



**GOVERNANCE AND THE EFFICIENCY  
OF ECONOMIC SYSTEMS  
GESY**

Discussion Paper No. 191

## Internet Peering as a Network of Relations

Steffen Lippert\*  
Giancarlo Spagnolo\*\*

November 2006

\*Steffen Lippert, Massey University, Department of Commerce, PB 102 904, NSMC, Auckland, New Zealand.  
Tel. +64 9 414 0800 Ext. 9283, s.lippert@massey.ac.nz

\*\*Giancarlo Spagnolo, Stockholm School of Economics, Department of Economics, P.O. Box 6501, SE-113 83  
Stockholm, Sweden, and Consip Research Unit, Via Isonzo, 19/E, I-00198 Rome, Italy.  
Tel +39 320 431 2186, giancaspagnolo@yahoo.com

Financial support from the Deutsche Forschungsgemeinschaft through SFB/TR 15 is gratefully acknowledged.

# Internet Peering as a Network of Relations\*

Steffen Lippert<sup>†</sup>  
Massey University  
Department of Commerce

Giancarlo Spagnolo<sup>‡</sup>  
Università degli Studi di Roma Tor Vergata,  
Stockholm School of Economics,  
Consip, and CEPR

November 27, 2006

## Abstract

We apply results from recent theoretical work on networks of relations to analyze optimal peering strategies for asymmetric ISPs. It is shown that - 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. We also explain why bilateral monetary transfers between asymmetric ISPs (Bilateral Paid Peering), though potentially good for bilateral peering, may have rather negative effects on the sustainability of the overall peering network.

---

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

<sup>†</sup>Massey University, Department of Commerce, PB 102 904, NSMC, Auckland, New Zealand. Phone +64 9 414 0800 Ext. 9283. Email [s.lippert@massey.ac.nz](mailto:s.lippert@massey.ac.nz).

<sup>‡</sup>Stockholm School of Economics, Department of Economics, P.O. Box 6501, SE-113 83 Stockholm, Sweden, and Consip Research Unit, Via Isonzo, 19/E, I-00198 Rome, Italy. Phone +39 320 431 2186. Email [giancaspagnolo@yahoo.com](mailto:giancaspagnolo@yahoo.com).

# 1 Introduction

Consumer welfare from Internet services, such as exchanging emails, accessing web sites, or concluding transactions with users located elsewhere, crucially depends on the ability of Internet Service Providers (ISPs) to connect their users with those of other ISPs anywhere in the world. There are two ways in which ISPs can be interconnected. The most common one is *transit*: ISPs buy connectivity from large Internet Backbone Providers (IBPs) that can carry traffic all over the Internet. The second way to interconnect is *peering agreements*, i.e. reciprocal relationships, where two ISPs establish and maintain a direct interconnection through which they exchange exclusively their own end-customers' traffic with each other (Giovannetti et al. [9]). 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, thus avoiding potential congestion on other paths and increasing overall traffic speed.

A frequently voiced problem with reciprocal peering is that traffic flows between two ISPs are often highly asymmetric, so that ISPs have asymmetric incentives to engage in peering and to maintain a high quality of interconnection, up to the point where too asymmetric situations have led to the most extreme form of interconnection quality degradation: complete depeering<sup>1</sup>. Monetary compensation to balance asymmetric allocations of costs and benefits of peering, so-called “paid peering”, is seen by many as a natural, though partial, solution to this problem (see e.g. Miller [20], Norton [21], or Jahn and Prüfer[14]).

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 insuring 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.

In this paper, we take a different road. We apply recent theoretical results on networks of long-term cooperative relationships 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. We show that taking an explicit network perspective – as appropriate, though opposed to the bilateral perspective typical of the discussion and literature on peering – asymmetric traffic flows need not lead to any asymmetric incentives to peer, if the network is “closed”. A closed network of peering relationships 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.

We proceed by clarifying why monetary transfers may indeed facilitate bilateral peering relationships; then showing that, when an appropriate network perspective is taken, *bilateral* monetary payments – bilateral paid peering – may

---

<sup>1</sup>Historical 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, or Level (3) vs Cogent.

actually end up harming the “ecosystem” of the peering networks 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 its ability to sustain a high quality of service by deterring opportunistic quality-reducing strategies with the threat of a “domino effect” on interconnection quality.

We conclude that large ISPs should adopt a global network approach when defining their interconnection strategies, valuing more asymmetric peering relationships with smaller ISPs when these close the network and speed up information transmission; 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 (e.g. 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

In this section, we will examine the peering decision of two ISPs in a purely bilateral setting.

**Peering vs. transit** Global connectivity is largely provided to the ISPs by Internet Backbone Providers (IBPs) in exchange for a payment to carry the traffic. IBPs invest in and maintain backbone capacity, i.e. typically in 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, which may be flat for the access, related to the *capacity* of the connection link, or usage dependent, where the traffic is metered and billed, or a combination. 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 [20]). The backbone services market is characterized by huge fixed investments, leading to a small number of IBPs, which are, in turn, able to exercise market power (see Giovannetti and Ristuccia [10], or Giovannetti et al. [8]), which often takes the form of discrimination in terms of connection quality (see Crémer et al. [6]), and which depends on the relative size of the networks exchanging data.

Another way of giving access to web pages hosted by an ISP different from that of the end user requesting the access is by means of *peering* agreements. Peering is a bilateral, reciprocal relationship in which two ISPs exchange exclu-

sively their own end-customers' traffic directly with each other (Giovannetti et al. [9]). Typically, these agreements are settlement free, i.e. traffic and access are not billed, even though Miller [20] reports that recently, some networks have started “charging for peering, [...] because the value proposition is *unbalanced* in some way.”




Figure 1: Two ways of interconnecting: transit and peering

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*: Two ISPs create a physical connection (a direct exchange point) between their networks, i.e. they run a cable directly from one ISP’s network to that of the other, the cost of which are typically shared equally among the two networks. The second option is *exchange-based*: Both ISPs connect a cable to a switch at the same Internet Exchange Point (IXP). An IXP is “a physical infrastructure that allows different Internet Service Providers (ISPs) to exchange Internet traffic between their networks”. The costs for operating an IXP are shared by its member ISPs<sup>2</sup>. Engaging in peering is, thus, *economically viable* only from a certain traffic size.

**The advantages of peering** Baake and Wichmann [2] (among other authors) argue 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, i.e. latency, is reduced. Second, 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, i.e. 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 each other. No congestion of a third network has an impact on connection quality. Finally, no upstream provider (Internet backbone) must be payed 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.

<sup>2</sup>More information can be found on [http://en.wikipedia.org/wiki/Internet\\_exchange\\_point](http://en.wikipedia.org/wiki/Internet_exchange_point)

**Cooperation in peering** However, the bilateral exchange of data within a peering agreement among 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. [9]). First, “only the amount of traffic exchanged between two networks can be measured, and not the paths followed by each packet”. This makes the quality of connections not contractible. Second, given that transmission within one’s own network is costly<sup>3</sup>, and given that ISPs often have multiple points of interconnection, it is common practice to pass traffic off to another network as quickly as possible, and thus use the other network for wide-area transit. This practice is called “Hot Potato Routing” (see also Laffont et al. [17] and [18] for “hot potato” behavior among Internet backbones). This phenomenon is described in figure 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.




Figure 2: Hot potato routing (adapted from Kende [15])

Third, 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 one’s own network as soon as possible, the most congested network on the path contributes the most to the delay. Each *single* ISP’s efforts aimed at avoiding congestion, i.e. maintenance of bandwidth capacity or of redundancy, benefits *both* peering ISPs, leading to free-riding and, thus, too low maintenance. 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 ISPs<sup>4</sup> through the threat of retaliation.

Constantiou and Courcoubetis [5] 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.

<sup>3</sup>This cost includes, for example, the maintenance of routers and long-distance cables and the avoidance of congestion (see Buccirosi et al. [4]).

<sup>4</sup>Indeed, given the necessity to cooperate, Giovannetti et al. [9] 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.

As a result,

- an ISP may not keep upgrading his network capacity after an interconnection agreement. This will result in poorer servicing of the partner’s traffic. As peering agreements currently are based on best effort services, such behavior cannot be easily verified;
- an ISP may actively discriminate against IP packets that enter into his 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 the result of intentional unilateral overbooking.

**Asymmetries and 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. [6]. The main reasons for a refusal to enter peering agreements are the “backbone free-riding” and “business stealing” effects.

Various studies (e.g. Norton [22]) highlight that most relationships between ISPs are asymmetric and – as emphasized by Giovannetti et al. [9] – 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 and Wichmann [2], or Kende [15] for the IBP market). Figure 3 describes the 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.




Figure 3: Backbone free-riding (adapted from Kende [15])

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, also the end users of the competing

ISP enjoy a low latency, thereby making customers more willing to move over. This effect has been described as a *business stealing effect* (Baake and Wichmann [2]), as it induces large ISPs to reduce the interconnection quality in peering agreements with smaller ISPs to steal their customers (see Crémer et al. [6]).

Both the backbone free-riding and the business stealing effect imply that the larger network has a higher incentive to renege on interconnection quality than the smaller one. As expressed by Huston [12]: “The larger provider often provides more traffic to a smaller attached provider than it receives from that provider [...]” and “[...] 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 emerged is *paid peering* (Miller [20], Norton [21]).

### 3 A simple model of cooperation in peering

Assume there to be two internet service providers, ISP 1 and ISP 2. Both 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 about their effort to provide interconnection quality. For simplicity, assume that each ISP can only choose either high effort,  $e_i^h$ , or low effort,  $e_i^l$ . If an ISP provides high (low) effort, his peering partner will receive a high (low) quality interconnection and a benefit  $B_i^h$  ( $B_i^l$ ), where  $B_i^l < B_i^h$ . Providing high (low) effort leads to a cost of  $C_i^h$  ( $C_i^l$ ), where  $C_i^h > C_i^l$ . The normal form game in figure 4 represents the strategic situation of the two ISPs.

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

Figure 4: Normal form representation of the effort decision

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

The theory of repeated games<sup>5</sup> tells us that if players maximize the present

---

<sup>5</sup>See Abreu [1]



value of their utility stream discounted with the relevant factor  $\delta$  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

$$\begin{aligned}\frac{(B_1^h - C_1^h)}{1 - \delta} &\geq (B_1^h - C_1^l) + \frac{\delta (B_1^l - C_1^l)}{1 - \delta}, \text{ and} \\ \frac{(B_2^h - C_2^h)}{1 - \delta} &\geq (B_2^h - C_2^l) + \frac{\delta (B_2^l - C_2^l)}{1 - \delta}.\end{aligned}$$

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, i.e.

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

This may lead to a situation where

$$\begin{aligned}\frac{(B_1^h - C_1^h)}{1 - \delta} &< (B_1^h - C_1^l) + \frac{\delta (B_1^l - C_1^l)}{1 - \delta}, \text{ while} \\ \frac{(B_2^h - C_2^h)}{1 - \delta} &\geq (B_2^h - C_2^l) + \frac{\delta (B_2^l - C_2^l)}{1 - \delta}.\end{aligned}$$

In this case, in a peering relation, low efforts would be provided and a peering agreement would only be concluded and sustained if the payoff from a peering agreement with low efforts,  $B_1^l - C_1^l$  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 peering relationships (e.g. Miller [20], Norton [21]). Indeed, standard game theoretic arguments show how the possibility of operating monetary transfers 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 an equilibrium if

$$\begin{aligned}\frac{(B_1^h - C_1^h) + T}{1 - \delta} &\geq (B_1^h - C_1^l) + T + \frac{\delta (B_1^l - C_1^l)}{1 - \delta}, \text{ and} \\ \frac{(B_2^h - C_2^h) - T}{1 - \delta} &\geq (B_2^h - C_2^l) + \frac{\delta (B_2^l - C_2^l)}{1 - \delta}.\end{aligned}$$

These two conditions imply the following proposition, the proof of which we relegate to the appendix.

**Proposition 1** *It is possible to reduce the discount factor necessary to sustain the high quality connection with a transfer  $T \geq 0$  if and only if*

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

Proposition 1 states that, for a given discount factor  $\delta$ , for which the big ISP has a negative net gain from the high quality connection, a positive *bilateral* monetary transfer (i.e. *bilateral* paid peering) helps sustain cooperative peering relationships if and only if the small ISP has enough to gain as compared to the big ISP. In sections 4.3 and 5.3, we will show that, when taking a network perspective, other considerations may run against this simple argument.

## 4 Pooling asymmetries

In this section, we will first briefly discuss some theoretical results in Lippert and Spagnolo<sup>6</sup> [19] (from here on LS) 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. We 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” from 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 relational network.

### 4.1 Pooling asymmetries in the theory of relational networks

LS [19] show that if players repeatedly interact in bilateral prisoners’ dilemmas with asymmetric 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.

Consider figure 5, which represents the normal form of a generic prisoners’ dilemma, in which  $c_i$  stands for the cooperation payoff,  $d_i$  for the payoff if both defect,  $b_i$  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, discounted with a common discount factor  $\delta$ , let

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

---

<sup>6</sup>There is a large and very interesting literature on network formation (for a survey of the literature, see Jackson [13]) 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, in this paper, we are not dealing with network *formation* but with the enforceability of implicit cooperative agreements in dynamic settings within an existing network. To our knowledge, for this purpose, the closest reference is Lippert and Spagnolo [19]

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

Figure 5: Normal form representation of a prisoners' dilemma

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. In order to sustain  $C_1, C_2$  forever as an equilibrium of the bilateral repeated prisoners' dilemma from figure 5, it is necessary and sufficient that *both* players have a positive net gain from cooperating<sup>7</sup>.

Assume that player 1 has a negative net gain, whereas player 2 has a positive one. 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 simplicity that each player observes the history of all other players in the network.

Denoting players as nodes of a network and the 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 figures 6 and 7 as an illustration. In these figures, the incoming arrow for player 2 means that, in his relation with player 1,  $\frac{1}{1-\delta}c_2 - b_2 - \frac{\delta}{1-\delta}d_2 \geq 0$ , i.e. 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,  $\frac{c_1}{1-\delta} - b_1 - \frac{\delta d_1}{1-\delta} < 0$ : player 1 has a negative net gain from cooperation with player 2.




Figure 6: Non-circular networks

In the networks in figure 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 in figure 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 agreement*

<sup>7</sup>Note that the terms in the net gains from cooperating correspond to the payoffs a player receives from playing Friedman's [7] grim trigger strategies.




Figure 7: Circular network

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 for ever otherwise.

Strategy profile 1 is an equilibrium as long the sum of the net gains of each player's for all his relations is positive. This implies the following result shown in LS [19]: For an environment with perfect transmission of information on the history of any player in the network, if each of the bilateral relations is such that due to an asymmetry in the payoffs, one player has a positive net gain from cooperation and the other has a negative one, 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” analogous to “end game effects” in finally repeated games;
2. cooperation in circular networks is 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. Let us examine an example where the asymmetries of ISPs come from 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, which have a small market share. Let us assume that a large part of the traffic originates and terminates in the same region.<sup>8</sup> The situation is depicted in figure 8.

Figure 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

<sup>8</sup>This is increasingly the case, as noted by Giovannetti and Ristuccia [10], since there is a growing cultural and linguistic differentiation of web contents, especially in Europe, and a simultaneous proliferation of regional European Internet Exchange Points.

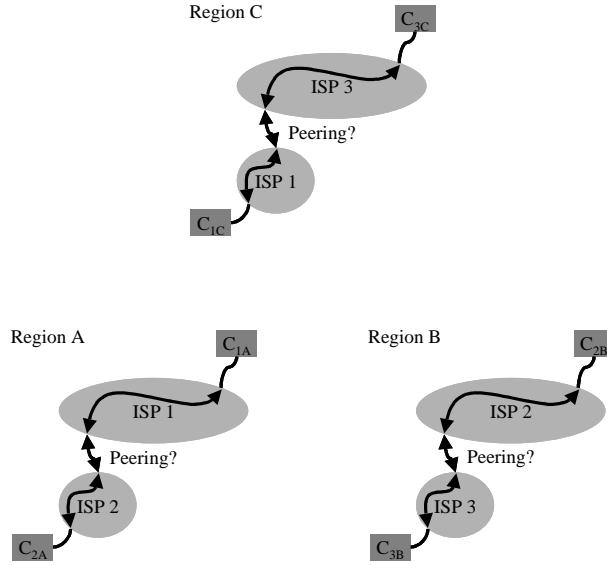


Figure 8: Three bilateral peering decisions and the backbone free-riding problem




Figure 9: Network Representation of the Bilateral (left) and the Multilateral (right) Peering Decision

net gains of cooperation, i.e. an ISP  $i$  has an incoming arrow iff  $\frac{(B_i^h - C_i^h)}{1 - \delta} \geq (B_i^l - C_i^l) + \frac{\delta(B_i^l - C_i^l)}{1 - \delta}$ . 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 from peering with ISP 2 and, bilaterally, peering would not be sustainable. However, taking into account that there is a region B and a region C and that the peering decisions with ISP 3 are to be taken, the ISPs might conclude the following multilateral agreement:

- The three pairs of ISPs conclude peering agreements.
- Each ISP provides high effort in his peering relations as long as he has

not observed too low a level of quality provided by any ISP, who is part of the multilateral agreement.

- Each ISP reverts to the low effort level in *both* peering relations (or just depeers) as soon as he observes too low a level of quality provided by any ISP, who is part of the multilateral agreement.

This multilateral agreement sustains the high quality connection if and only if for each ISP, the sum of the net gains from cooperation with both peering partners is non-negative. If any of the three providers starts providing low effort, there will be depeering in both markets, and he will also lose the beneficial high quality peering relationship in the foreign market. This multilateral mechanism is depicted in right-hand panel of figure 9. This mechanism can be refined in the spirit of Green and Porter [11] 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 his own traffic exchanges. The main insight stays the same and is summarized by proposition 2.

**Proposition 2** *Appropriate pooling of asymmetric peering relationships in multilateral agreements may enable ISPs to peer without transfers, as long as the sum of each ISPs net gains from peering is greater than zero, even if they were not able to peer bilaterally.*

### 4.3 Why pooling asymmetries may be a better idea than bilateral paid peering

Introducing bilateral paid peering may make a high effort in bilateral peering relations possible. However, as we will show now, it is only an *imperfect* substitute for a multilateral agreement that pools payoff asymmetries.

Consider the situation from figures 8 and 9 and assume the stage game between each pair of players to be given in figure 10, where L and R can take the values of 1, 2, and 3 and ISP L is “to the left of” ISP R, i.e. we are looking at ISP L’s home and ISP R’s foreign market.

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

Figure 10: Normal form representation of the effort decision

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

$$\begin{aligned}\frac{(B_L^h - C_L^h) + T}{1 - \delta} &\geq (B_L^h - C_L^l) + T + \frac{\delta (B_L^l - C_L^l)}{1 - \delta}, \text{ and} \\ \frac{(B_R^h - C_R^h) - T}{1 - \delta} &\geq (B_R^h - C_R^l) + \frac{\delta (B_R^l - C_R^l)}{1 - \delta}.\end{aligned}$$

However, there are 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

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

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

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

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 low effort. Proposition 3 summarizes this result.

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

However, if pooling of payoff asymmetries across peering relations is not possible, e.g. 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, we have concentrated on comparing bilateral paid peering with multilateral punishment mechanisms absent transfers. Having made the step toward a network view on peering agreements, it is only logical to conceive also a multilateral system of transfer payments  $\mathbf{T}$  among peering partners. Then, clearly any agreement implemented without transfer payments is equivalent to an agreement with transfers  $\mathbf{T} = (0, \dots, 0)$ . Thus, any feasible multilateral agreement involving positive transfers, if it is chosen unanimously over  $\mathbf{T} = (0, \dots, 0)$ , can only Pareto-improve on the situation 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 big ISPs<sup>9</sup>.

**Proposition 4** *A feasible multilateral agreement involving strictly positive transfers Pareto-improves on a multilateral agreement without transfers.*

## 5 Imperfect information transmission

A feature of peering relations is that the quality of the direct connection between two ISPs is observed privately by the two ISPs sharing it. The exchange of information, a frequently mentioned benefit from peering, may help overcome this problem. In this section, we will show that, once a network perspective is taken, this has implications for the peering decisions of large ISPs with smaller ISPs as well as for bilateral paid peering.

As in the previous section, also here we will first briefly discuss theoretical results in LS [19] and then apply them to shed light on the role of information exchange may play within peering relationships among ISPs if they were to make use of the network structure of the peering relationships. We will finally apply these insights to give 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 to be part of the *relational network* of peering relationships, which is formed by a multilateral agreement. Bilateral and multilateral agreements may co-exist.

### 5.1 Imperfect information transmission in the theory of relational networks

LS [19] show 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 know the histories of all other players, who are part of the multilateral agreement.

---

<sup>9</sup>We are grateful to an anonymous referee for pointing out this argument to us.






Figure 11: Circular network with four players

Consider the situation from figure 11, where each pair of players is interacting infinitely repeatedly in a prisoners' dilemma as the one in the normal form from figure 5 and where again an incoming arrow signifies that a player has a non-negative net gain from cooperation in a particular relation. Assume that every player only observes the history of his own interactions, i.e. 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 [19] show 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 for ever otherwise.

With strategy profile 2, if player 1 defects against player 2, he is punished by player 4 with defection forever two periods later<sup>10</sup>, and LS [19] show that cooperation in a multilateral agreement spanning a circular network, where each player has a relation with a positive and one with a negative net gain from cooperating, is sustainable if and only if for every player in the agreement the sum of the net gains from cooperation, discounted appropriately, is non-negative.

Consider, on the other hand, that each player can transmit information on the history of his play as well as of information received from other neighbors. Assume in addition that in every period of interaction, players can meet with more than one neighbor to transmit information. In this informational environment, LS [19] show that information will not be transmitted if it was made part of strategy profile 2 to immediately transmit information about defections to your neighbors. Given that there is no return to cooperation in the future, every player is better off by keeping that information for himself and reaping benefits of defection while his other neighbor cooperates for one last time.

<sup>10</sup>Given 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

LS [19] go on to define multilateral repentance strategies, which they show to 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 a  $T_j$  periods,
  - sends a message about the end of punishment and goes back to 2. thereafter.
4. If a neighbor  $k$  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 [19] show that with these strategies, information will be transmitted as (a) there is the prospect of a return to cooperation and (b) entering punishment benefits the punisher. They also show that players will be able to sustain more cooperation than with strategy profile 2, in fact, LS [19] show that repentance (strategy profile 3) is an optimal penal code because (a) it makes use of fast information transmission and (b), as the duration of the punishment phase  $T_j$  can be individualized, it gives an expected continuation payoff of a cheater in a punishment phase equivalent to his maximin payoff. In figure 11, player 4 could get to know about 1's betrayal against 2 already in the period of the betrayal, leading to his punishment against 1 after one period as compared to after two periods with the contagious strategies.

LS [19] finally study the impact of making bilaterally sustainable relations part of a multilateral agreement, when there is imperfect information transmission. In the network from figure 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 figure 12.

For this case, LS [19] show that if (1) the beliefs<sup>11</sup> of player 1, in case he observes player 4 deviate and player 2 cooperate, is such that he does not

---

<sup>11</sup>One needs to amend strategy profile 2 by beliefs as now there is not anymore an automatic incentive to punish. For reasons of brevity, we refer to LS [19] for a more detailed treatment of that matter.




Figure 12: Mixed network

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

## 5.2 Information transmission in peering

A frequently mentioned benefit from peering is 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 provided low effort and thus, provides higher incentives to provide high efforts.




Figure 13: Circular Peering Network

Consider figure 13, where four ISPs are active in a total of four markets, each of them 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 provided a high effort in his peering relation with ISP 2. He observes only the quality of service provided by ISP3 and ISP 1 in their peering agreement w.r.t. himself. 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 in 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. Figure 14 describes the situation.

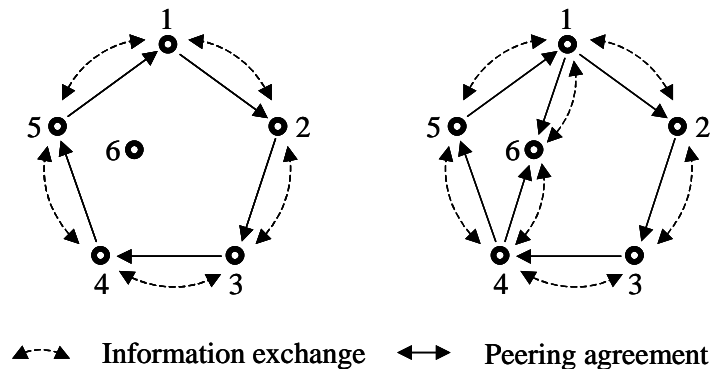


Figure 14: ISPs 1 and 4 peer with small ISP 6 for faster information transmission

If in the left-hand panel of figure 14 ISP 5 provides low efforts in his 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 his peering relation with ISP 5. If ISP 1 and ISP 4 peer with the small 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 within a multilateral peering agreement is not necessary for the result that appropriate pooling of asymmetric peering relationships in multilateral agreements may enable ISPs to peer without transfers.*

2. *If ISPs can exchange information on connection quality, they can implement high quality connections for lower discount factors with multilateral strategies, which use targeted, time-limited, harsh punishment of cheaters<sup>12</sup>, than with multilateral strategies, which punish cheating by de-peering<sup>13</sup>.*

<sup>12</sup>like strategy profile 3

<sup>13</sup>like strategy profile 2

### 5.3 Paid peering in peering networks

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

*Bilateral* paid peering may make the choice of high 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 and the peering network will be a mixed network with the result of a reduced interdependence of the individual elements of the network. If the punishment to low effort of a neighboring ISP is *permanent* service degradation (e.g. through depeering) toward both partner ISPs, this may prevent its 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 multilateral peering agreements.*

Note that our 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<sup>14</sup>.

## 6 Policy discussion and conclusions

Asymmetries between Internet Service Providers have been an obstacle to their ability to enter in bilateral cooperative peering agreements. Given that unless two ISPs are exchanging very little traffic, peering is the most efficient form of interconnection, this means that *desirable* agreements do not come into existence, also harming 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, we have 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, we find that ISPs may pool asymmetric incentives to engage in cooperative peering through multilateral

---

<sup>14</sup>We are grateful to an anonymous referee for pointing out this argument to us.

agreements. Asymmetric traffic flows need not lead to asymmetric incentives to peer, if the network is “closed” and information circulates in the network.

We 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 its ability to sustain a high quality of service by deterring opportunistic quality-reducing strategies with the threat of a “domino effect” on interconnection quality.

We argue 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 undervalue (at least from one side) asymmetric peering relations by not fully internalizing the “social” value of a closed network. However, public policies (e.g. subsidies) aimed at correcting this market failure could have the same negative consequence of monetary payments in terms of reduced interdependence and cooperation on quality maintenance in the network.

## A Appendix

### A.1 Proof of Proposition 1

**Proof.** The conditions for high quality to be an equilibrium,

$$\begin{aligned} \frac{(B_1^h - C_1^h) + T}{1 - \delta} &\geq (B_1^h - C_1^l) + T + \frac{\delta (B_1^l - C_1^l)}{1 - \delta} \Leftrightarrow B_1^l - B_1^h - \frac{C_1^l - C_1^h}{\delta} \geq T, \\ \frac{(B_2^h - C_2^h) - T}{1 - \delta} &\geq B_2^h - C_2^l + \frac{\delta (B_2^l - C_2^l)}{1 - \delta} \Leftrightarrow \delta (B_2^l - B_2^h) - (C_2^l - C_2^h) \leq T, \end{aligned}$$

can be rewritten as

$$\delta (B_2^l - B_2^h) - (C_2^l - C_2^h) \leq T \leq (B_1^l - B_1^h) - \frac{C_1^l - C_1^h}{\delta}. \quad (1)$$

One can find a  $T \geq 0$  such that condition (1) 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

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

■

## References

- [1] Abreu, D. (1988): “On the Theory of Infinitely Repeated Games with Discounting”, *Econometrica*, Vol. 56, pp. 383–396.
- [2] Baake, P. and T. Wichmann (1999): “On the Economics of Internet Peering”, *Netnomics*, Vol. 1, pp. 89–105.
- [3] Besen, S., P. Milgrom, B. Mitchell, and P. Srinagesh (2001): “Advances in Routing Technologies and Internet Peering Agreements”, *American Economic Review*, Vol. 91(2), pp. 292–296.
- [4] Buccirosi, P., L. Ferrari Bravo, and P. Siciliani (2005): “Competition in the Internet Backbone Market”, *World Competition*, Vol. 28(2), pp. 235–254.
- [5] Constantiou, I. and C. Courcoubetis (2002): “Information Asymmetry Models in the Internet Connectivity Market”, Proceedings of IEW 2001, Berlin, Germany.
- [6] Crémer, J., P. Rey, and J. Tirole (2000): “Connectivity in the Commercial Internet”, *The Journal of Industrial Economics*, Vol. 48(4), pp. 433–472.
- [7] Friedman, J. W. (1971): “A Non-Cooperative Equilibrium for Supergames”, *Review of Economic Studies*, Vol. 38, pp. 1–12.
- [8] Giovannetti, E., A. D’Ignazio, and J. Lepler (2005): “On the Peering/Transit Complementarity”, CoCombine Project Report D.17, <http://www.cocombine.org/pdf/D17final.pdf>.
- [9] Giovannetti, E., K. Neuhoff, and G. Spagnolo (2005): “Agglomeration in Internet Co-operation Peering Agreements”, CWPE 0505, University of Cambridge.
- [10] Giovannetti, E. and C. A. Ristuccia (2005): “Estimating Market Power in the Internet Backbone Using Band-X Data”, *Telecommunications Policy*, Vol. 29, pp. 269–284.
- [11] Green, E. and R. Porter (1984): “Non-cooperative collusion under imperfect Price Information”, *Econometrica*, Vol. 52, pp. 87–100.
- [12] Huston, G. (1999): “Interconnection, Peering, and Settlements”, *Proceedings of the INET99 Internet Society Conference*.
- [13] Jackson, M. (2004): “A Survey of Models of Network Formation: Stability and Efficiency,” Chapter 1 in *Group Formation in Economics; Networks, Clubs and Coalitions*, edited by Gabrielle Demange and Myrna Wooders, Cambridge University Press, Cambridge U.K., 2004.
- [14] Jahn, E. and J. Prüfer (2004): “Transit versus (Paid) Peering: Interconnection and Competition in the Internet Backbone Market”, Available at SSRN: <http://ssrn.com/abstract=590582>.

- [15] Kende, M. (2000): “The Digital Handshake: Connecting Internet Backbones”, OPP Working Paper No. 32, Federal Communications Commission, Washington DC.
- [16] Kranton, R. (1996): “Reciprocal Exchange: A Self-Sustaining System”, *American Economic Review*, Vol. 86(4), pp. 830–851.
- [17] Laffont, J.-J., S. Marcus, P. Rey, and J. Tirole (2001): “Internet Peering”, *American Economic Review*, Vol. 91(2), pp. 287–291.
- [18] Laffont, J.-J., S. Marcus, P. Rey, and J. Tirole (2003): “Internet interconnection and the off-net-cost pricing principle”, *RAND Journal of Economics*, Vol. 34(2), pp. 370–390.
- [19] Lippert, S. and G. Spagnolo (2005): “Networks of Relations and Social Capital”, CEPR Discussion Paper No. 5078.
- [20] Miller, R. (2002): “The Economics of Peering”, CarrierHotels.com.
- [21] Norton, W.B. (2002): “The Art of Peering: The Peering Playbook”, <http://www.xchangeport.net/info/wp20020625.pdf>.
- [22] Norton, W.B. (2003): “The Evolution of the U.S. Internet Peering Ecosystem”, <http://www.nanog.org/mtg-0405/pdf/norton.pdf>.