

Towards Network Games with Social Preferences

Petr Kuznetsov and Stefan Schmid

TU Berlin / Deutsche Telekom Laboratories, D-10587 Berlin, Germany

Abstract. Many distributed systems can be modeled as *network games*: a collection of *selfish* players that communicate in order to maximize their individual utilities. The performance of such games can be evaluated through the costs of the system *equilibria*: the system states in which no player can increase her utility by unilaterally changing her behavior. However, assuming that all players are selfish and in particular that all players have the same utility function may not always be appropriate. Hence, several extensions to incorporate also altruistic and malicious behavior in addition to selfishness have been proposed over the last years. In this paper, we seek to go one step further and study arbitrary relationships between participants. In particular, we introduce the notion of the *social range matrix* and explore the effects of the social range matrix on the equilibria in a network game. In order to derive concrete results, we propose a simplistic network creation game that captures the effect of social relationships among players.

1 Introduction

Many distributed systems have an open clientele and can only be understood when taking into account socio-economic aspects. A classic approach to gain insights into these systems is to assume that all players are selfish and seek to maximize their utility. Often, the simplifying assumption is made that all players have the same utility function. However, distributed systems are often “socially heterogeneous” whose participants run different clients and protocols, some of which may be selfish while others may even try to harm the system. Moreover, in a social network setting where members are not anonymous, some players may be friends and dislike certain other players. Thus, the state and evolution of the system depends on a plethora of different utility functions. Clearly, the more complex and heterogeneous the behavior of the different network participants, the more difficult it becomes to understand (or even predict) certain outcomes.

In this paper, we propose a more general approach to model the players’ utilities and introduce a social range matrix. This matrix specifies the *perceived* costs that are taken into account by the players when choosing a strategy. For example, a player who maliciously seeks to hamper the system performance has a perceived cost that consists of the negative costs of the other players. On the other hand, an altruistic player takes into account the costs of all other players

and strives for a socially optimal outcome. There are many more player types in-between that care about some players but dislike others.

In order to gain insights into the implications of different social ranges, we consider a novel network creation game that captures the willingness of a group of people to connect to each other. In this game, players do not incur infinite costs if they are not connected to some players. Rather, the utility of a player is given by the number of other players in her R -neighborhood, for some parameter R . For instance, in a game with $R = 1$, players can only collaborate with and benefit from their direct neighbors. Or imagine a peer-to-peer network like Gnutella where files are searched by flooding up to a certain radius (e.g., a time-to-live of $R = TTL = 5$), then a player is mainly interested in the data shared in her 5-hop neighborhood. Our motivation in using this model stems from its simplicity which allows to exemplify and quantify the effects of different social matrices.

1.1 Related Work

Over the last years, several models for distributed systems have been proposed that go beyond purely selfish settings. For instance, security and robustness related issues of distributed systems have been an active field of research, and malicious faults are studied intensively (e.g., [5,14]). To the best of our knowledge, the first paper to study equilibria with a malicious player is by Karakostas and Viglas [13] who consider a routing application where a single malicious player uses his flow through the network in an effort to cause the maximum possible damage. In order to evaluate the impact of such malicious behavior, a coordination ratio is introduced which compares the social costs of the worst *Wardrop equilibrium* to the social costs of the best *minimax saddle-point*. In [9], implementation problems are investigated with k faulty players in the population, but neither their number nor their identity is known. A planner's objective then is to design an equilibrium where the non-faulty players act according to her rules. Or in [1], the authors describe an asynchronous state machine replication protocol which tolerates Byzantine, Altruistic, and Rational behavior. Moscibroda et al. [17] discovered the existence of a so-called *fear factor* in the virus inoculation game where the presence of malicious players can improve the social welfare under certain circumstances. This windfall of malice has subsequently also been studied in the interesting work by Babaioff et al. [4] on congestion games.

There exists other work on game theoretic systems in which not every participating agent acts in a rational or malicious way. In the *Stackelberg theory* [19], for instance, the model consists of selfish players and players that are *controlled by a global leader*. The leader's goal is to devise a strategy that induces an optimal or near optimal so-called Stackelberg equilibrium. Researchers have recently also been interested in the effects of altruism that co-exists with selfishness [12,15]. For example, Meier et al. [15] have shown (for a specific game played on a social network) that friendship among players is always beneficial compared to purely selfish environments, but that the gain from social ties does not grow monotonically in the degree of friendship.

In contrast to the literature discussed above, we go one step further and initiate the study of games where players can be embedded in arbitrary social contexts and be selfish towards certain players, be friends with some other players, and even have enemies.

In particular, we apply our framework to a novel network creation game (for similar games, see the connection games described in Chapter 19.2 of [18]). Network creation has been a “hot topic” for several years. The seminal work by Fabrikant et al. [10] in 2003 seeks to shed light on the Internet’s architecture as built by economic agents, e.g., by Internet providers or *autonomous systems*. Recent subsequent work on network creation in various settings includes [2,3,6,7,8]. Moscibroda et al. [16] considered network creation games for peer-to-peer systems. The game proposed in our paper here can be motivated by peer-to-peer systems as well. However, in contrast to [16] where peers incur an infinite cost if they are not all connected to each other, we believe that our model is more appropriate for unstructured peer-to-peer systems.

The notion of interpersonal influence matrix, similar to our social range matrices, is used in sociology for understanding the dynamics of interpersonal agreement in a group of individuals (see, e.g., [11]).

1.2 Our Contributions

The main contribution of this paper is the introduction and initial study of the social range matrix which allows us to describe arbitrary social relationships between players. For instance, social range matrices can capture classic *anarchy* scenarios where each player is selfish, *monarchy* scenarios where players only care about one network entity, or *coalitions* that seek to support players within the same coalition but act selfishly or maliciously towards other coalitions. Despite this generality, we are able to derive interesting properties of such social matrices. For instance, we show that there are matrix transformations that do not affect the equilibria points (and the convergence behavior) of a game.

In addition, as a case study, we analyze a simplistic social network creation game where players can decide to which other players they want to connect. While a new connection comes at a certain cost, a player can also benefit from her neighborhood. That is, we assume that the players’ utility is given by the number of other players they are connected to up to a certain horizon, minus the cost of the links they have to pay for. For example, this game can be motivated by unstructured peer-to-peer systems where data is usually searched locally (in the peers’ neighborhood) and overall connectivity is not necessarily needed. We focus on this game due to its simplicity that allows us to study the main properties of the social matrix and exemplify the concepts. For example, in a social context where players can choose their neighbors, it is likely that players will connect to those players who they are friends with. We will show that this intuition is correct and that social relationships are indeed often reflected in the resulting network. As another example, we show that the social welfare of monarchic societies can be higher than that of anarchic societies if the price of establishing a connection is relatively low; otherwise, the welfare is lower.

Our new model and the network creation game open a large number of research directions. We understand our work as a first step in exploring the effect of social ranges on the performance of network games and use this paper to report on our first insights.

1.3 Paper Organization

The rest of the paper is organized as follows. We describe our model and formally introduce the social range matrix in Section 2. Section 3 presents our first insights on the properties of a social range matrix. We then report on our case study on social network formation (Section 4). The paper is concluded with a brief discussion and an outlook on future research directions in Section 5.

2 Social Range Matrices and Perceived Equilibria

In this section, we introduce the concept of a game theory where players are embedded in a social context; in particular, we define the social range matrix F describing for each player i how much she cares about every other player j .

We consider a set Π of n players (or nodes), $\Pi = \{0, \dots, n-1\}$. Let \mathcal{X}_i be the set of possible strategies player i can pursue in a given game \mathcal{G} . A *strategy profile* $s \in \mathcal{X}_0 \times \dots \times \mathcal{X}_{n-1}$ specifies a configuration, i.e., s is the vector of the strategies of all players.

The cost that actually arises at a player i in a given strategy profile s is described by its *actual* cost function $c_a(i, s)$. However, depending on the social context a player is situated in, it may experience a different *perceived cost* $c_p(i, s)$: While a purely selfish player may be happy with a certain situation, another player that cares about the actual costs of her friends may have a higher perceived cost and may want to change her strategy to a socially better one. (Note, however, that the distinction between “purely selfish players” and players that take into account the utility or cost of other players is artificial: Players whose action depends on other players’ utilities can be considered “purely selfish” as well, and simply have a different cost function.)

Formally, we model the perceived costs of a given player as a linear combination of the actual costs of all other players in the game. The *social range* of player i is a vector $f_i = (f_{i0}, \dots, f_{i(n-1)}) \in \mathbb{R}^n$. Intuitively, f_{ij} quantifies how much player i cares about player j , in both a positive (if $f_{ij} > 0$) and a negative way ($f_{ij} < 0$). $f_{ij} = 0$ means that i does not care about j . The social ranges of all the players constitute the *social range matrix* $F = \{f_{ij}\}$ of the game. We will later see (Lemma 1) that it is sufficient to focus on normalized matrices where $\forall i, j : -1 \leq f_{ij} \leq 1$ (rather than $f_{ij} \in \mathbb{R}$).

The perceived cost of player i in a strategy profile s is thus calculated as:

$$c_p(i, s) = \sum_j f_{ij} c_a(j, s).$$

In other words, the perceived cost of player i increases with the aggregate costs of i 's *friends* (players j with $f_{ij} > 0$) and decreases with the aggregate costs of

i 's *enemies* (players j with $f_{ij} < 0$). Note that we allow a player i to value other players' costs more than her own cost, i.e., f_{ii} can be smaller than some $|f_{ij}|$, $i \neq j$. This captures the effect of sacrificing one's own interests for the sake (or for the harm) of others.

Henceforth, a social matrix F with all 1's (resp., all -1 's, except for f_{ii}) is called *altruistic* (resp., *malicious*). Generally, a social matrix with a lot of zero or negligibly small elements describes a system with weak social ties. Some interesting social range matrices F are:

1. If F is the identity matrix, we are in the realm of classic game theory where each player is selfish.
2. A completely altruistic scenario is described by a social matrix F consisting of 1s only, i.e., $f_{ij} = 1$ ($\forall i, j$). Alternatively, we can also define an altruistic player that considers her own costs only to a small extent ($f_{ii} = \epsilon$, for some arbitrarily small $\epsilon > 0$).
3. In a situation where $\exists k$ such that $\forall i, j: f_{ij} = 0$ except for $f_{ik} = 1$, the players only care about a single individual. We will refer to this situation as a *monarchy scenario*. (Sometimes it makes sense to assume that players are at least a bit self-interested and $\forall i: f_{ii} = \epsilon$ for an arbitrarily small positive ϵ .)
4. If $\exists k$ such that $\forall i, j: f_{ij} = 0$ except for $f_{kj} = 1$ (and maybe $f_{kk} = \epsilon$), there is one benevolent player that seeks to maximize the utility of all players.
5. If $\exists k$ such that for all players $i: f_{ii} = 1$ and otherwise 0, and $f_{ki} = -1$, we have a selfish scenario with one malicious player k that seeks to minimize the utility of all the players. (Alternatively, we can also postulate that for a malicious player k , $f_{kk} = 1$.)
6. If $\exists j, k$ such that $\forall i: f_{ji} = f_{ki}$, then we say that players j and k *collude*: their incentives to deviate from a given strategy profile are identical. (We will show in Lemma 1 that j and k collude even if $\exists \lambda > 0: \forall i, f_{ji} = \lambda f_{ki}$.)

There are special player types to consider, e.g.:

Definition 1 (Ignorant and Ignored Players). *A player i is called ignorant if $f_{ij} = 0 \forall j$; the perceived cost of an ignorant player i does not depend on the actual costs. Now suppose that F contains a zero column: $f_{ji} = 0, \forall j$. In this case, no player cares about i 's actual cost, and we call i ignored.*

In game theory, (pure) *Nash equilibria* are an important solution concept to evaluate the outcomes of games. A Nash equilibrium is defined as a situation where no player can unilaterally reduce her cost by choosing another strategy given the other players' strategies. In our setting, where the happiness of a given player depends on her perceived costs, the equilibrium concept also needs to be expressed in terms of perceived costs. We formally define the *perceived Nash equilibrium* (PNE) as follows.

Definition 2 (Perceived Nash Equilibrium). *A strategy profile s is a perceived Nash equilibrium if for every s' that differs from s in exactly one position i , we have $c_p(i, s') \geq c_p(i, s)$.*

In order to evaluate the system performance, we study the social cost of an equilibrium. Note that the social cost is defined with respect to *actual* costs: the *social cost* of a strategy profile s is defined as $\text{Cost}(s) = \sum_j c_a(j, s)$. A strategy profile s is a *social optimum* if $\forall s': \text{Cost}(s') \geq \text{Cost}(s)$.

For a given game \mathcal{G} and a social matrix F , consider the ratio between the actual cost of the worst perceived Nash equilibrium and the cost of the social optimum. Comparing this ratio with the price of anarchy (the ratio computed with respect to actual Nash equilibria), we obtain the “effect of socialization” that captures the benefits or disadvantages that social relations contribute to the outcome of the game. Below we fix a game \mathcal{G} , and give some basic properties following immediately from the definitions.

3 Basic Properties of Social Range Matrices

We start our analysis by examining properties of the social range matrix. First, observe that F is invariant to row scaling.

Lemma 1. *Let F be a social matrix, and let $\lambda > 0$ be an arbitrary factor. Let F' be a social matrix obtained from F by multiplying a row of F by λ . Then s is a perceived Nash equilibrium w.r.t. F if and only if s is a perceived Nash equilibrium with F' .*

Proof. Let i be the player whose row is scaled. Since player i 's actual costs are not affected by multiplying f_i by λ , the perceived costs of all other players $j \neq i$ remain the same and hence, they still play their equilibrium strategy under F' . However, also player i will stick to her strategy in F' :

$$c_p(i, s) = \sum_j \lambda f_{ij} c_a(j, s) \leq c_p(i, s') = \sum_j \lambda f_{ij} c_a(j, s')$$

since we know that in F , $c_p(i, s) = \sum_j f_{ij} c_a(j, s) \leq c_p(i, s') = \sum_j f_{ij} c_a(j, s')$ for all s' that differ from s in i 's strategy. \square

In particular, Lemma 1 implies that we can normalize a social matrix F by $f'_{ij} = f_{ij} / \max_{\ell, k} |f_{\ell k}|$.¹ Therefore, in the following, we assume normalized matrices F for which $f_{ij} \in [-1, 1]$, $\forall i, j \in \{0, \dots, n-1\}$.

Lemma 2. *If $f_{ij} = 1 \ \forall i, j$, then every social optimum is a perceived Nash equilibrium. If $f_{ij} = -1 \ \forall i, j$, then every social minimum is a perceived Nash equilibrium.*

Proof. The proof is simple. By the definition of a social optimum s , $\sum_i c_a(i, s)$ is minimal, i.e., $\nexists s'$ with $\sum_i c_a(i, s') < \sum_i c_a(i, s)$. Thus, s is also an equilibrium if $f_{ij} = 1 \ \forall i, j$, as $\nexists s'$ for a given player j with $c_p(j, s') = \sum_i c_a(i, s') < c_p(j, s) = \sum_i c_a(i, s)$.

¹ Here we assume $\max_{\ell, k} |f_{\ell k}| > 0$; otherwise, every strategy is a perceived Nash equilibrium and the price of socialization is the worst possible.

Similarly for the minimum maximizing $\sum_i c_a(i, s)$ ($\nexists s'$ with $\sum_i c_a(i, s') > \sum_i c_a(i, s)$). Profile s is also a perceived equilibrium for $f_{ij} = -1 \forall i, j$, as $\nexists s'$ for a given player j with $c_p(j, s') = \sum_i c_a(i, s') > c_p(j, s) = \sum_i c_a(i, s)$. \square

Note however that the opposite direction is not true: there may be games with equilibria which are not optimal, even if all players are altruistic, namely if the game exhibits local optima.

Another special case that allows for general statements are ignorant and ignored players (see Definition 1). Note that neither ignorant nor ignored players can benefit from their unilateral actions: their perceived cost functions do not depend on their strategies. Moreover, no player's perceived cost depends on the actions of an ignored player. If s is a perceived equilibrium, then any strategy s' that differs from s only in position i , where i is an ignored player, is also a perceived equilibrium. In other words, it is sufficient to determine the set of equilibria PNE' with respect to the strategies of non-ignored players Π' .

Existing literature also provides interesting results on the properties and implications of certain types of social matrices. For instance, from the work by Babaioff et al. [4]—and even earlier, from the work by Karakostas and Viggas [13]!—we know that there are games where the presence of players who draw utility from the disutility of others, can lead to an *increase* of the social welfare; this however only holds for certain game classes that are characterized by some form of a generalized Braess paradox. Or from the work by Meier et al. [15], it follows that in a virus inoculation game where the social range matrix depends on the adjacency metrics of the social network, a society can only benefit from friendship (positive entries in the social range matrix), although *not always* in a monotonic manner.

Thus, in specific game classes, some “corner case” phenomena may be observed for certain types of social matrices. In order to focus on the principal properties of the social range, in the following we concentrate on our network creation game. It turns out that in games where choosing the neighbors can be a part of a player's strategy, there is a strong correlation between the social ties and the resulting network topology.

4 Case Study: Network Creation

In this section, we give a formal definition of our network creation game and investigate the implications of different social ranges on the formed topologies.

4.1 A Network Creation Game

As a use-case for employing our game-theoretic framework, we propose a novel simple network creation game where a node (or *player*) i can decide to which other nodes j she wants to connect in an undirected graph. Establishing a connection $\{i, j\}$ (or *edge*) entails a certain cost; we will assume that connections are undirected, and that one end has to pay for it. On the other hand, a player benefits from positive network externalities if it is connected to other players

(possibly in a multi-hop fashion). We assume that the gain or cost of a player depends on the number of players in her R -hop neighborhood, for some parameter $R \geq 0$. For instance, a network creation game with $R = 1$ describes a situation where players can only benefit from (or collaborate with) their direct neighbors. As motivation for larger radii, imagine an unstructured peer-to-peer network where searching is done by flooding up to radius R , and where the number of files that can be found increases monotonically in the number of players reached inside this radius.

Formally, the actual cost of player i is given by:

$$c_a(i, s) = \alpha \cdot s_i - g\left(\sum_{j=1}^R |\Gamma^j(i, s)|\right)$$

where parameter $\alpha \geq 0$ denotes the cost per connection, s_i is the number of connections player i pays for, and $|\Gamma^j(i, s)|$ specifies to how many nodes node i is connected with shortest hop-distance j in a graph incurred by strategy profile s . Moreover, $g : \mathbb{N}_n \rightarrow \mathbb{R}$ is a function that specifies the utility of being in a connected group of a given size (here $\mathbb{N}_n = \{0, \dots, n-1\}$). For example, $g(x) = x$ denotes that the utility grows linearly with the number of nodes within the given radius; a super-linear utility such as $g(x) = x^2$ may be meaningful in situations where the networking effects grow faster, and a sub-linear utility $g(x) = \sqrt{x}$ means that marginal utility of additional players declines with the size. By convention, we assume that $g(0) = 0$.

Finally, note that multiple strategy profiles (and hence perceived Nash equilibria) can describe the same network topology where the links are payed by different endpoints. Henceforth, for simplicity, we will sometimes say that a given topology *constitutes* (or *is*) a social optimum or an equilibrium if the corresponding profiles are irrelevant for the statement, are clear from the context, or if it holds for any strategy describing this network.

Given two network topologies of the same perceived costs but where one topology has some additional edges that need to be paid by a given player, this player is likely to prefer the other topology. That is, it often makes sense to assume that a player does not completely ignore the own actual cost, that is, $\forall i : f_{ii} = \epsilon$ for an arbitrarily small $\epsilon > 0$.

4.2 Social Optimum and Anarchy

First we describe the properties of the general network creation game in which players behave in a selfish manner. Social optima are characterized in the following lemma. It turns out that cliques and trees are the most efficient networks in our game.

Lemma 3. *Consider the network creation game where $\forall x \in \mathbb{N}_{n-1}$, $g(x+1) - g(x) > \alpha/2$. Then in the case $R = 1$, the only social optimum is the clique, and in the case $R > 1$, every social optimum is a tree of diameter at most $\min(R, n-1)$.*

Proof. Let s be any strategy profile. We say that an edge in s is *redundant* if in the strategy profile s' derived from s by dropping this edge, the R -neighborhood

of all nodes remains the same. Every non-redundant edge connecting a player with degree x to a player with degree y decreases the social cost by at least $g(x + 1) - g(x) + g(y + 1) - g(y) - \alpha > 0$. Naturally, every social optimum s will not have redundant edges. In the case $R = 1$, the clique has the most non-redundant edges, and thus is the only topology resulting from the social optimum.

In the case $R > 1$, suppose that the network described by s is not connected and does not contain redundant edges. Then every edge connecting the components of the graph decreases the social cost by a positive value. Hence, we can assume that the socially optimal topology is connected.

Now suppose that the network has diameter $R' > R$. Consider two nodes i and j such that j is at distance R' from i . Then adding an edge connecting i and j increases the R -neighborhood of each player by at least 1 and thus decreases the social cost. Therefore, the diameter of the social optimum topology is at most $\min(R, n - 1)$.

Finally, since over all connected graphs, trees have the least number of edges and hence the cost is minimized, every social optimum results in a tree. \square

In a selfish setting, players are less likely to connect to each other. Indeed, even for relatively small α , nodes remain isolated, resulting in a poor welfare.

Lemma 4. *In the network creation game, the set of isolated nodes is a Nash equilibrium if and only if $\forall x \in \mathbb{N}_n, g(x) \leq x\alpha$.*

Proof. Consider the strategy profile with no edges: $\forall j : s_j = 0$. If $\forall x \in \mathbb{N}_n, g(x) \leq x\alpha$, then unilaterally adding x edges may only increase the individual (actual) cost by at least $\alpha x - g(x)$, so no node has an incentive to deviate. On the other hand, if $\exists x \in \mathbb{N}_n, g(x) > x\alpha$, then every player has an incentive to add at least x edges, and thus the “isolated” strategy cannot be an equilibrium. \square

Lemmas 3 and 4 imply that in the case $1 < \alpha < 2$, the cost of the social optimum in the *linear* network creation game (when $g(x) = x$) is $n(n - 1)(\alpha/2 - 1)$ for $R = 1$ and $(n - 1)(\alpha - 2)$ for $R > 1$, while the cost of the worst Nash equilibrium is 0, i.e., selfishness may bring the system to a highly suboptimal state.

Below we describe the conditions under which certain topologies, like cliques and trees of bounded diameter, constitute Nash equilibria of the network creation game.

Lemma 5. *In the network creation game where $R = 1, \forall x \in \mathbb{N}_{\lfloor n/2 \rfloor}$, such that $\forall y \in \mathbb{N}_{n-x}: g(2x) - g(x + y) \geq \alpha(x - y)$, every $2x$ -regular graph constitutes a Nash equilibrium.*

Proof. Consider the strategy in which every player establishes x outgoing links so that the resulting topology is $2x$ -regular. Unilaterally establishing y (non-redundant) links instead of x (for any $y \in \mathbb{N}_{n-x}$), a player pays the cost $\alpha y - g(x + y) \geq \alpha x - g(2x)$, so no player has an incentive to deviate. \square

In the linear case with $R = 1$ and $\alpha < 1$, Lemma 5 implies that the clique is the only regular graph that results from an equilibrium: the only x that satisfies the

condition is $\lfloor n/2 \rfloor$. But in general, the resulting network may consist of up to $\lfloor n/2x \rfloor$ disconnected cliques of $2x$ players each.

Lemma 6. *In the network creation game with $R > 1$, where g is a monotonically increasing function on \mathbb{N}_n such that $\alpha < g(n-1)$, every tree of diameter at most $\min(R, n-1)$ corresponds to a Nash equilibrium.*

Proof. Consider the strategy in which every node but one maintains one edge so that the resulting graph is a tree of diameter at least $\min(R, n-1)$. Therefore, $n-1$ players have the cost $\alpha - g(n-1)$ and one player has the cost $-g(n-1)$. Every extra edge would be redundant, and dropping edges increases the cost by at least $g(n-1) - g(n)$. Thus, no player has an incentive to deviate. \square

Having described the classic setting with selfish players, we are ready to tackle social contexts.

4.3 Social Equilibria

We now turn our attention to more general matrices F , where player pairs i and j are embedded in a social context. For simplicity, we focus on values $f_{ij} \in \{-1, 0, \epsilon, 1\}$ where $f_{ij} = -1$ signifies that player i does not get along well with player j , $f_{ij} = 0$ signifies a neutral relation, and $f_{ij} = 1$ signifies friendship. We will sometimes assume that players care at least a little bit about their own cost, i.e., $\forall i : f_{ii} = \epsilon$ for some arbitrarily small positive ϵ . (This also implies that a player will not pay for a link which is already paid for by some other player.) We make two additional simplifications: we have investigated the network creation game where players can only profit from their direct neighbors (i.e., $R = 1$) in more detail, and assume a *linear* scenario where the utility of being connected to other players grows linearly in the number of contacts, that is, the marginal utility of connecting to an additional player is constant: we assume that $g(x) = x$.

Clearly, in this scenario, the cost $c_p(i, s)$ (and also $c_a(i, s)$) of a player i in a strategy profile s is independent of connections that are not incident to i . In this case, it holds that any social matrix F has a (pure) perceived equilibrium.

Lemma 7. *In the linear network creation game with $R = 1$, any social range matrix F has at least one pure perceived Nash equilibrium, for any α .*

Proof. Recall that in the $R = 1$ case, a player i can only benefit from her neighbors, that is, from a connection $\{i, j\}$ that either i or the corresponding neighbor j paid for. Player i will pay for the connection to player j if and only if the gain from this link is larger than the link cost α . We have that by establishing a new connection from player i to player j , the cost changes by $\Delta c_p(i) = f_{ii} \cdot (\alpha - 1) - f_{ij} \cdot 1$. If this cost is not larger than zero, it is worthwhile for player i to connect; otherwise it is not. On the other hand, player j will pay for a connection to player i iff $\Delta c_p(j) = f_{jj} \cdot (\alpha - 1) - f_{ji} \cdot 1 \geq 0$. As the decision of whether to connect to a given player or not does not depend on other connections, and as links cannot be canceled unilaterally, the resulting equilibrium network is unique assuming that the players will not change to a strategy of equal cost. \square

Observe that the equilibrium topology found in Lemma 7 is only unique if the cost inequalities Δc_p are strict. Moreover, a given equilibrium topology can result from different strategy profiles, namely if there are connections where both players have an incentive to pay for the connection to each other.

Intuitively, we would expect that the network formed by the players reflects the social context the players are embedded in. This can be exemplified in several ways. The following lemma shows that for the case of binary social matrices, there are situations where the social matrix translates directly into an equilibrium adjacency matrix.

Lemma 8. *In the linear network creation game with $R = 1$, assume a binary social matrix F where $\forall i, j : f_{ij} \in \{0, 1\}$ and where each player is aware of her own cost, i.e., $\forall i : f_{ii} > 0$. Then, for $1 < \alpha < 2$, there is an equilibrium topology that can be described by the adjacency matrix F' derived from F in the following manner: (1) $\forall i : f'_{ii} = 0$ and (2) if $f_{ij} = 1$ for some i, j , then $f'_{ij} = 1$ and $f'_{ji} = 1$.*

Proof. The claim follows from the simple observation that for $1 < \alpha < 2$, a player i is willing to pay for a connection to a player j if and only if $f_{ij} = 1$, as the cost difference is given by $\Delta c_p(i) = \alpha - f_{ii} - f_{ij}$: If $f_{ij} = 0$, player i only connects if $\alpha \leq 1$, and if $f_{ij} = 1$, it is worthwhile to pay for the connection as long as $\alpha \leq 2$. Therefore, as long as $1 < \alpha < 2$, one endpoint will pay for the link $\{i, j\}$ (and thus: $f'_{ij} = 1$ and $f'_{ji} = 1$) if $f_{ij} = 1$ or $f_{ji} = 1$. Clearly, it also holds that there are no loops ($\forall i : f'_{ii} = 0$). \square

Note that the condition that each player cares about her own cost is necessary for Lemma 8 to hold; otherwise, if $f_{ii} = 0$, a player could trivially connect to all players as this does not entail any connection costs. In this case, the social matrix still describes a valid equilibrium adjacency matrix, however, there are many other equilibria with additional edges.

4.4 Use Case: Anarchy vs. Monarchy

A natural question to investigate in the context of social ranges is the relationship between a completely selfish society (in game theory also known as an *anarchy*) and a society with an outstanding individual that unilaterally determines the cost of the players (henceforth referred to as a *monarchy*); as already mentioned, we assume that the players always care a little bit about their own actual costs, and hence in the monarchy, let $\forall i : f_{ii} = \epsilon$ for some arbitrarily small $\epsilon > 0$, and let $\forall i : f_{ij} = 1$ where player j is the monarch (we assume $f_{ji} = 0$ for all $i \neq j$).

Interestingly, while there are situations where a monarchy yields a higher social welfare, the opposite is also true as there are settings that are better for anarchies. The following result characterizes social optima, and Nash equilibria for anarchy and monarchy settings.

Lemma 9. *In the linear network creation game with $R = 1$, the social optimum cost is $(\alpha/2 - 1)n(n - 1)$ if $\alpha < 2$ and 0 otherwise, and the anarchy has social cost $(\alpha/2 - 1)n(n - 1)$ if $\alpha \leq 1$ and 0 otherwise. For the monarchy, there can be*

multiple equilibria (of the same cost): for any α , there is always an equilibrium with cost $(\alpha - 2)(n - 1)$; moreover, if $\alpha \leq 1$ there is an additional equilibrium with the same cost.

Proof. We consider the social optimum, the anarchy and the monarchy in turn. *Social optimum:* If $\alpha \leq 2$, then Lemma 3 implies that any social optimum implies the clique, with the cost $(\alpha/2 - 1)n(n - 1)$. If $\alpha > 2$, then every non-redundant link increases the social cost by $\alpha - 2$ and thus the set of isolated nodes has the minimal cost, 0.

Observe that the social cost is given by the total number of edges k in the network: k edges yield a connection cost of $k \cdot \alpha$, and the players are connected to $2k$ other players, thus $\text{Cost}(s) = k \cdot \alpha - 2k$. Using Lemma 3, for the social optimum we have:

$$\min_s \text{Cost}(s) = \min_k k \cdot \alpha - 2k = \begin{cases} (\alpha/2 - 1)n(n - 1) & , \text{ if } \alpha \leq 2 \text{ (clique)} \\ 0 & , \text{ otherwise (disconnected).} \end{cases}$$

Anarchy: In a purely selfish setting, a player connects to another player if and only if $\alpha \leq 1$. By Lemmas 3 and 4, if $\alpha \leq 1$, then the resulting equilibrium topology is the clique and the cost is thus $(\alpha/2 - 1)n(n - 1)$, and if $\alpha < 1$, then the resulting topology is the set of isolated nodes and the cost is 0.

$$\text{Cost(Nash equilibrium)} = \begin{cases} (\alpha/2 - 1)n(n - 1) & , \text{ if } \alpha \leq 1 \text{ (clique)} \\ 0 & , \text{ otherwise (disconnected).} \end{cases}$$

Monarchy: Let j denote the monarch and let $i \neq j$ denote any other player. Since the marginal utility of an additional neighbor of j is one while the connection cost is arbitrarily small ($\alpha\epsilon$), a player i will always connect (i.e., pay for the connection) to the monarch. On the other hand, the monarch will connect to a player if and only if $\alpha \leq 1$. The social cost of the network equilibrium is therefore always $(\alpha - 2)(n - 1)$ (up to the arbitrarily small ϵ components in the cost), for any α . \square

Using Lemma 9, we can compare the efficiency of the different social settings. For relatively low connection costs, a setting with a monarch gives stronger incentives for nodes to connect, and thus socially more preferable outcomes emerge. On the other hand, for high connection costs, due to the selfless behavior of the players ignoring their own connection prices, an anarchy is preferable. As a concrete example, according to Lemma 9, for $\alpha = 3/2$, the equilibrium network of the monarchic society is a star of utility $(n - 1)/2$ while in the anarchy nobody will connect, yielding a utility of zero. On the other hand, for $\alpha = 3$, the anarchy again has zero utility, while in the monarchy, players still connect which results in a negative overall utility of $-(n - 1)$. Thus, the following lemma holds.

Lemma 10. *There are situations where the social welfare of anarchy is higher than the welfare in a monarchy, and vice versa.*

4.5 Windfall of Friendship and Price of Ill-Will

An interesting property of our network creation game is that more friendship relations cannot worsen an equilibrium.

Lemma 11. *Consider a social range matrix F where $\forall i, j : f_{ij} \in \{0, 1\}$ and $f_{ii} = 1$. Let F' be a social range matrix derived from F where a non-empty set \mathcal{N} of 0-entries in F are flipped to 1. Then, for any equilibrium strategy s^F with respect to a social matrix F , there is an equilibrium strategy $s^{F'}$ with $\text{Cost}(s^{F'}) \leq \text{Cost}(s^F)$.*

Proof. We prove the claim by showing that for any equilibrium strategy s^F for F , there is an equilibrium strategy $s^{F'}$ for F' that has at least as many connections as s^F . Moreover, it holds that an equilibrium with more connections always implies a higher social welfare.

First, recall from Lemma 7 that an equilibrium s^F always exists. Now fix such an equilibrium s^F from which we will construct the equilibrium $s^{F'}$. If i and j are connected in s^F , then they are still connected in $s^{F'}$: as $R = 1$, whether or not a connection $\{i, j\}$ between two players i and j is established depends on the actual costs $c_a(i, \cdot)$ and $c_a(j, \cdot)$ only. If two players i and j are not connected in s^F , they have an incentive to connect in $s^{F'}$ if $f'_{ij} = 1$ and $\alpha \leq 2$. Thus, $s^{F'}$ contains a superset of the connections in s^F . Now let k be the number of edges in a given profile s . The social cost is then given by $\text{Cost}(s) = k\alpha - 2k$. For $\alpha \leq 2$, this function is monotonically decreasing, which implies the claim. On the other hand, for $\alpha > 2$, the set of isolated nodes constitutes the only equilibrium. \square

Lemma 11 implies that the best equilibrium with respect to F' cannot be worse than the best equilibrium with respect to F . On the other hand, it is easy to see that a similar claim also holds for the *worst* equilibrium: Consider the equilibrium $s^{F'}$ with the fewest connections k' . Then, there is an equilibrium s^F with $k \leq k'$ edges: either $s^F = s^{F'}$, or some edges can be removed. We have the following claim.

Corollary 1. *Consider a social range matrix F where $\forall i, j : f_{ij} \in \{0, 1\}$ and $f_{ii} = 1$. Let F' be a social range matrix that is derived from F by flipping one or several 0 entries to 1. Let s_w^F and s_b^F be the worst and the best equilibrium profile w.r.t. F , and let $s_w^{F'}$ and $s_b^{F'}$ be the worst and best equilibrium profile w.r.t. F' (maybe $s_w^F = s_b^F$ and/or $s_w^{F'} = s_b^{F'}$). It holds that $\text{Cost}(s_w^F) \geq \text{Cost}(s_w^{F'})$ and $\text{Cost}(s_b^F) \geq \text{Cost}(s_b^{F'})$.*

An analogous result can be obtained for settings where players dislike each other.

Lemma 12. *Consider a social range matrix F where $\forall i, j : f_{ij} \in \{-1, 0\}$ and $f_{ii} = 1$. Let F' be a social range matrix derived from F where a non-empty set \mathcal{N} of 0-entries in F are flipped to -1 s., where $|\mathcal{N}| \geq 1$. Then, for any equilibrium strategy s^F with respect to a social matrix F , there is an equilibrium strategy $s^{F'}$ with $\text{Cost}(s^F) \leq \text{Cost}(s^{F'})$.*

Proof. First recall from the proof of Lemma 11 that the social welfare increases with the total number of connections given that $\forall i : f_{ii} = 1$, and that it follows from Lemma 7 that an equilibrium s^F always exists. Fix an equilibrium s^F to construct the equilibrium $s^{F'}$. Similarly to the arguments used in the proof of Lemma 11, if i and j are not connected in s^F , then they are disconnected in $s^{F'}$ as well. On the other hand, if player i pays for the connection to player j in s^F , she has an incentive to disconnect in $s^{F'}$ if $f'_{ij} = -1$ and for any non-negative α . Thus, s^F contains a superset of the connections in $s^{F'}$. \square

5 Conclusions and Open Questions

We understand our work as a further step in the endeavor to shed light onto the socio-economic phenomena of today's distributed systems which typically consist of a highly heterogeneous population. In particular, this paper has initiated the study of economic games with more complex forms of social relationships. We introduced the concept of social range matrices and studied their properties. Moreover, we have proved the intuition right (under certain circumstances) that in our novel network creation game, the social relationships are reflected in the network topology.

This paper reported only on a small subset of the large number of questions opened by our model, and we believe that there remain many exciting directions for future research. For instance, it is interesting to study which conditions are necessary and sufficient for counter-intuitive phenomena such as the fear factor and the windfall of malice [4,17], or the non-monotonous relationship between welfare and friendship in social networks [15]. Another open question is the characterization of all topologies that correspond to a Nash equilibrium.

References

1. Aiyer, A., Alvisi, L., Clement, A., Dahlin, M., Martin, J.-P., Porth, C.: BAR Fault Tolerance for Cooperative Services. In: Proc. 20th ACM Symposium on Operating Systems Principles (SOSP), pp. 45–58 (2005)
2. Albers, S., Eilts, S., Even-Dar, E., Mansour, Y., Roditty, L.: On Nash Equilibria for a Network Creation Game. In: Proc. 17th ACM Symposium on Discrete Algorithms, SODA (2006)
3. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The Price of Stability for Network Design with Fair Cost Allocation. In: Proc. 45th Symposium on Foundations of Computer Science (FOCS), pp. 295–304 (2004)
4. Babaioff, M., Kleinberg, R., Papadimitriou, C.: Congestion Games with Malicious Players. In: Proc. ACM Conference on Electronic Commerce (EC), San Diego, CA, USA (2007)
5. Castro, M., Liskov, B.: Practical Byzantine Fault Tolerance. In: Proc. 3rd Symposium on Operating Systems Design and Implementation (OSDI), pp. 173–186 (1999)
6. Chen, H.-L., Roughgarden, T.: Network Design with Weighted Players. In: Proc. 18th ACM Symposium on Parallel Algorithms and Architectures (SPAA), pp. 29–38 (2006)

7. Corbo, J., Parkes, D.C.: The Price of Selfish Behavior in Bilateral Network Formation. In: Proc. 24th ACM Symposium on Principles of Distributed Computing (PODC), pp. 99–107 (2005)
8. Demaine, E.D., Hajiaghayi, M.T., Mahini, H., Zadimoghaddam, M.: The Price of Anarchy in Network Creation Games. In: Proc. 26th Annual Symposium on Principles of Distributed Computing (PODC) (2007)
9. Eliaz, K.: Fault Tolerant Implementation. *Review of Economic Studies* 69, 589–610 (2002)
10. Fabrikant, A., Luthra, A., Maneva, E., Papadimitriou, C.H., Shenker, S.: On a Network Creation Game. In: Proc. 22nd ACM Symposium on Principles of Distributed Computing (PODC), pp. 347–351 (2003)
11. Friedkin, N.E., Johnsen, E.C.: Social Influence Networks and Opinion Change. *Advances in Group Processes* 16 (1999)
12. Hoefer, M., Skopalik, A.: Altruism in Atomic Congestion Games. In: Proc. European Symposium on Algorithms (ESA) (2009)
13. Karakostas, G., Viglas, A.: Equilibria for Networks with Malicious Users. *Mathematical Programming A* 110(3), 591–613 (2007)
14. Li, H., Clement, A., Marchetti, M., Kapriltos, M., Robinson, L., Alvisi, L., Dahlin, M.: FlightPath: Obedience vs Choice in Cooperative Services. In: Proc. Symposium on Operating Systems Design and Implementation, OSDI (2008)
15. Meier, D., Oswald, Y.A., Schmid, S., Wattenhofer, R.: On the Windfall of Friendship: Inoculation Strategies on Social Networks. In: Proc. 9th ACM Conference on Electronic Commerce (EC) (2008)
16. Moscibroda, T., Schmid, S., Wattenhofer, R.: On the Topologies Formed by Selfish Peers. In: Proc. 25th Annual Symposium on Principles of Distributed Computing, PODC (2006)
17. Moscibroda, T., Schmid, S., Wattenhofer, R.: When Selfish Meets Evil: Byzantine Players in a Virus Inoculation Game. In: Proc. 25th Annual Symposium on Principles of Distributed Computing (PODC) (2006)
18. Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V.: *Algorithmic Game Theory*. Cambridge University Press, Cambridge (2007)
19. Roughgarden, T.: Stackelberg Scheduling Strategies. In: Proc. 33rd ACM Symposium on Theory of Computing (STOC), pp. 104–113 (2001)