

Implications of Selfish Neighbor Selection in Overlay Networks

NIKOLAOS LAOUTARIS^{†‡} GEORGIOS SMARAGDAKIS[†] AZER BESTAVROS[†] JOHN W. BYERS[†]
nlaout@cs.bu.edu gsmaragd@cs.bu.edu best@cs.bu.edu byers@cs.bu.edu

Abstract—In a typical overlay network for routing or content sharing, each node must select a fixed number of immediate overlay neighbors for routing traffic or content queries. A selfish node entering such a network would select neighbors so as to minimize the weighted sum of expected access costs to all its destinations. Previous work on selfish neighbor selection has built intuition with simple models where edges are undirected, access costs are modeled by hop-counts, and nodes have potentially unbounded degrees. However, in practice, important constraints not captured by these models lead to richer games with substantively and fundamentally different outcomes. Our work models neighbor selection as a game involving directed links, constraints on the number of allowed neighbors, and costs reflecting both network latency and node preference. We express a node’s “best response” wiring strategy as a k -median problem on asymmetric distance, and use this formulation to obtain pure Nash equilibria. We experimentally examine the properties of such stable wirings on synthetic topologies, as well as on real topologies and maps constructed from PlanetLab and AS-level Internet measurements. Our results indicate that selfish nodes can reap substantial performance benefits when connecting to overlay networks constructed by naive nodes. On the other hand, in overlays that are dominated by selfish nodes, the resulting stable wirings are optimized to such great extent that even uninformed newcomers can extract near-optimal performance through naive wiring strategies.

I. INTRODUCTION

Motivation: *Neighbor selection* is a key problem for a broad class of distributed services and applications that run atop large, amorphous overlay networks of autonomous nodes. For example, in an overlay routing or a peer-to-peer file sharing network, a new node must first select a relatively small number of direct neighbors before it can connect to the service. In these systems, and in many others, it is clear that the impact of the neighbor selection strategy is significant, as evidenced by the emerging body of work exploring network creation games and characterizing the equilibria of these games. To date, however, the bulk of the work (and main results) in this area have centered on games where edges are undirected, access costs are based on hop-counts, and nodes have potentially unbounded degrees [1], [2], [3]. While this existing body of work is extremely helpful for laying a theoretical foundation and for building intuition, it is not clear how or whether the guidance provided by this prior work generalizes to situations of practical interest, in which underlying assumptions in these prior studies are not satisfied. Another aspect not considered in previous work is the consideration of settings in which some

or even most players do not play optimally – a setting which we believe to be typical. Interesting questions along these lines include an assessment of the advantage to a player from employing an optimizing strategy, when most other players do not, or more broadly, whether employing an optimizing strategy by a relatively small number of players could be enough to achieve global efficiencies.

Scope and Contributions: In this paper, we formulate and answer such questions using a combination of modeling, analysis, and extensive simulations using synthetic and real datasets. Our starting point is the definition of a network creation game that is better suited for settings of P2P and overlay routing applications – settings that necessitate the relaxation and/or modification of some of the central modeling assumptions of prior work. In that regard, the central aspects of our model are:

(1) *Bounded Degree:* Most protocols used for implementing overlay routing or content sharing impose hard constraints on the maximum number of overlay neighbors. For example, in popular versions of BitTorrent a client may select up to 35 nodes from a neighbors’ list provided by the *Tracker* of a particular torrent file [4].¹ In overlay routing systems [8], the number of immediate nodes has to be kept small so as to reduce the monitoring and reporting overhead imposed by the link-state routing protocol implemented at the overlay layer. Motivated by these systems, we explicitly model such hard constraints on node degrees. Notice that in the prior studies cited above, node degrees were *implicitly bounded* (as opposed to *explicitly constrained*) by virtue of the trade-off between the additional cost of setting up more links and the decreased communication distance achieved through the addition of new links. We also note that some of these earlier network creation games were proposed in the context of physical communication networks. In such networks, the cost of acquiring a link is instrumental to the design and operation of a critical infrastructure. Such concerns do not apply in the case of overlay networks such as those we consider in this paper. Thus, we argue that models in which node degrees are outcomes of an underlying optimization process do not faithfully reflect the realities of systems and applications we consider.

(2) *Directed Edges:* Another important consideration in the

¹ KaZaA and FastTrack include neighbor constraints at multiple levels: ordinary nodes (ON) may select up to 5 super nodes (SN) from a larger list for establishing initial negotiation and then maintain connection with only one of these; SNs may connect to at most 50 other SNs (from a typical population of SNs ranging between 25K and 40K [5]) and accept between 55 to 70 (or 100 to 160) children ONs (depending on their provisioning). New versions of Gnutella and Limewire involve a similar two-level architecture [6] with associated constraints. Similarly, DHT routing protocols like Chord [7] impose hard constraints on the number of first hop neighbors.

[†] Computer Science Dept, Boston University, Boston, Massachusetts, USA.

[‡] Dept of Informatics and Telecommunications, University of Athens, Greece.

[§] N. Laoutaris is supported by a Marie Curie Outgoing International Fellowship of the EU MOIF-CT-2005-007230. A. Bestavros and J. Byers are supported in part by NSF CNS Cybertrust Award #0524477, CNS NeTS Award #0520166, CNS ITR Award #0205294, and EIA RI Award 0202067.

settings we envision for our work relates to link directionality. Prior models have generally assumed bi-directional (undirected) links. This is an acceptable assumption that fits naturally with the unbounded node degree assumption for models that target physical telecommunication networks because actual wire-line communication links are almost exclusively bidirectional. In overlay settings we consider, this assumption needs to be relaxed since the fact that node v forwards traffic or requests to node u does not mean that node u may also forward traffic or requests to v .

(3) *Non-uniform preference vectors*: In our model, we supply each node with a vector that captures its local preference for all other destinations. In overlay routing such preference may capture the percentage of locally generated traffic that a node routes to each destination, and then the aggregation of all preference vectors would amount to an origin/destination traffic matrix. In P2P overlays such preference may amount to speculations from the local node about the quality of, or interest in, the content held by other nodes. Other considerations may also include subjective criteria such as the perceived capacity of the node, its geographic location, or its availability profile.

(4) *Representative distance functions*: Although the initial models presented in this paper use assumptions made in several previous studies regarding equal unitary pair-wise distances for all one-hop overlay links, later in this paper, we relax this assumption by considering more representative distance models. As was done in [2], we consider synthetic distances obtained using topology generators. In addition, we consider more realistic settings in which topologies are obtained from real Internet settings – namely the PlanetLab overlay and actual AS-level maps – and in which associated distances are obtained through real measurements in these settings.

Our first technical contribution within this model is to express a node’s “best response” wiring strategy as a k -median problem on asymmetric distance [9], and use this observation to obtain pure Nash equilibria through iterative best response walks via local search. We then experimentally investigate the properties of stable wirings on synthetic topologies as we vary two key properties of interest: (i) the *edge density* of the graph and (ii) the *non-uniformity of popularity* of nodes within the topology.

Our experimental results then consider neighbor selection problems motivated and driven by measurements of PlanetLab and the AS-level topology with a realistic access cost model. Here, we find that selfish nodes can reap substantial performance benefits when connecting to overlay networks composed of non-selfish nodes. On the other hand, in overlays that are dominated by selfish nodes, the resulting stable wirings are already so highly optimized that even non-selfish newcomers can extract near-optimal performance through naive wiring strategies.

II. DEFINITIONS

Let $V = \{v_1, v_2, \dots, v_n\}$ denote a set of nodes. Associated with node v_i is a preference vector $p_i = \{p_{i1}, p_{i2}, \dots, p_{i(i-1)}, p_{i(i+1)}, \dots, p_{in}\}$, where $p_{ij} \in [0, 1]$ de-

notes the preference of v_i for v_j , $i \neq j$: $\sum_{j=1, j \neq i}^n p_{ij} = 1$. Node v_i establishes a *wiring* $s_i = \{v_{i_1}, v_{i_2}, \dots, v_{i_{k_i}}\}$ by creating links to k_i other nodes (we will use the terms link, wire, and edge interchangeably). Edges are *directed* and *weighted*, thus $e = (v_i, v_j)$ can only be crossed in the direction from v_i to v_j , and has cost d_{ij} . Going the opposite direction requires crossing edge (v_j, v_i) and incurring cost d_{ji} ($d_{ji} \neq d_{ij}$ in the general case). Let $S = \{s_1, s_2, \dots, s_n\}$ denote a *global wiring* between the nodes of V and let $d_S(v_i, v_j)$ denote the cost of a shortest directed path between v_i and v_j over this global wiring; $d_S(v_i, v_j) = M \gg n$ if there’s no directed path connecting the two nodes. For the overlay networks discussed here, the above definition of cost amounts to the incurred end-to-end delay when performing shortest-path routing along the overlay topology S , whose direct links have weights that capture the delay of crossing the underlying IP layer path that goes from the one end of the overlay link to the other. Let $C_i(S)$ denote the cost of v_i under the global wiring S , defined as the weighted (by preference) summation of its distances to all other nodes, i.e., $C_i(S) = \sum_{j=1, j \neq i}^n p_{ij} \cdot d_S(v_i, v_j)$.

Definition 1: (The SNS Game) The selfish neighbor selection game is defined by the tuple $\langle V, \{S_i\}, \{C_i\} \rangle$, where:

- V is the set of n players, which in this case are the nodes.
- $\{S_i\}$ is the set of strategies available to the individual players. S_i is the set of strategies available to v_i . Strategies correspond to wirings and, thus, player v_i has $\binom{n-1}{k_i}$ possible strategies $s_i \in S_i$.
- $\{C_i\}$ is the set of cost functions for the individual players. The cost of player v_i under an outcome S , which in this case is a global wiring, is $C_i(S)$.

The above definition amounts to a non-cooperative, non-zero sum, n -player game [10]. Let $S_{-i} = S - \{s_i\}$ denote the *residual wiring* obtained from S by taking away v_i ’s outgoing links.

Definition 2: (Best Response) Given a residual wiring S_{-i} , a best response for node v_i is a wiring $s_i \in S_i$ such that $C_i(S_{-i} + \{s_i\}) \leq C_i(S_{-i} + \{s'_i\})$, $\forall s'_i \neq s_i$.

Definition 3: (Stable Wiring) A global wiring S is stable iff it is composed of individual wirings that are best responses.

Therefore stable wirings are just pure Nash equilibria of the SNS game, i.e., they have the property that no node can re-wire unilaterally and reduce its cost. In [11] we have established that stable wirings always exist under uniform node popularity and overlay link weights.

III. DERIVING STABLE WIRINGS

In this section we start with a description of a general method for obtaining the best response of a node under general overlay link weights, which we then refine for the case that link weights are uniform. Next, we describe the iterative best response algorithm that we use for obtaining stable wirings. We conclude this section by presenting a simple lower bound for the social cost of a socially optimal solution — we later use this bound to evaluate the social cost of stable wirings.

A. The Best Response of a Node

A wiring for a node v_i can be defined using $n - 1$ binary unknowns Y_l , $1 \leq l \leq n$, $l \neq i$: $Y_l = 1$ iff v_i wires to v_l , and

0 otherwise. Define also the binary unknowns X_{lj} : $X_{lj} = 1$ iff v_i has v_l as a first-hop neighbor on a shortest path to v_j . A best response for v_i under residual wiring S_{-i} can be obtained by solving the following Integer Linear Program (ILP):

$$\text{Minimize:} \\ C_i(S_{-i}, X) = \sum_{j=1, j \neq i}^n p_{ij} \sum_{l=1, l \neq i}^n X_{lj} \cdot (d_{il} + d_{S_{-i}}(v_l, v_j)) \quad (1)$$

Subject to:

$$\sum_{l=1, l \neq i}^n X_{lj} = 1, \forall j \neq i \text{ and } \sum_{l=1, l \neq i}^n Y_l = k_i \text{ and } X_{lj} \leq Y_l, \forall l, j \neq i, \quad (2)$$

where d_{il} is the cost of a wire from v_i to v_l , and $d_{S_{-i}}(v_l, v_j)$ is the cost of a shortest path from v_l to v_j over the wiring S_{-i} .

B. Connection between the SNS Game and Facility Location

When all the wires have the same unitary weight, then the distances d_S are essentially ‘‘hop counts’’, in which case there is an interesting relationship between finding a node’s best response wiring and solving a k -median problem on asymmetric distance [9], [12]. The latter is defined as follows:

Definition 4: (Asymmetric k -median) Given a set of nodes V' , weight’s $w_j \forall v_j \in V'$, and an asymmetric distance function $d_{S'}$ (meaning that in general $d_{S'}(v, u) \neq d_{S'}(u, v)$), select up to k nodes to act as medians so as to minimize $C(V', k, w)$, defined as follows:

$$C(V', k, w) = \sum_{v_j \in V'} w_j \cdot d_{S'}(v_j, m(v_j)),$$

where $m(v_j)$ is the median that is closest to v_j .

Proposition 1: The best response of node v_i to S_{-i} under uniform link weights ($d_{ij} = 1, \forall i, j \in V$) can be obtained by solving an asymmetric k -median problem, in which:

- 1) $V' = V - \{v_i\}$
- 2) $k = k_i$
- 3) $w_j = p_{ij}, v_j \in V'$
- 4) $d_{S'}(u, w) = d_{S_{-i}}(w, u), u, w \in V'$,

Proof: Let s_i denote v_i ’s response to S_{-i} . The resulting cost will be:

$$\begin{aligned} C_i(S_{-i} + \{s_i\}) &= \sum_{v_j \in V'} p_{ij} d_{S_{-i} + \{s_i\}}(v_i, v_j) \\ &= \sum_{v_j \in V'} p_{ij} (d_{S_{-i} + \{s_i\}}(v_i, m(v_j)) + d_{S_{-i} + \{s_i\}}(m(v_j), v_j)) \\ &= \sum_{v_j \in V'} p_{ij} d_{S_{-i} + \{s_i\}}(v_i, m(v_j)) + \sum_{v_j \in V'} p_{ij} d_{S_{-i} + \{s_i\}}(m(v_j), v_j) \\ &= \sum_{v_j \in V'} p_{ij} + \sum_{v_j \in V'} p_{ij} d_{S_{-i} + \{s_i\}}(m(v_j), v_j) \\ &= \sum_{v_j \in V'} w_j + \sum_{v_j \in V'} w_j d_{S_{-i}}(m(v_j), v_j) \\ &= c + \sum_{v_j \in V'} w_j d_{S'}(v_j, m(v_j)) \end{aligned} \quad (3)$$

where c is a constant and $m(v_j)$ is v_i ’s next-hop neighbor on a shortest path to v_j under the global wiring $S_{-i} + \{s_i\}$. The transition from the third to the fourth line of Eq. (3) relies on the fact that all distances to first hop neighbors are equal to 1 under hop-count distance. Obtaining the best response requires minimizing $C_i(S_{-i} + \{s_i\})$. Equation (3) shows that

this is equivalent to minimizing $\sum_{v_j \in V'} w_j d_{S'}(v_j, m(v_j))$, which is exactly the objective function of the above mentioned asymmetric k -median problem. ■

Proposition 1 suggests that v_i ’s best response is to wire to the k_i medians of a distance function obtained by reversing the end-to-end distances of the residual wiring S_{-i} . Since even the metric version of k -median is NP-hard [12], so is its asymmetric version, and through Proposition 1, the best response of the SNS game as well. For the metric version of the k -median there exist several algorithms that provide constant-factor approximations of an exact solution [13], [14], [15], [16]. These guarantees do not hold for the asymmetric case. For the asymmetric k -median, Lin and Vitter [17] have given a bicriterion approximation that blows up the number of used medians by an $O(\log n)$ multiplicative factor to provide a cost that exceeds the optimal one by an additive factor. Archer [9] has shown that this is the best attainable approximation for this problem unless $NP \subseteq DTIME(n^{O(\log \log n)})$. Despite this negative result, simple heuristics like the p -swapping local search of Arya *et al.* [15] perform typically very well on the directed k -median (as also confirmed by our numerical results later in this paper).

C. Equilibrium Wirings through Iterative Best Responses

Definition 5: (Iterative best response) Given an initial global wiring $S^{(0)}$, start an iterative procedure where at the m -th iteration the nodes line up according to their ids (*i.e.*, v_1, v_2, \dots), and perform the following steps:

- 1) v_i computes its best response $s_i^{(m)}$ to $S_{-i}^{(m, i-1)}$, after v_{i-1} and before v_{i+1}
 - 2) $S^{(m, i)} = S_{-i}^{(m, i-1)} + \{s_i^{(m)}\}$
- $S^{(m, i-1)}$ is the global wiring at iteration m (after v_{i-1} ’s best response and prior to v_i ’s best response); $S_{-i}^{(m, i-1)}$ is the corresponding residual wiring with respect to v_i ($S_{-i}^{(m, 0)} = S^{(m-1, n)} - \{s_1^{(m-1)}\}$ and $S_{-i}^{(1, 0)} = S^{(0)} - \{s_1^{(0)}\}$). The iterative best response search stops and returns $S = S^{(\mathcal{M})}$ when at iteration \mathcal{M} : $s_i^{(\mathcal{M})} = s_i^{(\mathcal{M}-1)}, \forall v_i \in V$, *i.e.*, when no node can profit by re-wiring.

We use the iterative best response method to find stable wirings. In Sect. IV where we present synthetic results based on hop-count distance we take advantage of the connection established through Proposition 3, and employ exact (ILP) and approximate (p -swapping local search) solutions for the directed k -median in order to obtain best responses. In Sect. V we employ several real topologies in which distances are not hop-count and, therefore, employ the ILP formulation of Sect. III-A in order to obtain best responses.

D. A Lower Bound on the Cost of a Socially-Optimal Wiring

Let S^* denote a socially optimal (SO) wiring, *i.e.*, a global wiring that minimizes the *social cost* $C(S) = \sum_{v_i \in V} C_i(S)$. Let $S^{U, i}$ denote the *utopian* wiring for v_i , *i.e.*, the global wiring that minimizes $C_i(S)$ over all possible global wirings S (this should not be confused with a best response s_i that minimizes $C_i(S_{-i} + \{s_i\})$ granted a particular residual wiring S_{-i}). We can obtain a lower bound L on $C(S^*)$ by

summing the costs of the individual utopian solutions, *i.e.*, $L = \sum_{\forall v_i \in V} C_i(S^{U,i})$. We describe $S^{U,i}$ for some interesting cases below. Before that, let o_{-i}^j denote the node with the j th largest out-degree, excluding v_i — let this degree be denoted $k(o_{-i}^j)$.

Uniform node preference: When $p_i = p = \{1/n, \dots, 1/n\}$, $\forall v_i \in V$, it is easy to see that $S^{U,i}$ is a directed tree with downward pointing edges, where: (1) v_i is the root; (2) v_i connects to nodes $o_{-i}^1, o_{-i}^2, \dots, o_{-i}^{k_i}$ at level 1; (3) these nodes connect to the next $l_1 = \sum_{j=1}^{k_i} k(o_{-i}^j)$ nodes with highest degrees ($o_{-i}^{k_i+1}, o_{-i}^{k_i+2}, \dots, o_{-i}^{k_i+l_1}$) at level 2, and so on.

Uniform out-degree: When $k_i = k$, $\forall v_i \in V$, then $S^{U,i}$ is a directed regular k -ary tree with downward pointing edges, where (1) v_i is the root; (2) level l includes k^l nodes whose preference according to p_i ranks from $\left(\sum_{l'=1}^{l-1} k^{l'}\right) + 1$ to $\sum_{l'=1}^l k^{l'}$.

Uniform preference and out-degree: Combining the previous two cases results in a regular k -ary tree with l levels such that:

$$\sum_{l'=1}^l k^{l'} \geq n - 1 \Rightarrow k \frac{k^l - 1}{k - 1} \geq n - 1 \Rightarrow l \geq \log_k \left[\frac{(n-1)(k-1)}{k} + 1 \right]$$

The resulting (common) cost for all $v_i \in V$ is:

$$\begin{aligned} C_i(S^{U,i}) &= \left(\sum_{l'=1}^l l k^{l'} \right) - l \left(\sum_{l'=1}^l k^{l'} - (n-1) \right) \\ &= \frac{l(k-1)(n(k-1)+1) - k(k^l - 1)}{(k-1)^2} \end{aligned} \quad (4)$$

Later in this paper, we use the aforementioned bound to show numerically that the social cost of stable wirings is close to the social cost of socially optimal wirings.

IV. CHARACTERIZATION OF STABLE WIRINGS

In this section we assume that establishing a direct (overlay) link between any two nodes incurs unit cost and, therefore, the cost between any pair of nodes equals the number of hops along any shortest, directed path that connects these nodes at the overlay layer. Our goal will be to characterize the structure of stable wirings with respect to two key scaling parameters of interest. The first parameter, $\alpha \in [0, 1]$, reflects the non-uniformity (skew) in the popularity of different destinations. We create such non-uniformity by adopting a generalized power-law profile for node popularity with skewness α , meaning that the popularity of the i th most popular node is $q_i = \Lambda/i^\alpha$, where $\Lambda = \left(\sum_{k=1}^n \frac{1}{k^\alpha}\right)^{-1}$. We construct the preference vector p_i of node v_i by setting $p_{ij} = q_j/(1 - q_i), \forall v_j \in V : v_j \neq v_i$. High values of α mean that there are few highly-popular destinations among all the nodes, whereas low values mean that most destinations are equally popular.

The second parameter, $\beta \in [0, 1]$, determines the *link density* of a regular graph, which relates to the fanout (out-degree) of each node as follows: $k = \lceil n^\beta \rceil$.

For a given pair (α, β) we obtain the corresponding stable wiring by using the iterative best response method of Section III-C, where the best response amounts to a solution of a directed k -median problem. Here, it is worthwhile to notice

that different node orderings in the iterative best response search may lead to different stable wirings.² We have found that different stable wirings perform approximately the same and therefore it is of marginal value to look at the structure of different individual ones. To support this we note that it has been established analytically in [11] that provably existent stable wirings are guaranteed to perform approximately as well as socially optimal solutions under uniform node popularity. A similar conclusion is reached in the next section (albeit experimentally) for the case of non-uniform popularity.

A. Social Cost of Stable Wirings

We first consider the quality of stable wirings compared to the utopian wirings described in the previous section. As can be seen for the examples depicted in Figure 1 (a) and (b), which are representative of a much larger set of simulations we conducted, the gap between the stable solution and the Utopian solution is small, and this result holds across a wide range of settings for α and β , and for various values of n for which simulation was tractable. In terms of absolute values, the social cost decreases with both the skew in popularity and link density. In particular, a highly-skewed popularity profile ensures that shorter paths to the most popular destinations are realized, whereas higher link densities reduces the average length of shortest paths, and thus the social cost as well.

Since computing exact best response wirings is NP-hard, even under hop-count distance, it makes sense to study the performance of *approximate best responses* and corresponding *approximately stable wirings*. For this purpose, we used the Local Search (LS) heuristics described in [15] to solve the k -median problem, which yields the best response wiring by virtue of Proposition 1. We also considered ϵ -stable versions of the problem in which nodes do not re-wire unless they can reduce their current cost by *at least* a multiplicative factor ϵ (we combined ϵ -stability with both exact (ILP) and approximate (LS) best responses). As evident from Figure 1 (c), we found that ϵ -stable wirings have similar social costs.³

To summarize, stable wirings have performance close to the socially optimal wirings. Moreover, approximate best response wirings can be computed fast with LS and ϵ approximations.

B. Characterization of Stable Wirings

Next, we take a more in-depth look at the stable wirings that result for given values of α and β , as depicted in the set of graphs in Figure 2, where α varies from left to right and β varies from top to bottom.

The first interesting finding is evident from an examination of the structures that emerge when $\alpha = 0$ (*i.e.*, under uniform popularity – the leftmost column in the figure). Despite the equal popularity of nodes, the resulting stable wirings do not exhibit uniform in-degree node distributions. In particular, some nodes tend to be more desirable for other nodes. Had the links been bidirectional, the emergence of such “hubs”

² See [18] for a related discussion based on a different *object replication game*.

³ The results in Figure 1 (c) were obtained for $\epsilon = 0.05$, similar results (not shown) were obtained for $\epsilon \in [0.01, 0.1]$.




Fig. 1. **(a)** Comparison of the social cost $C(S)$ of stable wirings to a lower bound of the cost of a socially optimal solution (the Utopian solution of Section III-D) for $n = 15$. Stable wirings obtained using exact best responses based on an ILP formulation of the directed k -median problem of Section III-B. **(b)** same as (a) with $n = 50$. **(c)** Comparison of the social cost $C(S)$ of stable wirings obtained by using exact (ILP) and approximate (LS) best response and corresponding $\epsilon = 5\%$ versions. **(d)** Average path length for the stable graph obtained by using exact (ILP) best response for $n = 15$.

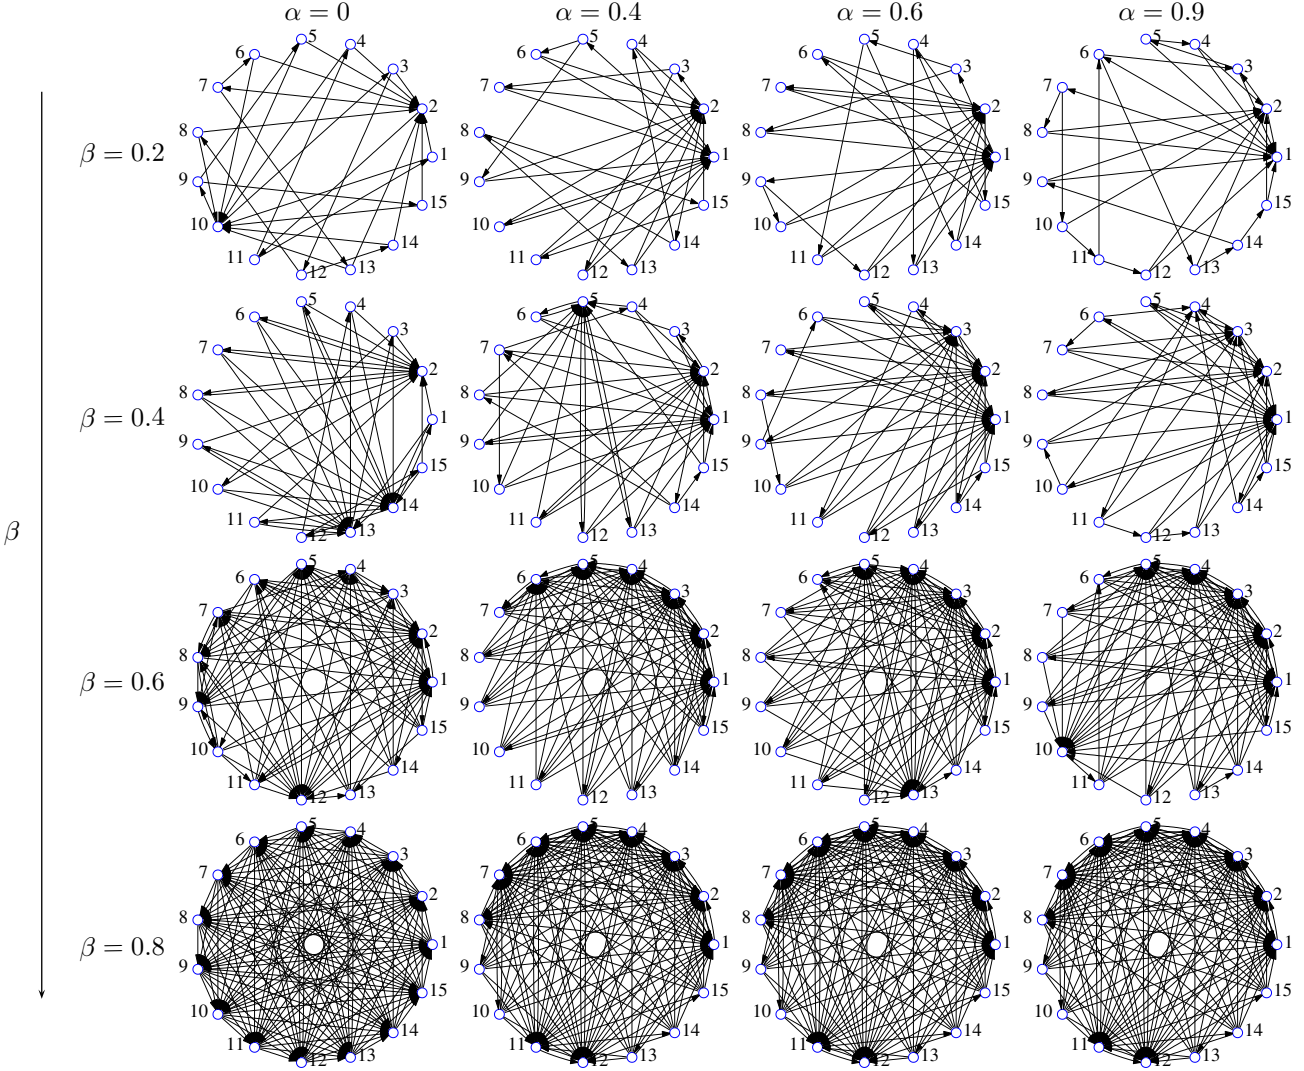


Fig. 2. Stable wiring motifs for $n = 15$ and different values of α and β .

could have been easily explained, by noting that they would be serving the purpose of providing short outgoing routes to many destinations. In our case, however, this cannot be the cause since these hubs have many *incoming* links, whereas their outgoing links are just as many as for the other nodes since all nodes have exactly k links, where k is controlled by the link density (β). Having made sure that these hubs did

not emerge due to bias in tie-breaking during the computation of best responses, we attribute this preferential attachment phenomenon to the *quality rather than the quantity of outgoing links of hub nodes*. In particular, the hubs are nodes that (by coincidence) managed to position their k outgoing links in such a way that is beneficial to others as well (despite the fact that the wiring has been decided solely based on selfish




Fig. 3. The social cost $C(S)$ of doubly capacitated stable wirings for $\beta = 0.2$ and different α, γ .

criteria).

Moving on to other larger values of α , where popularity is skewed, the hub creation process becomes a mix of the aforementioned phenomenon and the inherent preference for popular nodes. Nodes that are globally popular are natural candidates for becoming hubs. Even with relatively low skew ($\alpha = 0.4$), the most popular nodes are becoming hubs (node with $\text{id}=1$ is the most popular and that with $\text{id}=n$ is the least popular). We see this trend consistently for all values of β , as is to be expected. But interestingly, as β increases further, it is not simply a contiguous sequence of the most popular nodes that end up becoming hubs! For example, in the $\alpha = 0.6, \beta = 0.6$ case, several nodes in the “tail” end of the popularity distribution end up becoming hubs as well, facilitating relay shortcuts as in the uniform popularity case.

We also find that the average path length slowly increases with α for a given β (see Figure 1 (d)). This is to be expected since nodes prefer to be closer to the most popular nodes, and thus place less importance on the distance to much less popular nodes. Although this reduces the newcomers’ access costs, it increases some shortest paths, and the diameter.

C. Constraining the In-degree: A Doubly Constrained Overlay

We next examine the effects of constraining the maximum in-degree of nodes so that they never have more than ν incoming links, while maintaining also the constraint on the out-degree. We can enforce this constraint by including in the definition of $C_i(S)$ a large penalty for connecting to nodes that have more than $\nu - 1$ incoming links. We can define a scaling factor γ for the in-degree as done previously with β for the out-degree.

In Figure 3, we fix the out-degree scaling parameter to $\beta = 0.2$, and present the social cost for different values of the in-degree scaling parameter γ . Low values of γ increase the social cost under skewed popularity profiles, as in these cases, the highly-popular nodes quickly reach their maximum in-degree and thus, many nodes have to reach them indirectly through multi-hop paths. Note that without in-degree constraints most nodes would access them in a single hop by establishing a direct overlay link to them. When γ is low, e.g., $\gamma = 0.2$, the resulting graph (not shown here due to space constraints) looks much like a ν -regular graph. With large values of γ , i.e., γ approaching 1, the in-degree constraints become too loose and, thus, the corresponding stable graphs become similar to their unconstrained counterparts.

V. OVERLAY NEIGHBOR SELECTION: BEST-RESPONSE VS. k -RANDOM AND k -CLOSEST

In this section we take a closer look at the performance benefits from employing best response wiring instead of simpler wiring strategies. We also depart from the simplistic unit-distance model for the cost of direct links and instead use more realistic cost models on synthetic and measured topologies. Corresponding stable wirings are obtained by using the ILP model for a node’s best response under general distances as detailed in Section III-A.

A. Description and Design Methodology

In Section II, we defined the best response strategy for a node entering a given network. Now, we consider two other natural alternatives. Let d_{ij}^X denote the cost associated with creating a direct overlay link between nodes v_i and v_j under a model X for end-to-end IP layer distances. We say that a “newcomer” node v_i employs a k -Closest wiring strategy under the model X when it establishes a wiring s_i such that $d_{ij}^X \leq d_{i_j'}^X$ for all $v_j \in s_i, v_j' \notin s_i$. We say that a newcomer node v_i employs a k -Random wiring strategy when it chooses a wiring s_i uniformly at random from the space of all valid wirings of cardinality k_i .

To substantiate the benefits of best response, we consider the initial graph awaiting a “newcomer” upon its arrival. We assume that this initial graph has resulted from having its constituent nodes apply a specific wiring strategy.⁴ We refer to an instance of an n node graph for which each of the n nodes employed a k -Closest strategy as a k -Closest graph, and attribute similar meanings to a k -Random graph and a *Best Response (BR)* graph.

B. Description of the Datasets

In this section we describe the IP-layer end-to-end distance models X from which we obtain the d_{ij}^X ’s that are used as weights for direct overlay links between nodes v_i and v_j .⁵ The following three datasets are used:

BRITE: The first dataset is synthetically generated from the BRITE topology generator [19] following a Barabási-Albert [20] model with $N = 1000$ nodes and incremental growth parameter $\mu = 2$. The nodes were placed on the plane according to a heavy tail model that creates high density clusters. Based on the observation that the delay between two nodes in high speed networks is highly correlated to their physical distance [21], we assigned weights on the links at the physical layer by calculating the Euclidean distance between their two end nodes.

PlanetLab: PlanetLab is an overlay testbed network of approximately 700 nodes in more than 300 academic, industrial, and government sites around the world. We used a publicly available dataset [22] containing delays obtained using *pings* between all pairs of PlanetLab sites (inter-site delays are more representative than inter-node delays for overlay applications).

⁴ To guarantee connectivity, nodes that participate in a k -Random or a k -Closest graph, donate one link in order to create a ring. We note that a ring is a feature common to many other overlays, such as the Chord DHT [7].

⁵ Overlay nodes that do not have a direct link communicate through a shortest-path on the overlay topology.

AS-level map: As a third dataset, we use the relation-based AS topology map of the Internet from December 2001 (data available from [23]). This map was constructed by using the measurement methodology described in [24]. The dataset includes two kinds of relationships between ASes: (1) customer-provider: The customer is typically a smaller AS that pays a larger AS for access to the rest of the Internet. The provider may, in turn, be a customer of an even larger AS. A customer-provider relationship is modeled using a directed link from the provider to the customer. (2) Peer-Peer: Peer ASes are typically of comparable size and have mutual agreements for carrying each other’s traffic. Peer-peer relationships are modeled using undirected links. Overall the AS-level map includes 12779 unique ASes, of which 1076 are peers (joined by at least one peer-peer link), and the remaining 11703 are customers. These ASes are connected through 26387 directed and 1336 undirected links. We choose to present results based on the largest connected component of the dataset, which we found to include a substantial part of the total AS topology at the peer level: 497 peer ASes connected with 1012 links (we verified that this component contains all the top-20 larger peer ASes reported in [24]). The ASes that participate in this graph are responsible for routing the majority of the Internet traffic. We measured the hop-count distance between pairs of overlay nodes and used it as weight for a direct link between these two nodes at the overlay layer. To model the characteristics of IP routing (unique path), we broke ties by assigning each edge i a weight $1+\epsilon_i$ where ϵ_i is a zero-mean random noise as suggested in [25].

C. Comparison of Different Graphs

Using as input the weighted graphs from our three datasets, we obtained the social costs resulting from applying the various wiring strategies under consideration, for different values of β . The Best Response (BR) graph (resulting from having all nodes apply the best response wiring strategy) was by far the most optimized wiring, thus providing a lower-bound for the simpler k -Random and k -Closest strategies. Table I summarizes our results by providing the ratios of the social costs of the simple wiring strategies (k -Random, k -Closest) to that of the BR wiring. These results suggest that the premium provided by BR is highest for lower link densities (*i.e.*, when β is small). This is an intuitive result since in denser graphs, there is less of an opportunity for optimization.

The results in this section give us a baseline for the efficiency of the wirings that result from the adoption by all nodes in the graph of the same strategy (be it k -Random, k -Closest, or BR). This sets up the stage for our next set of questions: Given such an initial wiring, what is the marginal utility to a newcomer from executing each one of the three wiring strategies under consideration?

D. The Value of Best-Response

Given an initial wiring created (as described above) by having n overlay nodes follow one of our three wiring strategies, we quantify the benefit to a “newcomer” (*i.e.*, the $n + 1$ ’st node) from choosing its neighbors using one of the

three neighbor selection strategies. Nine possibilities exist for applying strategy S1 over a wiring obtained using S2, where S1 and S2 could be k -Random, k -Closest, or BR. We use $c(w|G(n))$ to denote the cost of a newcomer using wiring strategy w on a pre-existing graph G of n nodes, or simply $c(w)$ when the graph G is understood. For example, $c(k\text{-Random} - k\text{-Closest})$ denotes the cost of a newcomer using the k -Random wiring strategy to connect to a graph of n nodes, each of which employed the k -Closest wiring strategy to construct the initial graph (to which the newcomer will connect).

In the results presented below, we set $n = 50$ and evaluate the performance for 200 newcomers on the BRITE and AS dataset and 100 newcomers for the PlanetLab dataset (which is smaller). Our main results are shown in Figure 4, where each column corresponds to an underlying graph model, and each row corresponds to a strategy employed by the n newcomers. Within each plot, we vary the link density β along the x-axis, and plot the cost ratio of the $n+1$ ’st arrival for a given strategy versus the cost of the $n + 1$ ’st arrival if it were to use BR.

a) Connecting to a k -Random Graph: The plots in the top row of Figure 4 show the case in which the first n arrivals use k -Random, and thus the underlying graph is poorly optimized.

With such an initial graph, the k -Random wiring is a poor choice for the $(n + 1)$ st node, as it could lead to significantly higher costs (anywhere from 30% for the BRITE and AS datasets to 60% for the PlanetLab dataset) when compared to using BR. This performance gap closes, as one would expect, when β (and therefore k) becomes large. In fact this trend holds in all cases because finding a closer approximation to BR is easier when each node has more links — and therefore ample opportunity to make good connections, even when using simple strategies.

Using the k -Closest wiring, on the other hand, turns out to be a very reasonable choice, as it achieves a cost comparable to that achieved by BR (typically within 10% with small exceptions for the BRITE and PlanetLab datasets under low link densities). This finding suggests that in poorly optimized random graphs, simply connecting to your nearby neighbors (at low cost), is a good rule of thumb, especially when edge density is high.

b) Connecting to a k -Closest Graph: The plots in the middle row of Figure 4 show the case in which the first n arrivals use k -Closest, and thus the underlying graph consists mostly of local edges with few shortcuts. Here we see that it is considerably more important for newcomers to behave strategically. For example, on the BRITE topology, naively using k -Closest is a poor choice that perpetuates the lack of shortcuts in the underlying graph to the point that even using k -Random turns out to be a better choice! In the other topologies, k -Closest and k -Random are comparable, and the improvement in quality relative to BR as β increases is much more modest.

One conclusion from the results we obtained above for connection strategies to the k -Random and k -Closest graphs is of particular importance for P2P applications. In a P2P

	$\beta = 0.1$		$\beta = 0.2$		$\beta = 0.4$		$\beta = 0.6$		$\beta = 0.8$	
	k -Random/BR	k -Closest/BR	k -Random/BR	k -Closest/BR	k -Random/BR	k -Closest/BR	k -Random/BR	k -Closest/BR	k -Random/BR	k -Closest/BR
BRITE	1.44	1.53	1.52	1.84	1.38	2.07	1.28	1.46	1.09	1.16
PlanetLab	2.23	1.48	1.75	1.23	1.37	1.13	1.09	1.16	1.04	1.06
AS-level	2.04	1.90	1.83	1.61	1.58	1.39	1.24	1.23	1.12	1.16

TABLE I


SOCIAL COST RATIOS BETWEEN SIMPLE WIRING STRATEGIES (k -RANDOM, k -CLOSEST) AND BEST RESPONSE.

Fig. 4. The cost ratio between simple wiring (k -Random or k -Closest) and BR wiring for a newcomer node that connects to a pre-existing network of n nodes that was wired using k -Random, or k -Closest, or BR. We present the 25-,50-,75-quartiles for the aforementioned ratios using three different data sets (BRITE, PlanetLab, AS-level map) for obtaining the costs of establishing direct links.

network, nodes pick neighbors randomly from the list provided to them by a bootstrap server (*i.e.*, the initial graph to which a newcomer would connect is a k -random graph). Under such circumstances, *it pays to “cheat”*, by ping-pong the possible neighbors and connecting to the k -Closest ones. However, if the constituent nodes in the initial graph also cheat, (*i.e.*, the initial graph to which a newcomer would connect is a k -closest graph), then *it does not pay to cheat; it may even cost!*

c) Connecting to a Stable Graph: Finally, the plots in the bottom row of Figure 4 show the case in which the first n arrivals use BR, and thus the underlying graph ends up being highly optimized, prior to the arrival of newcomers. In this case, the graph is so much optimized for the newcomer that any reasonable strategy might well have good performance. Surprisingly, while the k -Closest strategy does indeed perform well for the newcomer across the three topologies, the alternative strategy of k -Random does not. This seemingly odd result could be explained by noting that given the very low overall

costs between nodes in the optimized initial graph, the cost to the newcomer from selecting its own neighbors (as opposed to the cost of reaching all nodes in the graph) could not be ignored. A poor choice of neighbors could backfire.⁶

d) General Observation: In conclusion, we find common trends across the three topologies with respect to strategic neighbor selection behavior. At the two extremes where the other players are playing completely at random or completely selfishly (top and bottom rows, respectively), the underlying graphs are either too poorly constructed, or too well constructed, for an uninformed newcomer to be at a significant disadvantage. In either of these two situations, the naive strategy of k -Closest is generally competitive to BR, especially under stable graphs. Picking links at random in these situations however, is unlikely to work well, unless the graph is already

⁶ Recall that the topologies we considered in this section feature non-unit link costs, and as such, selecting neighbors at random could put the newcomer at a disadvantage, especially if the initial graph was optimized, since the relative penalty from a bad random selection of neighbors would be high.

dense (large β).

But in the middle regime, in which all the other players adopt k -Closest, the newcomer must be much more careful. Here, there is much to be gained by the optimal shortcuts selected in BR, which neither k -Closest nor k -Random typically selects. Strikingly, our experimental results suggest that k -Closest is actually the worst of the possible strategies considered for the newcomer to adopt in this situation.

VI. RELATED WORK

Selfish neighbor selection for overlay networks was first mentioned by Feigenbaum and Shenker [26]. Fabrikant et al. [1] studied an unconstrained undirected version of the problem in which nodes can buy as many links as they want at a fixed per link price α . Chun et al. [2] studied experimental an extended version of the problem in which links prices need not be the same. Rocha et al. [3] is in the same spirit. All these works do not consider hard constraints on node degrees and thus are fundamentally different from ours.

Bindal et al. [4] propose a locality-enhanced version of BitTorrent in which only m out of the total k neighbors of a BitTorrent node are allowed to belong to a different ISP. Although the capacitated selection of neighbors is a central aspect of this work, their treatment is fundamentally different from ours in several regards: (i) there's no contention between selfish peers, (ii) the minimization objective is on inter-AS traffic therefore only two levels of communication distance are modeled, intra and inter-AS (we use finer topological information that includes exact inter-peer distances), and (iii) their "reachability" constraint amounts to asking for a similar level of data availability as the original one under the standard random neighbor selection mechanism of BitTorrent (we have fundamentally different reachability constraints, expressed as general preference functions over the potential overlay neighbors). Another recent work on neighbor selection is from Godfrey et al. [27]. It aims at selecting neighbors in a way that minimizes the effects of node churn (appearance of new nodes, graceful leaves and sudden malfunctions), but unlike our work, it does not focus on the impact of competing selfish nodes.

VII. SUMMARY AND CONCLUSION

Our experimental results on selfish neighbor selection, in a richer model that captures the nuances of overlay applications more faithfully than previous work, reveals numerous subtleties that are not apparent in simpler models. Among our most noteworthy findings is that it is typically in a newcomer's best interest (whether that newcomer is naive or sophisticated) to have had the prior arrivals behave selfishly, as the underlying "best response" graph is often highly optimized in favor of the newcomer. A corollary is that suboptimal behavior by a participant is often costly, not only to the individual, but to the population at large, i.e. suboptimal behavior leads to large negative externalities. Given that best-response neighbor selection has a significant performance advantage over other heuristics, our future work is to implement a practical and feasible realization.

REFERENCES

- [1] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker, "On a network creation game," in *Proc. of ACM PODC '03*. Boston, Massachusetts: ACM Press, 2003, pp. 347–351.
- [2] B.-G. Chun, R. Fonseca, I. Stoica, and J. Kubiatowicz, "Characterizing selfishly constructed overlay routing networks," in *Proc. of IEEE INFOCOM'04*, Hong Kong, 2004.
- [3] B. G. Rocha, V. Almeida, and D. Guedes, "Improving reliability of selfish overlay networks," in *Proc. of WWW'06*, Edinburgh, UK, 2006.
- [4] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, "Improving traffic locality in BitTorrent via biased neighbor selection," in *Proc. of IEEE ICDCS '06*, Lisbon, Portugal, 2006.
- [5] J. Liang, R. Kumar, and K. W. Ross, "The FastTrack overlay: A measurement study," *Computer Networks*, vol. 50, no. 6, pp. 842–858, 2006.
- [6] D. Stutzbach and R. Rejaie, "Characterizing the two-tier Gnutella topology," in *Proc. of ACM SIGMETRICS '05*, Banff, Alberta, Canada, 2005.
- [7] I. Stoica, R. Morris, D. Liben-Nowell, D. Karger, M. Kaashoek, F. Dabek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup protocol for internet applications," *IEEE/ACM ToN*, vol. 11, no. 1, pp. 17–32, 2003.
- [8] Y. Liu, H. Zhang, W. Gong, and D. F. Towsley, "On the interaction between overlay routing and underlay routing," in *Proc. of IEEE INFOCOM '05*, Miami, FL, 2005, pp. 2543–2553.
- [9] A. Archer, "Inapproximability of the asymmetric facility location and k -median problems," 2000, unpublished manuscript, available from the author's web-page.
- [10] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. MIT Press, 1994.
- [11] N. Laoutaris, R. Rajaraman, R. Sundaram, and S.-H. Teng, "A bounded-degree network formation game," 2007, arXiv-CoRR cs.GT/0701071.
- [12] P. Mirchandani and R. Francis, "Discrete location theory." John Wiley and Sons, 1990.
- [13] M. Charikar, S. Guha, D. B. Shmoys, and E. Tardos, "A constant factor approximation algorithm for the k -median problem," in *Proc. of ACM STOC '99*, Atlanta, GA, USA, 1999, pp. 1–10.
- [14] M. R. Korupolu, C. G. Plaxton, and R. Rajaraman, "Analysis of a local search heuristic for facility location problems," in *Proc. of ACM-SIAM SODA '98*, San Francisco, CA, USA, 1998, pp. 1–10.
- [15] V. Arya, N. Garg, R. Khandekar, K. Munagala, and V. Pandit, "Local search heuristic for k -median and facility location problems," in *Proc. of ACM STOC '01*, Hersonissos, Greece, 2001, pp. 21–29.
- [16] K. Jain and V. V. Vazirani, "Primal-dual approximation algorithms for metric facility location and k -median problems," in *Proc of IEEE FOCS '99*, New York City, NY, USA, 1999.
- [17] J.-H. Lin and J. S. Vitter, " ϵ -approximations with minimum packing constraint violation," in *Proc. of ACM STOC '92*, Victoria, British Columbia, Canada, 1992.
- [18] N. Laoutaris, O. Telelis, V. Zissimopoulos, and I. Stavrakakis, "Distributed selfish replication," *IEEE Transactions on Parallel and Distributed Systems*, vol. 17, no. 12, pp. 1401–1413, 2006.
- [19] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An Approach to Universal Topology Generation," in *Proc. of MASCOTS '01*, Cincinnati, OH, Aug 2001, pp. 346–354.
- [20] A.-L. Barabási and R. Albert, "Emergence of scaling in random networks," *Science*, vol. 286, no. 5439, pp. 509–512, 1999.
- [21] A. Ziviani, S. Fdida, J. F. de Rezende, and O. C. M. B. Duarte, "Toward a measurement-based geographic location service," in *Proc. of PAM'2004*, Antibes Juan-les-Pins, France, Apr. 2004, pp. 43–52.
- [22] C. Yoshikawa, <http://ping.eecs.uc.edu/ping/>, accessed on July 10, 2006.
- [23] <http://www.cc.gatech.edu/~mihail/ASdata.html>.
- [24] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the internet hierarchy from multiple vantage points," in *Proc. of IEEE INFOCOM '02*, New York City, NY, USA, 2002.
- [25] A. Lakhina, J. Byers, M. Crovella, and P. Xie, "Sampling biases in IP topology measurements," in *Proc. of IEEE INFOCOM'03*, Apr. 2003.
- [26] J. Feigenbaum and S. Shenker, "Distributed algorithmic mechanism design: Recent results and future directions," in *Proc. of the 6th International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*, 2002, pp. 1–13.
- [27] P. B. Godfrey, S. Shenker, and I. Stoica, "Minimizing churn in distributed systems," in *Proc. of ACM SIGCOMM '06*, Pisa, Italy, 2006.