

Loci of Competition for Future Internet Architectures

John Chuang
UC Berkeley
November 2010

ABSTRACT

Designing for competition is an important consideration for the design of future Internet architectures. Network architects should systematically consider the loci of competition in any proposed network architecture. To be economically sustainable, network architectures should encourage competition within each locus, anticipate and manage the interactions between the loci, and be adaptable to evolution in the loci. Given the longevity of network architectures relative to network technologies and applications, it is important to ensure that competition is not unnecessarily foreclosed at any particular locus of competition.

1. DESIGN FOR COMPETITION

In contemplating future Internet architectures, the networking community has identified economic viability as a key architectural requirement, along with other requirements such as security, scalability, and manageability. This is in recognition of the fact that any global-scale distributed communications infrastructure requires significant capital investments, and incentives must exist for the network owners to invest in new facilities and services in a sustainable fashion.

It is widely accepted that competition is important for promoting the long-term economic viability of the network. This is because competition imposes market discipline on the network operators and service providers, providing them with incentives for continual innovation and investment in their facilities. Conversely, given a lack of competition, a monopoly provider may become complacent and fail to invest in the long-term health of the network. However, competition does not occur automatically in a network. Therefore, “design for competition” must be an important design principle for any future Internet architecture.

“Designing for competition” can also promote the security of a network by encouraging diversity in all levels of the architecture. Robust competition pre-empts the vulnerabilities of a monoculture [1], including the increased likelihood of correlated failures or cascade failures. Diversity should be sought not just in hardware and software

levels, but also in the operators of networks, the providers of network services, and the policies for network management and control.

As a design principle, “design for competition” is similar to but different from the “design for choice” principle as articulated by Clark et al. in their influential tussles paper [2]. By designing for choice, Clark et al. refer to the ability of the architecture “to permit different players to express their preferences”. The “design for competition” principle extends this to the ability of the architecture to permit different players to express their preferences for services *by different providers*. For example, in thinking about architectural support for user-directed routing, the “design for choice” principle might suggest that route diversity be made available to end users by *a provider* over its network. The “design for competition” principle, on the other hand, might lead to a stronger requirement that the end users can choose routes offered by *multiple providers* and/or over multiple networks.

It should be clear that “design for competition” does not imply the mandating of competition. Instead, it is an argument against the unnecessary foreclosure of competition. Design for competition means the architecture should be designed to allow for the possibility of supporting multiple providers of a given service, even if the option of having multiple actual providers is not exercised from the start. This is because a network architecture usually outlives the network technologies and network applications. Sometimes, competition may not be feasible or desirable at a particular locus of the architecture. For example, the current generation of network technologies may be subject to strong economies of scale, such that the cost inefficiencies due to competition dominates the benefits obtained from market competition. However, the introduction of new technologies or the adoption of new applications may lead to a re-evaluation of the merits of actual competition at that locus at a future point in time. Put another way, a network architecture that is designed for competition would allow innovation, in the form of new entrants and/or new services, to be easily introduced into the network, to compete with existing offerings in the network.

2. LOCI OF COMPETITION

Designing for competition in a network architecture is not a straight-forward exercise. In facilitating a competitive networking environment, one needs to systematically think through the *what, where, who* (by whom and for whom), *when*, and *how* of supporting competition.

Network architects can begin by identifying the potential loci of competition in the architecture, i.e., candidate locations in a network architecture where a multiplicity of service providers can be supported. Based on the characteristics of prevailing technologies and uses, the architects can also gauge the level of competition that may arise organically for each locus. For example, loci subject to high fixed costs and low marginal costs (i.e., strong economies of scale), or to strong network effects, or to high switching costs, should be considered *competition-challenged*, i.e., they tend to face greater obstacles to competition.

In the event that a design calls for the preclusion of choice or competition at a particular locus, this decision needs to be clearly justified. In effect, a monopoly is being hard-coded into the architecture.

Recognizing that “loci are not silos,” network architects must also anticipate and understand the economic relationships that may exist between the loci of competition. For example, if a firm operates in two vertically related loci (i.e., vertical integration), its market power in one locus may influence the level of competition in the other locus. As another example, two adjacent loci may both be dominated by a small number of firms (i.e., bilateral oligopolies), leading to strategic behavior that may impede innovation on both sides. The network architects must define and develop interfaces to manage and lubricate both the technical and market transactions between the loci.

2.1 Loci Past, Present, and Future

In the early days of telephony in the United States, with AT&T controlling the long-lines, the local loops, and even the customer-premise equipment (CPE) leased to customers, there was effectively only a single locus of competition. A potential competitor would have had to construct an alternate nation-wide network capable for providing end-to-end service in order to compete effectively against AT&T. In reality, strong economies of scale meant it was economically infeasible for such an alternate network to be deployed. Hence, AT&T remained a dominant monopoly, and had to be held in check through regulation for many decades. However, starting in the 1960s and 1970s, supply-side and demand-side changes, in the form of improvements in communication technologies and increasing business demands for telephony services, set the stage for multiple loci of competition to emerge in the telephony network architecture. In particular, the divestiture of AT&T in 1984 made concrete the existence of three separate (but vertically related) loci: local telephony, long distance telephony, and customer premise equipment. Through the explicit intervention by the executive and judicial branches of the government,

firms can now compete in each locus individually, and the lowering of entry barriers led to significant innovations in each locus. Even so, the loci remain vertically related, and under the 1996 Telecommunications Act, firms that are dominant in one locus can compete in other loci. Therefore, government oversight remains necessary to this day.

For the current Internet architecture, multiple loci of competition can also be identified. In the topological dimension, the architecture includes long distance (backbone) networks, local access networks, and end hosts. In addition, the layering principle upon which the Internet is built means that separate loci of competition can also be identified for the physical, network, and application layers of the Internet. Further evolution of the architecture in support of video distribution, cloud computing, and other emerging applications may lead to additional loci of competition at datacenters, content distribution networks (CDNs), authentication authorities, and so on.

In contemplating future Internet architectures, it is not unreasonable to expect some loci of competition to look just like those of the current architecture, and some to be entirely different. In addition to devices, conduits, services, and applications, even the network architectures themselves can be potential loci of competition. The case has been made, on multiple occasions, that the incumbency of the current IPv4 network architecture has impeded innovation in the network itself [3]. Through network virtualization techniques, it is possible to allow multiple network architectures to be concurrently deployed over a shared physical infrastructure, to compete with one another. Then, end users can vote with their feet (or with their data packets) for their favorite, be it IPv4, IPv6, IPvN, or some other network architecture or protocol suite.

2.2 A 2x2 Example

As an exercise to identify loci of competition in a network architecture, let us consider a simplified organization of a network along two dimensions: topological and layering. In the topological dimension, we can organize the network into the Edge (E) and the Core (C). The Edge provides access to end hosts, and is commonly referred to as “the last mile” or “the local loop.” The Core provides connectivity between the edge networks over long distances, and is commonly referred to as “the backbone”. In the layering dimension, we can view the network as consisting of a Physical layer (P) that carries bits over a physical infrastructure, and a Logical layer (L) that sits on top of the Physical layer and provides network services to end users.

This organization of the network can be easily extended to include additional dimensions, as well as more than two tiers within the topological and layering dimensions. However, with this simple 2x2 matrix (Table 1), we can already articulate four different loci of competition.

	Edge (E)	Core (C)
Logical (L)	Number of access service providers	Number and coverage of transit service providers
Physical (P)	Number of wired and wireless “last-mile” conduits to each end-host	Number and coverage of wide-area physical conduits

Table 1. Loci of Competition in a Simple 2x2 Organization of a Network

In the Physical-Edge (P-E) cell of the matrix, competition is manifested in the form of multiple physical connections (wired or wireless) to the end hosts. Telecommunications regulators use the term “facilities-based competition” to describe competition in this locus. The build-out of new facilities involves high up-front costs for trenching, erection of towers, and/or acquisition of radio spectrum. Therefore the barriers to entry are high, and the number of parallel facilities is usually limited. However, once the facilities are built, the marginal cost of delivering data is very low, and so this part of the industry is subject to very strong economies of scale.

In the Logical-Edge (L-E) cell of the matrix, multiple firms may compete in providing network access services to end users. The providers may offer services over different physical networks, or over the same physical network. The providers may offer similar services (e.g., IPv4) and compete on quality and/or cost, or may offer different services for different segments of the market (e.g., VPN for enterprises). It is not uncommon to find a single firm operating in both the P-E and L-E cells of the matrix, i.e., a firm offers network service to end users over a physical infrastructure that it owns and operates.

In the Physical-Core (P-C) cell of the matrix, we can have multiple firms who deploy physical infrastructure for wide-area bit transport, such as fiber, submarine cables, terrestrial radio, and satellite links. To the extent that there is overlap in geographic coverage, the firms may compete for business with one another.

Finally, in the Logical-Core (L-C) cell of the matrix, the firms might compete with one another for transit customers while at the same time interconnect with one another to provide global connectivity for their customers. Once again,

it is possible, though not necessary, for a firm to be operating in the P-C and L-C cells of the matrix simultaneously. For that matter, it is even possible for a firm to be operating in all four cells of the matrix.

The presence of competition at these four loci directly translates into choice available to the end users. First, they may choose among different Physical-Edge (P-E) links based on requirements for mobility, bandwidth, and considerations of cost. End users may even choose to multi-home, i.e., connect to multiple physical networks at the same time, for availability or redundancy reasons. The end users can also choose among different Logical-Edge (L-E) providers, with the services offered over different physical networks, or over the same physical network. Furthermore, if the network architecture supports some form of user-directed routing, the end users may be able to exercise choice between different Logical-Core (L-C) and possibly even Physical-Core (P-C) options. In telephony, for example, consumers can select long-distance carriers on a per-call basis today by dialing additional prefixes. In sum, the end user can, in theory, exercise (or delegate) choice in all four cells of the matrix.

3. LOCI ARE NOT SILOS

While we can identify multiple loci of competition in a network architecture, it would be a mistake to think that we can design for competition one locus at a time, ignoring any interactions between the loci. In reality, there are many different ways through which competition in one locus can affect that in another locus.

3.1 Substitutes

First, consumer choice in one locus may sometimes serve as a substitute for choice in another locus. For example, if the network architecture seeks to promote route diversity, it could do so through the support of multi-homing (1st-hop diversity), and/or through the support of user-directed routing (ith-hop diversity). Multi-homing requires P-E (and possibly L-E) diversity, while user-directed routing is often targeted at L-C (and possibly P-C) diversity. To the extent that end users can access meaningful routing alternatives, multi-homing and user-directed routing are *imperfect substitutes*, and so a network architect can weigh the relative costs and benefits of supporting one or both mechanisms in his/her architectural design, and justify his/her design accordingly.

P-E competition (facilities-based competition) and L-E competition (e.g., through open-access regulation) is another good example of potential substitutes in the architecture. In the 1990s, when facilities-based competition was considered weak in the U.S., regulators and legislators pushed for L-E competition by requiring the physical network operators to open up access to their facilities (e.g., through unbundled network elements) to competing service providers. As a result, competitive local exchange carriers (CLECs) were able to offer alternate L-E service to consumers by gaining access to the existing local loop from the incumbent local exchange carriers (ILECs) such as the Baby Bells, rather than build out their own physical network. In response, the ILECs put up non-technical obstacles to access its facilities (e.g., switches in central offices, cell towers), in their efforts to frustrate competition in the L-E locus. In retrospect, one can view the CLEC experiment as a failure, i.e., open access may relieve the pressure for facilities-based competition, but it is not an effective substitute for P-E competition, as a monopoly P-E operator can still tilt the playing field in the L-E locus.

3.2 Vertical Integration

A network architecture should not preclude a single firm from operating in more than one locus of the competition matrix. A firm may be vertically integrated by operating in both the L-E and P-E loci (e.g., Comcast), in both the L-C and P-C loci (e.g., Level 3), or even in all four loci (e.g., AT&T). Once vertical integration comes into play, however, a network architect needs to be cognizant of the competition dynamics that arise.

The benefit of vertical integration derives from the fact that an integrated firm, e.g., one operating in both L-E and P-E, can provide service to consumers more cheaply than two independent firms operating in L-E and P-E separately. This *economies-of-scope* cost savings may come from a variety of sources, e.g., joint facilities operation, billing, customer service. For example, in Figure 1(b), the vertically integrated provider may be able to provide the same service at a lower cost, and therefore offer it at a lower price, than two separate providers at the logical and physical layers.




Figure 1. Examples of vertical industry structures. (a) No vertical integration -- users choose logical layer provider (black) and physical layer provider (white) separately; (b) With vertical integration, users can either choose logical and physical layer providers separately, or choose integrated provider; (c) All firms vertically integrated -- users choose one integrated provider; higher entry barrier for new entrants; (d) Competition at logical layer with monopoly at physical layer -- user has no choice in physical layer service; vertically integrated firm has control over whether or how user has choice of logical layer service.

To the extent that economies-of-scope are strong, vertical integration is socially desirable, and the industry structure may be dominated by vertically integrated firms, as illustrated in Figure 1(c). From a consumer's perspective, there is still effective choice in providers. However, this industry structure may discourage innovation at one of the loci. For example, a new firm with an innovation in the L-E locus may not be able to offer the new service unless it also enters the P-E market, which may have a high entry barrier.

Furthermore, vertical integration can be detrimental to competition, and therefore social welfare, if it allows a firm to use its dominant position (i.e., market power) in one locus to compete unfairly in another locus. Figure 1(d) offers an example where a vertically integrated firm enjoys a monopoly position in the P-E market, and has the ability to exercise its monopoly power to frustrate competition in the L-E market.

In our earlier example, the Baby Bells (e.g., Pacific Bell, Bell Atlantic) were vertically integrated ILECs that held monopoly positions in their respective P-E markets, and were accused of unfairly competing against the CLECs in the L-E market.

In an even earlier example, AT&T prior to the 1984 Divestiture was vertically integrated across the local and long-distance telephony markets, and was accused of using its monopoly power in the local telephony market to compete unfairly against MCI, Sprint, and other competitors in the long-distance telephony market. Even after the competing long-distance carriers successfully sued to gain access to interconnection with AT&T's local networks, AT&T was

still able to tilt the playing field in its favor by effecting hidden subsidies from its local service to its long-distance service. It eventually took anti-trust action by the U.S. Department of Justice to force a “vertical disintegration” of AT&T, in 1984, into separate firms for local and long-distance service.

The story continues in the 1990s and 2000s, when the digitization of voice and video, together with the widespread availability of mobile data networks, meant that local telephony operators no longer had a monopoly in the P-E market. Consequently, anti-trust objections to vertical integration no longer hold, and vertical re-integration occurred.

Given the historical evolution of the industry, it should be clear that the current market structure should not be assumed as the final equilibrium state of the industry. While multiple competitors occupy the P-E market today, continuing technological advances and/or changing consumer requirements may once again shift this locus towards a monopolistic market in the future.

In summary, the key takeaway is that vertical integration between two or more loci of competition can produce desirable cost savings due to economies-of-scope efficiencies. However, competition must be robust in each of the loci, or barriers to entry must be low. Otherwise, vertical integration can lead to anti-trust problems.

3.3 Tying

The designation of AT&T as the exclusive carrier in the U.S. for the iconic Apple iPhone is a notable recent example of tying. By selling one product or service only in conjunction with another product or service, tying can lead to competition dynamics in two loci. Similar to vertical integration, tying is not a major concern if consumers have choices in alternative products and services in both loci, or if potential entrants do not face high entry barriers into either locus. In the case of iPhone/AT&T, consumers have ample alternatives for both device and carrier, and so tying is not considered problematic from an anti-trust perspective. On the other hand, should one or both of the tying firms enjoy sufficient market power in their loci, the practice may in fact be enjoined by antitrust laws, e.g., the Sherman Antitrust Act and the Clayton Act in the U.S.

For the 700MHz spectrum auction in the U.S. in 2008, the Federal Communications Commission (FCC) adopted “open device” and “open application” rules for the highly sought after “C” Block. The rules effectively preempt

Verizon, the eventual auction winner, from engaging in tying of devices or applications to services offered over its 700MHz network. Given the ability of the winner to build out a high-quality nationwide network using the spectrum, the open rules were believed to encourage greater innovation in both the loci of mobile devices and applications.

3.4 Delegation, Gatekeeping, and Network Neutrality

Both “design for choice” and “design for competition” principles stipulate that users be permitted to express their preferences for services. In a general sense, this may be realized by an interface that allows users to make selections in each of the loci of competition in the architecture.

Some users or applications will clearly benefit from the ability to exercise full control over the selection of services in all loci. However, it is not clear that, even in a simple 2x2 architecture, most or many users will actually wish to make separate, explicit selections for each of the loci. Instead, it is more realistic to expect a typical user to (i) select a service provider in one of the loci, relying on the service provider to select services in the other loci, or (ii) delegate the service selection decisions to a third-party broker, who will compose a service package out of options in the various loci to match the preferences of the user.

Figure 2 illustrates, for a simple two-layer network architecture example, how delegation may be differently realized. Figure 2(a) shows, in the absence of delegation, how a user will have to explicitly select providers in both the Physical and Logical layers. In Figure 2(b), the user selects a third-party broker, and delegates the selection of Physical and Logical services to the broker. In Figure 2(c), the user selects a Physical layer provider, and relies on the latter to select Logical layer services on its behalf. In this case, the Physical layer provider serves as the *gatekeeper* to Logical layer services. Finally, in Figure 2(d), the user selects a Logical layer provider, and relies on the latter to select Physical layer services on its behalf. The Logical layer provider serves as the gatekeeper to the Physical layer services.

Given the complexity and dynamism of large distributed networks, service providers or third-party brokers are naturally in a much better position to collect and act upon low-level network information to make service composition decisions on behalf of their users. However, is there any difference, from a competition or innovation perspective, between the various forms of delegation?




Figure 2. Examples of Delegation and Gatekeeping. Selections made by user, broker, and gatekeeper are labeled as ‘U’, ‘B’, and ‘G’ respectively. (a) Mix and match – user selects logical (black) and physical (white) layer services; (b) Delegation – user delegates choice to third party broker; (c) Physical Gatekeeper – user selects physical layer service, which in turn selects logical layer services; (d) Logical Gatekeeper – user selects logical layer service, which in turn selects physical layer services.

The answer is yes. If there is a relative difference in strength of competition in the two loci, allowing the firms in the less competitive locus to serve as gatekeepers controlling access to services in the other locus will result in reduced competition and innovation in both loci [4]. Typically, competition is weaker in the physical layer because of the necessity of significant up-front capital investments. In this case, the physical gatekeeper model of delegation in Figure 2(c) would be the least desirable from an economic perspective. Architecturally, this means that it is important for the interface to allow end users or third party brokers to explicitly select logical layer services, independent of the physical layer provider. This will prevent a physical gatekeeper from becoming an obstacle between innovative logical layer services and the users who desire them. Architectural proposals like ROSE (‘Routing as a Service’) [5], CABO [6], and Cabernet [7] offer possible paths for the network to move away from a physical gatekeeper model, which currently dominates the Internet, to a model based on logical gatekeepers and/or delegation to third-party brokers.

Generalizing to more than two loci, we should ensure that the network architecture does not allow a competition-challenged locus to assume a gatekeeping role over any other locus. This is entirely consistent with, and perhaps provides a fundamental economic argument for, the principle of network neutrality. If the physical access network locus is the competition-challenged locus, then we must prevent the dominant firms in this locus from having the ability to discriminate against different offerings in the applications, services, content, and devices loci.

In the absence of explicit prohibition, architecturally or otherwise, we should expect firms in the competition-challenged loci to exercise their market power over service selection and become de-facto gatekeepers. At the same time, these firms have no incentives whatsoever to embrace changes in the architecture that may shift the power of service selection to other loci. Therefore, a network architecture cannot afford to be silent on this issue. One possible design choice is to foreclose, at the architectural level, the possibility of gatekeeping by the competition-challenged loci. This leaves open the possibility of gatekeeping by other more competitive loci. However, if we are uncertain of the long-run relative competitiveness of different loci in the architecture, an alternate approach may be to explicitly disallow gatekeeping by any locus, and to rely on delegation to independent third-party brokers.

4. EVOLUTION

Network architectures evolve over time, and so do the loci of competition.

The evolution can occur in several ways. First, the level of competition in a given locus may change over time. Consolidation via horizontal mergers may change a competitive locus to an oligopolistic or even a monopolistic one. Conversely, new entrants with novel technologies or business models may introduce competition to a locus. As we have seen in the preceding discussions, these changes can affect levels of competition in adjacent loci as well.

Second, new technologies and services may create new loci of competition that did not exist before, and/or destroy old ones. In the process, the adjacency of loci may also be rearranged. In the context of the current Internet architecture, we have seen an emergence of content distribution networks (CDNs) and large content providers (LCPs, e.g., Google) as major sources of inter-domain traffic [8]. We can consider the provision of these services as new loci of competition. Furthermore, the datacenters for these CDN/LCP servers are being deployed with direct connections to the consumer networks, bypassing the backbone networks altogether. This means that the vertical relationship between the oligopolistic local access locus and the competitive long-distance locus, traditionally subject to scrutiny by telecom economists and regulators, may be receding in importance. Instead, it is the bilateral oligopolistic relationship between the local access and the CDN/LCP loci that may become the new focus of attention. Traditionally, bilateral oligopolies are characterized by long-term contracts negotiated between players with strong market power. This has important implications to innovation in both of the loci. The “paid peering”

agreements that are currently being negotiated between local access networks and large content providers are precisely the type of long-term contracts that will shape the industry for years to come.

5. MINIMIZE SWITCHING COSTS

In addition to ensuring that a multiplicity of providers can be supported at each locus, there is another important role to be played by the architect. It is to influence how often, and how easily, the choice of providers can be exercised at each locus. Specifically, the lower the cost of switching providers, the easier it is for service consumers to try out new innovative services, and the more competitive a locus can be.

In the search engine market, for example, Hal Varian contends that Google has every incentive to continue to innovate because “competition is only one click away” [9]. So, should a network architecture allow choices to be made on a per-year basis, per-application basis, or per-packet basis?

The appropriate timescale of choice will likely differ for each locus, depending on the nature of service in question, and the amount of overhead incurred by the service provider each time a customer adds or drops the service. Extra attention should be given to those loci that are inherently more competition-challenged, e.g., due to strong scale economies. For example, number portability proved to be a particularly effective way to reduce switching costs in the local access locus of the telephony network architecture. Developments in software-defined radio technologies, as another example, offer the prospect of dramatically lowering switching costs for the wireless local access locus.

Importantly, the objective should be to minimize switching costs on both the customer side and the provider side. If we focus on only the former, but ignore the latter, the provider will have greater incentive to manage its own incurred costs by discouraging switching behavior altogether. In particular, the provider can create an artificial switching cost to the customers by employing long-term service contracts with early termination penalties. Such a strategy of contractual lock-in may serve the provider’s goal of churn-reduction, but it can also dampen competition significantly.

6. TAKEAWAYS

Let us close with a summary of main takeaways from this paper for network architects who are contemplating the design of future Internet architectures.

First, the architecture should ensure that a multiplicity of providers can be supported at each locus. We do not want to unnecessarily foreclose competition anywhere in the network.

Second, we need to recognize that architecture outlasts technologies and applications. The level of competition in each locus may change over time, and we cannot anticipate nor dictate the number of choices in each locus.

Third, in choosing designs that facilitate competition at each locus, we want to pay particular attention to those loci that are naturally competition-challenged due to stronger economies of scale, etc. In general, minimizing switching costs can be an effective way to promote competition in any locus.

Finally, remembering that “loci are not silos,” we cannot design for competition one locus at a time. Instead, the network architecture must facilitate robust competition in face of a wide range of possible strategic interactions between providers in different loci, as well as providers that straddle multiple loci. Interfaces must be carefully developed to manage and lubricate both the technical and market transactions between the loci.

7. REFERENCES

- [1] D. Geer, R. Bace, P. Gutmann, P. Metzger, C. Pfleeger, J. Quarterman, and B. Schneier, “Cyberinsecurity: The Cost of Monopoly,” Computer and Communications Industry Association Report, 2003.
- [2] D. Clark, J. Wroclawski, K. Sollins, and R. Braden, “Tussles in Cyberspace: Defining Tomorrow’s Internet,” *IEEE/ACM Transactions on Networking*, vol. 13, no. 3, Jun. 2005, pp. 462-475.
- [3] T. Anderson, L. Peterson, S. Shenker, and J. Turner, “Overcoming the Internet Impasse Through Virtualization,” *Computer*, vol. 38, no. 4, Apr. 2005, pp. 31-41.
- [4] P. Laskowski and J. Chuang, “Innovations and Upgrades in Virtualized Network Architectures,” *Proc. Workshop on the Economics of Networks, Systems, and Computation*, 2010.

- [5] K. Lakshminarayanan, I. Stoica, and S. Shenker, "Routing as a Service," UC Berkeley EECS Technical Report No. UCB/CSD-04-1327, 2004.
- [6] N. Feamster, L. Gao, and J. Rexford, "How to Lease the Internet in Your Spare Time," *ACM Computer Communication Review*, January 2006.
- [7] Y. Zhu, R. Zhang-Shen, S. Rangarajan, and J. Rexford, "Cabernet: Connectivity Architecture for Better Network Services," *Proc. Workshop on Rearchitecting the Internet*, 2008.
- [8] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-Domain Traffic," *Proc. ACM SIGCOMM*, 2010.
- [9] H. Varian, "The Economics of Internet Search," Angelo Costa Lecture delivered in Rome, Feb. 2007.