

Partitioning Friends Fairly

Lily Li, Evi Micha, Aleksandar Nikolov, Nisarg Shah

Department of Computer Science, University of Toronto
 {xinyuan, emicha, anikolov, nisarg}@cs.toronto.edu

Abstract

We consider the problem of partitioning n agents in an undirected social network into k almost equal in size (differing by at most one) groups, where the utility of an agent for a group is the number of her neighbors in the group. The core and envy-freeness are two compelling axiomatic fairness guarantees in such settings. The former demands that there be no coalition of agents such that each agent in the coalition has more utility for that coalition than for her own group, while the latter demands that no agent envy another agent for the group they are in. We provide (often tight) approximations to both fairness guarantees, and many of our positive results are obtained via efficient algorithms.

Introduction

The computer science department at University X is organizing a visit day for its newly admitted students. One of the most anticipated activity is the campus tour, during which the admitted students get to see the department they might one day join. Due to COVID-19 related capacity restrictions, the admitted students are divided into k separate tours. But more tours means more volunteers. Luckily, n current graduate students have agreed to help. We want to partition them almost equally between the k tours so that all the admitted students have equal opportunity to socialize with the current students. However, the current students have developed a social network and, to ensure a pleasant experience to encourage them to volunteer in the future, we would like to assign volunteers to tours with as many of their friends as possible.

In this paper, we introduce and study a model that captures such real-life applications. Specifically, we consider the problem of partitioning n agents into k almost equal-sized groups (each with size either $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$), when the agents are connected via an undirected social network indicating *friendships*. An agent's utility for being part of a group is the number of her friends who are in that group.

Formally, this model sits within the *hedonic games* formalism in cooperative game theory with nontransferable utilities (Aziz and Savani 2016). Two compelling axiomatic guarantees that have received significant attention in this literature are *the core* (Gillies 1953), which informally requires that there be no *deviating coalition* of agents such that

each agent in the coalition has strictly more utility for the coalition than for her group in the given partition, and *envy-freeness* (George and Marvin 1958), which informally requires that no agent receive strictly more utility when swapping places with another agent in the given partition. However, this literature typically does not impose any restriction on the partition (including on the number of groups it has). This would make our problem trivial because the grand coalition — a single group containing all agents — would trivially satisfy both the core and envy-freeness requirements. To study the core and envy-freeness meaningfully, this literature allows agents to have negative utility for other agents. We are interested in the case where the utilities are non-negative, but we require there to be exactly k groups and the groups to be of approximately equal sizes.¹

This problem can be viewed as a multi-dimensional generalization of the *stable roommates problem* (Irving 1985), in which the goal is to partition $2n$ agents between n rooms of capacity 2 each, and agents have preferences over who they wish to have as a roommate. The core becomes a notion of stability: if a pair of agents prefer each other to their assigned roommates, they may actually deviate and rent a room by themselves. But the core has also been studied in contexts where groups cannot really deviate, such as allocation of public resources (Aziz et al. 2017; Fain, Munagala, and Shah 2018; Conitzer et al. 2019) and clustering (Chen et al. 2019; Micha and Shah 2020). In such cases, it serves as a notion of *group fairness*, generalizing *proportionality* (Steinhaus 1948), because it posits that each set of $\lfloor n/k \rfloor$ or $\lceil n/k \rceil$ agents is entitled to be a group on its own and demands that the agents in the set be treated no worse than if they were their own group. Envy-freeness, on the other hand, serves as a notion of *individual fairness* and requires that no agent envy another agent for the group they belong to.

Our Results

For the core, we study bicriteria approximations of the form (α, β) -core, where a deviating coalition must improve the utility of each of its members by more than a multiplicative

¹In section *Beyond Balancedness*, we briefly consider imposing only the former restriction, allowing k arbitrarily-sized non-empty groups.

factor of α and an additive factor of β .²

We begin with the most well-studied case of $k = 2$. We show that the $(2, 0)$ -core is non-empty, and a 2-partition in the $(3, 0)$ -core can be found in polynomial time. For larger k , we note that a k -partition in the $(1, k)$ -core always exists when $n < k^2 + k$ and provide a lower bound proving that this guarantee is asymptotically the best possible with respect to an additive approximation. We show that a finite multiplicative approximation of the core is possible in general, but when $n \geq k^2 + k$, a min k -cut (a k -partition minimizing the so-called “cut size”) is in the $(2k - 1, 0)$ -core. While finding a min k -cut is known to be NP-hard, we present a polynomial-time algorithm that finds a (different) partition with the same approximation guarantee. We also show that min k -cuts cannot provide an asymptotically better approximation guarantee (i.e., our analysis is almost tight), and conjecture that no algorithm can.

For envy-freeness, we consider a similar additive approximation, EF- r , where an agent’s utility must not increase by more than r when swapping places with another agent. We make a connection to discrepancy theory (Chen et al. 2014) to show that a EF- $O(\sqrt{\frac{n}{k}} \cdot \log k)$ partition always exists, and it can be computed efficiently. We conjecture that a EF-2 partition may always exist for any k .

Finally, in section *Beyond Balancedness*, we consider a classical variation of our model, where the only requirement is to create k non-empty groups but the groups can be of arbitrary sizes. We study both the core and envy-freeness in this case and provide several tight approximation bounds.

Related Work

Our work can be viewed as a hedonic game with symmetric, binary, and additively separable preferences, and with the restriction that the partition produced have exactly k almost equal-sized parts. Hedonic games with restrictions on the number of groups have been studied before, but under other objectives, such as swap stability (Bild, Monaco, and Moscardelli 2022), Pareto optimality (Cseh, Fleiner, and Harján 2019). Moreover, Sless et al. (2018) consider a strictly weaker version of the core, without the size constraint, for which they provide positive results that do not carry over our imbalanced case. As noted in the introduction, our model is a generalization of the stable roommates problem of partitioning $2n$ agents into n pairs, where the widely studied notion of stability coincides with the core. In this problem, with asymmetric preferences a solution in the core does not always exist — unlike in the bipartite version, referred to as the *stable marriage problem*, in which it is guaranteed to exist (Gale and Shapley 1962) — but can be found in polynomial time when it does (Irving 1985). When preferences are symmetric, however, a solution in the core always exists and can be found efficiently; for instance, one can repeatedly match and remove a pair of agents with the highest utility. The three-dimensional version of this problem — partitioning $3n$ agents into groups of size 3 each — has also received significant attention. In this case, even with

²That is, an agent with utility u must receive utility more than $\alpha u + \beta$ after deviating.

symmetric additive preferences, a solution in the core may not exist (Arkin et al. 2009), and checking whether it does is NP-hard (Chen and Roy 2021). However, if we further restrict the preferences to be binary, then McKay and Manlove (2021) show that a solution in the core always exists and can be found efficiently. Our problem can be seen as a multi-dimensional generalization of the roommate problem with symmetric binary additive preferences.

Envy-freeness has been studied recently in the hedonic games literature (Barrot and Yokoo 2019), again with possibly negative utilities. Another concept similar to envy-freeness is *Nash-stability* (Bogomolnaia and Jackson 2002; Olsen, Bækgaard, and Tambo 2012), which requires that no agent be happier by *joining* another part (rather than by swapping places with an agent in another part).³ In our graph theoretic framework, this is equivalent to asking that each node have at least as many neighbors in its own part as in any other part. This has been studied extensively in graph theory using terms such as satisfactory partitions (Bazgan, Tuza, and Vanderpooten 2010), friendly partitions (Aharoni, Milner, and Prikry 1990), and internal partitions (Ban and Linial 2016), but under only the restriction that each part is non-empty. This problem is also studied in the case where the parts are required to be of almost the same size (Bazgan, Tuza, and Vanderpooten 2010). However, since such partitions do not always exist, this literature primarily focuses on the computational complexity of checking the existence of such partitions and approximating the most satisfactory partitions.

Instead, our focus is on providing worst-case guarantees on the necessary violation of envy-freeness, as is commonly done in the literature on fair resource allocation (Lipton et al. 2004; Caragiannis et al. 2019; Aziz et al. 2019). We make a connection to discrepancy theory (Chen et al. 2014) to establish an $O(\sqrt{n})$ bound. In discrepancy theory, the goal is to distribute each agent’s friends as evenly as possible between the parts, so that not only does an agent not have many more friends in another part than her own part, she also does not have many more friends in her own part than in any other part. The latter restriction, a flipped version of the satisfactory partition problem, has also been studied separately as the co-satisfactory or unfriendly partition problem (Aharoni, Milner, and Prikry 1990). Manurangsi and Suksompong (2021) use discrepancy theory in a similar problem with n agents partitioned into k groups, but with the agents having utilities over goods being allocated to the groups, not over the other agents.

Preliminaries

For $t \in \mathbb{N}$, let $[t] = \{0, \dots, t-1\}$. We consider a set $V = [n]$ of agents who are members of a social network. The network is represented by an undirected graph $G(V, E)$, where the agents are the nodes and an edge $(i, i') \in E$ indicates friendship between agents i and i' . This induces the utility function of agent i , denoted $u_i : V \rightarrow \{0, 1\}$, where $u_i(i') = 1$ if $(i, i') \in E$ and 0 otherwise. Let $N_G(i)$ denote

³The two differ only when the other part consists entirely of the agent’s friends.




Figure 1: An instance with $K_{n/2+2}$ (blue nodes) and $n/2 - 2$ isolated agents (white nodes), where partitions that achieve the best approximation with respect to core achieve the worst approximation with respect to envy-freeness, and vice versa.

the set of neighbors of agent i in G , i.e., $N_G(i) = \{i' \in V : (i, i') \in E\}$. We refer to $d_G(i) = |N_G(i)|$ as the degree of agent i . We omit G when it is clear from the context.

A k -partition of V is given by $X = (X_0, \dots, X_{k-1})$, where $X_j \cap X_{j'} = \emptyset$ for all distinct $j, j' \in [k]$; $X_j \neq \emptyset$ for all $j \in [k]$; and $\cup_{j \in [k]} X_j = V$. We may refer to an individual group X_j as a *part*. With slight abuse of notation, we denote by $X(i)$ the part X_j to which agent i belongs (i.e., $i \in X_j$). We assume that $n \geq k$, so a k -partition exists. A k -partition is called *balanced* if $\lfloor n/k \rfloor \leq |X_j| \leq \lceil n/k \rceil$ for all $j \in [k]$, and called *imbalanced* otherwise. Hereinafter, whenever we refer to a k -partition, we mean a balanced k -partition, unless explicitly specified otherwise. The utility of agent i for $S \subseteq V$ is denoted by, with slight abuse of notation, $u_i(S)$. We assume that utilities are additive, i.e., $u_i(S) = \sum_{i' \in S} u_i(i') = |S \cap N(i)|$.

In this work, we focus on two fairness criteria. The first one is the *core* which, informally, requires that there be no group of agents (coalition) of size $\lfloor n/k \rfloor \leq |S| \leq \lceil n/k \rceil$ such that every agent in the coalition prefers to be in that coalition than in her own part; such a coalition is called “blocking”.

Definition 1. Fix $\alpha \geq 1$ and $\beta \geq 0$. A coalition $S \subseteq V$ is called (α, β) -blocking for a k -partition X if

$$u_i(S) > \alpha \cdot u_i(X(i)) + \beta$$

for every $i \in S$. A k -partition X is said to be in the (α, β) -core if there is no (α, β) -blocking coalition S with $\lfloor n/k \rfloor \leq |S| \leq \lceil n/k \rceil$. When $\alpha = 1$ and $\beta = 0$, we simply use the terms blocking coalition, and core.

Another fairness criterion we focus on is envy-freeness.

Definition 2. For $r \geq 0$, a k -partition X is called *envy-free up to r* , denoted EF r or EF- r , if, for every pair of agents $i, i' \in V$, $u_i(X(i)) \geq u_i(X(i') \cup \{i\} \setminus \{i'\}) - r$. When $r = 0$, we simply refer to this as envy-freeness (EF).

For the proof techniques we plan to use, we need the following additional terminology. The *cut size* of a k -partition X , denoted $\text{cut}(X)$, is the number of edges between its different parts, i.e., $\text{cut}(X) = |\{(i, i') \in E : X(i) \neq X(i')\}|$. A k -partition with the smallest cut size is called a *min k -cut*. Note that

$$\text{cut}(X) = \sum_{i \in V} (|N(i)| - u_i(X(i))) = 2|E| - \sum_{i \in V} u_i(X(i)).$$

Hence, min k -cut also maximizes the social welfare among all k -partitions. Some of our results show that such solutions also satisfy good approximations of the core. Give disjoint sets of nodes A and B , $E(A, B)$ denotes the set of edges with one endpoint in A and the other in B .

Let us also introduce standard graph theory terminology. We denote by K_n , K_{n_1, n_2} and K_{n_1, n_2, n_3} the complete undirected graph of n vertices; the complete bipartite graph with n_1 and n_2 vertices on the two sides; and the complete tripartite graph with n_1 , n_2 , and n_3 vertices on the three sides, respectively.

Core vs. Envy-Freeness

While both core and envy-freeness are notions of fairness, they are quite different desiderata and it is easy to see that they are incompatible with each other in general. In fact, in our setting, there are instances in which every partition in the core achieves the worst possible approximation with respect to envy-freeness, and every partition that achieves the best approximations with respect to envy-freeness, achieves the worst possible approximation with respect to the core.

To see that, consider a graph that consists of the clique $K_{n/2+2}$ and $n/2 - 2$ isolated nodes, where n is a multiple of 4 and $k = 2$. The only 2-partitions in the core are the ones in which $n/2$ nodes from the clique form one group, say X_0 , and the remaining nodes form the other group, say X_1 . On the other hand, the 2-partitions that achieve the best approximation with respect to envy-freeness (in particular, EF1) are the ones in which each group contains $n/4 + 1$ nodes of the clique along with $n/4 - 1$ isolated nodes. These two types of partitions are illustrated in Figure 1.

Intuitively, the two partitions are complete opposites of each other: the former divides the clique as unequally as possible between the two groups, while the latter divides the clique exactly equally. The decision of which of the two notions (and correspondingly, partitions) is more desirable depends on the application at hand. Hence, we study both these notions separately in this paper.

Core

In this section, we study k -partitions in the (approximate) core. We start from the important case that $k = 2$, which has received particular attention in the literature on satisfactory

partitions (Bazgan, Tuza, and Vanderpooten 2010). First, we point out an interesting open question:

Open Question 1. *Does every graph admit a 2-partition in the core?*

While the existence of 2-partitions in the core is unsettled, we show that the $(2, 0)$ -core is always non-empty, and in particular, contains every min 2-cut.

Theorem 1. *For $k = 2$, a min 2-cut is in the $(2, 0)$ -core.*

Proof. Let $X = (X_0, X_1)$ be a min 2-cut. Suppose for contradiction that there exists a $(2, 0)$ -blocking coalition S of size $\lceil n/2 \rceil$ or $\lfloor n/2 \rfloor$. Let $X_0^* = X_0 \cap S$ and $X_1^* = X_1 \cap S$.

For each agent $i \in X_0^*$, $i \in S$ implies $u_i(S) > 2 \cdot u_i(X_0)$, which in turn implies $|N(i) \cap X_1^*| > 2 \cdot |N(i) \cap X_0 \setminus X_0^*|$. Summing over all $i \in X_0^*$, we obtain

$$E(X_0^*, X_1^*) > 2 \cdot E(X_0^*, X_0 \setminus X_0^*).$$

Similarly, for each agent $i \in X_1^*$, we have $|N(i) \cap X_0^*| > 2 \cdot |N(i) \cap X_1 \setminus X_1^*|$. Summing over all $i \in X_1^*$, we get

$$E(X_0^*, X_1^*) > 2 \cdot E(X_1^*, X_1 \setminus X_1^*).$$

Combining the two equations, we have

$$\begin{aligned} E(X_0^*, X_1^*) &> 2 \cdot \max\{E(X_0^*, X_0 \setminus X_0^*), \\ &\quad E(X_1^*, X_1 \setminus X_1^*)\} \\ &\geq E(X_0^*, X_0 \setminus X_0^*) + E(X_1^*, X_1 \setminus X_1^*). \quad (1) \end{aligned}$$

Now, consider the 2-partition $X' = (S, V \setminus S)$. We will show that $\text{cut}(X) > \text{cut}(X')$, which will contradict X being a min 2-cut. We have

$$\begin{aligned} \text{cut}(X) &= E(X_0, X_1) \\ &= E(X_0^*, X_1^*) + E(X_0^*, X_1 \setminus X_1^*) \\ &\quad + E(X_1^*, X_0 \setminus X_0^*) + E(X_0 \setminus X_0^*, X_1 \setminus X_1^*) \\ &\geq E(X_0^*, X_1^*) + E(X_0^*, X_1 \setminus X_1^*) \\ &\quad + E(X_1^*, X_0 \setminus X_0^*) \\ &> E(X_0^*, X_0 \setminus X_0^*) + E(X_1^*, X_1 \setminus X_1^*) \\ &\quad + E(X_0^*, X_1 \setminus X_1^*) + E(X_1^*, X_0 \setminus X_0^*) \\ &= \text{cut}(X'), \end{aligned}$$

where the strict inequality uses Equation (1). This is the desired contradiction. \square

While Theorem 1 is a strong existential result, it does not come with an efficient algorithm as finding a min 2-cut (also known as the minimum bisection problem) is NP-hard (Garey and Johnson 1979). This leads to our next open problem:

Open Question 2. *Can a 2-partition in the $(2, 0)$ -core be computed in polynomial time?*

As a slight remedy, we will later show that at least a 2-partition in the $(3, 0)$ -core can be computed efficiently.

Next, we focus on $k \geq 3$ and show that the core can be empty in this case.

Theorem 2. *When $k \geq 3$, there exists an instance in which no k -partition is in the $(\alpha, 0)$ -core for any $\alpha \geq 1$, and also there exists an instance in which no k -partition is in the $(1, \beta)$ -core for $\beta < k/2 - 2$.*

Proof. Fix $k \geq 3$. For the first claim, consider a cycle with $n = k + 1 \geq 4$ nodes. Fix an arbitrary k -partition X . Note that X must consist of one part with two nodes and $k - 1$ parts with a single node each. Without loss of generality, let X_0 be the part with $|X_0| = 2$. Note that in a cycle of length at least 4, the size of the smallest maximal matching is at least 2. Hence, there must exist agents $i, i' \notin X_0$ that are connected by an edge. Since the coalition $\{i, i'\}$ is allowed to deviate, they can both go from receiving utility 0 to receiving utility 1, implying that X is not in the $(\alpha, 0)$ -core for any $\alpha > 0$.

For the second claim, consider the graph G formed by $k + 1$ disjoint cliques of size $k - 1$ each, denoted C_0, \dots, C_k . Hence, $n = k^2 - 1$. Let X be any k -partition of G . First, we claim that there exists $\ell^* \in [k + 1]$ such that $|C_{\ell^*} \cap X_j| \leq (k + 1)/2$ for all $j \in [k]$. If this is not true, then for every $\ell \in [k + 1]$, there exists at least one $j_\ell \in [k]$ with $|C_\ell \cap X_{j_\ell}| > (k + 1)/2$. Note that such j_ℓ must be unique. Further, because $|X_j| \leq \lceil n/k \rceil \leq k + 1$ for all $j \in [k]$, we must have $j_\ell \neq j_{\ell'}$ for distinct $\ell, \ell' \in [k + 1]$. However, this is not possible as there are $k + 1$ cliques but only k parts. Now, for every agent $i \in C_{\ell^*}$, we have $u_i(X(i)) \leq (k - 1)/2$. On the other hand, C_{ℓ^*} is a feasible deviating coalition as $|C_{\ell^*}| = k - 1 = \lfloor n/k \rfloor$. Further, for every $i \in C_{\ell^*}$, we have $u_i(C_{\ell^*}) = k - 2 \geq u_i(X(i)) + (k - 3)/2$. Hence, C_{ℓ^*} is a $(1, k/2 - 2)$ -blocking coalition, as desired. \square

While above we show that the core can be empty when $k \geq 3$, these examples are somewhat unsatisfactory as they crucially rely on n not being divisible by k , which leads to another interesting open question:

Open Question 3. *Does every graph with n nodes admit a k -partition in the core, if k divides n ?*

When $n < k^2 + k$, note that any deviating coalition has size at most $\lceil n/k \rceil \leq k + 1$, which means that no agent can improve her utility by an additive factor of more than k when deviating. Hence, every k -partition is trivially in the $(1, k)$ -core.

Proposition 1. *When $n < k^2 + k$, every k -partition is in the $(1, k)$ -core.*

From Theorem 2, we know that one cannot obtain a purely multiplicative guarantee of the form $(\alpha, 0)$ -core for any $\alpha \geq 1$ and cannot obtain an additive guarantee of the form $(1, \beta)$ -core for any $\beta \leq k/2 - 2$. Thus, we conclude that the guarantee in Proposition 1 is asymptotically tight with respect to these two ways of approximation when $n < k^2 + k$.

Next, we consider the case of $n \geq k^2 + k$. Interestingly, while a purely multiplicative approximation of the core cannot be guaranteed in general, we show that this is possible when we know $n \geq k^2 + k$. Specifically, the next result shows that a k -partition in the $(2k - 1, 0)$ -core always exists and can be found in polynomial time, if $n \geq k^2 + k$. In this case, we in fact show that every min k -cut is in the $(2k - 1, 0)$ -core, but we can also use a local search procedure, presented as Algorithm 1, to obtain the same approximation guarantee efficiently. At a high level, the algorithm works as follows. It starts from an arbitrary k -partition and draws a directed edge from agent i to agent i' , with

Algorithm 1: Local Min-Cut

```

1:  $X \leftarrow$  an arbitrary  $k$ -partition
2: while true do
3:   Build a directed graph  $G' = (V', E')$  with  $V' = V$ 
   and  $E' = \{(i, i') : u_i(X(i')) > u_i(X(i)) + 1\}$ 
4:   if there is a cycle  $(i_0, i_1, \dots, i_{s-1}, i_0)$  in  $G'$  then
5:     for  $\ell \in [s]$  do
6:        $X(i_\ell) \leftarrow X(i_\ell) \setminus \{i_\ell\}$ 
7:        $X(i_{\{\ell+1 \bmod s\}}) \leftarrow X(i_{\{\ell+1 \bmod s\}}) \cup \{i_\ell\}$ 
8:     end for
9:   else if  $\exists(i, i')$  s.t.  $u_{i'}(X(i')) = 0$  and  $u_i(X(i')) >$ 
 $u_i(X(i))$  then
10:    if  $(i, i') \notin E$  or  $u_i(X(i')) > u_i(X(i)) + 1$  then
11:       $X(i) \leftarrow X(i) \cup \{i'\} \setminus \{i\}$ 
12:       $X(i') \leftarrow X(i') \cup \{i\} \setminus \{i'\}$ 
13:    end if
14:   else
15:     break
16:   end if
17: end while
18: return  $X$ 

```

$X(i) \neq X(i')$, if i envies i' by more than one. If there exists a directed cycle, say $(i_0, i_1, \dots, i_{s-1}, i_0)$, it shifts the agents along the cycle, i.e., i_0 is moved to $X(i_1)$, i_1 is moved to $X(i_2)$ and so on, while i_{s-1} is moved to $X(i_0)$. Then, it updates the directed edges and continues eliminating cycles in this fashion. When there are no cycles left, it searches for pairs of agents, i and i' , with $X(i) \neq X(i')$, such that i' has zero utility for her group and a positive utility for i 's group, i envies i' , and the envy is by more than one if i and i' are not neighbors. If such a pair exists, it swaps i and i' . Then, the algorithm updates the directed edges, search for cycles or such pairs, and eliminates them until until no cycles or such pairs are left. Throughout the procedure, the cut size strictly decreases, and we establish that the approximation guarantee holds at any local minimum.

Theorem 3. *When $n \geq k^2 + k$, every min k -cut is in the $(2k - 1, 0)$ -core, and Algorithm 1 returns a k -partition in the $(2k - 1, 0)$ -core in polynomial time.*

Proof. First, we show that Algorithm 1 terminates in polynomial time by arguing that $\text{cut}(X)$ strictly decreases in every iteration of the while loop. If we find a cycle in Line 4, then during the cyclic shift of nodes along this cycle, each node gains at least 1 utility. Since the social welfare strictly increases, $\text{cut}(X)$ strictly decreases. Similarly, if we find two agents i and i' such that i' has no neighbors in $X(i')$ but i has at least two more neighbors in $X(i')$ than in $X(i)$, then swapping i and i' also strictly decreases the cut size. Further, if i' is not a neighbor of i , then we only need i to have at least one more neighbor in $X(i')$ than in $X(i)$. Hence, in any case, $\text{cut}(X)$ strictly reduces in every iteration of the while loop, resulting in termination in polynomial time.

Let X be either a min k -cut or the output of Algorithm 1. Suppose for contradiction that there is a $(2k - 1, 0)$ -blocking coalition S of size $\lceil n/k \rceil$ or $\lfloor n/k \rfloor$. We first show the follow-

ing lemma.

Lemma 1. *For $i \in S$, if $u_i(S \cap X_j) \leq u_i(X(i)) + 1$ for each $j \in [k]$, then $u_i(X(i)) = 0$.*

Proof. Suppose there exists $i \in S$ with $u_i(S \cap X_j) \leq u_i(X(i)) + 1$ for each $j \in [k]$ but $u_i(X(i)) \geq 1$. Then,

$$\begin{aligned}
u_i(S) &= \sum_{j \in [k]} u_i(S \cap X_j) \\
&\leq (k-1)(u_i(X(i)) + 1) + u_i(X(i)) \\
&\leq 2(k-1) \cdot u_i(X(i)) + u_i(X(i)) \\
&= (2k-1) \cdot u_i(X(i)),
\end{aligned}$$

which contradicts S being a $(2k - 1, 0)$ -blocking coalition. \square

Suppose that there exists $i \in S$ such that $u_i(S \cap X_j) > u_i(X(i)) + 1$ for some $j \in [k]$. Let G' be the directed graph constructed from X according to Line 3 of Algorithm 1. Then, there must be an edge from i to every node in X_j in G' , as $u_i(X(i)) + 1 < u_i(S \cap X_j) \leq u_i(X_j)$. Further, since $u_i(S \cap X_j) > 0$, $S \cap X_j \neq \emptyset$. Hence, i has an edge to some node in S in G' . Note that there can be no cycle in G' : if X is the output of Algorithm 1, this would contradict the while loop terminating, and if X is a min k -cut, a cyclic shift of nodes like in Algorithm 1 would reduce the cut size, which would be a contradiction. Since there is no cycle in G' , consider the longest path in G' starting at i and only containing nodes in S . Suppose it is (i, i_1, \dots, i_t, i') . Then, i' must satisfy the condition of Lemma 1, otherwise by the same reasoning as before, there would exist $j' \in [k]$ such that $S \cap X_{j'} \neq \emptyset$ and i' has edges to all nodes in $X_{j'}$ in G' . This would either lead to a cycle or a longer path in G' starting at i and only containing nodes in S , which is a contradiction. Since i' satisfies the condition of Lemma 1, we have $u_{i'}(X(i')) = 0$. We also have $u_{i_t}(X(i')) > u_{i_t}(X(i_t)) + 1$. If X is returned by Algorithm 1, we get a contradiction because Algorithm 1 would have continued by swapping i_t and i' in Line 9. If X is a min k -cut, then swapping i_t and i' would reduce the cut size, which would be a contradiction.

We have established that all $i \in S$ satisfy the condition from Lemma 1. Hence, $u_i(X(i)) = 0$ for all $i \in S$. But, since $n \geq k^2 + k$, we have $|S| \geq \lfloor n/k \rfloor \geq k + 1$, which implies that there exist $i_1, i_2 \in S$ with $X(i_1) = X(i_2)$. If $(i_1, i_2) \in E$, then we have a contradiction with $u_{i_1}(X(i_1)) = u_{i_2}(X(i_2)) = 0$. Otherwise, there exists $i_3 \in S$, with $X(i_3) \neq X(i_1)$ and $(i_1, i_3) \in E$. But, then we would have that either $(i_2, i_3) \notin E$ or $u_{i_3}(X_2) > u_{i_3}(X_3) + 1$, which means i_2 and i_3 would have been swapped under Algorithm 1, and hence this partition would not be a min cut. Thus, there is no such $(2k - 1, 0)$ -blocking coalition S . \square

In the proof of Lemma 1, note that if we assumed the deviating coalition S to be a $(k, k - 1)$ -blocking coalition, then we would obtain a contradiction regardless of whether $u_i(X(i)) = 0$ or $u_i(X(i)) \geq 1$. Since the next part of the proof, which establishes that all $i \in S$ must satisfy the condition of Lemma 1, does not assume $n \geq k^2 + k$, we have

that Algorithm 1 always finds a solution in the $(k, k - 1)$ -core. In particular, for $k = 2$, we can efficiently guarantee $(3, 1)$ -core. Recall that Theorem 1 provides a slightly better guarantee of $(2, 0)$ -core, but not in polynomial time.

Corollary 1. *For $k = 2$, Algorithm 1 returns a 2-partition in the $(3, 1)$ -core in polynomial time.*

While it is an open question whether Algorithm 1 provides the best possible guarantee, we show, using an intricate instance, that our approximation analysis of min k -cuts in Theorem 3 is almost tight from a multiplicative point of view. Missing proofs can be found in the full version.

Theorem 4. *For $k \geq 3$, there exists an instance with $n \geq k^2 + k$ in which some min k -cut is not in the $(\alpha, 0)$ -core for $\alpha < 2k - 2$.*

Envy-Freeness

We now turn our attention to finding k -partitions that are (approximate) envy-free. We start by showing that EF1 cannot always be guaranteed.

Theorem 5. *Even when $k = 2$, a 2-partition that is EF1 does not always exist.*

Proof. Consider the $K_{3,3,3}$ graph which consists of three set of three nodes each, denoted by $C_1 = \{c_{11}, c_{12}, c_{13}\}$, $C_2 = \{c_{21}, c_{22}, c_{23}\}$ and $C_3 = \{c_{31}, c_{32}, c_{33}\}$, respectively, and every node of each set is adjacent to every node in the other two sets.

For the sake of contradiction, assume that $X = (X_0, X_1)$ is a partition of the graph that is EF1. Since the graph is 6-regular, we can see that $|X_0| \geq 4$ and $|X_1| \geq 4$, as if an agent i is assigned to a part with only at most two of its neighbours, then the other four of its neighbours are assigned to the other part along with an agent i' which is not neighbour of i , and then i envies i' for more than one agent. Without loss of generality, we assume that $|X_0| = 4$. If X_0 contains three nodes of the same set, then we can easily see that this partition is not EF1, as each of them is assigned to the same group with at most one of its neighbours. As there are three sets and X_0 contains four agents, we see that two agents of the same set, say c_{11} and c_{12} , are assigned to X_0 . Then these two agents are in the same part along with at most two of its neighbours, while all the remaining nodes are assigned to X_1 . Then, c_{11} and c_{12} envy c_{13} for more than one agents, which is a contradiction. \square

To obtain non-trivial approximations to envy-freeness for higher values of k , that too via partitions, we turn to the literature on discrepancy theory. Intuitively, we want to color the elements of a set using k colors such that each pre-determined subset has an approximately equal number of elements of each color. Formally, we are given a universe $\Omega = [n]$ and a set system $\mathcal{S} = \{S_0, \dots, S_{m-1}\}$, where $S_i \subseteq [n]$ for each $i \in [m]$. The k -color discrepancy of a coloring $\chi : \Omega \rightarrow [k]$ on the set system \mathcal{S} is defined as

$$\text{disc}_k(\mathcal{S}, \chi) = \max_{j \in [k], i \in [m]} \left| |\chi^{-1}(j) \cap S_i| - |S_i|/k \right|.$$

The k -discrepancy of \mathcal{S} is then the minimum k -color discrepancy over all χ : $\text{disc}_k(\mathcal{S}) = \min_{\chi: \Omega \rightarrow [k]} \text{disc}_k(\mathcal{S}, \chi)$.

A celebrated result from this literature is that $\text{disc}_k(\mathcal{S}) = O(\sqrt{\frac{n}{k} \ln(km/n)})$ for any set system \mathcal{S} and a k -coloring achieving this bound can be computed in polynomial time (Chen et al. 2014, Corollary 44).

In our setting, with $\Omega = V = [n]$, a k -coloring $\chi : \Omega \rightarrow [k]$ induces a k -partition X given by $X_j = \chi^{-1}(j)$ for all $j \in [k]$.⁴ Further, if we consider the set system \mathcal{S} where $S_i = N_G(i)$ for each $i \in [n]$ (i.e., with $m = n$), then we are guaranteed that agent i can have at most $2\text{disc}_k(\mathcal{S}, \chi)$ more neighbors in any other part than in her own part, implying $\text{EF-}2\text{disc}_k(\mathcal{S}, \chi)$. The above discrepancy bound then immediately yields the existence of a k -partition that is $\text{EF-}O(\sqrt{\frac{n}{k} \ln k})$. However, this may not be balanced.

To fix this, we add another set $S_n = V$ to our set system; we now have $m = n + 1$, which does not asymptotically change the discrepancy bound. Now, we obtain a k -partition X that is also approximately balanced: $||X_j| - |X_{j'}|| = O(\sqrt{\frac{n}{k} \ln k})$ for all $j, j' \in [k]$. By arbitrarily moving $O(\sqrt{\frac{n}{k} \ln k})$ agents between parts, we can make it perfectly balanced, while only increasing the EF approximation by $O(\sqrt{\frac{n}{k} \ln k})$. Thus, we get the following.

Theorem 6. *For any $k \geq 2$, a k -partition that is $\text{EF-}O(\sqrt{\frac{n}{k} \ln k})$ is guaranteed to exist and can be computed in polynomial time.*

For discrepancy, the aforementioned upper bound is known to be almost tight: there is a lower bound of $\Omega(\sqrt{n/k})$ (Chen et al. 2014, Theorem 61). However, for our “one-sided” envy-freeness guarantee, achieving a constant approximation remains an open question.

Open Question 4. *Does every graph admit a k -partition that is EF2, for all $k \geq 2$?*

Beyond Balancedness

An interesting variation of our problem is to drop the requirement of balancedness and simply seek k non-empty groups, i.e., imbalanced k -partitions. This variation was first introduced by Gerber and Kobler (2000) and, since then, it has been given much attention (Bazgan, Tuza, and Vanderpooten 2010) due to its importance to real-life applications such as clustering (Flake, Tarjan, and Tsioutsoulouklis 2004; Shafique 2004).

In this section, we briefly consider this case and study both the core and envy-freeness for imbalanced k -partitions. For the core, we provide matching upper and lower bounds. For envy-freeness, we provide a complete picture for $k = 2$ by making a connection to the literature on satisfactory partitions, and point out interesting open questions for $k \geq 3$.

Core

Recall that core requires that there be no group of agents (coalition) such that every agent in the coalition prefers to be in that coalition than in her own part. In general, there is no direct correlation between the size of a coalition and

⁴Technically, we also need to ensure $X_j \neq \emptyset$, but this is guaranteed due to the discrepancy bound.

the ease with which it can be blocking.⁵ Hence, in the imbalanced case, we impose the same restriction on the size of a deviating coalition as we have on the size of a part in an imbalanced k -partition. Note that all parts in an imbalanced k -partition X are required to be non-empty, which implies $1 \leq |X_j| \leq n - k + 1$ for all $j \in [k]$; hence, we require a deviating coalition S to also satisfy $1 \leq |S| \leq n - k + 1$. This gives rise to the following variant of the core.

Definition 3. Fix $\alpha \geq 1$ and $\beta \geq 0$. A coalition $S \subseteq V$ is called (α, β) -blocking for an imbalanced k -partition X if

$$u_i(S) > \alpha \cdot u_i(X(i)) + \beta$$

for every $i \in S$. An imbalanced k -partition X is said to be in the (α, β) -imbalanced core if there is no (α, β) -blocking coalition S with $1 \leq |S| \leq n - k + 1$. When $\alpha = 1$ and $\beta = 0$, we simply use the terms blocking coalition, and imbalanced core.

Note that the differing size restrictions on the deviating coalitions technically makes our results for the core under imbalanced k -partitions incomparable to our results for the core under balanced k -partitions.

We show that for any k , an imbalanced partition in the $(1, k - 2)$ -imbalanced core always exists and can be found in polynomial time.

Theorem 7. *When $k \geq 2$, we can find an imbalanced k -partition in the $(1, k - 2)$ -imbalanced core in polynomial time, and in particular, when $k = 2$, we can efficiently find an imbalanced 2-partition in the imbalanced core. Moreover, when $k \geq 3$, there exists an instance in which no imbalanced k -partition is in the $(1, \beta)$ -imbalanced core for any $\beta < k - 2$.*

Recall that in the proof of Theorem 2, we used an example with $n = k + 1$ to establish that there may not exist any (balanced) k -partitions in the $(\alpha, 0)$ -core for any $\alpha \geq 1$ or $(1, \beta)$ -core for any $\beta < k/2 - 2$. Because all imbalanced k -partitions are also balanced for $n = k + 1$, the imbalanced core becomes equivalent to the core. Hence, the negative results of Theorem 2 carry over to the imbalanced case, though the additive lower bound in Theorem 7 is better.

This shows that the guarantee in Theorem 7 is tight in two ways: one can hope for neither a purely multiplicative approximation of the form $(\alpha, 0)$ -imbalanced core for any $\alpha \geq 1$, nor a better additive approximation of the form $(1, \beta)$ -imbalanced core for any $\beta < k - 2$.

Envy-Freeness

Finally, we turn our attention to envy-freeness. First, we use the following result from the literature on satisfactory partitions, restated in our framework, to establish the existence of an EF-2 imbalanced partition when $k = 2$.

Theorem 8 (Stiebitz 1996, Bazgan, Tuza, and Vanderpooten 2007). *Given a graph $G = (V, E)$ and functions $a, b : V \rightarrow \mathbb{N}$ such that $d(i) \geq a(i) + b(i) + 1$ for every $i \in V$, there*

⁵Smaller coalitions have the advantage of only having to improve the utility of fewer agents, whereas larger coalitions can include more friends of their members.

exists an imbalanced 2-partition $X = (X_0, X_1)$ of V such that $u_i(X_0) \geq a(i)$ for each $i \in X_0$ and $u_i(X_1) \geq b(i)$ for all $i \in X_1$, and it can be computed in polynomial time.

In our case, we use functions $a(i) = b(i) = \lfloor (d(i) - 1)/2 \rfloor$ for all $i \in V$. Note that these satisfy the condition $d(i) \geq a(i) + b(i) + 1$. Hence, the above result allows us to efficiently compute a 2-partition X satisfying $u_i(X(i)) \geq \lfloor (d(i) - 1)/2 \rfloor$ for all $i \in V$. Since there are only two parts, this also implies that for all $i, i' \in V$,

$$\begin{aligned} u_i(X(i')) - u_i(X(i)) &\leq d(i) - 2 \cdot \lfloor (d(i) - 1)/2 \rfloor \\ &\leq d(i) - 2 \cdot (d(i) - 2)/2 = 2, \end{aligned}$$

which implies that X is EF-2.

Corollary 2. *An imbalanced 2-partition that is EF-2 always exists and can be computed in polynomial time.*

Theorem 8 admits an extension to $k > 2$ parts, but in our case, this only guarantees that $u_i(X(i)) \geq \lfloor (d(i) - k + 1)/k \rfloor$ for all $i \in V$ (Bazgan, Tuza, and Vanderpooten 2007). This does not meaningfully limit the number of neighbors that agent i has in another part and, therefore, fails to provide a non-trivial approximation to envy-freeness. That said, if one is interested in the slightly weaker guarantee of proportionality (Steinhaus 1948), which, in our setting, would require $u_i(X(i)) \geq d(i)/k$, then this would provide an additive 1-approximation.

For the satisfactory partition problem, where the goal is to indeed minimize $u_i(X(i')) - u_i(X(i))$, as in the equation above, it is easy to see that an additive error of 2 is the best possible. Consider dividing any clique with an odd number of nodes into two parts. An agent i in the smaller part will have at least two more neighbors in the larger part than in her own part. However, this does not hold for envy-freeness: if i envies i' by swapping places with an agent i' from the other part, then $X(i') \cup \{i\} \setminus \{i'\}$ will only contain one more neighbor of i than $X(i)$ does. Nonetheless, notice that the example that is used in the proof of Theorem 5 can also be used to show that EF-1 cannot always be guaranteed even in the imbalanced case when $k = 2$.

Discussion

In this paper, we considered the problem of partitioning n agents into k almost equal-sized groups, when the agents have binary preferences, induced by a social network. We designed algorithms which approximately satisfy two axiomatic fairness guarantees: the core and envy-freeness. Our work offers a number of exciting open questions. For example, is the core always non-empty when $k = 2$ or when k divides n ? Does an EF-2 partition always exist?

There are natural ways to extend our model. One can consider more general preferences than symmetric and binary. Symmetric weighted preferences are particularly interesting as while one can verify that our positive result of min 2-cut carries over this case, our guarantees of min k -cut for $k \geq 3$ are not easily expandable beyond the binary case. Moreover, if the agents are described by a number of attributes, the construction of fair and diverse groups is another interesting direction.

References

- Aharoni, R.; Milner, E. C.; and Prikry, K. 1990. Unfriendly partitions of a graph. *Journal of Combinatorial Theory, Series B*, 50(1): 1–10.
- Arkin, E. M.; Bae, S. W.; Efrat, A.; Okamoto, K.; Mitchell, J. S.; and Polishchuk, V. 2009. Geometric stable roommates. *Information Processing Letters*, 109(4): 219–224.
- Aziz, H.; Brill, M.; Conitzer, V.; Elkind, E.; Freeman, R.; and Walsh, T. 2017. Justified representation in approval-based committee voting. *Social Choice and Welfare*, 48: 461–485.
- Aziz, H.; Caragiannis, I.; Igarashi, A.; and Walsh, T. 2019. Fair Allocation of Indivisible Goods and Chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 53–59.
- Aziz, H.; and Savani, R. 2016. Hedonic Games. In Brandt, F.; Conitzer, V.; Endress, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 15. Cambridge University Press.
- Ban, A.; and Linial, N. 2016. Internal partitions of regular graphs. *Journal of Graph Theory*, 83(1): 5–18.
- Barrot, N.; and Yokoo, M. 2019. Stable and envy-free partitions in hedonic games. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 67–73.
- Bazgan, C.; Tuza, Z.; and Vanderpooten, D. 2007. Efficient algorithms for decomposing graphs under degree constraints. *Discrete Applied Mathematics*, 155(8): 979–988.
- Bazgan, C.; Tuza, Z.; and Vanderpooten, D. 2010. Satisfactory graph partition, variants, and generalizations. *European Journal of Operational Research*, 206(2): 271–280.
- Bilò, V.; Monaco, G.; and Moscardelli, L. 2022. Hedonic Games with Fixed-Size Coalitions. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, 9287–9295.
- Bogomolnaia, A.; and Jackson, M. 2002. Stability of Hedonic Coalition Structures. *Games and Economic Behavior*, 38(2): 201–230.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3): 1–32.
- Chen, J.; and Roy, S. 2021. Euclidean 3D Stable Roommates is NP-hard. arXiv:2108.03868.
- Chen, W.; Srivastav, A.; Travaglini, G.; et al. 2014. *A panorama of discrepancy theory*. Springer.
- Chen, X.; Fain, B.; Lyu, L.; and Munagala, K. 2019. Proportionally Fair Clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, volume 97, 1032–1041.
- Conitzer, V.; Freeman, R.; Shah, N.; and Wortman-Vaughan, J. 2019. Group Fairness for the Allocation of Indivisible Goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 1853–1860.
- Cseh, Á.; Fleiner, T.; and Harján, P. 2019. Pareto Optimal Coalitions of Fixed Size. *Mechanism and Institution Design*, 4(13): 87–108.
- Fain, B.; Munagala, K.; and Shah, N. 2018. Fair Allocation of Indivisible Public Goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 575–592.
- Flake, G. W.; Tarjan, R. E.; and Tsioutsoulouklis, K. 2004. Graph Clustering and Minimum Cut Trees. *Internet Mathematics*, 1(4): 385–408.
- Gale, D.; and Shapley, L. S. 1962. College Admissions and the Stability of Marriage. *American Mathematical Monthly*, 69(1): 9–15.
- Garey, M. R.; and Johnson, D. S. 1979. *Computers and Intractability: a Guide to the Theory of NP-Completeness*. W. H. Freeman and Company.
- George, G.; and Marvin, S. 1958. *Puzzle-math*. New York, NY, USA: Viking Press.
- Gerber, M. U.; and Kobler, D. 2000. Algorithmic approach to the satisfactory graph partitioning problem. *European Journal of Operational Research*, 125(2): 283–291.
- Gillies, D. B. 1953. *Some theorems on n-person games*. Princeton University.
- Irving, R. W. 1985. An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms*, 6(4): 577–595.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, 125–131.
- Manurangsi, P.; and Suksompong, W. 2021. Almost Envy-Freeness for Groups: Improved Bounds via Discrepancy Theory. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 335–341.
- McKay, M.; and Manlove, D. 2021. The Three-Dimensional Stable Roommates Problem with Additively Separable Preferences. In *Proceedings of the 14th International Symposium on Algorithmic Game Theory (SAGT)*, 266–280.
- Micha, E.; and Shah, N. 2020. Proportionally Fair Clustering Revisited. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, 85:1–85:16.
- Olsen, M.; Bækgaard, L.; and Tambo, T. 2012. On non-trivial Nash stable partitions in additive hedonic games with symmetric 0/1-utilities. *Information Processing Letters*, 112(23): 903–907.
- Shafique, K. H. 2004. *Partitioning a graph in alliances and its application to data clustering*. Ph.D. thesis, University of Central Florida.
- Sless, L.; Hazon, N.; Kraus, S.; and Wooldridge, M. 2018. Forming k coalitions and facilitating relationships in social networks. *Artificial Intelligence*, 259: 217–245.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16: 101–104.
- Stiebitz, M. 1996. Decomposing graphs under degree constraints. *Journal of Graph Theory*, 23(3): 321–324.