

Advances in Routing Technologies and Internet Peering Agreements

By STANLEY BESEN, PAUL MILGROM, BRIDGER MITCHELL, AND PADMANABHAN SRINAGESH*

Spurred by improvements in routing technology, the architecture of the Internet is evolving. Until recently, nearly all routing took place through parallel hierarchies, in which each *core* Internet service provider (ISP) at the top of its own hierarchy provided other core ISP's with routes to its own customers and customers of non-core ISP's in its hierarchy. Based partly on the ability and incentives of some core ISP's to deny or degrade service to others, antitrust authorities required the divestiture of internetMCI as a condition for the MCI-WorldCom merger and blocked the proposed merger of Sprint and MCI-WorldCom.

Recent changes in routing standards have enabled a wider range of routing arrangements, and these evidently reduce the market power of the core ISP's vis-à-vis their customers. We argue here that those new standards also reduce the incentives of core ISP's with large market shares to refuse or degrade service to ones with smaller market shares. Our analysis is based on a bargaining model, which provides a means to assess how the short-run bargaining positions of various core ISP's are affected by the new routing arrangements.

I. Background

Traditionally, there have been two common types of interconnection arrangements among Internet service providers. In a *transit arrangement*, the ISP selling transit services agrees to deliver traffic from its transit customers to any Internet destination and to deliver traffic from any Internet destination to its transit customers. In a *peering arrangement* between two ISP's, each ISP agrees to deliver to the other ISP only traffic that is destined to the other provider's

customers. Typically, peering arrangements do not require payments from one ISP to another: they are "bill-and-keep" arrangements. Greater detail can be found in Milgrom et al. (2000).

In the traditional "rigid" Internet hierarchy, a few core ISP's peer with one another to produce full routing capability among all Internet end users. Each non-core ISP purchases transit from a single core ISP; non-core ISP's are not directly connected to one another. End users purchase a single connection to a core or non-core ISP. Core ISP's compete in an upstream market to supply top-level backbone services; this upstream market is distinct from the downstream market for Internet access. Core and non-core ISP's compete in the downstream market for Internet access.

Recently, a new routing standard (BGP4) that supports alternative, less hierarchical, routing arrangements has become drastically cheaper (see Avi Freedman, 1999). Taking advantage of the cost reductions, non-core ISP's have increasingly entered into *secondary peering arrangements* in which the participating networks directly exchange traffic destined to each other's customers on a bill-and-keep basis, bypassing the core ISP's. Also, non-core ISP's and corporate end users have increasingly resorted to *multi-homing*—purchasing connections from multiple providers and routing traffic among them in real time. At the same time and with similar effect, publishers of web-based content are increasingly relying on intelligent caching strategies and on *content distribution services* (CDS) that replicate web pages at locations close to end users in order to deliver information in the most effective way.

II. A Simple Bargaining Model of Interconnection

We begin our analysis with a bargaining model of the traditional peering hierarchy. In this model, bill-and-keep arrangements are sometimes consistent with a bargaining equi-

* Besen: Charles River Associates, 1201 F Street, N.W., Suite 700, Washington, DC 20004-1204; Milgrom: Economics Department, Stanford University, Stanford, CA 94305 (e-mail: paul@milgrom.net); Mitchell and Srinagesh: Charles River Associates, 285 Hamilton Avenue, Suite 370, Palo Alto, CA 94301.

librium even when the larger network provider threatens to refuse interconnecting with its smaller peer. Thus, without extra conditions, a simple threat to refuse interconnection during a period of disagreement may not be sufficient for the effective exercise of market power.

Assume that there are N homogeneous customers in the market, served by n core ISP's. Each customer obtains service from only one ISP, and core ISP $_i$ serves a fraction α_i of the customers. When ISP $_i$ is not connected to any other core ISP, its representative customer enjoys a benefit or utility of $u(\alpha_i, N)$ per period and is willing to pay a corresponding amount for that connectivity. The presence of network externalities means that u is increasing in both α and N . Since we will be holding N fixed throughout this analysis, we use a less cumbersome notation by writing $f(\alpha) = u(\alpha, N)$ and conducting the analysis in terms of f .

Suppose that one core ISP initially serves a fraction α_1 of the customers and a second serves a fraction α_2 , and that these proportions are independent of the interconnection arrangements between the two ISP's. (This assumption implies that no ISP suffers a permanent loss of customers when it temporarily suffers a degradation of service.) Suppose further that both ISP's have obtained peering arrangements with all the other ISP's. The revenues of ISP $_1$ would be $N\alpha_1f(1 - \alpha_2)$ if it did not obtain a peering arrangement with ISP $_2$, and $N\alpha_1f(1)$ if it did obtain a peering arrangement. We assume for simplicity that there are no costs, so that revenues are equal to profits.

Assume that the lack of interconnection is sustained only temporarily during bargaining, until the parties reach a peering agreement. The outcome of negotiations according to the non-cooperative bargaining theory with short times between offers is approximately the same as that of the Nash bargaining model, provided the payoff each earns during the period of disrupted connection is treated as the Nash threat point (Kenneth Binmore et al., 1986). The total surplus to be divided in any agreement is $(N\alpha_1 + N\alpha_2)f(1)$, while the payoff pair during service disruption is $(N\alpha_1f(1 - \alpha_2), N\alpha_2f(1 - \alpha_1))$. At a noncooperative bargaining outcome, the two parties divide equally any gains relative to the threat point. The resulting bargaining payoff for ISP $_1$ is

$$(1) \quad \pi_1 = \frac{1}{2} [(N\alpha_1 + N\alpha_2)f(1) + N\alpha_1f(1 - \alpha_2) - N\alpha_2f(1 - \alpha_1)]$$

and that for ISP $_2$ is

$$(2) \quad \pi_2 = \frac{1}{2} [(N\alpha_1 + N\alpha_2)f(1) - N\alpha_1f(1 - \alpha_2) + N\alpha_2f(1 - \alpha_1)].$$

With full interconnection and a bill-and-keep arrangement, ISP $_1$ would be able to charge each of its customers a subscription fee of $f(1)$ and earn revenues (and profits) of $\rho_1 = N\alpha_1f(1)$. If the bargaining payoff π_1 is larger than ρ_1 , then the excess $\pi_1 - \rho_1$ may be interpreted as the negotiated net payment from ISP $_2$ to ISP $_1$. Since there are no costs in the formal model, such payments are not compensation for the costs imposed by one ISP on the other. Hence, a positive net payment represents a simple exercise of bargaining advantage due to market share. Doing the arithmetic, we find that

$$(3) \quad \pi_1 - \rho_1 = \frac{1}{2} N\alpha_1[f(1) - f(1 - \alpha_2)] - \frac{1}{2} N\alpha_2[f(1) - f(1 - \alpha_1)].$$

The first term is half the additional revenue that ISP $_1$ earns from its end users after it negotiates an interconnection arrangement with ISP $_2$, and the second is half the additional revenue for ISP $_2$; the bargaining solution splits the gains from agreement. Thus, when both ISP's gain equally from interconnection, neither party pays the other, and a bill-and-keep arrangement is the equilibrium outcome of the bargaining process. When the parties do not gain equally from interconnection, the ISP that gains more pays the other for interconnection.

Sufficient conditions for bill-and-keep interconnection arrangements are easily obtained. If either (i) the two ISP's are equally large ($\alpha_1 =$

α_2) or (ii) f is linear [$f(\alpha) = a + b\alpha$], then bill-and-keep is the outcome: $\pi_i = \rho_i$ for $i = 1, 2$. Intuitively, the second condition reflects the fact that, while the smaller ISP loses more per customer during a service interruption, the larger ISP suffers the loss over a larger customer base so, when f is linear, the total losses are equal.

Define $h(\alpha) \equiv [f(1) - f(1 - \alpha)]/\alpha$. Inspecting (3), one sees that ISP₁ receives a positive transfer if and only if, $h(\alpha_1) < h(\alpha_2)$. If h is decreasing, then the larger ISP receives a positive net transfer on account of its size; if h is increasing, then the predicted transfer is negative; and if h is constant (corresponding to f linear), then the transfer is zero and bill-and-keep is the prediction of the model. This last case, with linear f and constant h , corresponds to the frequently studied case of *isotropic* networks in which users connect to each other with equal probability, and each connection creates equal value.

III. Enriching the Account of Threats during Bargaining

The preceding conclusion is derived from a model of bargaining that is simplified in various respects. First, it assumes that, during any period of disagreement about transfers, the parties will suffer a service disruption. In practice, bargaining begins before current arrangements expire, and once current contracts expire bargaining over transfers might take place even while services continue to be delivered under old peering contracts. Endogenizing the choice to “get tough” by disrupting services leads to a much wider range of possibilities (Raquel Fernandez and Jacob Glaser, 1991), with roles for expectations, reputation-building, inefficient disagreements, and similar phenomena. In such models, bargaining theory does not lead to a unique prediction about the outcome.

Second, in the basic bargaining model, service disruptions are short enough that they do not induce the customer to switch ISP’s. While both ISP’s involved might be expected to lose customers during an extended degradation of service, if customer switching decisions are based on service quality thresholds, the smaller ISP is likely to suffer the larger loss.

In what follows, we modify our description of bargaining by making a simple specification

in which the disagreement payoff involves the permanent loss of some customers, rather than the loss of value to current customers that the ISP has to compensate. Specifically, suppose that for some $\gamma > 0$, ISP _{j} suffers losses due to customer switching at a total rate of $N\alpha_j k\alpha_j^\gamma$, for $i, j = 1, 2$. This loss rate is equal to the number $N\alpha_j$ of customers of ISP _{j} multiplied by a function of the degradation of service suffered during the period of disagreement. The parameter γ is interpreted as the elasticity of switching rates with respect to service degradation, while k parameterizes the speed of customer response to degradations in service. The model neglects second-order effects, in which customer switching itself affects the service quality during the period of degradation.

This specification maps neatly into the model of the previous section, provided that we set $f(\sigma) = 1 - k(1 - \sigma)^\gamma$. In this expression, we interpret $\sigma_i = 1 - \alpha_j$ as the index of service quality for the customers of ISP _{i} during disagreement. Using our earlier definition, $h(\alpha) = k\alpha^{\gamma-1}$. For $\gamma < 1$, this is decreasing, and hence the larger ISP is able to extract a positive payment from the smaller ISP in the bargaining. For $\gamma > 1$, the reverse advantage obtains, and $\gamma = 1$ is the linear case in which bill-and-keep obtains. We henceforth assume that $\gamma < 1$.

Suppose that ISP₁ has the larger market share, that is, $\alpha_1 > \alpha_2$, so that it is ISP₁ that can bargain for a positive transfer according to our model. In the developments below, we incorporate the effects of the multi-homing and secondary peering technologies into the service quality index, σ , while maintaining the assumption that net switching away from ISP _{i} during disagreement occurs at rate $N\alpha_i[f(1) - f(\sigma_i)]$. This allows us to assess the effects of changing technology on the bargaining outcomes among the core ISP’s.

IV. Competitive Effects of Secondary Peering and Multi-homing

Suppose some ISP’s that purchase transit only from core ISP₁ enter into secondary peering arrangements with some ISP’s that purchase transit only from core ISP₂. Assume that for each core ISP a proportion β of its subscribers can communicate over the secondary-peering interfaces. If the primary-peering interface between ISP₁ and ISP₂ is degraded, the fraction

$(1 - \beta)$ of customers of ISP_1 that are not connected to ISP 's with secondary peering arrangements will obtain high quality only on the fraction of their traffic destined for customers not served by ISP_2 (i.e., $1 - \alpha_2$). The remaining fraction of customers of ISP_1 , β , will also obtain high quality on the fraction of their traffic covered by the secondary peering arrangement (i.e., $\beta\alpha_2$). The average service quality for ISP_1 across all subscribers is $\sigma_1 = (1 - \beta)(1 - \alpha_2) + \beta(1 - \alpha_2 + \beta\alpha_2) = (1 - \alpha_2 + \beta^2\alpha_2)$. Similarly, customers of ISP_2 will on average obtain high quality on a fraction $\sigma_2 = (1 - \alpha_1 + \beta^2\alpha_1)$ of their traffic.

Notice that, if $\alpha_1 > \alpha_2$, increases in β have a larger effect in improving service quality for customers of ISP_2 than for those of ISP_1 ; that is, $\partial\sigma_2/\partial\beta = 2\alpha_1\beta > 2\alpha_2\beta = \partial\sigma_1/\partial\beta$. To see how this translates into an improved bargaining position for ISP_2 , we update formula (3). Conceptually, this involves two steps. First, since service quality is no longer represented by $1 - \alpha_i$, we replace each $f(1 - \alpha_i)$ term by $f(\sigma_j(\beta))$ to allow a more complex representation of service quality that depends on the network technology. Second, we use the formula $f(\sigma) = 1 - k(1 - \sigma)^\gamma$. The results are these expressions:

$$(4) \quad \pi_1 - \rho_1 = \frac{1}{2} Nk(\alpha_1[1 - \sigma_1(\beta)]^\gamma - \alpha_2[1 - \sigma_2(\beta)]^\gamma)$$

$$(5) \quad \frac{\partial}{\partial\beta} (\pi_1 - \rho_1) = -Nk\alpha_1\alpha_2\beta\gamma(1 - \sigma_1)^{\gamma-1} \times \left[1 - \left(\frac{1 - \sigma_2}{1 - \sigma_1} \right)^{\gamma-1} \right] < 0.$$

The calculation demonstrates that the bargaining advantage of the larger ISP is reduced by the introduction of secondary peering.

Increases in the extent of secondary peering are similar to reductions in the market share of a dominant core ISP. To implement this idea, let us assume that the customers lost during a service disruption between two ISP 's flow to the other core ISP 's. In that case, the impact of secondary peering on the competitiveness of top-level backbones can be represented by cal-

culating an *equivalent market share* (EMS) for ISP_1 , defined implicitly to be the market share that would provide the same relative service quality and hence bargaining position in the old "rigid hierarchy" without secondary peering or multi-homing: $(1 - \alpha_2)/(1 - EMS_1) = (1 - \alpha_2 - \beta^2\alpha_2)/(1 - \alpha_1 - \beta^2\alpha_1)$. Since $dEMS/d\beta$ is negative, increases in β are equivalent to a loss of market share, reducing the market power of the larger backbone.

The analysis of multi-homing is similar to that of secondary peering. Traffic that crosses a peering interface between two core ISP 's in a rigidly hierarchical Internet is able to take an alternative path from the origin to the destination after the secondary peering or multi-homing arrangement is implemented. The diverted traffic cannot be degraded by changes in the peering arrangement between the two core ISP 's and is therefore of high quality. Both networks experience high quality on the same volume of diverted traffic. This fixed traffic volume is a larger proportion of the smaller backbone's total traffic, leading to a greater proportionate increase for the smaller backbone in the proportion of traffic that is of high quality, and reducing the larger backbone's relative service quality σ_1/σ_2 which, as seen in (5), reduces its ability to extract bargaining concessions from its smaller partner. Increases in multi-homing, like increases in the extent of secondary peering, are equivalent for this analysis to a loss of market share for the larger backbone provider.

Specifically, if a proportion λ of customers of ISP_1 and ISP_2 are multi-homed, then $\lambda\alpha_1$ customers of ISP_1 can communicate with *all* customers of ISP_2 without sending traffic over a peering interface. Suppose that ISP_1 degrades its peering interface with ISP_2 . Multi-homed subscribers will continue to obtain high quality on all their traffic. Single-homed subscribers of ISP_1 will obtain high quality on traffic that is not bound for ISP_2 , and on all traffic to multi-homed customers of ISP_2 . The average quality of service obtained by all customers of ISP_1 is therefore $\lambda + (1 - \lambda)(1 - \alpha_2 + \lambda\alpha_2) = 1 - \alpha_2 - \lambda^2\alpha_2$. Similarly, customers of ISP_2 can obtain an average service quality of $1 - \alpha_1 - \lambda^2\alpha_1$. As with secondary peering, it can be shown that $\partial EMS/\partial\lambda$ is negative, implying that the incentives of a large core ISP to deny or degrade its peering relationships are reduced by

increases in multi-homing. The analysis of intelligent caching and CDS is similar.

V. Conclusions

The Internet's hierarchical structure has evolved from parallel hierarchies to a looser arrangement, in which messages can be exchanged among smaller ISP's without using the backbones of a core ISP. For the simple structures analyzed in this paper, indexes of "service quality" and "equivalent market share" can be used to quantify the effect of changing hierarchical structures on the incentives of larger backbone providers to degrade peering arrangements. The analysis suggests that, as the Internet develops a richer set of interconnection arrangements (such as secondary peering and multi-homing), the incentives of large backbone providers to refuse to enter into or degrade peering arrangements will be reduced.

REFERENCES

- Binmore, Kenneth; Rubinstein, Ariel and Wolinsky, Asher.** "The Nash Bargaining Solution in Economic Modeling." *Rand Journal of Economics*, Summer 1986, 17(2), pp. 176–88.
- Fernandez, Raquel and Glazer, Jacob.** "Striking for a Bargain between Two Completely Informed Agents." *American Economic Review*, March 1991, 81(1), pp. 240–52.
- Freedman, Avi.** "AVI RETURNS—Configuring CAR and CEF To Shape Traffic and Kill Smurfs." *Boardwatch Magazine*, July 1999, 13(7), pp. 96–100.
- Milgrom, Paul R.; Mitchell, Bridger M. and Srinagesh, Padmanabhan.** "Competitive Effects of Internet Peering Policies," in Benjamin Compaine and Ingo Vogelsang, eds., *The Internet upheaval*. Cambridge, MA: MIT Press, 2000, pp. 175–95.