A Game Theoretic Analysis of Network Design with Socially-Aware Users*

Jocelyne Elias[†], Fabio Martignon[‡], Konstantin Avrachenkov[§], Giovanni Neglia[§]

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.

*Preliminary results of this work have been presented in [1].

[†]Jocelyne Elias is with the Department of Electronics and Information, Politecnico di Milano. E-mail: elias@elet.polimi.it

[‡]Fabio Martignon is with the Department of Information Technology and Mathematical Methods, University of Bergamo, Dalmine (BG) 24044, Italy. E-mail: fabio.martignon@unibg.it

[§]Konstantin Avrachenkov and Giovanni Neglia are with the MAESTRO Project-team, INRIA Sophia Antipolis, 2004, route des Lucioles, B.P. 93 06902 Sophia Antipolis Cedex, France. E-mail: {Konstantin.Avratchenkov, Giovanni.Neglia}@sophia.inria.fr

1 Introduction

Network design with selfish users has been the focus of several recent works [2, 3, 4, 5, 6, 7], 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)[1]. Moreover, incentives could be introduced by some external authority (e.g., the overlay administrator) in order to increase the users' cooperation level.

In this work 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* [2] 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 [7]. 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 administrator) buys an appropriate subset of the network links (i.e., those belonging to the minimum cost generalized Steiner

[1] This observation motivates in [5] the study of *strong NE*, considering coalitions that could take decisions beneficial to all the members of the group.

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.

In summary, the main contributions of this paper are the following:

- the proposition of the Socially-Aware Network Design game, which combines both individual and social concerns in a unified and flexible manner.
- The determination of bounds on the Price of Anarchy, the Price of Stability and the Reachable Price of Anarchy of the proposed game.
- The proposition of a Stackelberg game where the network administrator leads the users to a system-wide efficient equilibrium by buying an appropriate subset of the network links.
- A thorough numerical evaluation of the proposed games in several realistic network scenarios, including real ISP topologies.

The paper is organized as follows: Section 2 discusses related work. Section 3 introduces the proposed Socially-Aware Network Design game. Section 4 proposes a greedy algorithm that implements best response dynamics, which permits to reach a Nash equilibrium in the proposed game. Section 5 provides precise bounds on the Price of Anarchy, the Price of Stability and the Reachable Price of Anarchy for the SAND game. Section 6 describes the proposed Stackelberg game (the Network Administrator-Driven SAND game), which enables very efficient Nash equilibria. Section 7 presents numerical results that demonstrate the effectiveness of the SAND and NAD-SAND games in several realistic network scenarios. Finally, Section 8 concludes this paper.

2 Related Work

Several recent works focused on network design with selfish users [2, 3, 4, 5, 6, 7].

The so-called *Shapley network design game* is proposed in [2]. In this game, each of the k players chooses a path from its source to its destination, and the overall network cost is shared among the players in the following way: each player pays for each edge a proportional share of the edge cost, i.e., the edge cost divided by the number of players that pass through such edge.

Of all the ways to share the social cost among the players, this proportional sharing method enjoys several desirable properties. First, it is budget balanced, in that it partitions the social cost among the players. Second, it can be derived from the Shapley value, and as a consequence is the unique cost-sharing method satisfying certain fairness axioms. Third, it admits pure strategy NEs. Specifically, Anshelevich et al.

showed in [2] that a pure-strategy Nash equilibrium always exists, that $PoA = k$ and $PoS = \mathcal{H}_k = \sum_{i=1}^{k} 1/i = O(\ln(k))$.

We observe that the Shapley network design game stimulates, to a certain extent, the cooperation among network users, since there is an incentive for players to share links in order to reduce their path cost. However, we will demonstrate in the Numerical Results section that introducing explicitly a socially-aware component in the users' objective functions (as we propose in this paper) increases consistently the level of cooperation among users, thus permitting to design more efficient and stable networks.

The network design model presented in [3] by Anshelevich et al. is general and does not admit pure Nash equilibria, even for very simple network instances. Briefly, in such model each player i has a set of terminal nodes that he must connect. A strategy of a player is a payment function p_i, where $p_i(e)$ is how much player i is offering to contribute to the cost of edge e. Any edge e such that $\sum_i p_i(e) \geq c_e$ is considered *bought* (c_e being the link cost). Each player i tries to minimize its total payments, $\sum_{e \in E} p_i(e)$, E being the edge set.

The authors in [4] have extended the network design model in [2], including weighted players. But while easy to define, this weighted network design game is challenging to analyze. In particular, it is shown in [4] that: (1) pure-strategy Nash equilibria exist in all weighted Shapley network design games with *two* players, and (2) there are no larger classes of weighted Shapley network design games that always possess pure-strategy Nash equilibria.

The works in [5, 6] study the existence of strong Nash equilibria (i.e., equilibria where no coalition can improve the cost of each of its members) in network design games under different cost sharing mechanisms. Strong Nash equilibria ensure stability against deviations by every conceivable coalition of agents. More specifically, the authors in [5] show that there are graphs that do not admit strong Nash equilibria, and then give sufficient conditions on the existence of approximate strong Nash equilibria.

Furthermore, the problem of designing a protocol that optimizes the equilibrium behavior of the induced network game is investigated in [7]. The authors study the design of optimal cost-sharing protocols for undirected and directed graphs, single-sink and multicommodity networks, different classes of cost-sharing methods, and different measures of the inefficiency of equilibria. Moreover, they provide upper and lower bounds on the best possible performance of non-uniform cost-sharing protocols.

Few works have considered individual and social concerns of user agents while dealing with different types of networking problems [9, 10].

A proposition that takes into account individual and social benefits has been considered in [9] in the general context of multi-agent systems, where individual and social concerns can conflict, leading to inefficient system performance. To address such problem, the authors have proposed a formal decision making framework, based on social welfare functions, that combines both social and individual perspectives.

An experimental investigation of the impact of cooperation in the context of routing games is conducted in [10]. The game is studied considering particular network topologies (i.e., parallel links and load balancing networks), shared by several users. Each user seeks to optimize either its own performance or some combination between its own performance and that of other users, by controlling the routing of its given flow demand.

However, unlike our work, none of the above papers [9, 10] applies these concepts to the network design game, nor provides a theoretical analysis of the efficiency of the achieved Nash equilibria.

3 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 $G = (V, E)$, with vertex set V and directed edge set E, where each edge e has a nonnegative cost c_e, and each player $i \in I = \{1, 2, \ldots, k\}$ 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 paths that go through 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:

$$J^i = \sum_{e\in S_i} \pi_e + \alpha \sum_{e\in\cup_j S_j} c_e. \tag{1}$$

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 (ILP) problem in order to optimize the routing choices for every change in the network (either in the topology, link costs, player locations etc ...), but can let the users solve the problem in a distributed way, converging to a good and stable solution. The social-aware cost can therefore be imposed, in practice, by the network administrator in order to stimulate the creation of cheaper networks, as we will detail in the application scenarios illustrated in Section 3.2.

Note that for $\alpha = 0$ objective function (1) corresponds to that of the Shapley network design game proposed in [2], 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.

The SAND game belongs to the class of *potential games*, which has been identified in [11]. Such games are characterized by a real-valued function, Φ, on the strategy vector (S) which measures exactly the difference in the cost that any player saves if he is the only player to deviate. Mathematically, if we define:

- $S = (S_1, S_2, \ldots, S_k)$ the vector of players' strategies,
- $S'_i \neq S_i$ an alternate strategy for some player i,
- S_{-i} the strategy played by all other players than i, and

- $S' = (S'_i, S_{-i})$ the vector of players' strategies where player i switches to alternate strategy S'_i, while all other players do not change their strategies S_{-i},

then a potential game with k players is characterized by a potential function, $\Phi(S)$, such that for any user i we have $\Phi(S) - \Phi(S') = J^i(S) - J^i(S')$.

Potential games have nice properties, such as existence of at least one pure Nash equilibrium, namely the strategy S that minimizes $\Phi(S)$ [12]. Furthermore, in such games, best response dynamics always converge to a Nash equilibrium.

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

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

We observe that the SAND game is a potential game without being at the same time a *congestion game* on the given graph G [13]. 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.)

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 5 shows that this is not necessarily the case. Before addressing this issue, we present in the following section a possible asynchronous users operation that implements best response dynamics and hence converges to a Nash equilibrium.

3.1 Comments

The SAND game can be easily extended to take into account performance metrics commonly used in routing problems in communication networks. For example, we can imagine that users are sensitive not only to link costs but also to the congestion level on each link e, which we can assume is directly proportional to the number of users (x_e) that pass through such link (like in *congestion games* [13]).

The extension can be performed adding a third term (a *congestion* cost) to cost function (1), which becomes as follows:

$$\widehat{J}^i = \sum_{e \in S_i} \pi_e + \alpha \sum_{e \in \cup_j S_j} c_e + \gamma \sum_{e \in S_i} \frac{x_e}{r_e} \tag{3}$$

where r_e is the maximum transmission rate of link e (i.e., the link's capacity), and γ is a weight that expresses the user's sensitivity to congestion costs. The rationale of adding $\gamma \sum_{e \in S_i} \frac{x_e}{r_e}$ to the cost function is that the congestion cost of using link e increases with the number of players that choose such link, since the allocated rate $(\frac{r_e}{x_e})$ decreases accordingly. Each player will therefore minimize its cost function (3), thus trying to avoid congested links.

The extended SAND game is still a potential game, with the following potential function:

$$\widehat{\Phi}(S) = \sum_{e \in E} \sum_{x=1}^{x_e} \left(\frac{c_e}{x} + \gamma \frac{x}{r_e} \right) + \alpha \sum_{e \in \cup_j S_j} c_e. \tag{4}$$

Therefore, the existence of pure Nash equilibria is guaranteed, as well as the convergence of best response dynamics.

We note that the best response algorithm illustrated in Section 4 can be extended straightforwardly to take into account the newly added congestion cost.

Finally, we observe that a natural extension of cost function (3) would be to consider users with different bandwidth demands (i.e., a sort of *weighted* game where players can have different requirements). However, it has been demonstrated in [14, 15] that with such extension the game is no longer a potential game, and as a consequence pure Nash equilibria may not exist. Indeed, in [14] the authors illustrate simple scenarios where no pure Nash equilibrium exists even with very few players and simple network topologies.

3.2 Application Scenarios

This section illustrates some notable application scenarios that can be envisaged for our proposed game. First of all, we underline that the cost sharing model proposed in the SAND game is general, and has received a significant attention in the networking community [2, 3, 4, 5, 6, 7]; as discussed in the Related Work section, these works have considered the specific cost sharing function also used in this paper.

Furthermore, our model directly applies to Service Overlay Networks (SONs) [16, 17, 18], which are application-layer networks built on top of the traditional IP-layer networks. In general, the SON is operated by a third-party overlay network operator that owns a set of overlay nodes residing in the underlying Internet Service Provider (ISP) domains. The SON operator must establish overlay links, purchasing them from the underlying ISPs.

In these networks, overlay operators provide a service to users through the creation of an overlay. A specific application, which will be detailed in the following, is selling connectivity with some specific guarantees.

In such scenario the overlay operator would buy bandwidth from the underlying ISPs and sell it to the users (e.g., companies) to connect their sites (for example, creating Virtual Private Networks). Then, the costs considered in the first term of cost function (1) are mainly the costs required for reserving some bandwidth from the underlying ISPs, increased by a given percentage to provide the overlay operator some revenues.

We can expect that negotiating with the underlying ISPs would have some serious constraints in terms of a) granularity of the purchased bandwidth [17] and b) the possibility to re-negotiate the contract on short time scales. Therefore, the overlay operator has interest to let the users share as much as possible the same overlay links in order to:

1. reduce the number of contracts it has to manage with the underlying ISPs, and hence the management costs,
2. take advantage of the multiplexing gain.

Therefore, instead of setting a fixed price for the connectivity service, the overlay operator could let users choose among different opportunities by making them pay a specific price that takes into account the real cost of the connection. A possibility is to make each user pay for each overlay link a share inversely proportional to the number of users that pass through such link. The interactions among selfish users in such scenario are captured by the Shapley network design game [2], which is the starting point of our work.

As we show in the Numerical Results section, however, the Shapley network design game can lead to relatively poor performance. Hence the overlay operator cannot find convenient to set prices according to this kind of allocation, but it can determine the price for each user including a given percentage (α) of the total network cost. This new pricing scheme is adopted in our game, whose formulation can correspond to partially socially-aware players (as the title of our paper suggests) or to selfish players that are trying to minimize a cost decided by the overlay operator as described above.

Note that the basic formulation of the Shapley network design game proposed in [2] (and also of our model that is built on it) does not capture a characteristic of the scenario we are considering. First, the cost of the overlay link is assumed to be constant, independently of the number of users, while probably the overlay operator should buy more bandwidth from the ISP the larger the number of users. However, due to the granularity issue we mentioned above, the overlay link cost is likely an increasing concave step function of the number of users. Hence, when a new user joins the system, it would probably not cause any cost increase if it uses existing links. Clearly, it would be nice to explicitly have in the model such a cost function, but this would make the general analysis more complex without bringing more insights in the basic problem.

We observe that we implicitly assumed that users are homogeneous in that they require the same bandwidth. This is not a real limitation: in fact, in a realistic scenario, also the overlay operator will provide the bandwidth with a given granularity (e.g., multiple of Mbps), even if much finer than the underlying ISPs. Hence, as a first approximation, a user requiring n Mbps can be considered as n users that independently try to choose n paths. In a large deployed system (with already $N \gg n$ users), there is not much difference in the decision of n users requiring 1 Mbps and one user requiring n Mbps.

Other application scenarios of our proposed model include large computer networks, like the Internet, which are built, operated and used by a large number of diverse and competitive entities. In this context, several independent service providers (Autonomous Systems or Internet Service Providers) seek to selfishly optimize the quality and cost of their own operation. Our work can provide a game theoretic framework for modeling such interests and the networks they generate, even when some sort of incentives to cooperation are introduced by external authorities (network administrators, governments etc ...).

Finally, note that the congestion component in equation (3), $\gamma \sum_{e \in S_i} \frac{x_e}{r_e}$, enables our extended SAND game to model independent routing decisions performed by selfish entities, which is an important aspect in large telecommunications networks (like the Internet) with no central authority.

4 Best Response Algorithm

We now describe a simple algorithm that implements best response dynamics, allowing each user to improve its cost function in the proposed SAND game. Such algorithm, detailed in the following, is the best response strategy for a user minimizing objective function (1), assuming other users are not changing their strategies.

We still consider a directed graph $G = (V, E)$, where each edge e has a nonnegative cost c_e. Let $S = (S_1, S_2, \ldots, S_k)$ be a vector of players' strategies, where $S_i \subset E$ is a set of edges that connect the source/sink pair (s_i, t_i). Let us denote by $L^{-i} = \cup_{j \neq i} S_j$ the set of edges used by all players except player i; as a consequence, the set $E - L^{-i}$

contains all links that are *not* used by the set of all players except i. Finally, let k_e denote the number of times that edge e lies in the paths chosen by such players (i.e., all players except player i).

Algorithm 1 allows user i to determine the *best response* to all other users' strategies. The rationale behind such algorithm is very simple: the link weights are set equal to the cost incurred by user i to choose them, according to objective function (1).

More specifically, if link e has already been selected by k_e players other than i (i.e., $e \in L^{-i}$, using the above notation), then the additional cost incurred by player i in using such link will be equal to $Cost(e) = \frac{c_e}{k_e+1}$, since now there will be $k_e + 1$ players that pass through such link, and therefore player i will be charged $\frac{c_e}{k_e+1}$ for using it. Since in this case link e is already selected by k_e other players, the choice of such link does not contribute to increase the second term of the cost perceived by user i ($\alpha \sum_{e \in \cup_j S_j} c_e$), which remains the same.

On the other hand, if link e is not selected by any player (i.e., $e \notin L^{-i}$), then if user i chooses such link, it will incur an additional cost equal to $\frac{c_e}{1} + \alpha c_e = c_e \cdot (1 + \alpha)$, since it will be the only player using link e, and hence it will pay for its whole cost c_e; furthermore, it will increase the overall network cost by c_e (and hence its perceived social cost by αc_e), since now such link will be included in the formed network.

Then, the shortest path computed in step 6 corresponds to the choice that minimizes the user's cost, and it represents therefore its best response. Note that, if multiple shortest paths exist with the same cost, the player chooses one of them randomly.

Algorithm 1 Pseudo-code specification for the best response dynamics for the SAND game

---

1: **if** $e \in L^{-i}$ **then**
2:     Set $Cost(e) = \frac{c_e}{k_e+1}$
3: **else**
4:     Set $Cost(e) = c_e \cdot (1 + \alpha)$
5: **end if**
6: Compute the least-cost path (using for example the Dijkstra algorithm) with link costs $Cost(e)$
7: Set S_i equal to the set of links contained in the least-cost path

---

5 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 [2] for the Shapley network design game. This allows us to determine the worst and best case performance of our proposed game.

5.1 Bound on the Price of Anarchy

We now 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:*

$$PoA \geq k(1+\alpha). \tag{5}$$

Proof: Let us consider the simple network scenario of Figure 1 with two parallel links, one of cost equal to 1, the other with an arbitrarily high cost equal to C. Each of the k players must connect the common source node s to the common destination node t. In the optimal outcome, each player chooses the lower link, with cost 1, and the cost of the formed network is obviously 1.

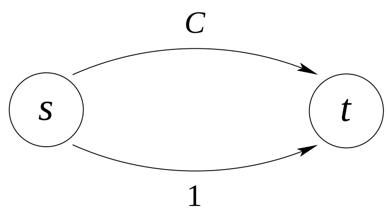

Figure 1: Two-link network topology: k players must connect the common source node s to destination node t.

However, this is not the only Nash equilibrium for the SAND game. Let us suppose that the initial network configuration sees all k users routed over the upper link, with cost C. It is easy to show that if $\alpha \geq \frac{C}{k} - 1$, no user has a gain to deviate and choose the link with cost equal to 1. Then, the cost of the network is C. Now if $C = k(1+\alpha)$, the above inequality is satisfied, and we obtain that the Price of Anarchy is at least $\frac{C}{1} = k(1+\alpha)$.

□

5.2 Bound on the Price of Stability

In the following 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:*

$$PoS \leq \frac{\mathcal{H}_k + \alpha}{1+\alpha}. \tag{6}$$

Proof: It is shown in [2, 12] that the function $\Psi(S) = \sum_{e \in E} \sum_{x=1}^{x_e} \frac{c_e}{x}$ is an exact potential function for the Shapley network design game revised in Section 2. Furthermore, it is shown in [12] that such function satisfies the following inequalities:

$$\text{cost}(S) \leq \Psi(S) \leq \mathcal{H}_k \text{cost}(S) \tag{7}$$

where $\text{cost}(S) = \sum_{e \in \cup_j S_j} c_e$ and $\mathcal{H}_k = \sum_{i=1}^{k} 1/i$ is the kth harmonic number.

The potential function of the SAND game, $\Phi(S)$, is directly related to that of the Shapley network design game, $\Psi(S)$. In fact, from expression (2) it can be easily seen that $\Phi(S) = \Psi(S) + \alpha \cdot \text{cost}(S)$. Hence, if we replace $\Psi(S) = \Phi(S) - \alpha \cdot cost(S)$ in expression (7), it follows that:

$$(1+\alpha)\text{cost}(S) \leq \Phi(S) \leq (\mathcal{H}_k + \alpha)\text{cost}(S). \tag{8}$$

Finally, Theorem 19.13 in [12] states that if a potential game with potential function $\Phi(S)$ satisfies the following inequality:

$$\frac{\text{cost}(S)}{A} \leq \Phi(S) \leq B\text{cost}(S) \tag{9}$$

with A and B positive constants, then the Price of Stability (PoS) is at most AB.

Therefore, in our problem, we obtain that the Price of Stability is at most $\frac{\mathcal{H}_k + \alpha}{1+\alpha}$.

□

5.3 Bound on the Reachable Price of Anarchy

We now derive a bound on the so-called Reachable Price of Anarchy ($RPoA$), a quantity defined in [7]. 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 for example the best response algorithm illustrated in the previous section). 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:*

$$RPoA < k\frac{\alpha+1}{\alpha+\frac{1}{k}}. \tag{10}$$

Proof: Let S_i^{opt} define the path from source node s_i to destination node t_i for player i in the optimal outcome (the least cost network), and let $\overline{S_i}$ denote the path chosen by player i in its first move in the game. Let $C(S_i^{opt}) = \sum_{e \in S_i^{opt}} c_e$ and $C(\overline{S_i}) = \sum_{e \in \overline{S_i}} c_e$ denote the costs of such two paths.

For the first player that enters the network, the following inequality holds:

$$\sum_{e \in \overline{S_1}} \pi_e + \alpha \sum_{e \in \overline{S_1}} c_e = (1+\alpha)C(\overline{S_1}) < (1+\alpha)C(S_1^{opt}).$$

For the second player, the following inequality holds:

$$\sum_{e \in \overline{S_2}} \pi_e + \alpha C(\overline{S_1} \cup \overline{S_2}) < \sum_{e \in S_2^{opt}} \pi_e + \alpha C(\overline{S_1} \cup S_2^{opt})$$

which can also be expressed as follows, observing that $C(\overline{S_1} \cup S_2^{opt}) < C(\overline{S_1}) + C(S_2^{opt})$, and that $\sum_{e \in S_2^{opt}} \pi_e \leq C(S_2^{opt})$

$$\sum_{e \in \overline{S_2}} \pi_e + \alpha C(\overline{S_1} \cup \overline{S_2}) < (1+\alpha) C(S_2^{opt}) + \alpha C(S_1^{opt}).$$

It can be further observed that:

$$\sum_{e \in \overline{S_i}} \pi_e \geq \frac{1}{k} C(\overline{S_i}),$$

since in the best outcome, each of the links belonging to the path chosen by player i is shared by all k players.

If we proceed iteratively, we finally obtain the following inequality:

$$(\alpha + \frac{1}{k}) C(\bigcup_{i=1}^{k} \overline{S_i}) < (1+\alpha) \sum_{i=1}^{k} C(S_i^{opt})$$

which can be rewritten as follows, observing that $\sum_{i=1}^{k} C(S_i^{opt}) \leq k C^{opt}$:

$$(\alpha + \frac{1}{k}) C(\bigcup_{i=1}^{k} \overline{S_i}) < (1+\alpha) k C^{opt}.$$

This allows us to obtain the following result:

$$\frac{C(\bigcup_{i=1}^{k} \overline{S_i})}{C^{opt}} < k \frac{\alpha + 1}{\alpha + \frac{1}{k}}. \tag{11}$$

Hence, the cost of the network reached after the first move of each of the k players, $C(\bigcup_{i=1}^{k} \overline{S_i})$, is no more than $k\frac{\alpha+1}{\alpha+\frac{1}{k}}$ times the optimal cost, C^{opt}. Since subsequent best response moves performed by players decrease the potential function value $\Phi(S)$ given by expression (2) (simulating a local search on it), the players will converge to a Nash equilibrium whose cost is still within $k\frac{\alpha+1}{\alpha+\frac{1}{k}}$ times the optimum.

Therefore, we can conclude that the Reachable Price of Anarchy is at most $k\frac{\alpha+1}{\alpha+\frac{1}{k}}$.

□

5.4 Comments

For $\alpha = 0$ the SAND game is equivalent to the original Shapley network design game. Indeed, expressions (5) and (6) confirm the results already demonstrated in [2] : $PoA = k$ 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 $\alpha \to \infty$ (then also $PoS \to 1$). This can be easily explained since for $\alpha \to \infty$ 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. This happens for example in the simple instance of Figure 1. 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 (10) that the $RPoA$ of the SAND game strictly decreases for increasing α values. For large α values (notably, for $\alpha \to \infty$), the $RPoA$ is less than k. Therefore, if the SAND game is played starting from an empty network, worst-case scenarios like that illustrated in Figure 1 are eliminated.

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 7.

6 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. This is the typical scenario of a Stackelberg game [19, 20] where the network administrator acts as a *leader*, which imposes his strategy on the self-optimizing users that behave as *followers*.

The NAD-SAND game occurs in the same scenario illustrated in Section 3 for the SAND game, i.e. in a directed graph $G = (V, E)$, with vertex set V and directed edge set E, where each edge e has a nonnegative cost c_e, and each player $i \in I = \{1, 2, \ldots, k\}$ is identified with a source/sink pair (s_i, t_i).

In the NAD-SAND game, the Network Administrator plays first, choosing a subset of network links (referred to as $E^{opt} \subseteq E$) for which he pays an equal share of their cost, thus providing an incentive for all other "ordinary" players to choose them. The goal is to stimulate all other players to build an efficient and stable network.

Then, each player plays exactly the SAND game as described in Section 3, picking a path from its source node s_i to its destination t_i, minimizing its objective function J^i, given in expression (1).

Since computing the optimal Stackelberg strategy for the Network Administrator 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 [21] as detailed below in Section 6.1, 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 3, 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 exactly as described in Section 3, 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, i.e. finding the least-cost network topology that connects all source/destination pairs (s_i, t_i). However, several efficient heuristics and approximation algorithms have been proposed to solve such problem in a reasonable computation time [21, 22, 23]. 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 illustrated in Section 6.1, and solving it with the CPLEX 11 solver [24].

Finally, we observe that as $\alpha \to \infty$, 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, $PoA \to 1$ as $\alpha \to \infty$.

6.1 Minimum cost generalized Steiner Tree: an exact ILP model

Hereafter we illustrate a simple ILP model that minimizes the total network cost while assuring full connection of all source/destination pairs $(s_i, t_i), \forall i \in I$.

Let d_n^i, $i \in I, n \in V$ be a parameter denoting source/destination nodes for player i:

$$d_n^i = \begin{cases} +1 & \text{if node } n \text{ is the source node of player } i \text{ (i.e., } n = s_i) \\ -1 & \text{if node } n \text{ is the destination node of player } i \text{ (i.e., } n = t_i) \\ 0 & \text{otherwise} \end{cases}$$

Decision variables of the problem include flow variables x_{nm}^i, $i \in I, (n,m) \in E$:

$$x_{nm}^i = \begin{cases} 1 & \text{if player } i \text{ chooses link } (n,m) \\ 0 & \text{otherwise} \end{cases}$$

and link utilization variables y_{nm}, $(n,m) \in E$:

$$y_{nm} = \begin{cases} 1 & \text{if at least one player passes through link } (n,m) \\ 0 & \text{otherwise} \end{cases}$$

Finally, let c_{nm} denote the cost of link (n,m).

Given the above definitions, the minimum cost generalized Steiner Tree problem can be stated as follows:

$$\min \sum_{(n,m)\in E} y_{nm} \cdot c_{nm} \tag{12}$$

$$s.t. \sum_{m\in V,(n,m)\in E} x^i_{nm} - \sum_{m\in V,(m,n)\in E} x^i_{mn} = d^i_n \qquad \forall n \in V, i \in I \tag{13}$$

$$y_{nm} \geq x^i_{nm} \qquad \forall (n,m) \in E, i \in I \tag{14}$$

$$x^i_{nm}, y_{nm} \in \{0,1\} \qquad \forall i \in I, (n,m) \in E \tag{15}$$

The objective function (12) minimizes the total network cost. Constraints (13) impose the flow conservation in node n for the traffic of each player i; these constraints are the same as those adopted for classical multicommodity flow problems. Constraints (14) are coherence constraints imposing that $y_{nm} = 1$ if at least one player uses link (n,m), i.e. if at least one of the x^i_{nm} variables is equal to 1. Finally, constraints (15) are the integrality constraints for the binary decision variables.

6.2 Bound on the Price of Anarchy

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:*

$$PoA \geq \frac{k}{2(1+\alpha)} \tag{16}$$

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

Proof: Let us consider the network scenario illustrated in Figure 2, where each of the k players must connect the source node s_i to the common destination node t. Link $\nu \to t$ has a cost equal to C, with $2 \leq C \leq k$.

The minimum cost generalized Steiner Tree is in this case formed by all links $s_i \to \nu$ (with cost zero) and by link $\nu \to t$, so that its cost is equal to C. Hence, the network administrator will choose all these links, which constitute the E^{opt} set.

In the optimal outcome, each player i chooses the $s_i \to \nu \to t$ path, and the cost of the formed network is C.

The NAD-SAND game, however, admits another Nash equilibrium. Let us suppose that initially all k users are routed over the direct path, $s_i \to t$, with total cost k. It can be easily seen that if $\alpha \leq \frac{C}{2} - 1$, no user has a gain to deviate and choose the $s_i \to \nu \to t$ path, with cost equal to C.

Then, the cost of the network is k. Now if $C = 2(1+\alpha)$, the above inequality is satisfied, and we obtain that the Price of Anarchy is at least $\frac{k}{C} = \frac{k}{2(1+\alpha)}$.

□

7 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 several

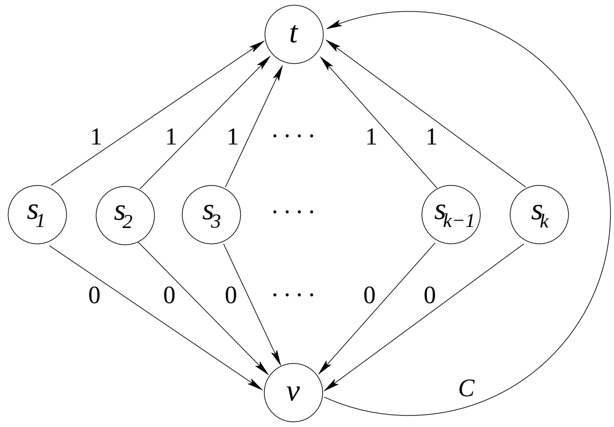

Figure 2: Parallel link topology: bound on the Price of Anarchy for the NAD-SAND game.

network scenarios, comparing them both to the Shapley network design game [2] and to the social optimum. To this end, we consider both randomly generated network instances and real ISP topologies mapped by the Rocketfuel tool [25, 26]. Random networks are obtained using a custom generator as well as a degree-based generator (BRITE [27, 28]) to create topologies with a power law distribution of node degrees.

In all cases, we start from the empty network and apply iteratively the best response algorithm illustrated in Section 4, 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 [21] with the ILP model illustrated in Section 6.1, using AMPL [29], and solving it with the CPLEX 11 solver [24].

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 5.

7.1 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 1. Furthermore, Figure 3 illustrates more precisely the total network cost achieved by the SAND and NAD-SAND games for $\alpha \in [0, 10]$, thus showing in detail the performance improvements obtained in such range.

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 ($\geq$ 5), the NAD-SAND game reaches exactly the optimum (i.e., the generalized Steiner Tree cost).

We observe that for large α values (i.e., for $\alpha \geq 50$) the NAD-SAND game reaches the optimum, while the SAND game achieves no further cost improvement. This behavior was observed in all the considered network scenarios, and is due to the fact that for such α values the selfish component of objective function (1), $\sum_{e \in S_i} \pi_e$, becomes negligible with respect to the socially-aware component, $\alpha \sum_{e \in \cup_j S_j} c_e$.

Table 1: Full-Mesh topology with 50 nodes randomly distributed on a 1000 $\times$ 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	

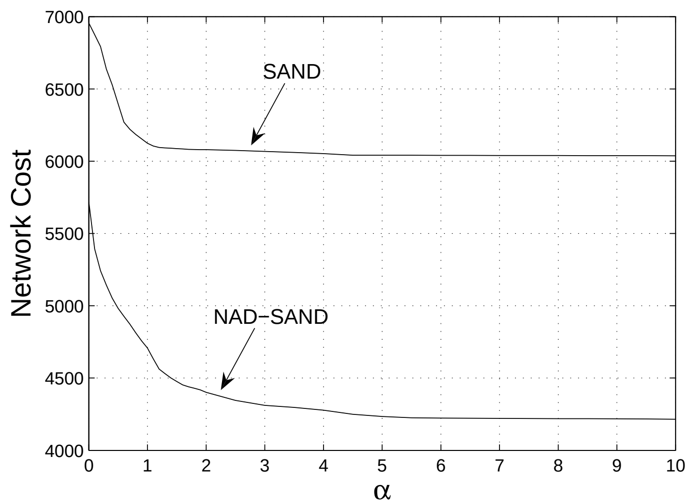

Figure 3: Full-Mesh topology with 50 nodes randomly distributed on a 1000 $\times$ 1000 area and 20 players: average network costs for the SAND and the NAD-SAND games (detail for $\alpha \in [0, 10]$).

7.2 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.

Figure 4 shows the networks resulting at the Nash equilibria with the SAND game in a random geometric graph scenario with $N = 50$ nodes, range $R = 500$, 20 source/destination pairs, and two different α values (viz., $\alpha = 0$ and $\alpha = 50$).

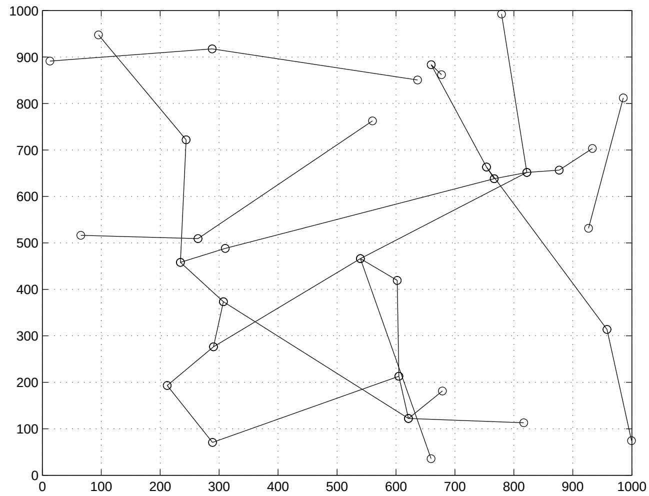

(a) $\alpha = 0$

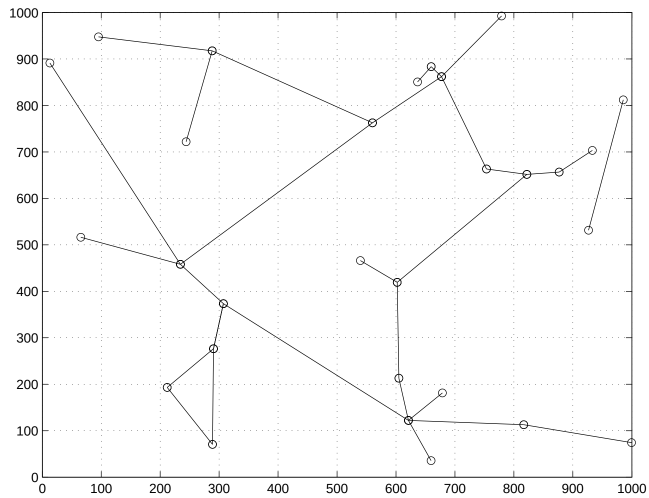

(b) $\alpha = 50$

Figure 4: Nash equilibria obtained by the SAND game in a random geometric network with 50 nodes, $R = 500$, 20 source/destination pairs, and different α values ($\alpha = 0$ and $\alpha = 50$).

We observe that for $\alpha = 50$ the topology is much closer to a tree-like topology than that obtained by the Shapley network design game ($\alpha = 0$). This is reflected in the total network cost, which is equal to 7046.3 for $\alpha = 0$ and to 5336.2 for $\alpha = 50$, thus resulting in a gain of more than 24%.

Table 2 illustrates the results obtained in the same random network scenario of Figure 4, 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.

Figure 5 shows the same numerical results, focusing on the $\alpha \in [0, 10]$ range, illustrating more precisely the performance improvements that occur owing to the socially-aware component. Also in this scenario it can be observed that the most notable cost reductions are mainly obtained for $\alpha \leq 5$.

Table 3 shows the numerical results obtained using Random network model "Uniform" in networks with 100 nodes, $L = 2000$ links, maximum link cost $C = 100$ and 20 source/destination pairs. The total network costs are illustrated in the table for both the SAND and NAD-SAND games, while the ILP column reports the optimal network cost. Although there is still room for improvement, in this scenario the SAND game approaches the ILP bound, and consistently improves the quality of network equilibria with respect to the $\alpha = 0$ case.

Also in this case, the NAD-SAND game performs better than the SAND game, lowering the overall network cost of more than 11%, reaching the socially optimal outcome for $\alpha \geq 50$.

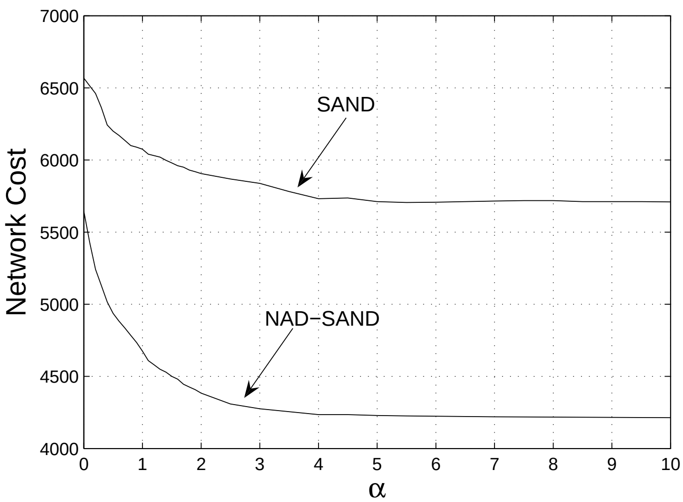

Figure 5: Random geometric graphs; random networks with 50 nodes, $R = 500$ and 20 players: average network costs for the SAND and the NAD-SAND games (detail for $\alpha \in [0, 10]$).

Table 2: 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 3: Random network model "Uniform"; random networks with 100 nodes, 2000 links, maximum link cost $C = 100$, 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	314.69	294.18	296.38	295.09	295.57	296.48	**263.25**
NAD-SAND	295.85	274.38	263.70	263.25	263.25	263.25	

7.3 Power Law topologies

We further considered BRITE, a degree-based topology generator [27, 28], which was used to generate power law topologies; the Barabási − Albert model [30] was adopted with default parameters provided by BRITE.

We generated two different power-law network scenarios: one with 100 nodes and average node degree equal to 3, which contained about 300 links; the other with 300 nodes and average node degree equal to 5, which contained about 1500 links. The cost of each link was set equal to its length.

Table 4 reports the corresponding numerical results, including the optimal network cost obtained solving the ILP model. In this case, the Nash equilibria reached by the Shapley network design game are already close to the optimum. This is mainly due to the fact that power law topologies, which result in short characteristic path lengths and heavy clustering, are the result of social concerns that eventually bring efficiency to the network design (even though not explicitly expressed in the underlying routing protocols design).

The SAND and NAD-SAND games, however, permit to improve the efficiency of such equilibria, finally obtaining the minimum cost network provided by the ILP model.

7.4 Real ISP topologies

Finally, we consider three real ISP topologies mapped using Rocketfuel [25, 26], listed in Table 5, with an increasing number of nodes and links. Note that Rocketfuel does not provide link capacities for these topologies, thus preventing the application of the extended SAND model illustrated in Section 3.1.

Table 6 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

Table 4: Power law topologies generated using BRITE: average network costs for the SAND and the NAD-SAND games. The optimal network cost is also reported.

	100 nodes						
Game	$\alpha = 0$	$\alpha = 1$	$\alpha = 10$	$\alpha = 50$	$\alpha = 100$	$\alpha = 1000$	**ILP**
SAND	15980.20	15684.28	15596.31	15593.69	15593.43	15602.21	**15314.14**
NAD-SAND	15479.98	15314.14	15314.14	15314.14	15314.14	15314.14	

	300 nodes						
Game	$\alpha = 0$	$\alpha = 1$	$\alpha = 10$	$\alpha = 50$	$\alpha = 100$	$\alpha = 1000$	**ILP**
SAND	22580.82	22312.27	22224.65	22210.97	22210.55	22210.55	**21927.41**
NAD-SAND	22052.06	21987.68	21927.41	21927.41	21927.41	21927.41	

Table 5: Rocketfuel-inferred ISP topologies: number of network nodes and links.

Network	Location	Nodes	Links
Telstra	AU	108	306
Sprintlink	US	141	748
Abovenet	US	315	1944

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.

Table 6: Rocketfuel topologies, 20 source/destination pairs: average network costs for the SAND and the NAD-SAND games. The optimal network costs are also reported.

Network	Game	$\alpha = 0$	$\alpha = 1$	$\alpha = 10$	$\alpha = 50$	$\alpha = 100$	$\alpha = 1000$	**ILP**
Telstra	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	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	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	

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.

8 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. 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] J. Elias, F. Martignon, K. Avrachenkov, and G. Neglia. Socially-Aware Network Design Games. In *Proc. of the 29th IEEE Conference on Computer Communications (INFOCOM'2010 mini-conference)*, San Diego, CA, USA, March 2010.

[2] 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.

[3] 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.

[4] H. L. 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.

[5] 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.

[6] 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.

[7] 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.

[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, March 4-6, 1999.

[9] L. Hogg and N. R. Jenning. Variable sociability in agent-based decision making. *Lecture Notes in Computer Science*, pages 305–318, vol. 1757, 1999.

[10] A. P. Azad, E. Altman, and R. El-Azouzi. From Altruism to Non-Cooperation in Routing Games. In *Proc. of Networking and Electronic Commerce Research Conference 2009 (NAEC 2009)*, Lake Garda, Italy, October 8-11, 2009.

[11] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, pages 124–143, vol. 14, no. 1, May 1996.

[12] 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.

[13] 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.

[14] D. Fotakis, S. Kontogiannis, and P. Spirakis. Selfish Unsplittable Flows. *Theoretical Computer Science*, pages 226–239, vol. 348, no. 2-3, December 2005.

[15] H. Ackermann, H. Roglin, and B. Vocking. Pure Nash Equilibria in Player-Specific and Weighted Congestion Games. *Lecture Notes in Computer Science*, pages 50–61, vol. 4286, November 2006.

[16] Z. Duan, Z.-L. Zhang, and Y.T. Hou. Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning. *IEEE/ACM Transactions on Networking*, pages 870–883, vol. 11, no. 6, December 2003.

[17] H.T. Tran and T. Ziegler. A design framework towards the profitable operation of service overlay networks. *Computer Networks*, pages 94–113, vol. 51, 2007.

[18] A. Capone, J. Elias, and F. Martignon. Models and Algorithms for the Design of Service Overlay Networks,. *IEEE Transactions on Network and Service Management*, 5(3):143–156, September 2008.

[19] Y. A. Korilis, A. A. Lazar, and A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, pages 161–173, vol. 5, no. 1, February 1997.

[20] R. B. Myerson. Game Theory: Analysis of Conflict. Cambridge, Massachussets: Harvard University Press, 1997.

[21] 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.

[22] 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.

[23] M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. In *Proc. of the 3rd Annual ACM-SIAM Symposium on Discrete algorithms*, pages 307–316, Orlando, Florida, USA, January 1992.

[24] ILOG Optimization Products. ILOG CPLEX. http://www.ilog.com/products/cplex/.

[25] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proc. of ACM SIGCOMM 2002*, Pittsburgh, PA, USA, August 2002.

[26] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson. Inferring link weights using end-to-end measurements. In *Proc. of the 2nd Internet Measurement Workshop*, Marseille, France, November 2002.

[27] A. Medina, A. Lakhina, I. Matta, and J. Byers. BRITE: an approach to universal topology generation. In *Proc. of the 9th International Workshop on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2001)*, Cincinnati, OH, USA, August 2001.

[28] A. Medina, I. Matta, and J. Byers. On the origin of power-laws in Internet topologies. *ACM Communications Review*, pages 18–28, vol. 30, no. 2, April 2000.

[29] *AMPL: a modeling language for mathematical programming*. Available at http://www.ampl.com.

[30] A. L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, pages 509–512, vol. 286, no. 5439, 1999.