

Socially-Aware Network Design Games

Jocelyne Elias
Politecnico di Milano
elias@elet.polimi.it

Fabio Martignon
University of Bergamo
fabio.martignon@unibg.it

Konstantin Avrachenkov, Giovanni Neglia
INRIA Sophia Antipolis
{Konstantin.Avrachenkov, Giovanni.Neglia}@sophia.inria.fr

Abstract—In many scenarios network design is not enforced by a central authority, but arises from the interactions of several self-interested agents. This is the case of the Internet, where connectivity is due to Autonomous Systems’ choices, but also of overlay networks, where each user client can decide the set of connections to establish.

Recent works have used game theory, and in particular the concept of Nash Equilibrium, to characterize stable networks created by a set of selfish agents. The majority of these works assume that users are completely non-cooperative, leading, in most cases, to inefficient equilibria.

To improve efficiency, in this paper we propose two novel socially-aware network design games. In the first game we incorporate a socially-aware component in the users’ utility functions, while in the second game we use additionally a Stackelberg (leader-follower) approach, where a leader (e.g., the network administrator) architects the desired network buying an appropriate subset of network’s links, driving in this way the users to overall efficient Nash equilibria.

We provide bounds on the Price of Anarchy and other efficiency measures, and study the performance of the proposed schemes in several network scenarios, including realistic topologies where players build an overlay on top of real Internet Service Provider networks. Numerical results demonstrate that (1) introducing some incentives to make users more socially-aware is an effective solution to achieve stable and efficient networks in a distributed way, and (2) the proposed Stackelberg approach permits to achieve dramatic performance improvements, designing almost always the socially optimal network.

Index Terms: - Network Design, Social Awareness, Game Theory, Nash Equilibrium, Stackelberg Game.

I. INTRODUCTION

Network design with selfish users has been the focus of several recent works [1], [2], [3], [4], [5], [6], which have modeled how independent selfish agents can build or maintain a large network by paying for possible edges. Each user’s goal is to connect a given set of terminals with the minimum possible cost. Game theory is the natural framework to address the interaction of such self-interested users (or players). A *Nash Equilibrium* (NE) is a set of users choices, such that none of them has an incentive to deviate unilaterally. For this reason the corresponding networks are said to be *stable*.

However, Nash equilibria in network design games can be much more expensive than the optimal, centralized solution. This is mainly due to the lack of cooperation among network users, which leads to design costly networks.

Actually, the majority of existing works assume that users are completely non-cooperative. However, this assumption could be not entirely realistic, for example when network design involves long-term decisions (e.g., in the case of

Autonomous Systems peering relations). Moreover, incentives could be introduced by some external authority (e.g., the overlay administrator) in order to increase the users’ cooperation level.

In this work¹ we overcome this limitation by first proposing a novel network design game, the Socially-Aware Network Design (SAND) game, where users are characterized by an objective function that combines both *individual* and *social* concerns in a unified and flexible manner. More specifically, the cost function of each user is a combination of its own path cost (the *selfish* component) and the overall network cost, which represents the *social* component. A parameter (α) weights the relative importance of the network cost with respect to the user path cost. Changing the value of α permits to take into account different levels of social awareness or user cooperation.

We investigate systematically the impact of cooperation among network agents on the system performance, through the determination of bounds on the *Price of Anarchy* (PoA), the *Price of Stability* (PoS) and the *Reachable Price of Anarchy* ($RPoA$) of the proposed game. They all quantify the loss of efficiency as the ratio between the cost of a specific stable network and the cost of the optimal network, which could be designed by a central authority. In particular the PoA , first introduced in [8], considers the worst stable network (that with the highest cost), while the PoS [1] considers the best stable network (that with the lowest cost); finally, the $RPoA$ considers only Nash equilibria reachable via best response dynamics from the empty solution [6]. Hence, PoA and $RPoA$ indicate the maximum degradation due to distributed users decisions (anarchy), while the PoS indicates the minimum cost to pay to have a solution robust to unilateral deviations. Our analytical results show that as α increases, i.e., when users are more sensitive to the social cost, the PoS converges to 1, i.e., the best stable network is more efficient, as expected. Surprisingly, an opposite result holds for the worst case. Indeed, for large α values (highly socially-aware users) the worst stable network can be much more expensive than the networks designed by purely selfish users.

For this reason, we further propose a Stackelberg approach, the Network Administrator-Driven SAND game (NAD-SAND), which enables very efficient Nash equilibria, avoiding worst-case scenarios: a leader (e.g., the network

¹An extended version of this paper is available as INRIA technical report [7].

administrator) buys an appropriate subset of the network links (i.e., those belonging to the minimum cost generalized Steiner tree covering all source/destination pairs), inducing the followers (the network users) to reach an efficient Nash equilibrium.

We measured the performance of the proposed games in several network topologies, including realistic scenarios where players build an overlay on top of real Internet Service Provider networks, and we observed that socially-aware users always generate better networks. Furthermore, we observed that the proposed Stackelberg approach achieves dramatic performance improvements in all the considered scenarios, even for small α values, since it leads most of the times to the optimal (least cost) network. Hence, introducing some incentives to make users more socially-aware could be an effective solution to achieve stable and efficient networks in a distributed way.

The paper is organized as follows: Section II introduces the proposed Socially-Aware Network Design game. Section III provides precise bounds on the Price of Anarchy, the Price of Stability and the Reachable Price of Anarchy for the SAND game. Section IV describes the proposed Stackelberg game (the Network Administrator-Driven SAND game), which enables very efficient Nash equilibria. Section V presents numerical results that demonstrate the effectiveness of the SAND and NAD-SAND games in several realistic network scenarios. Finally, Section VI concludes this paper.

II. THE SOCIALLY-AWARE NETWORK DESIGN GAME

In this Section we illustrate the proposed Socially-Aware Network Design (SAND) game, motivating the reason to introduce such model.

The SAND game occurs in a directed graph $G = (V, E)$, where each edge e has a nonnegative cost c_e , and each player $i \in I = \{1, 2, \dots, k\}$ is identified with a source-sink pair (s_i, t_i) . Every player i picks a path S_i from its source to its destination, thereby creating the network $(V, \cup_i S_i)$ with total cost equal to $\sum_{e \in \cup_i S_i} c_e$, which will be referred to as *social* cost. We will refer to the path S_i also as the strategy chosen by player i .

This social cost is assumed to be shared among the players in the following way: let x_e denote the number of overlay paths that go through overlay link e ; if edge e lies in x_e of the chosen paths, then each player choosing such an edge pays a proportional share $\pi_e = \frac{c_e}{x_e}$ of the cost.

The objective function J^i that user i wants to minimize is therefore given by:

$$J^i = \sum_{e \in S_i} \pi_e + \alpha \sum_{e \in \cup_j S_j} c_e \quad (1)$$

The first term takes into account the *selfish* nature of each player, since it is the cost for user i to buy the edges belonging to the chosen path, S_i ; on the other hand, the second term represents the total network cost (i.e., the *social* cost), α being a parameter that permits to give more weight to one component with respect to the other.

An alternative interpretation of objective function (1) is also possible: we can think that users are completely selfish, but the social-aware term in cost function (1), $\alpha \sum_{e \in \cup_j S_j} c_e$, is imposed by the network operator. The advantage of such approach is that the operator does not need to solve a large-scale Integer Linear Programming problem in order to optimize the routing choices for every change in the network (either in the topology, links cost, players location etc ...), but can let the users solve the problem in a distributed way, converging to a good and stable solution.

Note that for $\alpha = 0$ objective function (1) corresponds to that of the Shapley Network Design Game proposed in [1], which represents a particular case of our proposed game. Furthermore, in our game, users need only a limited amount of information, which is exactly equal to that of the Shapley network design game.

We observe that the SAND game is a *potential game* without being at the same time a *congestion game* on the given graph G [9]. In fact, the cost incurred by player i , J^i , is not simply the sum of the costs incurred by such player from each link in the chosen path S_i , but it includes a further term, the total network cost (multiplied by the parameter α).

It is easy to verify that the SAND game is characterized by the following potential function:

$$\Phi(S) = \sum_{e \in E} \sum_{x=1}^{x_e} \frac{c_e}{x} + \alpha \sum_{e \in \cup_j S_j} c_e \quad (2)$$

$S = (S_1, S_2 \dots S_k)$ being the vector of players' strategies.

As a consequence, such game has at least one pure Nash equilibrium, namely the strategy S that minimizes $\Phi(S)$ [10]. Furthermore, in such a game, best response dynamics always converge to a Nash equilibrium.

Since we have introduced a term proportional to the network cost in the objective function of each player, we expect that players should design better networks, even though Section III shows that this is not necessarily the case.

III. BOUNDS ON THE PRICE OF ANARCHY, PRICE OF STABILITY AND REACHABLE PRICE OF ANARCHY FOR THE SAND GAME

In this Section, we derive bounds on the Price of Anarchy (*POA*), the Price of Stability (*POS*) and the so-called Reachable Price of Anarchy (*RPOA*) for the Socially-Aware Network Design game, and compare them with the results presented in [1] for the Shapley Network Design game. This allows us to determine the worst and best case performance of our proposed game.

We first establish a lower bound on the Price of Anarchy for the SAND game, which is defined as the ratio between the cost of the worst stable network (that with the highest cost), and the cost of the optimal network.

Proposition 1: In the SAND game, a lower bound on the Price of Anarchy (*POA*) is given by the following expression:

$$POA \geq k(1 + \alpha). \quad (3)$$

Then, we compute an upper bound on the Price of Stability for the SAND game, which is defined as the ratio between the cost of the best stable network (that with the lowest cost), and the cost of the optimal network.

Proposition 2: In the SAND game, the Price of Stability (PoS) is upper bounded by the following expression:

$$PoS \leq \frac{\mathcal{H}_k + \alpha}{1 + \alpha}. \quad (4)$$

Finally, we derive a bound on the so-called Reachable Price of Anarchy ($RPoA$), a quantity defined in [6]. The denominator of this ratio is the cost of the socially optimal network, which we denote by C^{opt} ; the numerator is the largest cost of an equilibrium reachable via the following process: the k players enter the game one-by-one in an arbitrary order, and each picks a path of minimum cost (according to objective function (1)), given the choices of previous players. After all players have entered, the game proceeds exactly as for the SAND game, with each player re-optimizing its path, given the current strategies of all other players (using a best response algorithm). When the process reaches a Nash equilibrium (as it must, since it is a potential game), it stops.

Proposition 3: In the SAND game, the Reachable Price of Anarchy ($RPoA$) is upper bounded by the following expression:

$$RPoA < k \frac{\alpha + 1}{\alpha + \frac{1}{k}}. \quad (5)$$

Comments

For $\alpha = 0$ the SAND game is equivalent to the original Shapley network design game. Indeed, expressions (3) and (4) confirm the results already demonstrated in [1]: $PoA = k$ and $PoS = \mathcal{H}_k$. In the SAND game, for increasing α values, the upper bound on the PoS decreases, and tends to 1 for $\alpha \rightarrow \infty$ (then also $PoS \rightarrow 1$). This can be easily explained since for $\alpha \rightarrow \infty$ the *social* component of objective function (1) is predominant. In this limiting behavior, all users share the same utility function, which is the second term of expression (1); in such a situation, the social network optimum (i.e., the network with the minimum total cost) is obviously also a Nash equilibrium, since the objective function of each single player coincides with the social network cost (multiplied by α).

On the other hand, our lower bound on the PoA increases when α increases. This corresponds to the (quite counter-intuitive) fact that, in some cases, more socially-aware users can design less efficient networks. Such inefficiency is due to the myopic decision criterion: each player only considers the effect of its own choice (i.e., changing the selected path), without considering eventual future decisions from the other players.

However, it can be observed from expression (5) that the $RPoA$ of the SAND game strictly decreases for increasing α values. For large α values (notably, for $\alpha \rightarrow \infty$), the $RPoA$ is less than k .

Finally, our simulations show that the Nash equilibria reached in the SAND game are, in average, consistently better

than those achieved by the Shapley game, as we will show in Section V.

IV. THE NETWORK ADMINISTRATOR-DRIVEN SOCIALLY-AWARE NETWORK DESIGN GAME

We now illustrate a variation of the SAND game, named the Network Administrator-Driven SAND (NAD-SAND) game, where a network administrator plays before the users, and his aim is to drive them to the best Nash equilibrium possible.

Since computing the optimal Stackelberg strategy is NP-hard, we present in this paper a simple strategy that achieves consistent performance improvements. Such approach is implemented via the following heuristic:

- 1) Given the network topology, the network administrator solves a generalized Steiner Tree problem [11], determining the minimum-cost subnetwork such that the source/destination nodes of each player are connected by a path. Let E^{opt} be the set of edges belonging to such optimal subnetwork.
- 2) The network administrator chooses all links belonging to E^{opt} , thus offering to share eventually their cost with the other players. Therefore, using the notation introduced in Section II, after this step we have $x_e = 1, \forall e \in E^{opt}$ (that is, the network administrator has already chosen all links that are *optimal* from a social point of view).
- 3) At this point, all the k users play the SAND game, each trying to optimize its own objective function, which is the same of expression (1).

The rationale behind the proposed NAD-SAND game is the following: the network administrator tries to motivate all players to use the links that belong to the socially optimal solution by sharing their cost with network users. We will show in the next Section that such heuristic is very effective, and permits to obtain dramatic performance improvements with respect to the SAND game.

We observe that the first step of the NAD-SAND game involves solving an NP-Complete problem. However, several efficient heuristics and approximation algorithms have been proposed to solve such problem in a reasonable computation time [11], [12]. In the numerical results presented in the next Section we were able to compute exactly the minimum cost generalized Steiner tree using a simple ILP formulation, and solving it with the CPLEX 11 solver.

Finally, we observe that as $\alpha \rightarrow \infty$, the NAD-SAND game always reaches the minimum cost network since for each player the cost of choosing any link that does not belong to the minimum-cost subnetwork (i.e., to E^{opt}) has an exceedingly large cost. As a consequence, $PoA \rightarrow 1$ as $\alpha \rightarrow \infty$.

We now derive a lower bound on the Price of Anarchy for the NAD-SAND game.

Proposition 4: In the NAD-SAND game, a lower bound on the Price of Anarchy (PoA) is given by the following expression:

$$PoA \geq \frac{k}{2(1 + \alpha)} \quad (6)$$

for $\alpha \leq \frac{k}{2} - 1$.

V. NUMERICAL RESULTS

In this Section, we report the results obtained by the proposed Socially-Aware Network Design (SAND) and Network Administrator-Driven SAND (NAD-SAND) games in random network topologies as well as in real ISP topologies, comparing them both to the Shapley network design game [1] and to the social optimum. To this end, we consider both randomly generated network instances and real ISP topologies mapped by the Rocketfuel tool [13].

In all cases, we start from the empty network and we apply iteratively a best response algorithm until a Nash equilibrium is reached.

For each scenario we report the total network cost obtained by the proposed games for different α values, as well as the optimal network cost (the ILP column), obtained formulating the generalized Steiner Tree problem [11] with an Integer Linear Program (ILP), using AMPL [14], and solving it with the CPLEX solver [15]. Solving such problem provides the least-cost network topology that connects all source/destination pairs, thus representing a term of comparison for the efficiency of the equilibria reached by our proposed games.

Finally, we evaluated numerically the Price of Anarchy and the Price of Stability in all our simulation settings, and we found that they are, respectively, always less than 1.78 and 1.56, hence consistently lower than the bounds provided in Section III.

Full-Mesh Topologies

We first consider full-mesh network topologies with 50 nodes randomly distributed on a 1000×1000 square area and 20 players (source/destination pairs). The cost of each link is equal to its length, and the numerical results, averaged over 20 random extractions, are reported in Table I.

The SAND game permits to obtain network equilibria with a cost significantly lower (approximately 13%) than that obtained with the Shapley network design game (i.e., the SAND row with $\alpha = 0$), even though there is still room for improvements, as demonstrated by the gap existing between the SAND game and the optimal cost (the ILP column).

This gap is filled by the NAD-SAND game, which achieves consistently cheaper Nash equilibria for all α values (up to more than 30%), including the $\alpha = 0$ case, thus representing a very effective approach also when applied to the original Shapley network design game. Furthermore, when α is sufficiently large (≥ 10), the NAD-SAND game reaches exactly the optimum (i.e., the generalized Steiner Tree cost).

Random Topologies

To generate random network instances, we have implemented a topology generator which considers a square area with edge equal to 1000, and randomly extracts the position of N nodes, uniformly distributed on the square area.

As for the network links, which can be bought by players to connect their endpoints, we consider two alternative choices:

- Random geometric graphs: links exist between any two nodes located within a range R . The link cost is set equal to its length.
- Random network model “Uniform”: a given number of links, L , is generated between randomly extracted couples of nodes. The cost of each link is uniformly distributed in the 0 to C range (C being the maximum link cost).

Given a random network, 20 random selections of $k = 20$ source/destination couples are considered.

Table II illustrates the results obtained in a random geometric graph scenario with 50 nodes, range $R = 500$ and 20 source/destination pairs (players). The results are averaged on 20 source/destination random selections, and also on 5 random topologies. The Table reports the total cost of the network planned by the SAND and NAD-SAND games, as well as the optimal network cost. Also in this scenario, the SAND game achieves improved equilibria with respect to the Shapley network design game (up to 13.2%). The NAD-SAND game outperforms the SAND game, further decreasing the planned network cost, and reaches, for $\alpha \geq 10$, the social optimum.

Similar results were obtained using Random network model “Uniform”, and they are reported in [7].

Real ISP topologies

Finally, we consider three real ISP topologies mapped using Rocketfuel [13], listed in Table III, with an increasing number of nodes and links. The same Table shows the results obtained in such topologies, that is, the total network costs as well as the optimal network costs obtained solving the ILP model. The link costs are those provided by Rocketfuel, and we performed 20 random selections of 20 source/destination pairs.

In the small-size Telstra topology, the equilibria found by the SAND game (as well as those of the Shapley network formation game) are very close to the optimum, which is reached by the NAD-SAND game for $\alpha \geq 50$. As for the other ISP topologies, the SAND game diminishes the total network costs, and approaches the optimal ILP solution, while the NAD-SAND game always plans cheaper networks, achieving the optimal outcome already for relatively small α values.

A final interesting point that we would like to mention is the resilience of the proposed network design game solutions to link failures. In particular, for the three ISP scenarios described above we individuated the most congested link (i.e., the one used by the largest number of players), dropped it and recomputed new equilibria. It turns out that the new network costs in all the cases are at most 2% higher than the costs of the original networks.

VI. CONCLUSION

In this paper we proposed the Socially-Aware Network Design game, a novel network formation game where users are characterized by an objective function that combines both social and individual concerns in a unified and flexible manner.

We studied the efficiency of the Nash equilibria achieved by the proposed game, providing bounds on the Price of Anarchy, the Price of Stability and the Reachable Price of Anarchy.

TABLE I

FULL-MESH TOPOLOGY WITH 50 NODES RANDOMLY DISTRIBUTED ON A 1000×1000 AREA AND 20 PLAYERS: AVERAGE NETWORK COSTS FOR THE SAND AND THE NAD-SAND GAMES. THE OPTIMAL NETWORK COST IS ALSO REPORTED. THE COST OF EACH LINK IS SET EQUAL TO ITS LENGTH.

Game	$\alpha = 0$	$\alpha = 1$	$\alpha = 10$	$\alpha = 50$	$\alpha = 100$	$\alpha = 1000$	ILP
SAND	6956.17	6124.62	6037.53	6043.82	6050.94	6046.66	4214.63
NAD-SAND	5718.86	4707.48	4214.63	4214.63	4214.63	4214.63	

TABLE II

RANDOM GEOMETRIC GRAPHS; RANDOM NETWORKS WITH 50 NODES, $R = 500$ AND 20 PLAYERS: AVERAGE NETWORK COSTS FOR THE SAND AND THE NAD-SAND GAMES. THE OPTIMAL NETWORK COST IS ALSO REPORTED.

Game	$\alpha = 0$	$\alpha = 1$	$\alpha = 10$	$\alpha = 50$	$\alpha = 100$	$\alpha = 1000$	ILP
SAND	6567.73	6074.57	5708.95	5724.18	5736.17	5706.09	4213.82
NAD-SAND	5645.44	4675.13	4213.82	4213.82	4213.82	4213.82	

TABLE III

ROCKETFUEL TOPOLOGIES, 20 SOURCE/DESTINATION PAIRS: AVERAGE NETWORK COSTS FOR THE SAND AND THE NAD-SAND GAMES. THE OPTIMAL NETWORK COST IS ALSO REPORTED.

Network	Nodes	Links	Game	$\alpha = 0$	$\alpha = 1$	$\alpha = 10$	$\alpha = 50$	$\alpha = 100$	$\alpha = 1000$	ILP
Telstra (AU)	108	306	SAND	165.55	164.55	163.15	163.10	163.10	163.10	157.95
			NAD-SAND	163.25	160.70	158.55	157.95	157.95	157.95	
Sprintlink (US)	141	748	SAND	286.25	282.60	280.90	280.30	280.30	280.30	255.60
			NAD-SAND	270.30	263.35	256.85	255.65	255.60	255.60	
Abovenet (US)	315	1944	SAND	301.30	292.65	287.45	287.45	287.05	287.05	261.55
			NAD-SAND	282.95	272.50	262.10	261.55	261.55	261.55	

Our analytical results show that when users are more sensitive to the social cost, the best stable network is more efficient. However, an opposite result holds for the worst case, since highly socially-aware users can design stable networks that are much more expensive than the networks designed by purely selfish users.

For this reason, we further proposed the Network Administrator-Driven SAND game, where the network administrator implements a link building strategy that drives users to the “best” Nash equilibrium in terms of system performance, thus architecting the desired network.

We measured the performance of our proposed games in several network scenarios, including real ISP topologies, and we showed how they outperform classical network formation games (like the Shapley network design game), often obtaining the socially optimal outcome.

Such results demonstrate that introducing incentives to make users more socially-aware can be a very effective solution to achieve stable and efficient networks in a distributed way. Furthermore, if incentives are deployed, the intervention of a network administrator can lead to dramatic performance improvements, as demonstrated by our proposed Stackelberg game.

ACKNOWLEDGMENTS

This work was partially supported by MIUR in the framework of the PRIN SESAME project.

REFERENCES

- [1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, Rome, Italy, October 17–19, 2004.
- [2] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proc. of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 511–520, San Diego, CA, USA, 2003.
- [3] Ho-Lin Chen and T. Roughgarden. Network design with weighted players. In *Proc. of the 18th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA’06)*, Cambridge, MA, USA, July 30–August 2, 2006.
- [4] S. Albers. On the value of coordination in network design. In *Proc. of the 19th annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 294–303, San Francisco, CA, USA, 2008.
- [5] A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. In *Proc. of the 8th ACM conference on Electronic commerce*, pages 84–92, San Diego, CA, USA, June 2007.
- [6] H. L. Chen, T. Roughgarden, and G. Valiant. Designing networks with good equilibria. In *Proc. of the 19th annual ACM-SIAM symposium on Discrete algorithms*, pages 854–863, 2008.
- [7] J. Elias, F. Martignon, K. Avrachenkov, and G. Neglia. Socially-Aware Network Design Games. *INRIA Research Report no.7141*, December 2009. Available at <http://hal.inria.fr/inria-00439687>.
- [8] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, Trier, Germany, 4–6 March 1999.
- [9] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, pages 65–67, vol. 2, no. 1, 1973.
- [10] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. Algorithmic game theory. *Chapter 19, Network Formation Games*, pages 487–516, Cambridge University Press, 2007.
- [11] B. Awerbuch, Y. Azar, and Y. Bartal. On-line generalized Steiner problem. *Theoretical Computer Science*, pages 313–324, vol. 324, no. 2–3, September 2004.
- [12] A. Agrawal, P. Klein, and R. Ravi. When Trees Collide: An Approximation Algorithm for the Generalized Steiner Problem on Networks. *SIAM Journal on Computing*, pages 440–456, vol. 24, no. 3, June 1995.
- [13] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proc. of ACM SIGCOMM 2002*, Pittsburgh, PA, USA, August 2002.
- [14] *AMPL: a modeling language for mathematical programming*. Available at <http://www.ampl.com>.
- [15] ILOG Optimization Products. ILOG CPLEX. <http://www.ilog.com/products/cplex/>.