

Evolution of Internet Infrastructure in the 21st century: The Role of Private Interconnection Agreements

Rajiv Dewan*, Marshall Freimer, and Pavan Gundepudi
{dewan, freimer, gundepudpa}@ssb.rochester.edu
Simon School, University of Rochester, Rochester, NY 14627.

Keywords: AM Economic Theory, Internet, public peering, network interconnection, BA0212
Communications Industry.

Abstract:

The infrastructure of the Internet is evolving from that of a public network to interconnected private networks that are selective in terms and conditions for connectivity to others. The connectivity, reach and bandwidth of the Internet in the next century will depend on its infrastructure, which in turn depends on the economics and forms of the interconnection agreements that create the global network. We show that the benefits of interconnection are unevenly distributed and that smaller networks gain more. Further, the smaller networks gain even more as the public peering points get more congested. This makes larger networks less willing to peer with others. This raises public policy issues about network mergers and consolidations that result in some networks being much larger than others.

1 Introduction

The Internet and its infrastructure are undergoing rapid evolution. On the surface, from the perspective of an end-user, the Internet provides a seamless and public wide area network that connects users and resources across the globe. Underneath, however, the way in which the Internet achieves this global connectivity is undergoing rapid change with consolidations, mergers and private interconnection agreements among the firms that provide these services. In this paper we seek to shed light on the evolving infrastructure of the Internet as we head into the 21st century.

The National Science Foundation owned the Internet backbone network prior to 1994 [Leiner, et. al., 1998]. The traffic was primarily from universities and other research sites that were supported by NSF. As the Internet became more commercial, NSF's backbone was replaced by a number of backbone networks owned by different companies connected at a few locations called *public peering points*. NSF initially set up three public peering points, in Chicago, Palo Alto, and Pennsauken, NJ. It later added two more industry run sites, called Metropolitan Access Exchanges, East in Vienna, VA and West in San Jose, CA. In the absence of any private agreements, the networks exchange traffic bound for each other's destinations at these sites.

Figure 1 describes the networking aspects of the Internet. A user connected to her ISP requests a page. The request is sent to the backbone network to which her ISP is connected. As shown in the diagram, the web site may not be directly connected to the user's network. In this case, the request is transferred to the network that hosts the web site at an interconnection and then eventually reaches the server. The requested resource takes a similar route back to the user.

Public peering, though successful at tying the Internet together, has not been a successful solution. For starters, there are just a few public peering or traffic exchange points and they may not match the reach of the networks. For example, most of the major networks have points of presence in cities such as New York City and Boston. And yet, the closest public peering site is in Pennsauken, New Jersey. Consequently,

traffic from a network site in Boston to another network's site in the same city may have to be exchanged in New Jersey. This network disadvantage is further exacerbated by the bad economics of public peering.

No single network provider has the incentive to invest in public peering points and yet they all benefit from them. A network provider investing in a high capacity line into a public peering point gets all the costs but only some of the benefits. This creates a disincentive to invest in higher capacities for public peering [Winkleman, 1998]. Further, because of a lack of pricing of traffic at public peering points, the network providers tend to overuse these to the point where the congestion becomes excessive [Barrett, 1998]. So we have the convergence of two negative economic phenomena at these public exchanges: lack of willingness to invest in capacity and overuse. Both of these make public peering unsatisfactory for most large networks.

Network providers get around the public peering or exchange bottlenecks by forming private interconnections that are governed by bilateral agreements. These agreements vary considerably [Bailey, 1997], [Srinagesh, 1997]. While some are on a quid pro quo basis, permitting free exchange of traffic, others specify a charge for net flow of traffic. Yet others lay down complex conditions that the provider with a smaller network must satisfy before the large network provider is even willing to consider a private connection agreement. Some examples of conditions follow:

- Fee and Quality of Service: PSInet requires quality of service and reliability assurance while charging a fee for transit traffic [Schrader, 1998].
- Geographic conditions: IBM (backbone now owned by AT&T) would only peer with other network providers in specific geographic regions (possibly to regions where it did not have substantial presence.) [IBM Global Services]
- Quality of service and comparable size: Fiber-Network Solutions, and many others, require the potential peer to be of comparable size and to maintain a full time network management center.
- Presence at multiple interconnection points: Worldcom's UUnet requires a potential peer to have a national network that will connect with it at four or more locations with superfast DS-3 lines [Yang, 1998].

In this paper we focus on private interconnection agreements between firms that provide Internet backbone services. We find considerable asymmetries in benefits that depend on size. When the size discrepancy is very large, i.e., the smaller network is very small compared to the larger one, the benefits largely accrue to the smaller network provider. This explains the size requirements for peering that were described above. Surprisingly, the distribution of gains from private interconnection switches as the networks become comparable sized. The somewhat larger network does better. This implies that these networks will tend to form "clubs" where they willingly peer with comparable size networks while shunning smaller ones.

Public peering points have gotten increasingly congested and the quality of service has fallen well below that of privately managed switch centers [Barrett, 1998]. This creates delay and other problems for smaller networks that have to rely on public peering points for reaching the global Internet resources. Larger networks are less dependent on public peering points for getting to Internet resources. Consequently, as the public peering points get more congested, the smaller networks gain far more than the larger networks. Surprisingly, the larger network providers prefers a small amount of peering point congestion as it keeps the traffic from the smaller network away from it. However, as the delay gets larger, the larger network provider also does worse.

This further exacerbates the size disadvantage of smaller networks and makes private interconnection agreements even more vital. This explains the changes in reported contracts.

Internet backbone consolidations make the study of private interconnection agreements important for social reasons. Four of the top six backbone network providers have undergone consolidation within the last year. Cable and Wireless acquired the MCI backbone, Bell Atlantic is acquiring GTE's backbone network, AT&T has bought all of IBM's global backbone assets, and Worldcom's UUnet acquired AOL's and CompuServe's backbone networks. These consolidations create a lopsided distribution of network sizes and more hurdles and conditions for peering by smaller networks. On the other hand, the larger networks also benefit users as they have more resources located on the home network. These competing effects demand a more detailed analysis of social implications of network consolidation.

In the next section we model the economics of interconnected networks. In section 3 we examine public peering. In section 4 we analyze private interconnection and the impact of differing network sizes and peering point congestion.

2 Modeling Interconnected Networks

To simplify the analysis, we consider two networks owned by ISP-1 and ISP-2, respectively, that are in geographically distinct, but possibly adjacent, areas with population sizes N_1, N_2 , respectively. Without loss of generality, we assume that $N_1 \geq N_2$. We consider two scenarios:

1. Public peering with no private interconnection. In this case the firms do not enter into a private agreement but instead use public peering points or exchanges to transmit cross traffic. This is the default connection that is available in the absence of any private agreement. In this situation, the customers experience less delay when they access a web page hosted on their own network than on the other network. Accessing a page on the other network requires the use of the public peering point that is a bottleneck and adds to the delay.
2. Private Network Interconnection Agreement. In this case the firms agree, possibly for a fee, to exchange traffic at more advantageous locations and at higher speeds. We assume that once the networks are interconnected, a customer sees no difference in delay in accessing a resource hosted by a site in its network or the other network.

In each scenario, the firms pick prices, p_1 and p_2 , respectively, to charge customers for ISP services. Given these prices, some customers get ISP services. Let m_1 be the number of customers, out of N_1 total, that choose to get Internet services from ISP-1 and m_2 be the corresponding number for ISP-2.

Customers get value from accessing web resources but at a decreasing marginal rate. We assume that this demand curve is linear: $a - b q$, where $a, b > 0$ and q is the number of web pages visited by the customer in a given period. Potential customers differ in their demand curves with a being uniformly distributed in $[0, A]$. This difference may arise from the different value that customers place on other Internet services such as e-mail, chat rooms, etc.

These web pages may be in the customer's network or on the other network. We assume that an equal proportion of customers in each network create web pages and that the probability of using any one is equal. The expected cost for a customer of ISP_{*i*} (ISP_{*j*} being the other ISP) of accessing a web page is taken to be:

$$c_i = \begin{cases} \frac{m_i}{m_i + m_j} c_L + \frac{m_j}{m_i + m_j} c_H, & \text{no private interconnection} \\ c_L, & \text{private interconnection} \end{cases} \quad (1)$$

where c_L and c_H are the costs to any user of accessing a web resource in his own and another network, respectively. Here, $m_i / (m_i + m_j)$ is the probability that a customer of ISP_{*i*} will access a web page on network *i* and $m_j / (m_i + m_j)$ is the probability that a customer of ISP_{*i*} will access a web page on network *j*.

Each customer is seeking to maximize his net surplus by equating the marginal value and cost, i.e., $a - bq = c_i$. For a customer of ISP_{*i*} with demand curve $a - bq$, the optimal number of web resources to use per period is:

$$q_i^*(a) = \frac{1}{b}(a - c_i) \quad (2)$$

The optimal expected net value to a customer of type *a* joining network *i* is

$$V_i^*(a) = \int_0^{q_i^*(a)} (a - bq - c_i) dq - p_i = \frac{1}{2}b(q_i^*(a))^2 - p_i \quad (3)$$

A summary of the notation introduced in this section is given next, following which we turn to the networks in each of the two scenarios.

Summary of Notation

Customers:

<i>a</i>	Customer type, uniformly distributed in $[0, A]$
<i>b</i>	Reduction in customer value from accessing a web page
<i>q</i>	Number of web pages visited by a customer in a given period
$q_i^*(a)$	Optimal number of web pages visited by customer of type <i>a</i> when joining network <i>i</i>
$V_i^*(a)$	Optimal expected net value to a customer of type <i>a</i> when joining network <i>i</i>
a_{im}	Type of marginal customer joining network <i>i</i>

ISPs:

N_i	Potential customer base of network <i>i</i>
m_i	Actual customer base of network <i>i</i> in a given period
p_i	Price set by network <i>i</i> in a given period
π_i	Profit to network <i>i</i> in a given period

3 Public Peering Only

In this scenario, the interconnection between the two networks occurs at the public peering points and the users experience greater delay in accessing pages and resources that are hosted on the other network.

Each ISP sets its price and only customers who get a non-negative net value, determined by equation (3), join the network. The marginal customer has a net value of zero. Setting V_i^* to zero and solving for a we get a description of the marginal customer:

$$a_{im} = c_i + \sqrt{2bp_i} = \frac{m_i c_L + m_j c_H}{m_i + m_j} + \sqrt{2bp_i} \quad (4)$$

Thus, customers in region i will get ISP services if their a is greater than a_{im} . Now, since m_i is taken to be the customer base of network i , it follows that:

$$m_i = N_i \Pr[a \geq a_{im}] = N_i \frac{A - a_{im}}{A} = \frac{N_i}{A} \left(A - \frac{m_i c_L + m_j c_H}{m_i + m_j} - \sqrt{2bp_i} \right) \quad (5)$$

Note that the equation for m_1 (and m_2) depends on both m_1 and m_2 . So we have two simultaneous quadratic equations in m_1 and m_2 . Solving these gives us the number of customers who join the two networks as a function of each other's prices.

We are now ready to determine the profit of the ISPs as a function of their prices and behavior of the customers that is described above.

The profit, before any fixed costs associated with interconnection, to ISP $_i$ is $p_i m_i^*$ where m_i^* is obtained by solving (5). Using equations (4) and (5), we get:

$$p_i = p_i \frac{N_i}{A} \left(A - \frac{m_i^* c_L + m_j^* c_H}{m_i^* + m_j^*} - \sqrt{2bp_i} \right) - (\text{Interconnection Fees and Costs}) \quad (6)$$

Since the number of customers that join a particular ISP depends on its and the other ISP's prices, the profits for each ISP also depend on each other's prices. Hence the two ISPs are in an economic game where each of them is trying to maximize its profit given the other's strategy. The prices that emerge from this strategic game are said to be in simultaneous Nash equilibrium. This is described in Figures 2 and 3.

Figure 2 describes the pricing strategy for ISP-1. The figure is drawn for $A = 1$, $b = 0.5$, $c_L = 0$, $c_H = 0.5$, and $N_2 / N_1 = 0.5$. The profit for ISP-1 is plotted against prices set by the ISPs. For any given price set by ISP-2, ISP-1 will set a price to stay on the ridge of the profit surface. In the figure this is delineated by a bold line. ISP-2 has a similar strategy choice. Each ISP wants to stay on its profit ridge. The combination of prices at which both profits are maximized is at the intersection of the two ridges. This is shown in Figure 3.

Now let us turn to the second scenario of private interconnection.

4 Private Network Interconnection

Private network interconnections provide high speed links between the two networks. The terms and fees for these are negotiated by the firms that own the networks. We assume that the private interconnection is

of a quality that all web sites, whether on the home network of a customer or not, have the same delay for the customer. Hence, the expected cost of accessing any web site is $c_i = c_L$. Substituting into (3) and solving as before we get the optimal price and profit, which are $2(A - c_L)^2 / 9b$ and $2(A - c_L)^3 N_i / 27Ab$ for ISP_{*i*}, respectively.

4.1 Value of Private Network Interconnection: Impact of network sizes

While each of the two firms gain from interconnection, the gains are not evenly distributed. In Figure 4 the gain from interconnection is plotted against the ratio of network sizes. Note that when one of the network is much smaller, it sees a greater increase in profit from a private interconnection agreement. This changes as the ratio gets larger.

The asymmetry in gains from interconnection arises from two competing effects:

1. Consider the customers of the two networks. To simplify the explanation, let c_L , the cost to a customer of accessing a page on his own network, be zero. The expected cost of accessing a web page is $m_L c_H / (m_L + m_S)$ for the smaller network and $m_S c_H / (m_L + m_S)$ for the larger network, where c_H is the delay cost of accessing a page through a public peering point, and m_L and m_S , the number of customers getting services in the larger and smaller network, respectively. Hence, the impact of a decrease in c_H to c_L from private interconnection is larger for the customers of the smaller network.
2. Customers of both the network providers benefit from reduced delay offered by a private network interconnection. There are more customers in the larger network and hence for a given change in delay costs, the impact on absolute profit of the larger network is larger.

The impact of these competing effects is illustrated by the crossing of the gain curves for the two network providers in Figure 4. This figure is plotted for parameters: $A = 10$, $b = 1$, $c_H = 6$, N_1 (size of the larger network) = 1,000,000, and $c_L = 0$. When the size of the smaller network is smaller than 220,000, the smaller network gains more from interconnection. This reverses when the size of the smaller network gets larger.

This asymmetry in gains from interconnection is reflected in the reports of interconnection agreements that were described in an earlier section. While most interconnection agreements require non-disclosure of terms, reports in trade press and web sites indicate that most network providers are only willing to peer with other network providers of comparable size. Smaller networks are “depeered” and treated as customers with connection and traffic fees [Dunlap, 1997], [Gerwig, 1998].

4.2 Impact of public peering point congestion on private interconnection

To assess the impact of increase in public peering point congestion, consider the customers of the two networks. As discussed in point 1 in the previous section, the impact of a change in c_H is greater on the customers of the smaller network. They experience a much greater increase in expected cost of accessing a web page from an increase in congestion. Even though the network owner compensates by lowering prices, the market coverage in the smaller network decreases. This is illustrated by the dashed curve in Figure 5 in which the market coverages (percentage of potential customers who get Internet services) for the two networks are plotted against public peering delay cost. As mentioned earlier, the price set by the smaller network firm decreases also, reflecting the increased cost for customers using web pages. This is illustrated by the dashed curve in Figure 6. Both of these figures are for the case with $A = 10$, $b = 1$, $N_2 / N_1 = 0.1$, $c_L = 0$. With a decrease in number of customers and in prices, the profit for the smaller network firm decreases rapidly with increasing public peering point congestion. So at least from the smaller firm’s

perspective, a private network interconnection that allows its customer to access resources on the larger network without having to go through the public peering point gets increasingly attractive with increase in public peering point congestion. This is reflected by the dashed curve in Figure 7 in which the gains (increase in profit) from private interconnection are plotted against public peering delay costs. Note that the gain to the smaller network increases monotonically with increasing congestion.

Interestingly, the impact of increasing public peering congestion on the larger network is quite different from that on the smaller network. Firstly, the impact is not that large since only a small fraction of pages are on a foreign network for customers of the larger network. Secondly, the effect is non-monotonic. This arises from two competing reasons:

1. An increase in congestion costs, c_H , increases the expected cost of accessing a web page for the customers.
2. Increases in congestion cost, c_H , is accompanied by a decrease in number of customers in the other, smaller, network. This was discussed above. A decrease in number of non-home network customers reduces the probability that the customer has to access foreign pages through the public peering point. This reduces the expected cost of accessing a web page.

The non-monotonic but smaller impact of congestion costs on customers of the larger network is reflected in Figures 5, 6 and 7. The market coverage and prices are plotted by solid curves in Figures 5 and 6, respectively.

The gain from private interconnection for the firm with the larger network differs markedly from that for the smaller firm. While the gain for the smaller firm is monotonically increasing, the gain for the larger firm is non-monotonic and goes to zero as the congestion costs get very high. The latter requires some explanation. The smaller network has no customers when the congestion costs are too high and consequently there are no resources on the smaller network that a customer from the larger network may want. This implies that the expected cost of accessing a web page is the lowest possible, c_L . This is also the case with a private network interconnection. Hence the customers of the larger network are indifferent between these two states and the large network firm makes the same profit. Hence the gain from interconnection when the congestion cost is too high is zero for the larger firm. The gain is also zero for $c_H = c_L$. In between, the larger network benefits from interconnection. These effects are illustrated by inverted U-shape of the gain curve for the large network which is plotted as a solid curve in Figure 7.

5 Conclusions

Private interconnection between network providers considerably changes the economics of Internet services. A customer placing web resources increases the value for other customers on its network and may reduce the value to customers in other networks as they may experience greater delay in getting resources across the interconnection which typically is the bottleneck. Consequently, the profits of firms that provide Internet services also depend on the interconnections. In this paper we have shown that distribution of gains from interconnection depends on a number of factors that include relative sizes, service levels, and congestion at the public peering points.

Smaller firms may gain more from an interconnection agreement than larger firms. Consequently, larger firms would rather treat the smaller network providers as customers, charging them for connectivity and traffic, than treat them as peers. This disparity gets larger as network providers consolidate and public peering points get more congested. This raises important social policy issues in managing the Internet infrastructure in the next century.

6 References

1. Bailey, J.P., "The economics of Internet interconnection agreements" in *Internet Economics* by McKnight, L.W. and Bailey, J.P., Cambridge, MIT Press, 1997.
2. Barrett, R., "Peering into the future," *Interactive Week*, Dec. 7, 1998.
3. Dunlap, C., "UUnet defends peer pricing," *Computer Reseller News*, May 2, 1997.
4. Gerwig, K., "Peering standoff must come to an end," *InternetWeek*, September 07, 1998a.
5. IBM Global Services, "Peering with IBM Global Services," www.ibm.com/globalnetwork/peering.htm
6. Leiner, B.M., Cerf, V.G., et. al, "A brief history of the Internet," www.isoc.org, 1998.
7. Schrader, W., "Peering into the future," Presentation at ISPCON, slides are available at www.psinet.com, 1998.
8. Srinagesh, P., "Internet cost structures and interconnection agreements" in *Internet Economics* by McKnight, L.W. and Bailey, J.P., Cambridge, MIT Press, 1997.
9. Winkleman, J., "Getting connected," *America's Network*, Aug. 15, 1998.
10. Yang, C., "How the Internet works: All you need to know," *BusinessWeek*, July 1998.

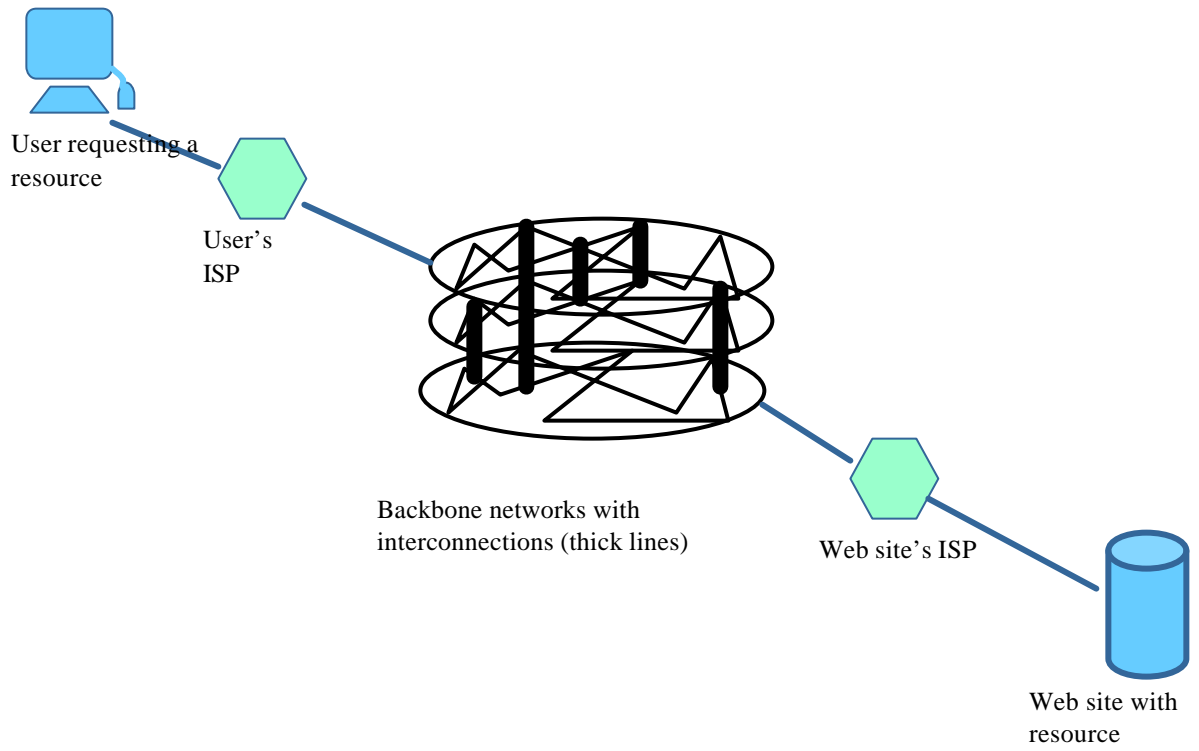


Figure 1: The Internet

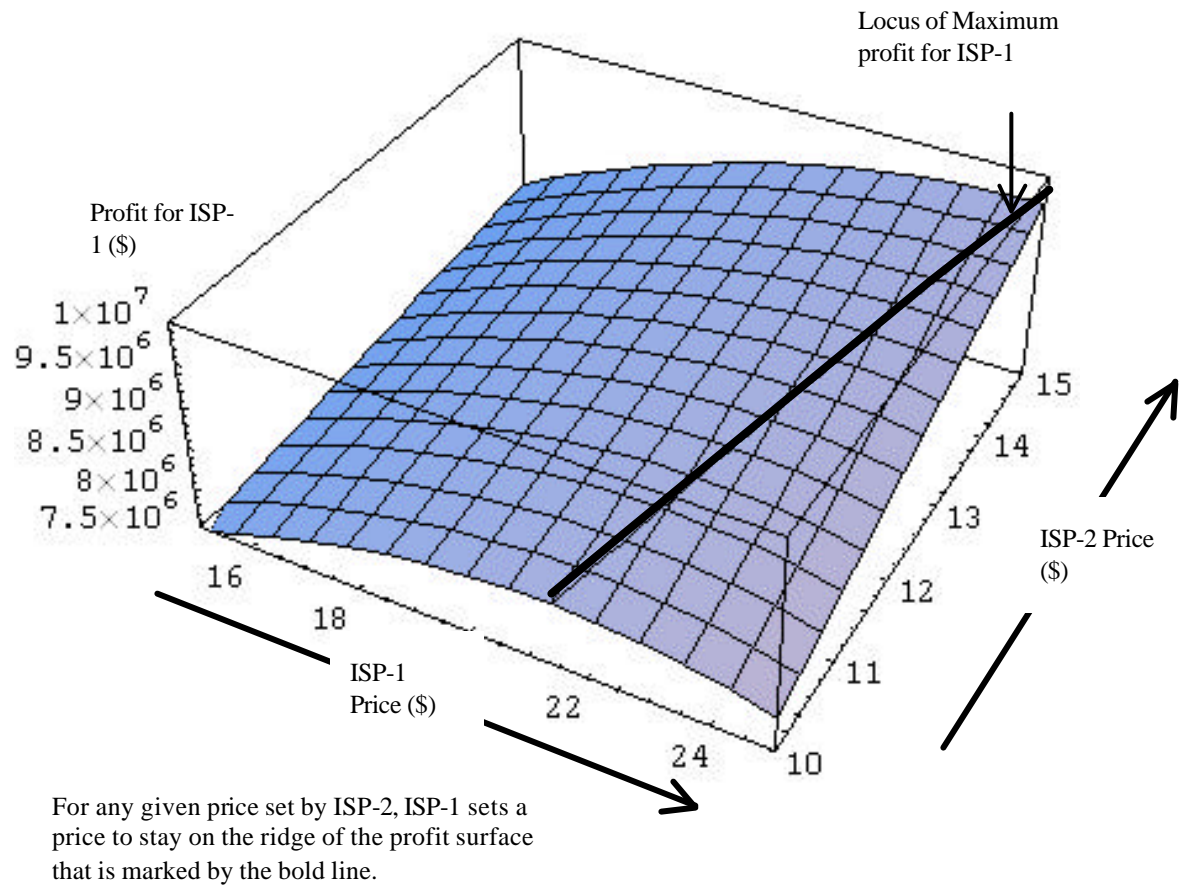


Figure 2: Strategy for ISP-1

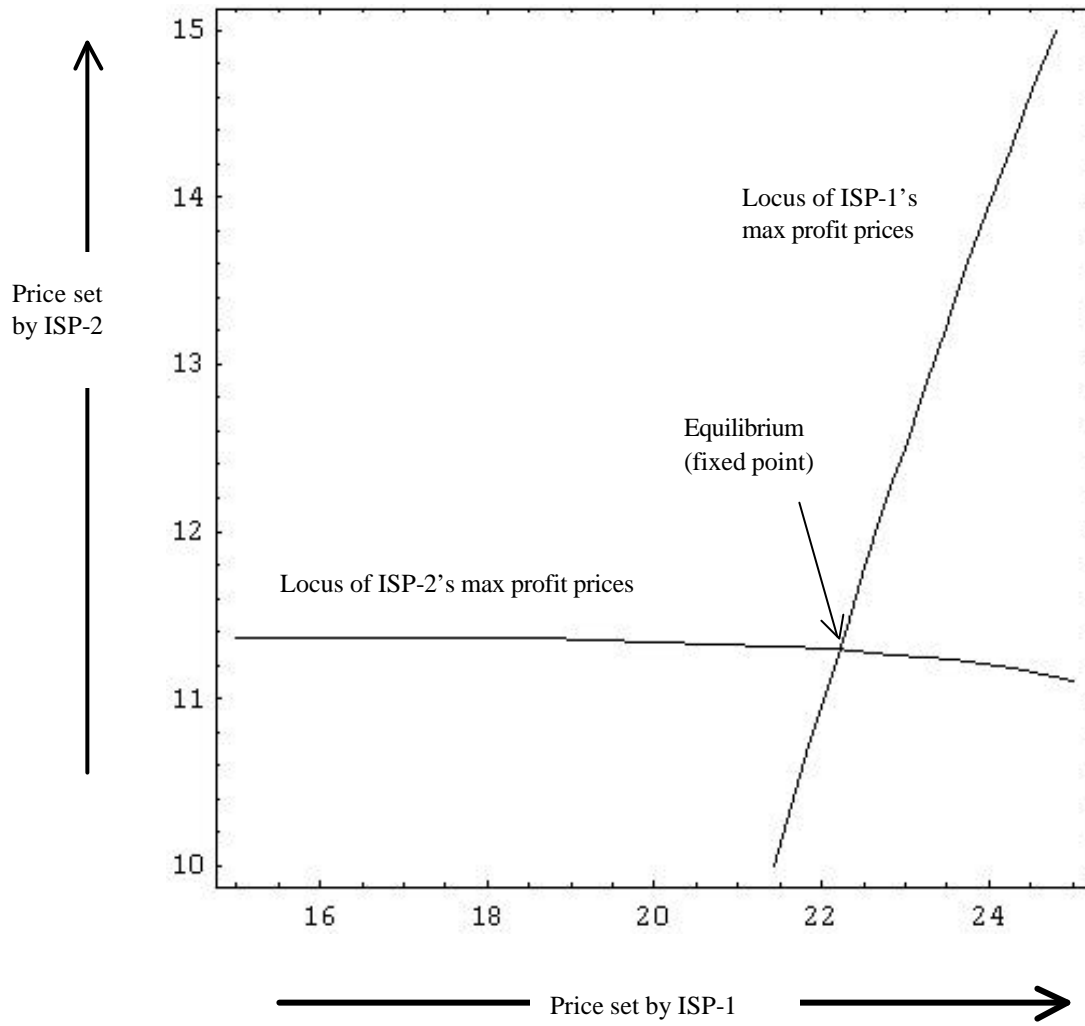


Figure 3: Equilibrium between ISPs

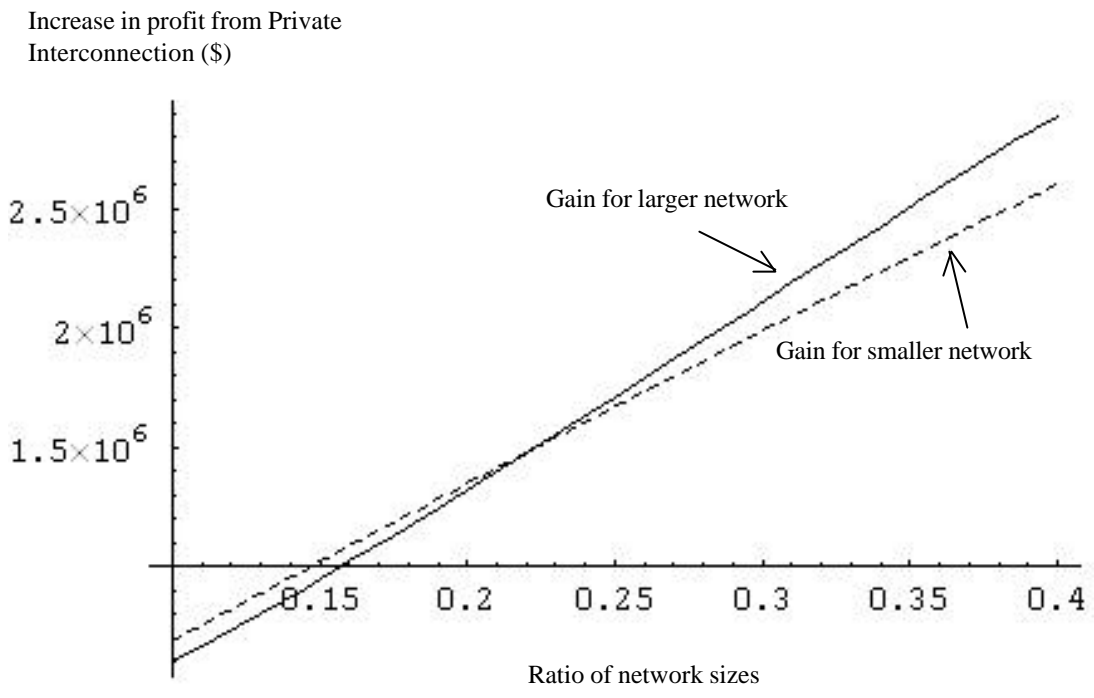


Figure 4: Increase in profits from Private Interconnection agreements Vs. Relative network sizes

Market
Coverage (%)

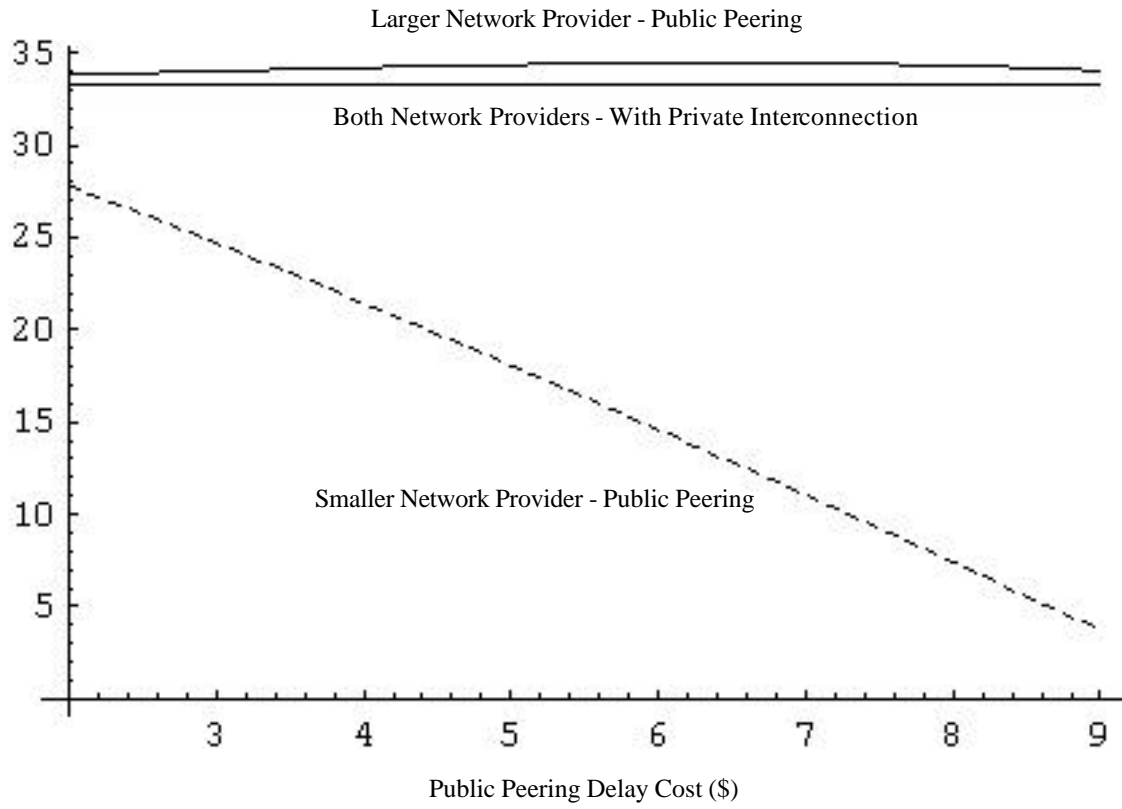


Figure 5: Impact of Peering Point Congestion on Market Coverage

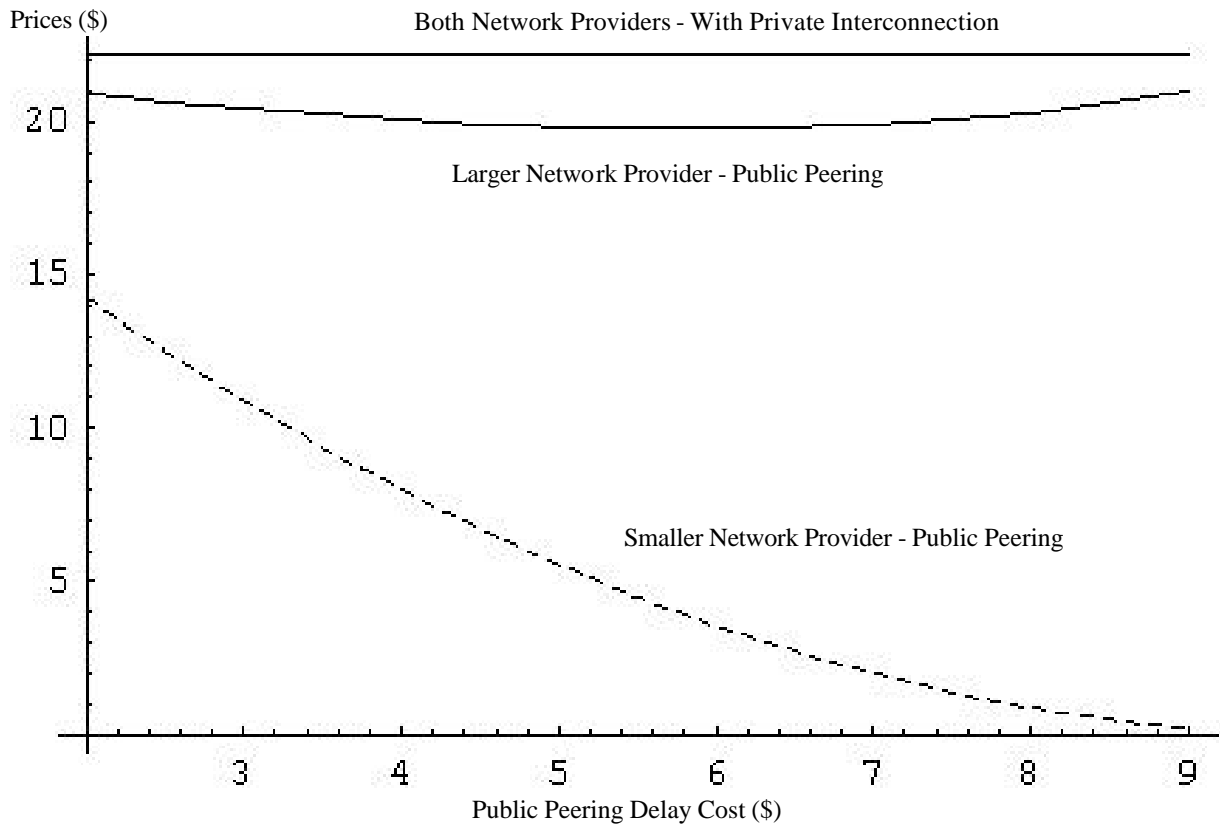


Figure 6: Impact of Public Peering Congestion on ISP Prices

Increase in profit from
Private Interconnection (\$)

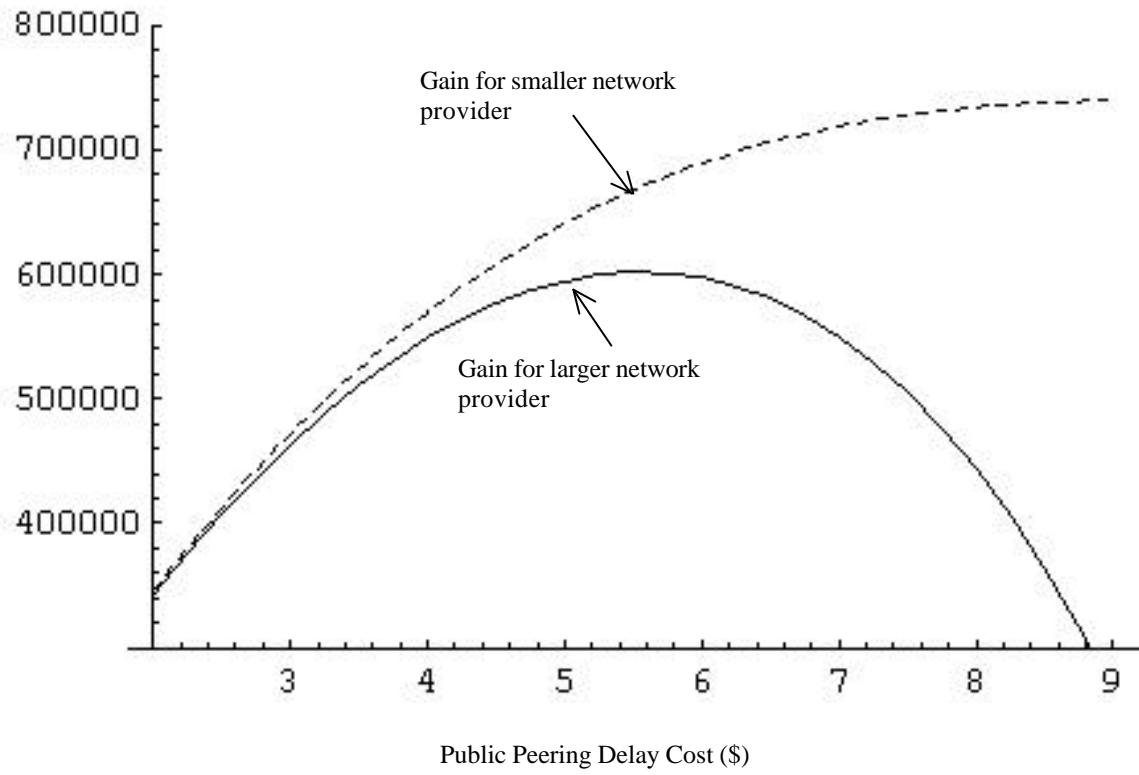


Figure 7: Impact of Public Peering Point Congestion