A Unified Framework for Strong Price of Anarchy in Clustering Games*

Michal Feldman and Ophir Friedler

Blavatnik School of Computer Science Tel Aviv University

Abstract. We devise a unified framework for quantifying the inefficiency of equilibria in clustering games on networks. This class of games has two properties exhibited by many real-life social and economic settings: (a) an agent's utility is affected only by the behavior of her direct neighbors rather than that of the entire society, and (b) an agent's utility does not depend on the actual strategies chosen by agents, but rather by whether or not other agents selected the same strategy. Our framework is sufficiently general to account for unilateral versus coordinated deviations by coalitions of different sizes, different types of relationships between agents, and different structures of strategy spaces. Many settings that have been recently studied are special cases of clustering games on networks. Using our framework: (1) We recover previous results for special cases and provide extended and improved results in a unified way. (2) We identify new settings that fall into the class of clustering games on networks and establish price of anarchy and strong price of anarchy bounds for them.

1 Introduction

Suppose that mobile phone provider offer a significant discount for calls between their subscribers. In such a case, users selecting a provider would benefit the most by subscribing to the provider of the friends with whom they talk most. Alternatively, consider radio stations selecting radio frequencies on which to broadcast. Because nearby stations that select the same frequency incur interference, each station would favor a frequency that is used the least by its nearby stations. Finally, in some opinion formation settings, an agent forming an opinion aims to have the same opinion as similar agents and an opinion that is different from dissimilar agents.

At a first glance the different settings described above seem different in their nature. In the first example, people want to make similar choices to others, while in the second example they wish to differentiate themselves from others, and in the third example their choice depends on the type of relationship they have with others. Indeed, different models for these real-life settings have been

* This work was partially supported by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement number 337122.

devised and studied in the algorithmic game theory literature [15, 13, 14, 18, 2]. For example, in coordination games on graphs [2], each agent has her own set of feasible options, and an agent benefits by sharing her option with as many friends as possible. In Max-Cut games [15], all agents choose one of two options, and they all seek to distinguish themselves from others. In 2-NAE-SAT games [13], CNF clauses are satisfied when their literals have different values. Agents correspond to literals that select truth values, where agents derive utility from the satisfied clauses of which they are part of. In this case, an agent may wish to share her truth value with some (e.g. in a clause where she is negated, and the other literal is not), while distinguishing herself from others.

Inefficiency of Equilibria. It is well known that settings in which individual agents follow their self interests may exhibit economic inefficiencies. A great deal of the work in the algorithmic game theory literature has attempted to quantify the inefficiencies that may arise as a result of selfish behavior. The most common measure that has been used is the *price of anarchy* (PoA) [17], which is defined as the ratio of the welfare obtained in the worst Nash equilibrium (NE) and the unconstrained optimal welfare. But because of some weaknesses of the NE solution concept, researchers have turned to additional solution concepts, such as *strong equilibrium* (SE), and q*-strong equilibrium* (q-SE). Whereas an NE is an outcome that is resilient against unilateral deviations, a q-SE is resilient against coordinated deviations of coalitions of size at most q. Therefore, a 1-SE is a NE, and an n-SE is resilient against *any* coordinated coalitional deviation, i.e., an n-SE is an SE. In settings in which agents can coordinate their actions (in particular their deviations), the SE concept is more suitable. One can then study the *strong price of anarchy* (SPoA) [1], which is the worst case ratio of an SE and the social optimum, or the q-SPoA, in which the SE is replaced by the q-SE.

The inefficiency of equilibria, as measured by the PoA (and in some cases also the SPoA and q-SPoA), has been studied in all of the above-mentioned games [15, 13, 14, 18, 2]. Because these models correspond to different settings, each work developed its own tools and techniques to compute the (q-)strong price of anarchy.

Clustering. The starting point of this paper is the observation that all the above-mentioned settings are special cases of a very well known setting — that of *clustering* [5]. In this setting, a network is given, with each edge labeled as either a *coordination* or an *anti-coordination* edge. The objective is to find a clustering (i.e., a partition of the nodes into clusters) that maximizes the number of edges (or the total weight of edges in the weighted case) that are *satisfied*, where a coordination edge is satisfied if its adjacent nodes are in the same cluster, and an anti-coordination edge is satisfied if its adjacent nodes are in different clusters.

The above description represents the traditional perspective on the problem, but one can also consider its game-theoretic variant. In the game-theoretic setting each node corresponds to a strategic agent whose strategy space is a feasible subset of all clusters. Each strategic agent chooses a cluster, and given the selected clusters of all the agents (i.e. an *outcome* of the game), the payoff of an agent is the number (or total weight, in the weighted case) of the satisfied edges to which she is incident. We refer to the class of games that arises from this setting as *clustering games on networks*.

Let us recall some of the classes mentioned above, and restate them as special cases of clustering games on networks. In coordination games on graphs [2], each agent has a specified subset of clusters she can join, and all edges are coordination edges. In Max-k-Cut games [15] all the edges are anti-coordination edges, and all agents can join any of the k existing clusters. In 2-NAE-SAT games [13], each agent can join one of two clusters, and edges can be either coordination or anti-coordination edges.

As mentioned above, each case has its own analysis. In light of the observation that all of these settings are special cases of clustering games on networks, the objective of the persent paper is the following:

> *Construct a unified recipe for quantifying the degradation of social welfare (i.e., PoA, SPoA, q-SPoA) in various settings that fall into the class of clustering games on networks.*

The class of clustering games on networks is a rich class that lies at the intersection of two important classes of games: hedonic [11] and graphical games [16]. In hedonic games the utility of each player is fully determined by the set of players that selected the same strategy. They serve as a fundamental model in the study of coalition formation, which has implications for many social, economic, and political scenarios [11, 4, 9, 10, 8, 6]. In graphical games (see [16, 20] and references therein) the agents are nodes in a graph, and the payoff of each agent depends strictly on the strategies of her neighbors. Such games attempt to capture the local nature of interactions in a network.

Our Contribution is four-fold:

1. We provide a unified framework for computing the PoA, SPoA, and q-SPoA in clustering games on graphs.
2. We use our framework to recover previous results on special cases.
3. We use our framework to establish new PoA, SPoA, and q-SPoA bounds on previously studied games.
4. We identify new settings that fall into the class of clustering games on networks and establish PoA, SPoA, and q-SPoA bounds for them.

A summary of our results appears in Fig. 1. The different rows correspond to different games, whose description is given in the three left columns. For each class of games, we provide results on the PoA, SPoA, and q-SPoA. For q-SPoA, we find it convenient to present our results using the *coordination factor* $z(q) = \frac{q-1}{n-1}$, which is a real number in $[0, 1]$ that equals 0 in the case of PoA (i.e., $q = 1$) and 1 in the case of SPoA (i.e., $q = n$). Because of lack of space we elaborate below on two types of games that appear in the table. For the other games, we refer the reader to the corresponding sections.

Symmetric Coordination Games on Graphs (SCGGs). We introduce the class of SCGGs, in which all the relationships are coordination, and all players have the same k strategies. We use our framework to establish tight bounds of $1/k$ for the PoA, and $\frac{k}{2k-1}$ for the SPoA, and provide lower bounds on the q-SPoA.

Max-Cut games. It is part of folklore that the PoA of Max-Cut games is exactly $1/2$. Recently, [13] showed that the SPoA is exactly $2/3$, and that the q-SPoA remains roughly $1/2$ even for $q = O(n^{\frac{1}{2}-\epsilon})$ for any $\epsilon > 0$. Using our framework, we provide more refined results on the q-SPoA, specifically, we show that the q-SPoA is at least $\frac{2}{4-z(q)}$ for every q. A direct corollary of this result is that when the coordination factor is at least some constant $\alpha \in [0, 1]$ (i.e., q is roughly αn), the q-SPoA strictly improves to $\frac{2}{4-\alpha}$ (in particular, recovering the PoA and SPoA results, which correspond to the two extremes, $\alpha = 0$ and $\alpha = 1$, respectively). On the negative side, we construct an instance in which the q-SPoA remains roughly $1/2$ even for $q = O(n^{1-\epsilon})$ (and even slightly more relaxed values of q), thus further tightening the previous negative results. Referred to Fig. 1 for a complete list of our results.

Class Name	Case Description: +/−	Case Description: # of Str.	Case Description: Sym	Result: PoA	Result: SPoA	Result: q-SPoA
Max-Cut	−	2	√	1/2	2/3	$\frac{2}{4-z(q)}$ ⋆
2-NAE-SAT	+/−	2	√	1/2	2/3	$\frac{2}{4-z(q)}$ ⋆
Max-k-Cut	−	k	√	$\frac{k-1}{k}$	$\frac{k-1}{k-\frac{1}{2(k-1)}}$ ⋆	$\frac{k-1}{k-\frac{1}{2(k-1)}\cdot z(q)}$ ⋆
SCGGs	+	k	√	$1/k$ ⋆	$\frac{k}{2k-1}$ ⋆	$\frac{2+(k-2)\cdot z(q)}{2k-z(q)}$ ⋆
CGGs	+	k	×	0	1/2	$\frac{z(q)}{2}$
SCGs	+/−	k	√	$1/k$	$\frac{1}{2-\frac{1}{k(k-1)}}$ ⋆	$\frac{2+(k-2)\cdot z(q)}{2k-\frac{1}{k-1}\cdot z(q)}$ ⋆
CGs	+/−	k	×	0	1/2	$\frac{z(q)}{2}$ ⋆

Fig. 1. In the first column: SCGs = Symmetric Clustering Games on Networks, SCGGs = Symmetric Coordination Games on Graphs, CGGs = Coordination games on graphs, CGs = Clustering Games on Networks. Each class is fully described using three attributes: (a) the "+/−" column states the type of relationships on the edges (+ for coordination, − for anti-coordination); (b) the "# of Str." column states how many strategies exist in the game; And (c) the "Sym" (i.e., symmetry) column states whether all players have the same strategies or not. The *coordination factor* $z(q)$ is $\frac{q-1}{n-1}$. New results are marked with a '⋆', the rest are recovered by our framework. PoA results are tight. SPoA results for Max-Cut, 2-NAE-SAT, SCGGs, CGGs and CGs are tight. The q-SPoA results still have a gap between lower and upper bounds.

Existence of Equilibria. While there exist clustering games on networks that do not possess any SE (see, e.g., [2]), for some special cases, it has been shown that an SE always exists if the size of the strategy space is 2 [13, 2]. Theorem 1 extends this result to every clustering game with two strategies.

Our techniques. In our analysis, we utilize equilibrium properties to order the agents of a coalition so that each agent does not benefit when deviating together with the agents following her. We obtain a lower bound on the total welfare of a coalition, with respect to a different outcome (typically an optimal one), and utilize the potential function to translate it into an expression which we break into combinatorial objects such as cuts – edges between disjoint sets of agents (e.g. the coalition and its complement), and interiors – edges between agents in the same set (e.g. edges between agents in the coalition). The obtained expression for the lower bound is generic enough to encompass all clustering games on networks, yet expressive enough to give each special case its unique treatment. For symmetric games, rather than providing a lower bound with respect to a single (optimal) outcome, we generalize the method from [13, 14] and consider all the optimal outcomes that can be obtained by permuting the agents strategies. We combine all these lower bounds to achieve an improved lower bound which is quantified for each special case separately.

Related Work. PoA analysis was initiated by [17], and continued in a long line of studies in the algorithmic game theory literature. PoA with respect to SE and q-SE was first considered in [1]. The notion of *coalitional smoothness* was introduced by [3]. Specifically, if a game is (λ, μ)-coalitionally smooth, the SPoA is at least $\frac{\lambda}{1+\mu}$, and the same bound extends to more general solution concepts (see [3] and references therein for details on strong correlated and strong coarse correlated equilibria). In clustering games on networks, half the social welfare is a *potential function* [19], therefore a result from [3] implies that clustering games on networks are $(1/2, 0)$-coalitionally smooth, which implies that the SPoA is at least $1/2$ (which is tight for clustering games on networks). We improve the lower bound for SPoA in most of the special cases, and provide q-SPoA bounds, but our technique does not extend to the more general solution concepts.

Many clustering models have been studied from a game theoretic perspective. In [12] two variants are considered, and in both the weights of the edges are derived from a metric. In one the utilities are computed differently, and in the other the number of clusters is not pre-defined. Moreover, only NE is studied. In [7], all edges are coordination edges and the utility of each player is the utility of clustering games on networks, divided by the number of players that selected the same strategy, therefore each player derives her utility not only from her neighbours but from everyone in her cluster. Finally, clustering, i.e. the partitioning of objects with respect to similarity measures, is an area of research with an explosive amount of studies (e.g. [5]).

2 Model and Preliminaries.

A Clustering Game (CG) is a tuple $\{G = (V, E), (w_e)_{e \in E}, (b_e \in \{0, 1\})_{e \in E}, (\Sigma_i)_{i \in V}\}$ where G is an undirected graph with no self-loops. Each node corresponds to a player, and Σ_i is the (finite) strategy space of player i. Each edge e has a weight $w_e \in \mathbb{R}_{\geq 0}$, and a type $b_e \in \{0, 1\}$, where 0 implies that e is an *anti-coordination* edge and 1 implies that e is a *coordination* edge. Let $|V| = n$, $|E| = m$. Each element $\sigma = (\sigma_1, \ldots, \sigma_n) \in \times_i \Sigma_i$ is an *outcome* of the game. For each edge $e = \{i, j\}$, if e is a coordination edge, then it is *satisfied* if and only if $\sigma_i = \sigma_j$; if e is an anti-coordination edge, then it is satisfied if and only if $\sigma_i \neq \sigma_j$. Let 1_e^σ equal 1 if the edge e is satisfied in outcome σ and 0 otherwise. The utility of player i is the weight of satisfied edges she is incident to, i.e., $u_i(\sigma) = \sum_{e:i\in e} w_e \cdot 1_e^\sigma$. A CG is *symmetric* if $\Sigma_i = \{1, \ldots, k\}$ for every player i. In this case we abuse notation and denote the strategy space by k.

We identify several games from the literature as special cases of clustering games on networks and introduce *symmetric coordination games on graphs* as another special case. Due to space limitations, the proof of the following proposition as well as most proofs are deferred to the full version.

Proposition 1. *The following games are special cases of clustering games.*

1. *Max-Cut games [13] are CGs of the form:* $\langle G, (w_e)_{e\in E}, (0)_{e\in E}, 2\rangle$
2. *2-NAE-SAT games [13] are CGs of the form:* $\langle G, (w_e)_{e\in E}, (b_e \in \{0,1\})_{e\in E}, 2\rangle$
3. *Max-k-Cut games [14] are CGs of the form:* $\langle G, (w_e)_{e\in E}, (0)_{e\in E}, k\rangle$
4. *Coordination games on graphs [2] are CGs of the form:* $\langle G, (w_e)_{e\in E}, (1)_{e\in E}, (\Sigma_i)_{i\in V}\rangle$
5. *Symmetric coordination games on graphs are CGs of the form:* $\langle G = (V, E), (w_e)_{e\in E}, (1)_{e\in E}, k\rangle$.

Given an outcome σ, $(\sigma_i^*, \sigma_{-i})$ denotes the outcome where σ_i is replaced by $\sigma_i^* \in \Sigma_i$, e.g., $\sigma = (\sigma_i, \sigma_{-i})$. Given outcomes σ, σ^* and a set of players $A \subseteq V$, $(\sigma_A^*, \sigma_{-A})$ denotes the outcome where all players in A play by σ^*, and all players in A^c play by σ, where $A^c = V \setminus A$ (and so $\sigma = (\sigma_A, \sigma_{-A})$).

Definition 1. *A* Nash equilibrium *(NE)*[1] *is an outcome σ such that no player can strictly increase her utility by deviating unilaterally, i.e., if for every player i and strategy $a \in \Sigma_i$, $u_i(\sigma_i, \sigma_{-i}) \geq u_i(a, \sigma_{-i})$. σ is a* q-strong equilibrium *(q-SE) if there is no coalition A of size at most q and outcome σ_A^* of A's members such that $u_i(\sigma_A^*, \sigma_{-A}) > u_i(\sigma_A, \sigma_{-A})$ for all $i \in A$. For $q = n$, the outcome σ is called a* strong equilibrium *(SE).*

Throughout this paper the quality of an outcome σ is measured by its *social welfare*, i.e., the sum: $SW(\sigma) = \sum_{i\in V} u_i(\sigma)$. In addition, for every set of players $A \subseteq V$, denote the total welfare of the players in A by $SW_A(\sigma) = \sum_{i\in A} u_i(\sigma)$.

Theorem 1. *For every clustering game with two strategies, any optimal outcome is an SE.*

The degradation of social welfare is commonly quantified as follows: Consider a solution concept (e.g. a NE), and quantify the ratio between the social welfare of the worst solution and that of an optimal outcome. When the solution concept is a q-SE the ratio is called the *q-strong price of anarchy*, which is the measure of interest in this paper. Note that for $q = 1$ the ratio is the *price of anarchy* (PoA), and for $q = n$ it is the *strong price of anarchy* (SPoA) which was defined in [1]. Formally, given a class of games Γ, let $\mathcal{G} \in \Gamma$, and let q-$SE(\mathcal{G})$ be the set of q-SE in game $\mathcal{G}$, and let σ^* be an optimal outcome of $\mathcal{G}$ (i.e. $\sigma^* \in \arg\max_\sigma SW(\sigma)$):

$$q\text{-SPoA} = \min_{\mathcal{G}\in\Gamma} \frac{\min\{SW(\sigma) : \sigma \in q\text{-}SE(\mathcal{G})\}}{SW(\sigma^*)} \quad (1)$$

[1] In this paper we restrict attention to pure-strategy equilibrium (i.e., no randomization is used in an agent's behavior).

By definition, the q-SPoA ranges from 0 to 1, where 1 corresponds to full efficiency. In this terminology lower bounds are positive results on the efficiency and upper bounds are negative results. Note that in order to lower bound the q-SPoA, it is enough to bound the term $SW(\sigma)/SW(\sigma^*)$ for every q-SE σ.

2.1 Welfare Guarantees in Equilibrium

In this section we establish the lemmas required to analyse the efficiency loss quantified by the q-SPoA in various special cases of clustering games on networks.

For two m-vectors $v = (v_1, \dots v_m)$, $u = (u_1, \dots u_m)$, let $v \cdot u = (v_1 \cdot u_1, \dots, v_m \cdot u_m)$, and denote by $\langle v \rangle$ the inner product with the edge weights, i.e., $\langle v \rangle = \sum_{e \in E} w_e \cdot v_e$. Observe that the operator $\langle \cdot \rangle$ is linear, i.e., given vectors $v^1, \dots, v^k$, it holds that: $\sum_{i=1}^k \langle v^i \rangle = \sum_{i=1}^k \sum_{e \in E} w_e \cdot v_e^i = \sum_{e \in E} w_e \cdot \sum_{i=1}^k v_e^i = \langle \sum_{i=1}^k v^i \rangle$.

Edge Partition. Let $1^\sigma = (1_e^\sigma)_{e \in E}$. Consider a pair of outcomes σ and σ^*. Let $\mathcal{B}$ be the characteristic vector of edges that are satisfied *both* by σ and σ^*, i.e., $\mathcal{B} = 1^\sigma \cdot 1^{\sigma^*}$. Let $\mathcal{E}$ be the characteristic vector of edges that are satisfied in σ but not satisfied in σ^*, i.e., $\mathcal{E} = 1^\sigma \cdot \overline{1^{\sigma^*}}$. Let $\mathcal{O}$ be the characteristic vector of edges that are satisfied in σ^* but not satisfied in σ, i.e., $\mathcal{O} = \overline{1^\sigma} \cdot 1^{\sigma^*}$. The choice of these notations will become clear later, regardless, observe that $1^\sigma = \mathcal{B} + \mathcal{E}$ and $1^{\sigma^*} = \mathcal{B} + \mathcal{O}$. Sometimes we abuse notation and let a characteristic vector be the set of edges it represents (e.g. $e \in 1^\sigma$ if and only if e is satisfied in σ).

Consider a set of players A. Let 1_e^A equal 1 if $e \cap A \neq \emptyset$ and 0 otherwise. Let $1^A = (1_e^A)_{e \in E}$. The *interior of* A is the set $\{e \in E : e \subseteq A\}$. Observe that $\mathcal{I}^A = \overline{1^{A^c}}$ is the characteristic vector of the interior of A, since $e \subseteq A$ if and only if $e \cap A^c = \emptyset$. Moreover, for two disjoint sets of players $A, B \subseteq V$, the vector $\delta^{A,B} = 1^A \cdot 1^B$ is the characteristic vector of the A–B cut, since an edge is in the cut if and only if it intersects with A and B.

Lemma 1. *For every set of players $A \subseteq V$ it holds that:*
$SW_A(\sigma) = \langle (2 \cdot \mathcal{I}^A + \delta^{A,A^c}) \cdot 1^\sigma \rangle$. *As a result, if $A = V$ then $SW(\sigma) = 2\langle 1^\sigma \rangle$*

Clustering games admit a *potential function.* A function Φ is a potential if it encodes the difference in utility of a player when deviating from one strategy to another, i.e., $\forall i, \forall a, b \in \Sigma_i, \forall \sigma : u_i(a, \sigma_{-i}) - u_i(b, \sigma_{-i}) = \Phi(a, \sigma_{-i}) - \Phi(b, \sigma_{-i})$.

Theorem 2. *The function $\Phi(\sigma) = \langle 1^\sigma \rangle$ is a potential for every clustering game.*

A direct corollary of Theorem 2 and Lemma 1 is that for every outcome σ, $\Phi(\sigma) = \frac{1}{2} SW(\sigma)$, therefore the problem of bounding the q-SPoA reduces to bounding the ratio $\Phi(\sigma)/\Phi(\sigma^*)$ for every q-SE σ and optimal outcome σ^*. By rearranging the terms of the potential function property, we get:

$$u_i(a, \sigma_{-i}) = \Phi(a, \sigma_{-i}) - \Phi(b, \sigma_{-i}) + u_i(b, \sigma_{-i}) \tag{2}$$

The Renaming Procedure. Let σ be a q-SE, σ^* an optimal outcome, and $K \subseteq V$ a set of players of size at most q. If all players in K deviate together from σ to σ^*, then by the definition of a q-SE, there is some player $i \in K$ so that $u_i(\sigma) \geq u_i(\sigma_K^*, \sigma_{-K})$. Rename the players so that this is player 1. Similarly, there is some player i in $K' = K \setminus \{1\}$ so that $u_i(\sigma) \geq u_i(\sigma^*_{K'}, \sigma_{-K'})$. Rename the players in K' so that this is player 2. Iterate this argument to rename all players in K and conclude that for each $i \in K$ it holds that $u_i(\sigma) \geq u_i(\sigma^*_{\{i\ldots|K|\}}, \sigma_{-\{i\ldots|K|\}})$. To simplify notation, for each $i \in K$ let $p^{\sigma,\sigma^*}_{K,i} = (\sigma^*_{\{i\ldots|K|\}}, \sigma_{-\{i\ldots|K|\}})$. When clear in the context, we omit the outcomes and denote it by $p_{K,i}$. Note that $p_{K,1} = (\sigma^*_K, \sigma_{-K})$ and $p_{K,|K|+1} = \sigma$. By this notation, we get $u_i(\sigma) \geq u_i(p_{K,i})$ for all $i \in K$. Therefore, it holds that $SW_K(\sigma) = \sum_{i\in K} u_i(\sigma) \geq \sum_{i\in K} u_i(p_{K,i})$. For each $i \in K$ we apply (2) to $u_i(p_{K,i})$ and $u_i(p_{K,i+1})$ and we get that $\sum_{i\in K} u_i(p_{K,i}) = \sum_{i\in K} \left(\Phi(p_{K,i}) - \Phi(p_{K,i+1}) + u_i(p_{K,i+1})\right)$. Observe that the sum on the potential function telescopes. We conclude that

$$SW_K(\sigma) \geq \Phi(\sigma^*_K, \sigma_{-K}) - \Phi(\sigma) + \sum_{i\in K} u_i(p_{K,i+1}) \tag{3}$$

Given an ordering o on a set of players K, for every edge $e = \{i, j\}$ from the interior of K, if $o(i) < o(j)$ then let $[K]^{\sigma,\sigma^*}_e = 1^{(\sigma_i,\sigma^*_{-i})}_e$. For each edge $e \nsubseteq K$, let $[K]^{\sigma,\sigma^*}_e = 0$. Let $[K]^{\sigma,\sigma^*} = ([K]^{\sigma,\sigma^*}_e)_e$. The essence of $[K]^{\sigma,\sigma^*}_e$ is that the *first* player in the edge e plays according to her strategy in σ, and the *second* player plays according to her strategy in σ^*. Therefore, when considering edges from $[K]^{\sigma,\sigma^*}$, we are guaranteed that *exactly* one player changes color (from σ to σ^*) on each edge in the interior of K. We prove in the full version that the right-most sum of (3) can be expressed as follows:

$$\sum_{i\in K} u_i(p_{K,i+1}) = \langle \delta^{K,K^c} \cdot 1^\sigma \rangle + \langle \mathcal{I}^K \cdot \left(1^\sigma + [K]^{\sigma,\sigma^*}\right)\rangle \tag{4}$$

The object $[K]^{\sigma,\sigma^*}$ formally describes the result of the renaming procedure, i.e., it encodes the fact that we are considering edges from the interior of K such that *one* player plays according to σ, and *the other* plays according to σ^*. Equation (3) together with (4) provides a welfare guarantee for a set of players of size q. To provide a welfare guarantee for a larger set of players A, $|A| > q$, apply (3) for each subset $K \subseteq A$, and sum all inequalities. The resulting equation is:

$$\begin{aligned} SW_A(\sigma) \geq \frac{q-1}{|A|-1}\langle \mathcal{I}^A \cdot (\mathcal{B} + \mathcal{O})\rangle + \\ \binom{|A|-1}{q-1}^{-1} \sum_{K\subseteq A, |K|=q} \langle \left(\mathcal{I}^K \cdot [K]^{\sigma,\sigma^*} + \delta^{K,K^c} \cdot 1^{(\sigma^*_K,\sigma_{-K})}\right) \cdot (\mathcal{B} + \mathcal{O} + \mathcal{E})\rangle \end{aligned} \tag{5}$$

Let Π be the set of all permutations $\pi : \{1, \ldots, k\} \to \{1, \ldots, k\}$. Given a permutation π, let σ^*_π be the outcome in which every player i plays $\pi(\sigma^*_i)$, i.e., $\sigma^*_\pi = (\pi(\sigma^*_i))_{i\in V}$. For every permutation π, $1^{\sigma^*} = 1^{\sigma^*_\pi}$. Therefore, $\mathcal{B}$, $\mathcal{O}$ and $\mathcal{E}$ are *permutation invariant*, i.e., for every π, $\mathcal{B} = 1^\sigma \cdot 1^{\sigma^*_\pi}$, and similarly for $\mathcal{O}$ and $\mathcal{E}$. Permutation invariance was previously considered for Max-k-Cut games [14]. Let $D_\pi(\sigma, \sigma^*)$ be the set of players with *different* strategies in σ and σ^*_π, i.e., $D_\pi(\sigma, \sigma^*) = \{i : (\sigma^*_\pi)_i \neq \sigma_i\}$. When clear in the context we omit the outcomes and write D_π. Let $q_\pi = \min\{q, |D_\pi|\}$. Lemma 2 establishes our general inequality for symmetric clustering games.

Lemma 2. *For every q-SE σ, and optimal outcome σ^*, it holds that:*

$$\sum_{\pi\in\Pi} SW_{D_\pi}(\sigma) \geq \sum_{\pi\in\Pi}\Bigg(\langle\Big(\frac{q_\pi-1}{|D_\pi|-1}\mathcal{I}^{D_\pi}+\delta^{D_\pi,D_\pi^c}\Big)\cdot(\mathcal{B}+\mathcal{O})\rangle+$$
$$\binom{|D_\pi|-1}{q_\pi-1}^{-1}\sum_{K\subseteq D_\pi,|K|=q_\pi}\langle\Big(\mathcal{I}^K\cdot[K]^{\sigma,\sigma^*_\pi}+\delta^{K,D_\pi\setminus K}\cdot 1^{((\sigma^*_\pi)_K,\sigma_{-K})}\Big)\cdot(\mathcal{B}+\mathcal{O}+\mathcal{E})\rangle\Bigg)$$

Lemma 3. *Let σ and σ^* be two outcomes. Then it holds that:*
$\sum_{\pi\in\Pi} SW_{D_\pi}(\sigma) = (k-1)(k-1)!\, SW(\sigma)$.

Game Specific Analysis. For each special case of clustering games, given a set K, the values of $[K]_e^{\sigma,\sigma^*_\pi}$ and $1_e^{((\sigma^*_\pi)_K,\sigma_{-K})}$ in the right-hand side of Lemma 2 may be interpreted differently by making arguments that are specific to the special case. While nothing is assumed on the order of the players in any set K, since K is a subset of D_π, it is guaranteed that for each edge $e\in\mathcal{I}^K$, the outcome that is considered in $[K]_e^{\sigma,\sigma^*_\pi}$ is such that one player plays her strategy in σ, while the other player plays her strategy in σ^*_π *which is different* from her strategy in σ. Similarly, in the outcome $((\sigma^*_\pi)_K,\sigma_{-K})$, all the players in K play their strategies in σ^*_π *which are different* from their strategies in σ.

3 Strong Price of Anarchy Bounds

In this section we show how to utilize the proposed framework to establish bounds on the q-strong price of anarchy. Define the *coordination factor* to be $z(q)=\frac{q-1}{n-1}$, with $z(1)=0$ and $z(n)=1$, which correspond to the PoA and SPoA, respectively. Intuitively, it measures the amount of coordination that is assumed to exist in the regarded solution concept. Our first result generalizes the lower bound in [2].

Theorem 3. *The q-SPoA of every clustering game is at least* $\frac{z(q)}{2}$

Proof. Let σ be a q-SE and σ^* an optimal outcome. Substitute V for A in (5) and omit all terms with $[K]_e^{\sigma,\sigma^*}$ and $1^{(\sigma^*_K,\sigma_{-K})}$ (since they are non-negative). Observe that $1^V=(1,\dots,1)$ and recall that $\langle 1^{\sigma^*}\rangle=\frac{1}{2}SW(\sigma^*)$. Therefore,

$$SW(\sigma)=SW_V(\sigma)\geq\frac{q-1}{n-1}\langle\mathcal{B}+\mathcal{O}\rangle=z(q)\cdot\langle 1^{\sigma^*}\rangle=z(q)\cdot\frac{1}{2}\cdot SW(\sigma^*)$$

□

A tight upper bound of 1/2 for $q=n$, and an upper bound of $\frac{z(q)}{2-z(q)}$ for every q was shown in [2] for coordination games on graphs. For $q=n$, the same lower bound is also achieved by Theorem 10 in [3]. Indeed, clustering games on networks are utility-maximization potential games with only positive externalities, and $\Phi(\sigma)=\frac{1}{2}SW(\sigma)$. Therefore, the game is $(1/2,0)$-*coalitionally smooth*[2] which implies that the SPoA is at least 1/2. This bound extends to more general solution concepts (see [3] for details regarding mixed, correlated and coarse-correlated strong equilibria).

[2] The reader is referred to [3] for the exact definition.

3.1 Symmetric Coordination Games on Graphs (SCGGs)

When all edges are coordination edges, and all players have the same strategy space $\{1, \ldots, k\}$, the class that is obtained is SCGGs with k strategies.

Theorem 4. *The SPoA of SCGGs with k strategies is at least $\frac{k}{2k-1}$.*

Proof. (sketch.) Let σ be an SE and σ^* an optimal outcome. Consider Lemma 2 in the context of SCGGs. Since all edges are coordination edges, any edge e that is satisfied in σ or σ^*_π is certainly unsatisfied when exactly one of its adjacent nodes changes strategy, i.e. $[D_\pi]_e^{\sigma,\sigma^*_\pi} = 0$. Therefore, for $q = n$ we get:

$$\sum_{\pi \in \Pi} SW_{D_\pi}(\sigma) \geq \sum_{\pi \in \Pi} \langle \left(\mathcal{I}^{D_\pi} + \delta^{D_\pi, D^c_\pi}\right) \cdot (\mathcal{B} + \mathcal{O}) \rangle \tag{6}$$

By Lemma 3, in the left-hand side we get $(k-1)(k-1)!\, SW(\sigma)$. In the right-hand side, since $\mathcal{B}$ and $\mathcal{O}$ are permutation invariant, we get four expressions:

$$\langle (\sum_\pi \mathcal{I}^{D_\pi}) \cdot \mathcal{B} \rangle + \langle (\sum_\pi \delta^{D_\pi, D^c_\pi}) \cdot \mathcal{B} \rangle + \langle (\sum_\pi \mathcal{I}^{D_\pi}) \cdot \mathcal{O} \rangle + \langle (\sum_\pi \delta^{D_\pi, D^c_\pi}) \cdot \mathcal{O} \rangle \tag{7}$$

We use simple combinatorial arguments to compute the expressions above.

First expression: $e = \{i, j\} \in \mathcal{B}$ if and only if $\sigma_i = \sigma_j$ and $\sigma^*_i = \sigma^*_j$. In such a case, for every permutation π we get that $e \in \mathcal{I}^{D_\pi}$ if and only if $\pi(\sigma^*_i) \neq \sigma_i$. There are $(k-1)$ options to fix $\pi(\sigma^*_i)$, and for each such option there are $(k-1)!$ options to set the other $(k-1)$ values of π (as there are no additional restrictions). Therefore, $\langle \left(\sum_\pi \mathcal{I}^{D_\pi}\right) \cdot \mathcal{B} \rangle = (k-1)(k-1)!\, \langle \mathcal{B} \rangle$. Second expression: Since e cannot be in the D_π–D^c_π cut, it holds that $\langle (\sum_\pi \delta^{D_\pi, D^c_\pi}) \cdot \mathcal{B} \rangle = 0$. Third expression: $e = \{i, j\} \in \mathcal{O}$ if and only if $\sigma_i \neq \sigma_j$ and $\sigma^*_i = \sigma^*_j$. In such a case $e \in \mathcal{I}^{D_\pi}$ if and only if $\{\pi(\sigma^*_i)\} \cap \{\sigma_i, \sigma_j\} = \emptyset$. There are $(k-2)$ options to fix $\pi(\sigma^*_i)$, and for each such option there are $(k-1)!$ options to set the other $(k-1)$ values of π (as there are no additional restrictions). Therefore, $\langle \left(\sum_\pi \mathcal{I}^{D_\pi}\right) \cdot \mathcal{O} \rangle = (k-2)(k-1)!\, \langle \mathcal{O} \rangle$. Fourth expression: $e \in \delta^{D_\pi, D^c_\pi}$ in exactly two disjoint events: $\pi(\sigma^*_i) = \sigma_i$ or $\pi(\sigma^*_j) = \sigma_j$, therefore, $\langle \left(\sum_\pi \delta^{D_\pi, D^c_\pi}\right) \cdot \mathcal{O} \rangle = 2(k-1)!\, \langle \mathcal{O} \rangle$.

In total, we conclude that (7) is at least $k!\, \Phi(\sigma^*) - (k-1)!\, \Phi(\sigma)$. Recall that $SW(\sigma) = 2 \cdot \Phi(\sigma)$. Divide both sides by $(k-1)!$ and reorganize terms to get $(2k-1) \cdot \Phi(\sigma) \geq k \cdot \Phi(\sigma^*)$, as desired. □

The following proposition shows that Theorem 4 is tight.

Proposition 2. *The symmetric coordination game on a line graph with $2k$ nodes and k strategies for each player has a SPoA of $\frac{k}{2k-1}$.*

Theorem 4 extends to q-SPoA case as follows:

Theorem 5. *The q-SPoA of SCGGs with k strategies is at least $\frac{2+z(q)\cdot(k-2)}{2k-z(q)}$.*

Proposition 3 shows that the PoA (i.e., for $q = 1$) bound is tight.

Proposition 3. *There exists a SCGG with k strategies, with a PoA $= 1/k$.*

3.2 Symmetric Anti-Coordination Games on Graphs

When all edges are anti-coordination edges, and all players have the same strategy space $\{1,\ldots,k\}$, the class that is obtained is Max-k-Cut [14]. Previous work [15, 18] showed that the PoA exactly $(k-1)/k$. For $k=2$, [13] showed that the SPoA is exactly 2/3. In [14] an upper bound of $\frac{2k-2}{2k-1}$ was established[3]. Theorem 6 uses the framework to establish lower bounds on the q-SPoA for any k and q.

Theorem 6. *The q-SPoA of Max-k-Cut games is at least:* $\frac{k-1}{k-\frac{1}{2(k-1)}\cdot z(q)}$.

3.3 Symmetric Clustering Games on Networks (SCGs)

In this class, edges are either coordination or anti-coordination edges, and all players have k strategies. For $k=2$, the class coincides with 2-NAE-SAT, for which a tight SPoA bound of 2/3 was shown [13]. We establish a bound on the q-SPoA for every q and k:

Theorem 7. *The q-SPoA of SCGs with k strategies is at least* $\frac{2+z(q)\cdot(k-2)}{2k-z(q)\cdot\frac{1}{k-1}}$.

A special case of Theorem 7 is that the SPoA of SCGs is at least $\frac{1}{2-\frac{1}{k(k-1)}}$, for which a simplified proof is provided in the full version. For the class of 2-NAE-SAT ($k=2$), it was shown in [13] that for $q=O(n^{\frac{1}{2}-\epsilon})$ (for any $\epsilon>0$), the q-SPoA is 1/2. Theorem 8 shows this is true even for $q=O(n^{1-\epsilon})$.

Theorem 8. *For any* $\epsilon>0$ *and* $q=O(n^{1-\epsilon})$, *the q-SPoA of Max-Cut is* 1/2.

On the other hand, the positive result that is implied by Theorem 7, is that for any $\alpha\in(0,1]$ and $z(q)\geq\alpha$, (roughly, when $q\geq\alpha n$), the q-SPoA is $\frac{2}{4-\alpha}$.

Corollary 1. *The q-SPoA of 2-NAE-SAT games is at least* $\frac{2}{4-z(q)}$.

4 Future Directions

Our results and analysis suggest interesting directions for future research. First, the existence of q-SE in clustering games on networks is only partially understood. For example, it is conjectured in [14] that every Max-k-Cut game admits an SE. This is clearly an interesting open problem, as well as the more general problem of whether all symmetric clustering games on networks possess an SE.

A full characterization of q-SE existence for clustering games on networks is an interesting open problem.

Second, our analysis leaves a gap in the q-SPoA for Max-k-Cut games. We note that the analysis of this gap gives rise to a combinatorial problem that is of independent interest.

[3] In [14], the authors also presented a matching lower bound of $\frac{2k-2}{2k-1}$. However, one of the derivations in the analysis contained an error. This was verified by personal communication with the authors. Their corrected analysis leads to a lower bound that matches the lower bound, established in this paper in Theorem 6 for $z(q)=1$.

Third, our proof techniques provide q-SPoA results only with respect to pure equilibria. It is desired to extend this analysis to handle other solution concepts such as mixed, correlated and coarse correlated equilibria. Finally, it would be interesting to explore ways in which our analysis can shed light on coalitional dynamics in clustering games on networks.

References

1. N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *SODA*, pages 189–198, 2007.
2. K. R. Apt, M. Rahn, G. Schäfer, and S. Simon. Coordination games on graphs. In *WINE*, pages 441–446. Springer, 2014.
3. Y. Bachrach, V. Syrgkanis, É. Tardos, and M. Vojnović. Strong price of anarchy, utility games and coalitional dynamics. In *SAGT*, pages 218–230. Springer, 2014.
4. S. Banerjee, H. Konishi, and T. Sönmez. Core in a simple coalition formation game. *Social Choice and Welfare*, 18(1):135–153, 2001.
5. N. Bansal, A. Blum, and S. Chawla. Correlation clustering. *Machine Learning*, 56(1-3):89–113, 2004.
6. S. Barberà and A. Gerber. On coalition formation: durable coalition structures. *Mathematical Social Sciences*, 45(2):185–203, 2003.
7. V. Bilò, A. Fanelli, M. Flammini, G. Monaco, and L. Moscardelli. Nash stability in fractional hedonic games. In *WINE*, pages 486–491. Springer, 2014.
8. F. Bloch and E. Diamantoudi. Noncooperative formation of coalitions in hedonic games. *International Journal of Game Theory*, 40(2):263–280, 2011.
9. A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
10. E. Diamantoudi and L. Xue. Farsighted stability in hedonic games. *Social Choice and Welfare*, 21(1):39–61, 2003.
11. J. H. Dreze and J. Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica: Journal of the Econometric Society*, pages 987–1003, 1980.
12. M. Feldman, L. Lewin-Eytan, and J. Naor. Hedonic clustering games. In *SPAA*, pages 267–276. ACM, 2012.
13. L. Gourvès and J. Monnot. On strong equilibria in the max cut game. In *WINE*, pages 608–615, 2009.
14. L. Gourvès and J. Monnot. The max k-cut game and its strong equilibria. In *TAMC*, pages 234–246. Springer, 2010.
15. M. Hoefer. *Cost sharing and clustering under distributed competition*. PhD thesis, University of Konstanz, 2007.
16. M. Kearns, M. L. Littman, and S. Singh. Graphical models for game theory. In *UAI*, pages 253–260. Morgan Kaufmann Publishers Inc., 2001.
17. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413. Springer, 1999.
18. J. Kun, B. Powers, and L. Reyzin. Anti-coordination games and stable graph colorings. In *SAGT*, pages 122–133. Springer, 2013.
19. D. Monderer and L. Shapley. Potential games. *Games and economic behavior*, 14(1):124–143, 1996.
20. N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani. *Algorithmic game theory*, chapter 7. Cambridge University Press, 2007.