ISP is *permanent* service degradation (or depeering) towards both partner ISPs, then bilateral paid peering may prevent the network's ability to sustain a high quality of service by deterring opportunistic quality-reducing strategies with the threat of a domino effect on interconnection quality.

Proposition 6. *Uncoordinated bilateral paid peering reduces the interdependence of individual ISPs in the peering network and may, thus, prevent more efficient multilateral peering agreements.*

Note that this paper's simple model neither takes into account any other costs of the ISPs than those of maintaining the connection nor does it deal with network formation or equilibrium selection. However, if one was to write down a more complicated model, one might be able to come up with a strategic use of bilateral paid peering: As it makes it more difficult for other ISPs to peer, using a network of relations to enforce the high quality connection, it: (a) raises rivals' costs and (b) reduces rival's service quality. Both may be reasons for potential entrants not to enter a market in the first place (see footnote¹¹).

6. Policy discussion and conclusions

Asymmetries between ISPs have been an obstacle to the formation of bilateral peering agreements. Given that, unless two ISPs are exchanging very little traffic, peering is the most efficient form of interconnection, *desirable* agreements do not come into existence. This also harms consumer welfare. Monetary compensation to balance asymmetric allocations of costs and benefits of peering, the so-called *paid peering*, is seen by many as a natural solution to this problem.

In this paper, it has been argued that ISPs should take a network perspective in their decision of whether to peer and whether to engage in paid peering. Likewise, policy makers should adopt a network perspective in their advocacy of peering and, eventually, paid peering.

Applying the theory of networks of relations, it has been shown that ISPs may pool asymmetric incentives to engage in cooperative peering through multilateral agreements. Asymmetric traffic flows need not lead to asymmetric incentives to peer, if the network is *closed* and information circulates in the network.

Moreover, it has been clarified why bilateral paid peering, even though it may be good in a bilateral peering relation, may actually be bad for the peering ecosystem by making more bilateral peering relations sustainable on their own, independent of other peering relationships. This reduced interdependence of individual elements of the network may prevent the network from sustaining a high quality of service, by deterring opportunistic quality-reducing strategies with the threat of a *domino effect* on interconnection quality.

¹⁴Like Strategy Profile 3.

¹⁵Like Strategy Profile 2.

Finally, it has been argued that large ISPs should adopt a global network approach when defining their interconnection strategies, valuing asymmetric peering relationships with smaller ISPs when these close the network and speed up information transmission. Moreover, even if they take a network perspective, ISPs will still be likely to undervalue (at least from one side) asymmetric peering relations by not fully internalizing the *social* value of a closed network. Nevertheless, public policies (e.g. subsidies) aimed at correcting this market failure could have the same negative consequence as would monetary payments, in terms of reduced interdependence and cooperation on quality maintenance in the network.

Acknowledgments

The authors would like to thank Emanuele Giovannetti for encouraging them to write this paper as well as Simona Fabrizi and conference participants at the University of Cambridge 2006 ASSET meeting for their helpful comments. They are especially grateful to two anonymous referees for excellent suggestions which significantly improved the paper. Steffen Lippert gratefully acknowledges funds from the CEPR RTN ‘Competition Policy in International Markets’ (HPRN-CT-2002-00224) and research funds from Massey University College of Business.

Appendix A. Proof of Proposition 1

The conditions for high quality to be an equilibrium

$$\frac{(B_1^h - C_1^h) + T}{1 - \delta} \geq (B_1^l - C_1^l) + T + \frac{\delta(B_1^l - C_1^l)}{1 - \delta} \Leftrightarrow B_1^l - B_1^h - \frac{C_1^l - C_1^h}{\delta} \geq T,$$

$$\frac{(B_2^h - C_2^h) - T}{1 - \delta} \geq B_2^h - C_2^l + \frac{\delta(B_2^l - C_2^l)}{1 - \delta} \Leftrightarrow \delta(B_2^l - B_2^h) - (C_2^l - C_2^h) \leq T,$$

can be rewritten as

$$\delta(B_2^l - B_2^h) - (C_2^l - C_2^h) \leq T \leq (B_1^l - B_1^h) - \frac{C_1^l - C_1^h}{\delta}.$$

One can find a $T \geq 0$ such that condition () holds if, and only if,

$$\delta(B_2^l - B_2^h) - (C_2^l - C_2^h) \leq (B_1^l - B_1^h) - \frac{C_1^l - C_1^h}{\delta},$$

which simplifies to

$$\delta \left((B_2^l - B_2^h) - \frac{(C_2^h - C_2^l)}{\delta} \right) - \left((B_1^h - B_1^l) - \frac{(C_1^h - C_1^l)}{\delta} \right) \geq 0.$$

References

- Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica*, 56, 383–396.
- Baake, P., & Wichmann, T. (1999). On the economics of Internet peering. *Netnomics*, 1, 89–105.
- Buccirosi, P., Ferrari Bravo, L., & Siciliani, P. (2005). Competition in the Internet backbone market. *World Competition*, 28(2), 235–254.
- Constantiou, I., & Courcoubetis, C. (2002). Information asymmetry models in the Internet connectivity market. In *Proceedings of IEW 2001, Berlin, Germany*.
- Crémer, J., Rey, P., & Tirole, J. (2000). Connectivity in the commercial Internet. *Journal of Industrial Economics*, 48(4), 433–472.
- Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38, 1–12.
- Giovannetti, E., Neuhoff, K., & Spagnolo, G. (2005). Agglomeration in Internet co-operation peering agreements. *CWPE 0505*, University of Cambridge.
- Giovannetti, E., & Ristuccia, C. A. (2005). Estimating market power in the Internet backbone using Band-X data. *Telecommunications Policy*, 29, 269–284.
- Green, E., & Porter, R. (1984). Non-cooperative collusion under imperfect price information. *Econometrica*, 52, 87–100.
- Huston, G. (1999). Interconnection, peering, and settlements. In *Proceedings of the INET99 Internet society conference*.

- Jackson, M. (2004). A survey of models of network formation: Stability and efficiency. In G. Demange, & M. Wooders (Eds.), *Group formation in economics: Networks, clubs and coalitions* (pp. 11–57). Cambridge: Cambridge University Press.
- Jahn, E., & Prüfer, J. (2004). Transit versus (paid) peering: Interconnection and competition in the Internet backbone market. Available at SSRN: <<http://ssrn.com/abstract=590582>>.
- Kende, M. (2000). The digital handshake: Connecting Internet backbones. *OPP Working Paper No. 32*. Washington, DC: Federal Communications Commission.
- Laffont, J.-J., Marcus, S., Rey, P., & Tirole, J. (2001). Internet peering. *American Economic Review*, 91(2), 287–291.
- Laffont, J.-J., Marcus, S., Rey, P., & Tirole, J. (2003). Internet interconnection and the off-net-cost pricing principle. *RAND Journal of Economics*, 34(2), 370–390.
- Lippert, S., & Spagnolo, G. (2005). Networks of relations and social capital. *CEPR Discussion Paper No. 5078*.
- Miller, R. (2002). The economics of peering. CarrierHotels.com.
- Norton, W. B. (2002). The art of peering: The peering playbook. <<http://www.xchangepoint.net/info/wp20020625.pdf>>.
- Norton, W. B. (2003). The evolution of the US Internet peering ecosystem. <<http://www.nanog.org/mtg-0405/pdf/norton.pdf>>.