

Schelling Games on Graphs

Edith Elkind¹, Jiarui Gan¹, Ayumi Igarashi²,
Warut Suksompong¹ and Alexandros A. Voudouris¹

¹Department of Computer Science, University of Oxford

²Department of Mathematical Informatics, University of Tokyo

{edith.elkind, jiarui.gan, warut.suksompong, alexandros.voudouris}@cs.ox.ac.uk,
igarashi@mist.i.u-tokyo.ac.jp

Abstract

We consider strategic games that are inspired by Schelling’s model of residential segregation. In our model, the agents are partitioned into k types and need to select locations on an undirected graph. Agents can be either stubborn, in which case they will always choose their preferred location, or strategic, in which case they aim to maximize the fraction of agents of their own type in their neighborhood. We investigate the existence of equilibria in these games, study the complexity of finding an equilibrium outcome or an outcome with high social welfare, and also provide upper and lower bounds on the price of anarchy and stability. Some of our results extend to the setting where the preferences of the agents over their neighbors are defined by a social network rather than a partition into types.

1 Introduction

In 2015, African Americans constituted 83% of the population of the City of Detroit. At the same time, the neighboring Oakland County was 77% white, and in the city of Dearborn in Detroit metropolitan area about 30% of the residents were Arab Americans. Similar phenomena can be observed in many other major metropolitan areas around the globe. In the developed world, the leading cause of such population patterns is not direct discrimination, which is typically illegal; rather, it is the residents themselves who tend to select neighborhoods where their ethnic or social group is well-represented. Schelling [1969; 1971] proposed the following stylized model of this phenomenon: Agents of two different types are placed on a line or on a grid, and are assumed to be happy if at least a fraction τ of the agents within distance w from them are of the same type, for some parameters τ and w ; unhappy agents can either jump to empty positions or swap positions with other agents. Using simple experiments, Schelling showed that, even in cases where the agents are not opposed to integration ($\tau < 1/2$), this behavior leads to almost complete segregation.

In the 50 years since Schelling’s pioneering paper, this segregation model attracted the attention of many researchers, mostly in sociology and economics [Alba and Logan, 1993;

Benard and Willer, 2007; Benenson *et al.*, 2009; Clark and Fossett, 2008; Pancs and Vriend, 2007; Young, 2001; Zhang, 2004a; Zhang, 2004b], but recently also in computer science [Barmpalias *et al.*, 2014; Barmpalias *et al.*, 2015; Brandt *et al.*, 2012; Immorlica *et al.*, 2017]. While the early work in this area was mainly empirical, the more recent papers have provided theoretical analysis. In particular, it was proved that the local behavior of unhappy agents is likely to create very large regions consisting of agents of the same type, even when τ is small, i.e., even when the agents themselves are tolerant towards having neighbors of the other type. The vast majority of this work was based on Schelling’s original model, where agents’ behavior was explained by a simple stochastic model rather than strategic considerations.

An alternative approach is to assume that the behavior of each agent is *strategic*, and exploit tools and techniques from non-cooperative game theory to analyze the induced games. To the best of our knowledge, there are only two papers in the literature that pursue this agenda. Specifically, Zhang [2004b] considered a model with transferable utility where agents prefer to be in a balanced neighborhood. More recently, Chauhan *et al.* [2018] investigated a setting that is closer to Schelling’s motivating scenario and incorporates the idea that, in addition to preferences over the composition of their neighborhood, agents may also have preferences over locations. In the model of Chauhan *et al.* [2018], there are two types of agents, and an agent i ’s *happiness ratio* is defined as the fraction of agents of i ’s type among i ’s neighbors. Each agent has two further parameters: a tolerance threshold $\tau \in (0, 1)$ and a preferred location. An agent’s primary goal is to find a location where her happiness ratio exceeds the tolerance threshold; if no such location is available, she aims to maximize her happiness ratio. An agent’s secondary goal is to minimize the distance to her preferred location. To achieve these goals, agents can either swap locations (swap games) or jump to unoccupied locations (jump games). The main contribution of the paper is to identify conditions under which agents are guaranteed to converge to an equilibrium; for instance, the authors establish that in jump games, convergence is guaranteed if agents have no preferred locations and the underlying network is a ring.

Our Contribution

The model of Chauhan *et al.* [2018] makes an important contribution to the literature by enriching Schelling’s model with

two additional components: agents who are fully strategic, and location preferences. However, the resulting model of agents’ preferences is quite complex, and, consequently, not easy to analyze: the positive results in the paper are limited to special cases of the utility function and highly regular networks. In this paper, we propose a simpler model that aims to capture the same phenomena and is more amenable to formal analysis.

Specifically, in our basic model the agents are partitioned into k types and, just as in the work of Chauhan *et al.* [2018], the set of available locations is represented by an undirected graph, which we refer to as the *topology*. We also incorporate location preferences in our model, but instead of assuming that optimizing the distance to the preferred location is the secondary goal of every agent, we assume that agents are either *stubborn*, in which case they stay at their chosen location irrespective of their surroundings, or *strategic*, in which case they aim to maximize their happiness ratio by jumping to an unoccupied location (we do not consider swaps in this paper). Our model captures the fact that, in practice, many residents (such as older people or those with underwater mortgages) are unwilling to move even if they are no longer satisfied with the composition of their neighborhood. Importantly, unlike Chauhan *et al.* [2018] or Schelling in his original work, we do not assume that agents have tolerance thresholds; rather, a strategic agent is willing to move as long as there exists another location with a better happiness ratio. Towards the end of the paper (Section 6), we also discuss several variants of this basic model. In particular, we show that some of our positive results extend to the setting where there are no types, but rather the agents are connected by a social network and care about the fraction of their friends (i.e., their neighbors in the social network) among their neighbors in the topology; we refer to the resulting class of games as *social Schelling games*.

The rest of the paper is organized as follows. We define our model in Section 2. Then, in Section 3, we show that for some classes of topologies, such as stars and graphs of maximum degree two, our games always admit a pure Nash equilibrium, i.e., the strategic agents can be assigned to the nodes of the topology so that none of them wants to move to a different location; this result holds even for social Schelling games. In contrast, an equilibrium may fail to exist even if the topology is acyclic and has maximum degree four. In Section 4, we complement this result by presenting a dynamic programming algorithm that decides whether an equilibrium exists on a tree topology; this algorithm runs in polynomial time if the number of types is bounded by a constant. For more general topologies, we prove that deciding whether an equilibrium exists is an NP-complete problem. Similar hardness and easiness results hold for the problem of maximizing the social welfare (the total utility of all strategic agents). In Section 5, we study the effect of the strategic behavior on the social welfare, by bounding the price of anarchy [Koutsoupias and Papadimitriou, 1999] and the price of stability [Anshelevich *et al.*, 2008]. In particular, we show that even in the absence of stubborn agents it may be impossible to achieve the maximum social welfare in equilibrium. In Section 6 we discuss several variants and extensions of our model and es-

tablish some preliminary results for these new models, as well as outline directions for future work.

Other Related Work

Due to space constraints, we omit the survey of the literature on non-strategic variants of the Schelling model; for an accessible introduction, see chapter 4 in the book of Easley and Kleinberg [2010].

Besides the work of Chauhan *et al.* [2018], which was discussed in detail earlier, our model shares a number of properties with hedonic games [Drèze and Greenberg, 1980; Bogomolnaia and Jackson, 2002]; these are games where agents split into coalitions, and each agent’s utility is determined by the composition of her coalition. Specifically, in *fractional hedonic games* [Aziz *et al.*, 2017] the relationships among the agents are described by a weighted directed graph, where the weight of an edge (i, j) is the value that agent i assigns to agent j , and an agent’s utility for a coalition is her average value for the other members in the coalition. If the graph is undirected and all edge weights take values in $\{0, 1\}$, it can be interpreted as a friendship relation; then an agent’s utility in a coalition is computed as the fraction of her friends among the coalition members, which is very similar to how utilities are defined in social Schelling games. On the other hand, the type-based model is closely related to the Bakers and Millers game discussed by Aziz *et al.* [2017]. This connection between Schelling games and hedonic games motivates much of the discussion in Section 6. Of course, a fundamental difference between hedonic games and our setting is that in the former agents derive their utilities from pairwise disjoint coalitions, whereas in our model utilities are derived from (overlapping) neighborhoods.

2 The Model

Let $N = \{1, \dots, n\}$ be a set of $n \geq 2$ agents. The agents are partitioned into $k \geq 2$ different types T_1, \dots, T_k so that $\cup_{j=1, \dots, k} T_j = N$; we write $\mathcal{T} = (T_1, \dots, T_k)$. We say that two agents $i, j \in N$, $i \neq j$, are *friends* if $i, j \in T_\ell$ for some $\ell \in [k]$; otherwise we say that i and j are *enemies*. For each $i \in N$, we denote the set of all friends of agent i by $F(i)$.

A *topology* is an undirected graph $G = (V, E)$ with no self-loops. Each agent in N has to select a node of this graph so that there are no collisions. The agents are classified as either *strategic* or *stubborn*; let R and S denote these sets of agents so that $R \cup S = N$. Each stubborn agent occupies some node of the topology and never moves away. We describe the locations of the stubborn agents by an injective mapping $\lambda : S \rightarrow V$ such that $\lambda(i)$ is the node that agent $i \in S$ occupies. In contrast, strategic agents do not care about their location per se, but want to be in a neighborhood that has a large proportion of their friends, and are willing to move to a currently unoccupied node in order to increase their utility.

A tuple $I = (R, S, \mathcal{T}, G, \lambda)$, where R is the set of strategic agents, S is the set of stubborn agents, $\mathcal{T} = (T_1, \dots, T_k)$ is a list of types, $G = (V, E)$ is a topology that satisfies $|V| > |R| + |S|$, and λ is an injective mapping from S to V , is called a *k-typed Schelling game* or *k-typed instance*. Let \mathcal{I} be the set of all possible games.

An *assignment* is a vector $\mathbf{v} = (v_1, \dots, v_n) \in V^n$ such that (1) $v_i = \lambda(i)$ for each $i \in S$ and (2) $v_i \neq v_j$ for all $i, j \in N$ such that $i \neq j$; here, v_i is the node of the topology where agent i is positioned. For a given assignment \mathbf{v} and an agent $i \in N$, let $N_i(\mathbf{v}) = \{j \in N : \{v_i, v_j\} \in E\}$ be the set of neighbors of agent i . Let $f_i(\mathbf{v}) = |N_i(\mathbf{v}) \cap F(i)|$ be the number of neighbors of i in \mathbf{v} who are her friends. Similarly, let $e_i(\mathbf{v}) = |N_i(\mathbf{v})| - f_i(\mathbf{v})$ be the number of neighbors of i in \mathbf{v} who are her enemies. Following Chauhan *et al.* [2018], we define the utility $u_i(\mathbf{v})$ of an agent $i \in R$ in \mathbf{v} to be 0 if $f_i(\mathbf{v}) = 0$; otherwise, her utility is defined as the fraction of her friends among the agents in the neighborhood:

$$u_i(\mathbf{v}) = \frac{f_i(\mathbf{v})}{f_i(\mathbf{v}) + e_i(\mathbf{v})}.$$

We say that an assignment \mathbf{v} is a *pure Nash equilibrium* (or, simply, *equilibrium*) of an instance I if no strategic agent i has an incentive to unilaterally deviate to an empty node z of G in order to increase her utility, i.e., for every $i \in R$ and for every node $z \in V$ such that $z \neq v_j$ for all $j \in R \cup S$ it holds that $u_i(\mathbf{v}) \geq u_i(z, \mathbf{v}_{-i})$, where (z, \mathbf{v}_{-i}) is the assignment obtained by changing the i -th entry of \mathbf{v} to z . Let $\text{EQ}(I)$ denote the set of all equilibria of game I .

The *social welfare* of an assignment \mathbf{v} is defined as the total utility of all strategic agents:

$$\text{SW}(\mathbf{v}) = \sum_{i \in R} u_i(\mathbf{v}).$$

Let $\mathbf{v}^*(I)$ be an assignment that maximizes the social welfare for a given game I ; we refer to it as an *optimal* assignment.

The *price of anarchy* (PoA) of game I with at least one equilibrium is the ratio between the optimal social welfare and the social welfare of the *worst* equilibrium; its *price of stability* (PoS) is defined as the ratio between the optimal social welfare and the social welfare of the *best* equilibrium:

$$\text{PoA}(I) = \sup_{\mathbf{v} \in \text{EQ}(I)} \frac{\text{SW}(\mathbf{v}^*(I))}{\text{SW}(\mathbf{v})},$$

$$\text{PoS}(I) = \inf_{\mathbf{v} \in \text{EQ}(I)} \frac{\text{SW}(\mathbf{v}^*(I))}{\text{SW}(\mathbf{v})}.$$

The price of anarchy and the price of stability are the suprema of $\text{PoA}(I)$ and $\text{PoS}(I)$ over all $I \in \mathcal{I}$ such that $\text{EQ}(I) \neq \emptyset$, respectively.

All omitted proofs can be found in the full version of this paper [Elkind *et al.*, 2019].

3 Existence of Equilibria

In this section, we focus on the existence of equilibria. To start, we observe that for highly structured topologies such as paths, rings, and stars, there is always at least one equilibrium assignment, and some such assignment can be computed efficiently. This can be shown directly, and also follows from a more general result established in Section 6 (Theorem 6.1).

Theorem 3.1. *Every k -typed Schelling game where the topology is a star or a graph of maximum degree 2 admits at least one equilibrium assignment, which can be computed in polynomial time.*




Figure 1: Example of the topology used in the proof of Theorem 3.2 for $k = 2$.

However, in general, an equilibrium may fail to exist; this holds even if the topology is acyclic and there are no stubborn agents.

Theorem 3.2. *For every $k \geq 2$ there exists a k -typed instance $(R, S, \mathcal{T}, G, \lambda)$ where $S = \emptyset$ and G is a tree that does not admit an equilibrium.*

Proof. Given $k \geq 2$, we construct an instance with $2k + 1$ agents per type; the total number of agents is $n = k(2k + 1)$. The topology $G = (V, E)$ is a tree that consists of $|V| = n + 1$ nodes, which are distributed over four layers. Specifically, the tree has a root α , which has one child β . Node β has $2k - 1$ children; we denote the set of its children by Γ . Each node in Γ has k children, which are the leaves of the tree; we denote the set of all leaves by Δ . Figure 1 depicts the topology for $k = 2$. Now, assume that there is an equilibrium assignment; note that exactly one node is left empty. We consider four cases depending on the location of the empty node.

Node α is empty. Assume that the agent occupying node β is of type T . Then, since there are $2k$ other agents of type T , there must exist some subtree rooted at a node in Γ that contains both agents of type T and agents that belong to other types. Then an agent of type T from this subtree has an incentive to deviate to α .

Node β is empty. Assume that the agent occupying node α is of type T ; note that her utility is 0. If she does not have an incentive to deviate to β , it follows that no agent of type T occupies a node in Γ . But then there is an agent of type T who occupies a node in Δ ; as her parent is not of type T , her utility is 0, and she can increase it by moving to β .

Some node $\gamma \in \Gamma$ is empty. Consider the agents occupying the children of γ ; note that their utility is 0. If at least two of them have the same type, each of them has an incentive to deviate to γ in order to increase her utility to at least $\frac{1}{k}$. If all of them have different types, then there is exactly one agent of each type in this set. In particular, there is an agent i who has the same type as the agent occupying β ; then i can move to γ to increase her utility.

Some node $\delta \in \Delta$ is empty. Let γ denote the parent of this node, and suppose that γ is occupied by an agent i of

type T . We say that an agent j of type T is *hungry* if $j \neq i$ and j is adjacent to at least one agent of a different type; note that a hungry agent has an incentive to deviate to δ . We claim that at least one agent is hungry. Indeed, if β is occupied by an agent j of type T , then either j is hungry or every agent in $\Gamma \setminus \{\gamma\}$ is hungry. If the agent in β is not of type T and there is an agent ℓ of type T in $\Gamma \setminus \{\gamma\}$, then ℓ is hungry. Finally, if no agent in $\Gamma \setminus \{\gamma\}$ is of type T , there exists a leaf node not in γ 's subtree that is occupied by an agent r of type T ; r is then hungry.

The proof is complete. \square

4 Computational Complexity

We now turn our attention to the computational complexity of k -typed Schelling games. The main result of this section is that finding an equilibrium assignment is computationally intractable.

Theorem 4.1. *For every $k \geq 2$, given a k -typed Schelling game I , it is NP-complete to decide whether I admits an equilibrium assignment. The hardness result holds even if all strategic agents belong to the same type.*

Proof. We give a proof for $k = 2$; it is straightforward to extend it to $k \geq 2$, by adding isolated stubborn agents of additional types. We will use a reduction from the CLIQUE problem. An instance of this problem is an undirected graph $H = (X, Y)$ and an integer s ; it is a yes-instance if H has a complete subgraph of size s . Given an instance $\langle H, s \rangle$ of CLIQUE, we assume without loss of generality that $s \geq 5$ and construct an instance of our problem as follows:

- There are two agent types: red and blue.
- There are s strategic red agents; all remaining agents are stubborn. We will describe the stubborn agents and their locations when defining the topology.
- The topology $G = (V, E)$ consists of three disjoint components G_1, G_2 , and G_3 such that
 - $G_1 = (V_1, E_1)$, where $V_1 = X \cup W$, $|W| = s - 2$, $E_1 = Y \cup \{\{v, w\} : v \in X, w \in W\}$. There is a stubborn blue agent at each node $w \in W$;
 - G_2 is a complete bipartite graph with parts L and R , $|L| = s - 2$, $|R| = 4s$. Of the $4s$ nodes in R , $2s + 1$ nodes are occupied by red agents and $2s - 1$ nodes are occupied by blue agents;
 - G_3 has three empty nodes, denoted x, y , and z , and 121 nodes — 41 red and 80 blue — occupied by stubborn agents. There is an edge between nodes x and y ; also, x is connected to 1 red agent and 2 blue agents; y is connected to 41 red agents and 80 blue agents, and z is connected to 5 red agents and 7 blue agents.

Note that a strategic red agent obtains a utility of $\frac{2s+1}{4s} = \frac{1}{2} + \frac{1}{4s}$ by choosing an available node in G_2 and a utility of $\frac{5}{12}$ by choosing z . If she chooses x , her utility is $\frac{1}{3}$ if y is unoccupied and $\frac{1}{2}$ otherwise. Similarly, if she chooses y , her utility is $\frac{41}{121}$ if x is unoccupied and $\frac{42}{122}$ otherwise; note that $\frac{1}{3} < \frac{41}{121} < \frac{42}{122} < \frac{5}{12}$.

Now, suppose that H contains a clique of size s . If strategic red agents occupy the nodes of that clique, the utility of each such agent is $\frac{s-1}{(s-1)+(s-2)} = \frac{1}{2} + \frac{1}{4s-6}$. Thus, by our choice of parameters, no agent has a profitable deviation.

On the other hand, suppose that H does not contain a clique of size s . Assume for the sake of contradiction that there is an equilibrium assignment v .

Suppose first that in v some strategic agents are located in G_1 . It cannot be the case that each of them is adjacent to $s - 1$ friends, as this would mean that their locations form a clique of size s . Hence, at least one of these agents is adjacent to at most $s - 2$ friends. As this agent is also adjacent to the $s - 2$ stubborn blue agents in W , her utility is at most $\frac{1}{2}$. By our choice of parameters, all unoccupied nodes of G_2 offer a higher utility, namely, $\frac{1}{2} + \frac{1}{4s}$. Thus, if there are strategic agents in G_1 , all $s - 2$ nodes of G_2 that are available to strategic agents must be occupied. But then, there are at most two strategic agents in G_1 , which means that their utility is at most $\frac{1}{s-1} < \frac{1}{3}$ (recall that we assume that $s \geq 5$). This leads to a contradiction, as these strategic agents would be better off moving to G_3 where their utility would be at least $\frac{1}{3}$.

Therefore, in equilibrium no strategic agent can be located at a node of G_1 . Further, since all unoccupied nodes of G_2 always offer more utility than any unoccupied nodes of G_3 can offer, in equilibrium all nodes of G_2 are occupied, and the two remaining strategic agents must be in G_3 , with one of x, y , and z left empty.

Suppose that z is empty. Then the agent located at y can increase her utility from $\frac{42}{122}$ to $\frac{5}{12}$ by moving to z , a contradiction. If y is empty, the agent located at x can increase her utility from $\frac{1}{3}$ to $\frac{41}{121}$ by moving to y , a contradiction. Finally, if x is empty, the agent located at z can increase her utility from $\frac{5}{12}$ to $\frac{1}{2}$ by moving to x , a contradiction. As we have exhausted all possibilities, it follows that if G does not have a clique of size s , then there is no equilibrium assignment. \square

The proof of Theorem 4.1 can be adapted to show that maximizing social welfare in Schelling games is NP-hard as well.

Theorem 4.2. *For every $k \geq 2$, given a k -typed Schelling game I and a rational value s , it is NP-complete to decide whether I admits an assignment with social welfare at least s . The hardness result holds even if $k = 2$, all strategic agents belong to one type, and the other type consists of a single stubborn agent.*

On the positive side, for small k we can efficiently decide whether an equilibrium exists if the topology G is a tree. Our algorithm is based on dynamic programming: it selects an arbitrary node of G to be the root, and then for every node v of G , it fills out a multidimensional table whose dimension is linear in the number of types, proceeding from the leaves to the root. It decides whether the given instance admits an equilibrium by scanning the table at the root node.

Theorem 4.3. *Given a k -typed Schelling game I with n agents, where the topology G is a tree, we can decide whether I admits an equilibrium (and compute one if it exists) in time $\text{poly}(n^k)$, i.e., this problem lies in the complexity class XP with respect to the number of types k .*




Figure 2: The topology used in the proof of Theorem 5.1 for $k = 2$.

By slightly modifying our algorithm, we can compute an assignment that maximizes the social welfare, either among all assignments or among equilibria.

Corollary 4.4. *Given a k -typed Schelling game, where G is a tree, the problems of computing an equilibrium with maximum social welfare or a socially optimal assignment are in XP with respect to k .*

We have not been able to determine whether the problem of computing an equilibrium assignment is fixed-parameter tractable with respect to the number of types; we leave this question for future work.

5 Price of Anarchy and Stability

In this section, we investigate the loss in social welfare caused by strategic behavior, as measured by the price of anarchy and the price of stability; unless otherwise specified, the topology is assumed to be a connected graph.¹ We start by establishing bounds on the PoA for instances with no stubborn agents.

Theorem 5.1. *For k -typed Schelling games with no stubborn agents and n strategic agents, the PoA*

- can be unbounded for each $k \geq 2$;
- is $\Theta(n)$ when there are at least two agents per type;
- is $k + o(1)$ if each type has the same number of agents.

Proof. Due to space constraints, we only prove the last claim.

For the lower bound, fix $k \geq 2$, let $\ell \geq 2$ be a parameter and consider an instance with $k(\ell + 1)$ agents per type; altogether there are $n = k^2(\ell + 1)$ agents. The topology consists of $n + 1$ nodes and is defined as follows. There are k cliques V_1, \dots, V_k of size $k\ell$ each, and a node y . In each clique V_i there is a special node w_i that is connected to y . Also, for each $i \in [k]$ there are k auxiliary nodes $z_{i,1}, \dots, z_{i,k}$; each of these nodes is connected to a distinct set of ℓ nodes in V_i . Let $z_{i,i}$ be the auxiliary node that is connected to w_i . Figure 2 illustrates this topology for $k = 2$.

There is an optimal assignment where all $k(\ell + 1)$ agents of type T_i are placed at the nodes of clique V_i and the corresponding auxiliary nodes, so that all agents are connected only to agents of the same type and have maximum utility 1. Therefore, the optimal social welfare is $k^2(\ell + 1)$.

¹To see that the PoA can be unbounded if the topology is not connected, consider a topology that contains a connected component W of size exactly n , as well as some other connected components: any assignment of agents to nodes in W is an equilibrium, even if its social welfare is 0.

In contrast, consider the following equilibrium assignment: node y is empty, and for each $i, j \in [k]$ all ℓ nodes in V_i that are connected to the auxiliary node $z_{i,j}$ as well as $z_{i,j}$ itself are occupied by agents of type T_j . Since node y is connected to k nodes that are occupied by agents of different types, any agent would get utility $1/k$ by deviating there. No agent occupying an auxiliary node has an incentive to deviate since she is connected only to agents of her type. For every clique, each agent is connected to exactly ℓ agents of the same type ($\ell - 1$ of whom occupy nodes of the clique and one that occupies the corresponding auxiliary node) and $(k - 1)\ell$ agents of different type; thus, her utility is $1/k$. Consequently, no agent has an incentive to deviate, and the social welfare is $k \cdot k\ell \cdot \frac{1}{k} + k^2 = k(\ell + k)$. Hence, the PoA is at least $\frac{k\ell + k}{\ell + k}$; this expression becomes arbitrarily close to k as ℓ grows.

For the upper bound, consider an arbitrary instance with n agents and $k \geq 2$ types, where there are n/k agents per type. We will show that the social welfare of any equilibrium assignment is at least $n/k - 1$. The bound on the PoA then follows, since the optimal social welfare is at most n .

Recall that we assume that the number of available nodes exceeds the number of agents and the topology is connected, so there must exist some empty node v with at least one non-empty neighbor. Suppose that v is connected to x_i agents of type T_i , for $i \in [k]$, and let $s = \sum_{i \in [k]} x_i$. Consider an agent of type T_i . A deviation to v would give her utility $\frac{x_i}{s}$ if she is not connected to v , and utility $\frac{x_i - 1}{s - 1}$ otherwise (for readability we use the convention that $\frac{0}{0} = 0$). Since at equilibrium no agent has any incentive to deviate, her utility is at least the utility she would get by deviating to v . Therefore, the social welfare at equilibrium is at least

$$\begin{aligned} \text{SW}(\mathbf{v}) &\geq \sum_{i \in [k]} \left(\left(\frac{n}{k} - x_i \right) \frac{x_i}{s} + x_i \frac{x_i - 1}{s - 1} \right) \\ &\geq \frac{1}{s} \sum_{i \in [k]} \left(\left(\frac{n}{k} - x_i \right) x_i + x_i(x_i - 1) \right) = \frac{n}{k} - 1. \end{aligned}$$

The proof is complete. \square

In the setting considered in Theorem 5.1 the PoA improves significantly if each type has the same number of agents. In the presence of stubborn agents, to ensure that the PoA does not depend on n , we additionally require that this constraint holds both for strategic and for stubborn agents.

Theorem 5.2. *For k -typed Schelling games with n agents the PoA*

- is $\Omega(n)$ for each $k \geq 2$ even if each type has the same number of agents;
- is $k + o(1)$ if each type has the same number of strategic agents and the same number of stubborn agents.

Finally, we show that even the best equilibrium need not be socially optimal, even if all agents are strategic.

Theorem 5.3. *For k -typed Schelling games the PoS*

- can be unbounded for each $k \geq 2$;
- is at least 3 for each even $k \geq 2$, if there is the same number of stubborn agents per type;

- is at least $34/33$ for each $k \geq 2$, even in the absence of stubborn agents.

6 Variants and Extensions

Throughout this paper, we focused on a setting where agents are classified into k types and their utilities are defined by the proportion of their friends among their neighbors. In this section, we introduce three variants of this model and briefly discuss some preliminary results; a more thorough investigation of these alternative models is left for future work.

Schelling Games with Social Networks

In k -typed Schelling games, the friendship relation is defined by types: an agent’s set of friends consists of all agents of the same type. One can also consider a more general friendship relation, defined by an arbitrary undirected graph \mathcal{G} with vertex set N , which we refer to as the *social network*: the set of friends of agent i consists of all neighbors of i in \mathcal{G} . We refer to the resulting class of games as *social Schelling games*.

By definition, k -typed Schelling games form a subclass of social Schelling games: a k -typed game corresponds to a social network consisting of k cliques. Hence, our next theorem implies Theorem 3.1 in Section 3.

Theorem 6.1. *Every social Schelling game where the topology is a star or a graph of maximum degree 2 admits at least one equilibrium, which can be computed in polynomial time.*

Conversely, all our non-existence results (Theorem 3.2), hardness results (Theorems 4.1 and 4.2) and lower bounds on the PoA and PoS (Section 5) apply to social Schelling games as well. In fact, maximizing the social welfare in social Schelling games is NP-hard even if all agents are strategic (whereas our hardness reduction for k -typed games uses stubborn agents). Moreover, this hardness result holds even if G is a graph of maximum degree 2, i.e., social welfare maximization may be hard even when finding equilibria is easy.

Theorem 6.2. *Given a social Schelling game I and a rational value s , it is NP-complete to decide whether I admits an assignment with social welfare at least s . The hardness result holds even if all agents are strategic and even if G is a graph of maximum degree 2.*

Identifying special classes of social Schelling games that allow for good upper bounds on the price of anarchy and the price of stability is an interesting research direction. We note that the upper bounds in Section 5 only apply to k -typed instances with further restrictions on the structure of each type, so they cannot be extended to the social setting.

Schelling Games with Enemy Aversion

In our model, if an agent is not adjacent to any friends, it does not matter how many enemies she is adjacent to. This is also the case in fractional hedonic games: agents are indifferent between being alone and being in coalitions consisting of their enemies. This assumption makes sense when the “enemies” of an agent are simply agents that do not contribute to her welfare. However, an agent may prefer being alone to being in a group of enemies. In the context of hedonic games, such preferences are modeled by *modified*

fractional hedonic games [Olsen, 2012; Elkind *et al.*, 2016; Monaco *et al.*, 2018], where the utility of an agent in a coalition with f friends and e enemies is $\frac{f+1}{f+e+1}$, i.e., the agent herself is included in the set of her friends.

Many of our results extend to this definition of utility. For example, we can construct instances without equilibria even for 2-typed games, using ideas similar to those in the reduction of Theorem 4.1. Further, for k -typed games with a tree topology and a constant number of types, equilibrium existence can be decided in polynomial time, by adapting the proof of Theorem 4.3. However, it remains an open question if instances with no stubborn agents always admit an equilibrium in this model.

Schelling Games with Linear Utilities

Throughout the paper we assume that an agent’s utility is determined by the fraction of her friends among her neighbors. Alternatively, an agent may simply care about the number of friends in her neighborhood or the difference between the number of friends f_i and the number of enemies e_i ; more broadly, her utility may be an arbitrary linear function of f_i and e_i (in the context of hedonic games, this model corresponds to a subclass of *additively separable hedonic games*; see, e.g., the survey by Aziz and Savani [2016]). It turns out that games of this form are potential games and therefore have at least one equilibrium; also, in the absence of stubborn agents there is always an equilibrium that is socially optimal.

Theorem 6.3. *Consider a variant of the (social) Schelling model where the utility of each agent i , who is adjacent to f_i friends and e_i enemies, is $\alpha f_i - \beta e_i$ for some $\alpha, \beta \geq 0$. Then, every instance has an equilibrium assignment, which can be computed in polynomial time. Moreover, if no agent is stubborn, the price of stability is 1.*

7 Conclusions

In this paper, we investigated Schelling games on graphs, from the perspective of equilibrium analysis and the perspective of social welfare. Concerning equilibrium existence, our positive results are rather limited in scope: while an equilibrium always exists for simple topologies, like stars and paths, it may fail to exist even if the topology does not contain cycles. It would be interesting to obtain a complete characterization of topologies that guarantee existence of equilibria.

For welfare maximization, a natural question is whether one can efficiently compute assignments with nearly optimal social welfare. We note that our NP-hardness reductions are not approximation preserving, so they do not rule out this possibility. Another interesting algorithmic question is whether the problem of computing equilibria in k -typed games remains hard in the absence of stubborn agents; we conjecture that this is indeed the case, but were unable to prove it.

Acknowledgements

This work is supported by the European Research Council (ERC) under grant number 639945 (ACCORD), by the EP-SRC International Doctoral Scholars Grant EP/N509711/1, and by the KAKENHI Grant-in-Aid for JSPS Fellows number 18J00997.

References

- [Alba and Logan, 1993] Richard D. Alba and John R. Logan. Minority proximity to whites in suburbs: An individual-level analysis of segregation. *American Journal of Sociology*, 98(6):1388–1427, 1993.
- [Anshelevich *et al.*, 2008] Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [Aziz and Savani, 2016] Haris Aziz and Rahul Savani. Hedonic games. In *Handbook of Computational Social Choice*, pages 356–376. 2016.
- [Aziz *et al.*, 2017] Haris Aziz, Florian Brandl, Felix Brandt, Paul Harrenstein, Martin Olsen, and Dominik Peters. Fractional hedonic games. *CoRR*, abs/1705.10116, 2017.
- [Barmpalias *et al.*, 2014] George Barmpalias, Richard Elwes, and Andrew Lewis-Pye. Digital morphogenesis via Schelling segregation. In *Proceedings of the 55th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 156–165, 2014.
- [Barmpalias *et al.*, 2015] George Barmpalias, Richard Elwes, and Andrew Lewis-Pye. From randomness to order: unperturbed Schelling segregation in two or three dimensions. *CoRR*, abs/1504.03809, 2015.
- [Benard and Willer, 2007] Stephen Benard and Robb Willer. A wealth and status-based model of residential segregation. *Journal of Mathematical Sociology*, 31(2):149–174, 2007.
- [Benenson *et al.*, 2009] Itzhak Benenson, Erez Hatna, and Ehud Or. From Schelling to spatially explicit modeling of urban ethnic and economic residential dynamics. *Sociological Methods and Research*, 37(4):463–497, 2009.
- [Bogomolnaia and Jackson, 2002] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- [Brandt *et al.*, 2012] Christina Brandt, Nicole Immorlica, Gautam Kamath, and Robert Kleinberg. An analysis of one-dimensional Schelling segregation. In *Proceedings of the 44th Symposium on Theory of Computing Conference (STOC)*, pages 789–804, 2012.
- [Chauhan *et al.*, 2018] Ankit Chauhan, Pascal Lenzner, and Louise Molitor. Schelling segregation with strategic agents. In *Proceedings of the 11th International Symposium on Algorithmic Game Theory (SAGT)*, pages 137–149, 2018.
- [Clark and Fossett, 2008] William Clark and Mark Fossett. Understanding the social context of the Schelling segregation model. *Proceedings of the National Academy of Sciences*, 105(11):4109–4114, 2008.
- [Drèze and Greenberg, 1980] Jacques H. Drèze and Joseph Greenberg. Hedonic coalitions: optimality and stability. *Econometrica*, 48(4):987–1003, 1980.
- [Easley and Kleinberg, 2010] David A. Easley and Jon M. Kleinberg. *Networks, Crowds, and Markets – Reasoning about a Highly Connected World*. Cambridge University Press, 2010.
- [Elkind *et al.*, 2016] Edith Elkind, Angelo Fanelli, and Michele Flammini. Price of Pareto optimality in hedonic games. In *Proceedings of the 20th AAAI Conference on Artificial Intelligence (AAAI)*, pages 475–481, 2016.
- [Elkind *et al.*, 2019] Edith Elkind, Jiarui Gan, Ayumi Igarashi, Warut Suksompong, and Alexandros A. Voudouris. Schelling games on graphs. *CoRR*, abs/1902.07937, 2019.
- [Immorlica *et al.*, 2017] Nicole Immorlica, Robert Kleinberg, Brendan Lucier, and Morteza Zadimoghaddam. Exponential segregation in a two-dimensional Schelling model with tolerant individuals. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 984–993, 2017.
- [Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.
- [Monaco *et al.*, 2018] Gianpiero Monaco, Luca Moscardelli, and Yllka Velaj. Stable outcomes in modified fractional hedonic games. In *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 937–945, 2018.
- [Olsen, 2012] Martin Olsen. On defining and computing communities. In *Proceedings of the 18th Computing: Australasian Theory Symposium (CATS)*, pages 97–102, 2012.
- [Pancs and Vriend, 2007] Romans Pancs and Nicolaas J. Vriend. Schelling’s spatial proximity model of segregation revisited. *Journal of Public Economics*, 91(1–2):1–24, 2007.
- [Schelling, 1969] Thomas C. Schelling. Models of segregation. *American Economic Review*, 59(2):488–493, 1969.
- [Schelling, 1971] Thomas C. Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186, 1971.
- [Young, 2001] H. Peyton Young. *Individual Strategy and Social Structure: an Evolutionary Theory of Institutions*. Princeton University Press, 2001.
- [Zhang, 2004a] Junfu Zhang. A dynamic model of residential segregation. *Journal of Mathematical Sociology*, 28(3):147–170, 2004.
- [Zhang, 2004b] Junfu Zhang. Residential segregation in an all-integrationist world. *Journal of Economic Behavior and Organization*, 54(4):533–550, 2004.