A

Hedonic Clustering Games

Moran Feldman, CS Dept., Technion
Liane Lewin-Eytan, Yahoo Labs, Haifa, Israel
Joseph (Seffi) Naor, CS Dept., Technion

Clustering, the partitioning of objects with respect to a similarity measure, has been extensively studied as a global optimization problem. We investigate clustering from a game-theoretic approach, and consider the class of *hedonic clustering games*. Here, a *self-organized* clustering is obtained via decisions made by independent players, corresponding to the elements clustered. Being a hedonic setting, the utility of each player is determined by the identity of the other members of her cluster. This class of games seems to be quite robust, as it fits with rather different, yet commonly used, clustering criteria. Specifically, we investigate hedonic clustering games in two different models: *fixed clustering*, which subdivides into k-median and k-center, and *correlation clustering*. We provide a thorough analysis of these games, characterizing Nash equilibria, and proving upper and lower bounds on the price of anarchy and price of stability. For fixed clustering we focus on the existence of a Nash equilibrium, as it is a rather non-trivial issue in this setting. We study it both for general metrics and special cases, such as line and tree metrics. In the correlation clustering model, we study both minimization and maximization variants, and provide almost tight bounds on both the price of anarchy and price of stability.

Categories and Subject Descriptors: C.2.4 [**Computer Systems Organization**]: Computer-Communication Networks—*Distributed Systems*

General Terms: Theory, Performance

Additional Key Words and Phrases: Clustering Games, Hedonic Games, Price of Anarchy, Price of Stability

ACM Reference Format:
Moran Feldman, Liane Lewin-Eytan, and Joseph (Seffi) Naor. 2014. Hedonic Clustering Games. *ACM Trans. Parallel Comput.* V, N, Article A (January YYYY), 48 pages.
DOI:http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

Clustering is the partitioning of objects or elements with respect to a similarity measure. The greater the similarity of elements belonging to a cluster, or the distance between elements belonging to different clusters, the "better" the clustering. Clustering has been extensively treated as a global optimization problem, employing a variety of optimization methods. We adopt here a novel game-theoretic approach, and consider a setting in which a *self-organized* clustering is obtained from decisions taken by in-

---

Work supported in part by the Technion-Micorsoft Electronic Commerce Research Center and by ISF grant 954/11. A preliminary version of this work appeared in *Proceedings of the* 24^{th} *ACM Symposium on Parallelism in Algorithms and Architectures (SPAA12)*, pages: 267-276.
Authors addresses: (Current Address) Moran Feldman, EPFL-IC-THL2, Station 14, 1015 Lausanne, Switzerland; e-mail: moran.feldman@epfl.ch. Liane Lewin-Eytan, Yahoo Labs, Haifa, Israel; e-mail: liane@yahoo-inc.com. Joseph (Seffi) Naor, CS Dept., Technion, Haifa, Israel; e-mail: naor@cs.technion.ac.il.

DOI:http://dx.doi.org/10.1145/0000000.0000000

dependent players. We assume that the players correspond to the elements clustered, and their goal is to maximize their own utility functions. From the perspective of a single player, the quality, or the utility of a clustering, depends on the player's similarity to elements in her own cluster and perhaps on dissimilarity to elements in other clusters.

Our clustering games belong to the well known class of *hedonic* games, introduced in the Economics literature as a model of coalition formation. In a hedonic game, the utility of a player is solely determined by the identity of the players belonging to her coalition, and is independent of the partition of the other players into coalitions. Hedonic games were first introduced and analyzed by Dreze and Greenberg [1980] in the context of cooperative games, and were motivated by situations in which individuals carry out joint activities as coalitions. Examples of such situations are individuals organizing themselves in groups for consumption or production purposes, or individuals relying upon local communities for the provisioning of public goods. Thus, hedonic games can be used to model settings arising in a wide variety of social, economic, and political problems, ranging from communication and trade to legislative voting. See [Bogomolnaia and Jackson 2002] for a discussion of several real-life situations fitting the hedonic model. The notion of stability in hedonic games has been investigated both from cooperative, as well as non-cooperative, aspects [Banerjee et al. 2001; Bogomolnaia and Jackson 2002; Bloch and Diamantoudi 2010]. The non-cooperative framework makes sense in environments lacking a social planner, or if the cost of coordination is high. We note that most work on hedonic games has mainly focused on the existence of stable coalition partitions, whether core stable or individually stable, and on the complexity of finding such outcomes.

We investigate non-cooperative hedonic clustering games, in which elements are independent selfish players. Each player joins a group maximizing her utility, and the resulting clustering is the outcome of the choices of all players. We present a case study of two different well-known clustering models, with commonly used utility functions. The first model is *fixed clustering*, in which the number of clusters is fixed, and each cluster has a *centroid* whose position is determined by the identity of the cluster members. A player's utility depends on the location of the *centroid* of her cluster. The second model is *correlation clustering*, in which a player's utility depends on her similarity to other elements in her cluster as well as on her dissimilarity to elements in other clusters. In general, various settings in which players form clusters, and then each cluster provides a public good, or a service from a set of available alternatives, is captured by hedonic clustering games. Following are two motivating examples coming from different application areas.

In an ad-hoc (or sensor) network there is a large number of *autonomous devices* which are spread over a geographic area and wish to communicate with each other. In order to establish communication, devices invest transmission power which depends on the physical distance between them. Power is a critical resource for battery-limited devices, and thus, the goal of each device is to minimize its transmission power and save on battery time. Fixed clustering is a proven method for enhancing energy efficiency and lifetime of large ad-hoc networks, and has been extensively studied in this context [Abbasi and Younis. 2007; Bandyopadhyay and Coyle 2003]. Proposed clustering protocols organize the devices in data aggregation clusters to reduce network traffic. Each cluster has a center that receives data from other devices in the cluster, and sends it beyond the cluster limits, possibly after aggregating the received data and reducing its volume. A device will then join a cluster having the closest center to minimize the power needed for transmission, thus leading to a game-theoretic setting. We note that clustering in ad-hoc networks has been studied from a game-theoretic perspective in [Koltsidas and Pavlidou 2011], yet their game definition is completely different from ours.

In online web advertising, publishers wish to join advertising services. Publishers are partitioned into clusters, and each cluster provides a different type of advertising service to its members. The type of service a cluster offers is derived from the attributes of its members, where possible attributes are, *e.g.*, fields of specialization, geographical area, organization size, types of product and annual budget. Publishers join or leave a cluster depending on the advertising service that the cluster offers and on the attributes of its members. For example, a new and relatively small business would prefer not to be coupled with a well-known large company specializing in a similar field. Thus, the utility of each player (publisher) depends on her similarity to the cluster, *i.e.*, how close are her advertising needs to the service provided by the cluster, and on her dissimilarity to players in other clusters (in the latter case, small and large businesses would be considered dissimilar). This is precisely the type of utility captured by correlation clustering.

Despite extensive work on clustering, not much work has been done from a game-theoretic perspective, and we believe that this work contributes in that direction. We emphasize that the focus of our paper is not on a specific setting, but rather on the study of the general game-theoretic framework in the context of hedonic clustering games.

1.1. Our Model

Our clustering problems are defined on a set of n points lying in a metric space with a distance function $d(\cdot,\cdot)$. The points correspond to selfish, non-cooperative, players (or users) moving between clusters at will. Players within a cluster are provided a service depending on the set of players belonging to it. A player achieves a utility from being a member of a cluster, and will naturally join the one maximizing her utility (or minimizing her cost). The notion of *social welfare* (or *social cost*) corresponds to the overall utility achieved by the system (or overall cost). The strategies of a player in a clustering game correspond to the set of clusters to which she can belong. Every choice of strategies by the players partitions them into clusters, and is called a *clustering configuration*. A Nash equilibrium[1] of the clustering game corresponds to a clustering configuration in which no user can unilaterally increase her utility (reduce her cost) by changing clusters. We investigate the clustering game in two different models: *fixed clustering* and *correlation clustering*.

Fixed Clustering. In a fixed clustering game, the number of clusters is known beforehand and is denoted by k. Each cluster C has a *centroid*, $c(C)$, defined to be the element minimizing the cost of the cluster. We consider two well-known definitions in the clustering literature, known as k-*median* and k-*center*. In the k-median clustering problem, the cost of a cluster is defined as the sum of distances between all members of the cluster and its centroid. The centroid is thus defined as

$$c(C) = \arg\min_{u\in C}\left\{\sum_{v\in C} d(u,v)\right\} \ .$$

In the k-center clustering problem, the cost of a cluster is defined by its radius, *i.e.*, the maximum distance between an element in the cluster and its center. Hence, the centroid is

$$c(C) = \arg\min_{u\in C}\left\{\max_{v\in C} d(u,v)\right\} \ .$$

[1]We consider in this paper only pure Nash equilibria.

We note that in both models the choice of a centroid might not be unique, and therefore a tie-breaking rule is needed. We elaborate on this issue later on.

In both models, the strategy space of a player is defined by the k clusters. Let C^{+v} denote a cluster C with the addition of player v. A clustering $C_1, C_2, \ldots, C_k$ is a pure Nash equilibrium if for all C_i, C_j, and $v \in C_i$, we have $d(v, c(C_j^{+v})) \geq d(v, c(C_i))$, i.e., node v cannot unilaterally reduce its cost by changing clusters from C_i to C_j. Note that following v's addition to C, the centroid of C might change, *i.e.*, it might be that $c(C) \neq c(C^{+v})$. For k-median, the social cost is defined to be the sum of costs of all the clusters, whereas for k-center the social cost is defined to be the maximum cost of a cluster.

In fixed clustering, the service offered by a cluster is represented by its centroid. The example of ad-hoc networks fits this model, since the centroid is the node to which transmissions within a cluster are sent, and transmission costs depend on the distances to the centroid.

Correlation Clustering. In settings where only the relationship among objects is known, correlation clustering is a natural approach. Unlike most clustering formulations, specifying the number of clusters as a separate parameter is not necessary in this formulation. We assume that the similarity metric is captured by a distance metric $d(\cdot,\cdot) \in [0,1]$. If $d(u,v) \approx 0$, then u and v are very similar, and if $d(u,v) = 1$, then they are highly distinctive, unrelated elements. Each element v has a weight w_v denoting its "measure of influence" on other elements. Elements wish to be clustered with similar elements of high weight, and to be partitioned away from unrelated elements of high weight. Since the number of clusters is not fixed, the possible strategies of a player are either to join an existing cluster, or to create a new cluster and become its sole member.

Given an element v, denote by C_v its cluster in a given configuration. Typically, two variants are studied in correlation clustering. In the *minimization variant*, the objective of each element v is to minimize its cost

$$\sum_{u \in C_v} w_u \cdot d(u,v) + \sum_{u \notin C_v} w_u \cdot (1 - d(u,v)) \ \ ,$$

i.e., an element pays for being in a cluster with unrelated nodes of high weight, and for being partitioned away from similar elements of high weight. The *social cost* is defined as the sum of the costs paid by all elements. In the *maximization* variant, the objective of each node is to maximize its utility

$$\sum_{u \in C_v} w_u \cdot (1 - d(u,v)) + \sum_{u \notin C_v} w_u \cdot d(u,v) \ \ ,$$

i.e., an element achieves utility from being in a cluster with similar nodes of high weight and from being partitioned away from unrelated nodes of high weight. Again, the *social welfare* is the sum of the utilities achieved by all elements.

Correlation clustering essentially models scenarios in which the objective is to either minimize the difference or maximize the similarity among objects within clusters. This type of clustering depends on the relationship among elements. The advertising example fits this model if distances between objects represent willingness to be clients of the same advertising service, and weights represent market influence.

1.2. Our Contribution

We provide a thorough analysis of hedonic clustering games under several models, characterizing Nash equilibria, and proving upper and lower bounds on the prices of anarchy and stability[2]. Our study covers a broad subclass of hedonic games which seems to lack previous investigation from a game-theoretic perspective. This subclass captures clustering as a self-organizing process governed by game theoretic considerations. We note that it is important to study clustering games in several models, since it is a diverse subject, and cannot be captured by a single framework [Blum 2009]. Our models seem to have quite a robust definition, as they fit well with rather different, yet commonly used, clustering criteria.

1.2.1. Fixed Clustering. The first clustering model we consider is fixed clustering. For a general metric, we show that a Nash equilibrium does not necessarily exist. Clearly, imposing high enough penalties on players for changing the location of a centroid (when moving to a different cluster) would guarantee the existence of a Nash equilibrium for both k-median and k-center models. We prove that setting the penalty to be equal to the distance traveled by the centroid suffices. This choice of penalty is very natural, and can be thought of as a fee imposed by a cluster on nodes joining it, in order to cover the incurred expenses. This choice of penalty also resembles the way costs are determined by the VCG mechanism [Groves 1973].

Generally speaking, the issue of improving on the overall system performance even in the face of selfish behavior has been considered extensively in the game theory literature, and designing mechanisms to improve the coordination of selfish agents is a well known idea, e.g., mechanism design and coordination mechanisms. However, in these mechanisms, the system is designed once and for all. In contrast, in our setting, penalties are dynamic and determined by the state of the system.

Since a Nash equilibrium does not necessarily exist in general, we study the fixed clustering game in specific metrics, *i.e.*, tree and line metrics. The strict definition of hedonic games requires that the members of a cluster uniquely determine the location of its centroid. To achieve that one needs to specify some (possibly arbitrary) static tie-breaking rules. However, there exist instances for which no Nash equilibrium exists under certain choices of a static tie-breaking rules, even for line metrics. We circumvent this issue by using tie-breaking rules that are *history dependent*; when the choice of a centroid is not unique, it depends on the previous states of the system.[3] In this respect, our work on fixed clustering deviates from the class of hedonic games. However, we emphasize that even if fixed clustering games were strictly hedonic, our results would not have followed from existing literature. We also point out that many of our results have implications also for static tie-breaking rules. For example, when we prove an instance has a Nash equilibrium under history dependent tie-breaking rules, it implies that it also has a Nash equilibrium under some static tie-breaking rule.

The proof of existence of a Nash equilibrium in a tree metric for the k-median model is rather involved. It is based on a characterization of a centroid, definition of a potential function, and a judiciously chosen schedule of moves of players resulting in an equilibrium. We note that for the k-center model, the proof of existence of a Nash equilibrium under tree metrics requires allowing centroids to be located in an arbitrary location between the two end points of an edge. We believe that this relaxation is not necessary (as is the case for line metric), but we have not managed to prove that. For both the k-median and k-center models, we show that the price of stability is 1, while the price of anarchy is unbounded.

[2]The price of stability is defined as the ratio between the social welfare/cost of the *best* Nash equilibrium and the social optimal solution, while the price of anarchy is defined as the ratio between the social welfare/cost of the *worst* Nash equilibrium and the social optimal solution.
[3]The history of the system is one of the parameters for our tie-breaking rules. For a one-shot game, these rules reduce to static tie-breaking rules.

Going back to the example of ad-hoc networks, from a designer's perspective, our work implies that simple greedy-like algorithms are sufficient for reaching a Nash equilibrium, and thus, the devices do not need to run more sophisticated protocols.

We summarize our results for the fixed clustering model as follows.

(1) The k-median model:
- For general metrics, a Nash equilibrium does not necessarily exist, even in the presence of history-dependent tie breaking rules.
- For line and tree metrics, a Nash equilibrium always exists: the price of stability is 1 and the price of anarchy is unbounded. Moreover, best response dynamics always converge into a Nash equilibrium for line metrics.
- With naturally chosen penalties, a Nash equilibrium exists: the price of stability is 1 and the price of anarchy is unbounded.

(2) The k-center model:
- For general metrics, a Nash equilibrium does not necessarily exist.
- For a line metric, a Nash equilibrium always exists and the price of anarchy is unbounded.
- For tree metrics, a Nash equilibrium exists if the centroids can be placed at any location along edges. In this case, the price of stability is 1 and the price of anarchy is unbounded.
- With naturally chosen penalties, a Nash equilibrium exists: the price of stability is 1 and the price of anarchy is unbounded.

1.2.2. Correlation Clustering. The second clustering model we consider is correlation clustering, for which we obtain results for both the minimization and maximization variants. This model is closely related to additively-separable hedonic games [Gairing and Savani 2010], and techniques used for this class of games can be extended to show that Nash equilibria always exist for correlation clustering games, but finding them is PLS-complete.

Our main results for correlation clustering games are lower and upper bounds on both the price of anarchy and price of stability. The bounds are proved by characterizing the distance between nodes belonging to the same cluster vs. the distance between nodes belonging to different clusters. The specific bounds are:

- In the special case of identical node weights, the price of stability is 1 for both the minimization and maximization variants. For arbitrary weights, the price of stability is strictly larger than 1.
- For the minimization variant, an upper bound on the price of anarchy is $O(n^2)$ and the corresponding lower bound is $n-1$ (for even n). For the special case of identical node weights, the lower bound still holds, and there is an improved upper bound of $n-1$ on the price of anarchy. We note that the lower bound of $n-1$ holds even for a line metric with identical node weights.
- For the maximization variant, the price of anarchy is $\Theta(\sqrt{n})$. This bound holds even for a tree metric with identical node weights. In the case of a line metric, the price of anarchy is $\Theta(n^{1/3})$, even for identical node weights.

For both the fixed clustering and correlation clustering models, an intriguing question is what kind of mechanisms can be used to reduce the price of anarchy. We make a first step in this direction by showing that the price of anarchy of the maximization variant of correlation clustering can be bounded by k (for $k \geq 2$) by limiting the game to k clusters. Moreover, if all node weights are identical, this constraint can be removed after the game reaches a Nash equilibrium, and the price of anarchy still remains at most k.

Previous work. Clustering is a vast area of research with abundant results, and therefore we mention only a few, directly related, results. Fixed clustering is the classic approach to clustering data, and the goal of the optimization problem is to find a minimum cost partitioning of the nodes into k clusters. The k-center problem was considered by Gonzalez [1985], who gave the first 2-approximation algorithm (see also the 2-approximation algorithm of Hochbaum and Shmoys [1985].) For the k-median problem, the first constant-factor algorithm was given by Charikar et al. [1999]. The approximation factor has since been improved in a sequence of papers (see, *e.g.*, [Vazirani 2001]). We note that in case of a tree metric, k-median can be solved optimally in polynomial time [Tamir 1996].

Correlation clustering was first defined by Bansal et al. [2004]. They considered the version where the edges of a complete graph are labeled as either "$+$" (similar) or "$-$" (different), and the goal is to find a partition of the nodes into clusters that agrees as much as possible with the edge labels. They considered both maximizing agreements and minimizing disagreements, obtaining a constant-factor approximation for the former and a PTAS for the latter. These results were generalized for real-valued edge weights; Demaine and Immorlica [2003], Emanuel and Fiat [2003] and Charikar et al. [2005] obtained a logarithmic approximation algorithm for the minimization version, while Charikar et al. [2005] obtained a constant (greater than $1/2$) factor approximation for the maximization version.

From game-theoretic perspectives, in addition to being a hedonic game, our correlation clustering game falls into a class known as *polymatrix games*, introduced by Yanovskaya [1968]. Few other games that fall into both classes were also considered. We mention only those that are most closely related to our model. Hoefer [2007] considered a game called "MaxAgree" which is equivalent to the maximization variant of our correlation clustering game with identical node weights and a limit ℓ on the number of possible clusters. For this game Hoefer [2007] shows that best response dynamics converge in polynomial time to a Nash equilibrium (which does not seem likely in our model), and gives bounds on the price of anarchy. Another example is the game version of Max-Cut, considered by multiple works, *e.g.*, Hoefer [2007] and [Gourvès and Monnot 2009]. In a sense, this game represents the inverse of our model.

A different game-theoretic representation of clustering is given by Pelillo [2009] and Bulò [2009], however, their approach is completely different from ours. In their framework, clusters are derived from a competition among players (who bear no relation to the objects being clustered). Each player simultaneously selects an object, and after having revealed his choice, receives a payoff according to the similarity of the selected object to those of the opponents. It turns out that the competition induces the players to learn in an unsupervised manner a common notion of a clustering by reaching an equilibrium.

In competitive facility location models, facilities compete against one another attempting to attract as many customers as possible in order to maximize market share. Competitive facility location was first introduced by Hotelling [1929] who considered competition on a segment (such as "main street"). In its most simple form, two competitors (providers) choose a location for their facilities (servers), and users can choose among the providers. This model can be generalized to show that in many markets it is rational for producers to make their products as similar as possible. Hotelling's game is also an example of a self-organizing clustering. However, in our terminology, the centroids (the providers) in Hotelling's game are the players and their choice of location determines the clustering; whereas in our game, the customers are the players. Competitive location models have been generalized and extensively studied beyond the work of Hotelling to tree networks and general networks. Excellent surveys on these models appear in Mirchandani and Francis [1990] and Drezner [1995].

Table I. Static tie-breaking rule resulting in no Nash equilibrium.

Cluster Nodes	Centroid Location
$\{A, B\}$	A
$\{B, C\}$	B
$\{A, C\}$	C

2. FIXED CLUSTERING - PRELIMINARIES

We remind the reader that in the fixed clustering game the number of clusters is known beforehand, and is denoted by k. Sections 3 and 4 investigate two variants of this model: k-median and k-center. In both variants, given the set of nodes in a cluster, the choice of a centroid may not be unique. As already mentioned, in order to fit perfectly into the hedonic model, a static tie-breaking rule for choosing a unique centroid is required. However, such tie-breaking rules may have a negative impact on our results, as demonstrated by Proposition 2.1.

PROPOSITION 2.1. *With (static) tie-breaking rules, there may not exist a Nash equilibrium for both the k-median and k-center variants, even for line metrics with three nodes.*

PROOF. Consider a line with three unique nodes A, B and C and assume $k = 2$. For clusters of 2 nodes, either node can be the centroid, and thus, tie-breaking rules can choose the centroid's location according to Table I. In Nash equilibrium there cannot be an empty cluster (as each node that is not a centroid will gain by joining it), and thus, there must be a cluster C_1 with two nodes, and a cluster C_2 containing the remaining node. The above tie-breaking rules imply that the node of C_1 which is not the centroid will become a centroid by moving to C_2, which will decrease its cost. Hence, there is no stable configuration. □

We circumvent this issue by analyzing game dynamics that allow tie-breaking rules which are *history dependent*. Initially, an arbitrary static tie-breaking rule R is used. (Rule R also applies to a one shot game.) During the dynamics, whenever a player performs a move and changes her strategy, each centroid remains at the same node if it is still a valid location for a centroid; otherwise, the static tie-breaking rule R is used to relocate it. The use of a history dependent rule implies that the cost observed by a player does not solely depend on the identity of the other players in her cluster, but also on the history of the game. We can thus consider the fixed clustering model as a hedonic game with an additional attribute.

We note that although our results are proved under history dependent tie-breaking rules, they often have implications for static tie-breaking rules. For example, an instance that has a Nash equilibrium under history dependent tie-breaking rules also has a Nash equilibrium under some static tie-breaking rule. On the other hand, our proofs that some instances have no Nash equilibria under history dependent tie-breaking rules apply also under static tie-breaking rule.

To avoid confusion, we introduce a notation for referring to the different cases of fixed clustering we analyze. Each case is described by notation of the form $A|B|C$, where:

- A represents the model – it can take two values *median* and *center*.
- B represents the constraint assumed on the metric – it can take three values *general*, *tree*, and *line*.
- C: sometimes, we consider a slight modification of the problems presented above: either by imposing penalties on deviating nodes or allowing centroids to be located along edges. If we consider such a modification, then C represents it – it can take two

values: *penalties* and *edges*. Whenever we consider such a modification, the details are explained beforehand. When considering a case with no modification, we omit C (*i.e.*, we write an expression of the form $A|B$).

3. THE K-MEDIAN MODEL

In the k-median model, the cost of a cluster C is defined as the sum of the distances between all members of the cluster and its centroid, and the centroid $c(C)$ is defined as the node minimizing the cost of the cluster, that is, $c(C) = \min_{u \in C} \left\{ \sum_{v \in C} d(u, v) \right\}$. We denote by $D(u, C)$ the sum of distances between u and the other nodes in C. Thus, $c(C) = \min_{u \in C} D(u, C)$. For ease of notation, we denote by $D(C)$ the cost of a cluster C, that is, $D(C) = D(c(C), C)$.

The cost of a node v in the k-median clustering game is defined as its distance from the centroid of its cluster, $d(v, c(C_v))$. A clustering configuration of the k-median clustering game is a Nash equilibrium if no player can reduce its cost by choosing a different cluster, assuming the other players stay in their cluster. We assume a node v changes its strategy from cluster C_1 to C_2 only in case it strictly decreases its distance to the centroid, that is $d(v, c(C_1)) > d(v, c(C_2^{+v}))$. As mentioned, we assume that following a move performed by a player, the centroid of a cluster will not change its location unless forced (*i.e.*, only in case the sum of distances of the points from the new location of the centroid is strictly lower than from its previous location).

We first consider the general metric case, and prove that a Nash equilibrium does not necessarily exist. Moreover, we show that this is the case even if centroids are allowed to be located at any location along edges, rather than only at a node of their own cluster. Then, we notice that by imposing a high enough penalty on a node whose move to a cluster changes the location of its centroid, a Nash equilibrium is guaranteed to exist. Moreover, we show that it is enough to set the penalty to be equal to the distance traveled by the centroid. Motivated by these results, we study the game in line and tree metrics, and show that in these cases a Nash equilibrium always exists, with no further assumptions such as penalties. Note that the existence of a Nash equilibrium in the line case is implied by our result for the tree case. Still, we consider line metrics separately, since we can show that for such metrics best-response dynamics converge to an equilibrium.

3.1. The General Metric Case

For a general metric there is no guarantee that Nash equilibrium exists.

THEOREM 3.1. *There exists an instance of* median|general *with* 9 *nodes having no Nash equilibria for two clusters.*

Figure 1 displays a graph representing an instance having the properties guaranteed by Theorem 3.1. The numbers on the edges of the graph represent the distances. The proof that no Nash equilibrium exists for this graph is done via a case analysis, showing that any selection of two nodes to be the centroids of the clusters results in a configuration where either a cluster has a non optimal centroid choice, or there is a node wishing to deviate. The proof itself is long and technical, and is, therefore, deferred to Appendix A.

A natural way to try to guarantee the existence of Nash equilibrium is by allowing the centroid of a cluster to be located at any location along an edge of the graph. Formally, every edge (u, v) and $\lambda \in [0, 1]$ represents a possible location for a centroid. The distance of this centroid from a node w is $\min\{d(u, w) + \lambda \cdot d(u, v), d(v, w) + (1 - \lambda) \cdot d(u, v)\}$. The centroid of the cluster is placed at the location minimizing the sum of distances from the nodes of the cluster. Unfortunately, this generalization fails to guarantee the

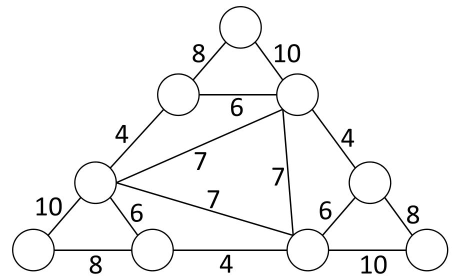

Fig. 1. Graph with no Nash equilibrium (assuming two clusters).

existence of a Nash equilibrium. The graph in Figure 1 represents a counterexample for this case as well. In order to establish the counterexample based on this graph, we use the notion of weak Nash equilibria.

Our original definition of Nash equilibrium states that a centroid will not change location unless forced. If the centroid is forced to change location, it might have multiple locations to which it can move. The choice among these possible new locations is based on some arbitrary tie-breaking rule. On the other hand, a configuration is said to be in *weak Nash equilibrium* if there is no node u wishing to deviate to any cluster $C \neq C_u$, given that following u's deviation, C's centroid will move to the location farthest away from u among its possible locations. In other words, if there are several possible locations for the centroid of C after u's move, then the centroid will move to the worst possible location from the point of view of u. Moreover, this move occurs even if the original location of the centroid is still a possible location for it.

Clearly, every Nash equilibrium is also a weak Nash equilibrium, though the reverse is not necessarily true. We would like to show that the graph in Figure 1 allows no weak Nash equilibria with two clusters; and therefore, it also does not allow any Nash equilibria (assuming two clusters). The proof consists, again, of a case analysis. Before doing the case analysis, we need the next lemma. Lemma 3.2 shows that it is sufficient to consider only a finite set of cases because we can assume the centroids of a weak Nash equilibrium are half integral locations on edges. The lemma also shows that checking whether a given configuration is a weak Nash equilibrium is a finite task.

LEMMA 3.2. *An instance of* median|general|edges *with integral edge distances has a weak Nash equilibrium if and only if it has a weak Nash equilibrium configuration in which the centroids are placed in half integral locations on edges. Moreover, given a node u and the set P of potential locations for the centroid of u's cluster, there exists a point in P which is as far as any other point in P from u and is located in a half integral location on an edge.*

PROOF. Consider an instance of median|general|edges and a cluster C with centroid $c(C)$ located on some edge (u, v). Let α be the fractional part of the distance between $c(C)$ and u, and assume, w.l.o.g, that $\alpha \in [0, 1/2]$. Assume we move $c(C)$ by a small fraction and change α in the range $[0, 1/2]$. It can easily be noted that for every node $w \in C$, its distance from $c(C)$ is given by one of two formulas: $\delta_w + \alpha$ or $\delta_w + (1 - \alpha)$, where δ_w is a w-specific integer constant. Both formulas are linear in α, hence, the derivative of $D(C)$ by α is constant in the range (0, 1/2). Since a centroid is placed in the location which minimizes $D(C)$, one of the two following cases must hold:

— The derivative of $D(C)$ by α in the range $(0, 1/2)$ is either positive or negative, and thus, the original value of α was either 0 or $1/2$, *i.e.*, the centroid is in a half integral location.

— The derivative of $D(C)$ by α is 0. In this case, every value of $\alpha \in [0, 1/2]$ gives an equally good location for the centroid of C.

The second part of the lemma follows immediately from the above discussion since the location of $c(C)$ which is most further away from u is given by either $\alpha = 0$ or $\alpha = 1/2$ (among the locations corresponding to $\alpha \in [0, 1/2]$). Thus, for every point in P not located in a half integral location on an edge, one can find another point in P which is located in a half integral location on an edge and is at least as far away from u as the original location.

Consider now a weak Nash equilibrium configuration $\mathcal{E}$, and let $\mathcal{E}'$ be the same clustering configuration with every centroid kept at the same location if it is half integral, or else moved to the closest half integral location. Following the discussion above, every centroid that changed its location still minimizes the cost of its cluster. We next show that $\mathcal{E}'$ is a weak Nash equilibrium.

Consider a node v that is a member of cluster C_1 in $\mathcal{E}$. Let $\mathcal{E}_{v \to C_2}$ (respectively, $\mathcal{E}'_{v \to C_2}$) denote the clustering that results from $\mathcal{E}$ (respectively, $\mathcal{E}'$) after moving v to a cluster C_2, and relocating the center of C_2 to the location farthest away from v among the possible new locations. Following our observations, the centroid of C_2 in $\mathcal{E}_{v \to C_2}$ and $\mathcal{E}'_{v \to C_2}$ is in a half integral location. Clustering $\mathcal{E}$ is a weak Nash equilibrium, hence, v has no incentive to deviate from C_1. Formally, if we denote by $d(v, C_1, \mathcal{E})$ the distance of v from the center of cluster C_1 in $\mathcal{E}$, then for every cluster $C_2 \neq C_1$ we have:

$$d(v, C_1, \mathcal{E}) \leq d(v, C_2, \mathcal{E}_{v \to C_2}) \ .$$

Notice also that C_2 contains the same set of nodes in $\mathcal{E}_{v \to C_2}$ and $\mathcal{E}'_{v \to C_2}$, and therefore, $d(v, C_2, \mathcal{E}_{v \to C_2}) = d(v, C_2, \mathcal{E}'_{v \to C_2})$. We now distinguish between two cases:

— The location of the centroid of C_1 was not changed in the process of creating $\mathcal{E}'$. Hence, $d(v, C_1, \mathcal{E}') = d(v, C_1, \mathcal{E}) \leq d(v, C_2, \mathcal{E}_{v \to C_2}) = d(v, C_2, \mathcal{E}'_{v \to C_2})$.
— The location of the centroid of C_1 is different between $\mathcal{E}$ and $\mathcal{E}'$. In this case, it changed by at most $1/4$, and thus:

$$\begin{aligned} d(v, C_1, \mathcal{E}) \leq d(v, C_2, \mathcal{E}_{v \to C_2}) &\Rightarrow d(v, C_1, \mathcal{E}') - 1/4 \leq d(v, C_2, \mathcal{E}'_{v \to C_2}) \\ &\Rightarrow d(v, C_1, \mathcal{E}') - d(v, C_2, \mathcal{E}'_{v \to C_2}) \leq 1/4 \ . \end{aligned}$$

As all centroids in $\mathcal{E}'$ and $\mathcal{E}'_{v \to C_2}$ are located in half integral locations, it follows that the distances are also half integral, and therefore, $d(v, C_1, \mathcal{E}') \leq d(v, C_2, \mathcal{E}'_{v \to C_2})$.

We conclude that in both cases, v has no incentive to deviate in $\mathcal{E}'$, and therefore, $\mathcal{E}'$ is a weak Nash equilibrium. □

To verify that Figure 1 allows no weak Nash equilibria with two clusters, we do the following:

(1) Enumerate every partition of the nodes into two clusters.
(2) Enumerate for each partition all the optimal centroid locations which are half integral locations on edges.
(3) Show that no choice of partition and and two centroids induces a weak Nash equilibrium.

Notice that the last step can be performed due to the second part of Lemma 3.2. Due to the huge number of cases to check, we do the check using a computer program.[4] The next theorem follows by checking all the cases.

[4]The source code of the program (in VB.net) can be found at: http://theory.epfl.ch/moranfe/Resources/Journals/ACM%20Transactions%20on%20Parallel%20Computing%202014/NoNashValidation.vb.

THEOREM 3.3. *There exists an instance of* median|general|edges *with* 9 *nodes having no weak Nash equilibria (and thus, also no Nash equilibria) for two clusters.*

In order to guarantee the existence of a Nash equilibrium, we add a rule penalizing a node whose move to a cluster C changes the location of the centroid of C. That is, the cost of a node u performing an improvement move will consist of two values:

— The distance from the updated centroid of the new cluster, $d(u, c(C^{+u}))$.
— A cost equal to the distance between the original and new centroids of C, *i.e.*, $d(c(C), c(C^{+u}))$.

Intuitively, one can think of this choice of a penalty as a fee imposed by a cluster on nodes joining it, in order to recover the incurred expenses. Note that the total cost of the nodes in the cluster deserted by u can only decrease.

LEMMA 3.4. *In* median|general|penalties*, a node* u *will only deviate from cluster* C_1 *to cluster* C_2 *if its distance from the original location of the centroid* $c(C_2)$ *is shorter than its distance from* $c(C_1)$*.*

PROOF. Before deviating, u pays for its distance from $c(C_1)$. If u deviates, it will have to pay both for its distance to $c(C_2^{+u})$ and for the distance between $c(C_2)$ and $c(C_2^{+u})$. By the triangle inequality, the sum of these two distances is at least the distance from u to the original location of the centroid $c(C_2)$. □

Lemma 3.4 implies that the above penalizing rule prevents a node from deviating unless the deviation results in a decreased social cost. Since the social cost cannot decrease forever, the game must converge to Nash equilibrium. The following theorem formalizes this argument.

THEOREM 3.5. *In* median|general|penalties*, every best-response move strictly decreases the social cost. Hence, there always exists a Nash equilibrium, and it is guaranteed to be reached by best-response dynamics.*

PROOF. Think of a best-response move as a two steps process. In the first step, the node moves from one cluster to the other, and in the second step the centroids of the two clusters change locations according to their new set of nodes. By Lemma 3.4, the first step is guaranteed to reduce the objective function (social cost). The second step cannot increase the social cost as the new location selected for the centroid of each cluster is the one minimizing its cost. Thus, the social cost must decrease whenever a best response move occurs. As the number of configurations is finite, Nash equilibrium is guaranteed to be reached by a best-response dynamics. □

COROLLARY 3.6. *The price of stability of* median|general|penalties *is* 1*.*

PROOF. By Theorem 3.5, the social cost is guaranteed to decrease with every best response move. Thus, any optimal solution must be a Nash equilibrium. □

The next proposition shows that the price of anarchy of the k-median clustering game without penalties is unbounded. Note that a Nash equilibrium in the setting without penalties is also valid with penalties (as penalties can only decrease the benefit of a move), thus, the proposition also holds when there are penalties.

PROPOSITION 3.7. *The price of anarchy of the* k*-median clustering game is unbounded, even when restricted to* median|line*.*

PROOF. Consider a line with three nodes at locations 0, 1 and M, for some large M, and assume $k = 2$. There are two clusters $\{0\}$ and $\{1, M\}$, and the centroid of the second cluster is the node 1. Clearly this is a Nash equilibrium, since the only node which is

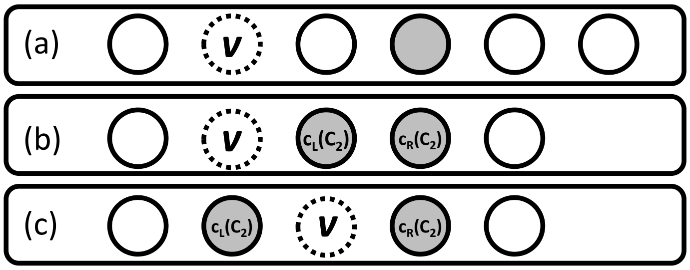

Fig. 2. Three cases from the proof of Lemma 3.9. The circles in each case represent the nodes of C_2, except for the dotted circle which represent v. The dark shaded circles in each case represent the centroids of C_2. The centroid of the new cluster C_2^{+v} is: (a) unchanged, (b) $c_L(C_2)$ and (c) v.

not a centroid will gain nothing by deviating. However, the cost of this equilibrium is $M-1$. On the other hand, the optimal solution has the following two clusters: $\{0,1\}$ and $\{M\}$, yielding a cost of 1. The price of anarchy of this instance is thus $M-1$. □

3.2. The Line Metric Case

We prove that if players are nodes on a line, a Nash equilibrium always exists (without penalties, and allowing a centroid to be placed only at a node). We begin by characterizing the centroid of a cluster on a line. A median node of a cluster is a cluster node for which the number of cluster nodes to its left and to its right differ by at most 1.

LEMMA 3.8. *The centroid of a cluster in* median|line *is a median of that cluster.*

The proof of Lemma 3.8 is quite standard, and is, therefore, deferred to Appendix B. In case the number of nodes in a cluster is odd, there exists a single median node. Thus, this node is the single optional centroid node, and divides the other nodes in the cluster to right and left sets where each set contains $(m-1)/2$ nodes. In case the number of nodes in a cluster is even, there are two optional centroid nodes. These are two consecutive nodes on the line, each dividing the nodes to right and left sets such that one of the sets contains $\lceil \frac{m-1}{2} \rceil$ nodes and the other set contains $\lfloor \frac{m-1}{2} \rfloor$ nodes.

Given a clustering configuration $\mathcal{F}$, we define a potential function Φ, and show that it strictly decreases with each strategy change performed by a player. We assume a player changes her strategy only if it strictly decreases her distance to the center. The potential function is equal to the social cost function, that is

$$\Phi(\mathcal{F}) = \sum_{i=1}^{k} \sum_{v \in C_i} d(v, c(C_i)) \ . \tag{1}$$

The following lemma is proved using the characterization of the centroid given by Lemma 3.8.

LEMMA 3.9. *In* median|line*, the potential function* Φ *strictly decreases during the clustering game's natural dynamics.*

PROOF. A strategy change performed by a node v consists of moving from a cluster C_1 to a different cluster C_2 in order to lower its distance to the center. We denote by C_1^{-v} cluster C_1 after v's leave and by C_2^{+v} cluster C_2 after v joins. We consider the nodes that are influenced by the move of v, and show that the sum of their distances to their centers strictly decreases:

— Nodes in C_1: following node v's move, the center of C_1^{-v} may change. As C_1^{-v} consists now of a subset of the nodes that were in C_1 (all the previous nodes of C_1 except v), by definition of the centroid it holds that $D(C_1^{-v}) \leq D(C_1) - d(c(C_1), v)$.
— Node v: as v is the node performing a change of strategy, its distance to the center strictly decreases, that is, $d(v, c(C_2^{+v})) < d(v, c(C_1))$.
— Nodes in C_2: we consider two cases, according to the number of nodes in C_2.
  (1) The number of nodes in C_2 is odd (see part a of Figure 2). In this case, after v joins C_2, the number of nodes in C_2^{+v} becomes even, and its old centroid remains one of the two optional new centroids. As the centroid doesn't change place unless forced, $c(C_2) = c(C_2^{+v})$, and the distance of nodes in C_2 from their center remain unchanged. Note that we assume here, that if a centroid is not forced to move, it remains at its current location.
  (2) The number of nodes in C_2 is even. In this case, there are two optional centroids, that is two consecutive nodes $c_L(C_2)$ and $c_R(C_2)$ of C_2, where $c_L(C_2)$ is the left median node of C_2 and $c_R(C_2)$ is the right median node of C_2. Note that $D(c_L(C_2), C_2) = D(c_R(C_2), C_2)$. If v is located to the left of $c_L(C_2)$ or to the right of $c_R(C_2)$ (see part b of Figure 2) then one of them is also the centroid of $c(C_2^{+v})$, and therefore, the total distance of the nodes of C_2 from their centroid remains unchanged. If v is located between $c_L(C_2)$ and $c_R(C_2)$ (see part c of Figure 2), then v himself is the sole centroid of $c(C_2^{+v})$. Notice that:

$$
\begin{aligned}
&D(v, C_2) = \sum_{u \in C_2} d(v, u) \\
&= \sum_{u \in C_2} \left[\frac{d(c_L(C_2), v)}{d(c_L(C_2), c_R(C_2))} \cdot d(v, c_L(C_2)) + \frac{d(v, c_L(C_1))}{d(c_L(C_2), c_R(C_2))} \cdot d(v, c_R(C_2)) \right] \\
&= \frac{d(c_L(C_2), v)}{d(c_L(C_2), c_R(C_2))} \cdot D(c_L(C_2), C_2) + \frac{d(v, c_L(C_1))}{d(c_L(C_2), c_R(C_2))} \cdot D(c_R(C_2, C_2)) \ .
\end{aligned}
$$

  Since $D(c_L(C_2), C_2) = D(c_R(C_2), C_2)$, the rightmost hand side of the above equality is also equal to them. □

As the strategy space of all nodes is finite, the potential function will decrease following each move performed by a node until reaching a local (or global) minimum, corresponding to a Nash equilibrium.

COROLLARY 3.10. *A Nash equilibrium always exists for* median|line*, and such an equilibrium is guaranteed to be reached by best-response dynamics.*

3.3. The Tree Metric Case

In this section we assume that the distances are defined by a tree metric, and the players are exactly the nodes of the tree. We prove that in this case, a Nash equilibrium always exists. First, we characterize the centroid node of a cluster in the tree case. To this end we define the *median node of a tree* as follows.

Definition 3.11. Given a tree T with $|V|$ nodes, a node $v^* \in T$ is called a *median node* if its removal partitions T into connected components of size at most $\left\lceil \frac{|V|-1}{2} \right\rceil$ each.

The following lemma is well known; its proof appears, for completeness, in Appendix B.

LEMMA 3.12. *There are at most* 2 *median nodes in a tree.*

In order to define the relation between a median and the centroid of a cluster, we use the next definition.

Definition 3.13. A cluster C has the *closure property* if all nodes on each path between two nodes of C belong to C as well.

LEMMA 3.14. *In* median|tree*, the centroid of a cluster with the closure property is always a median node of the cluster.*

PROOF. Given a cluster C of size $|C|$, assume by way of contradiction that the centroid is a node v that is not a median node. We consider the connected components of the subtree C after v's removal. As v is not a median node, the maximal connected component Q_{max} contains at least $\left\lceil \frac{|C|-1}{2} \right\rceil + 1$ nodes. Thus, by changing the location of the centroid from v to $v_{in}(Q_{max})$ (the sole neighbor of v in the connected component Q_{max}), we can lower the sum of all nodes distance to the center, as the distance of $\left\lceil \frac{|C|-1}{2} \right\rceil + 1$ nodes will decrease by $d(v, v_{in}(Q_{max}))$, while the distance of the other nodes in C will increase by the same value. This is a contradiction to the assumption that v is the centroid of C. $\square$

In order to prove that the k-median clustering game always has a Nash equilibrium, we use the potential function of Equation (1), and describe a schedule that converges to an equilibrium. Unlike in the line metric case, we cannot simply use an arbitrary schedule of a best-response dynamics because the closure property can easily be violated by a best-response move. Instead, we need to define a set of moves that has two properties: it keeps the closure property, which allows us to use the median characterization of centroids, and it is guaranteed to strictly decrease the potential function.

We describe the convergence schedule. It consists of iterations, where each iteration starts and ends with a clustering configuration in which all clusters have the closure property. We call such a configuration a *closed configuration*. Starting from a closed configuration, consider a node v in cluster C_1 wishing to move to cluster C_2. All nodes on the path between v and $c(C_2)$, denoted by $\delta(v, c(C_2))$, are either already in C_2 or wish to move to C_2 as well (if there are other clusters with equal distance to centroid, we choose C_2). This follows since nodes in $\delta(v, c(C_2))$ are closer than v is to $c(C_2^{+v})$. Note that the center of C_2 will move in the same way given that any node of $\delta(v, c(C_2))$ moves to C_2. In addition, if there is a node $w \in \delta(v, c(C_2))$ having a better cluster C_m such that $d(w, c(C_m^{+w})) < d(w, c(C_2^{+w}))$, then the same holds for v as well since: $d(v, c(C_m^{+v})) \leq d(w, c(C_m^{+w})) + d(w, v) < d(w, c(C_2^{+w})) + d(w, v) = d(v, C_2^{+v})$.

Thus, when starting a new iteration of the schedule, we choose a node u adjacent to C_2, and make it perform a best-response move from C_1 to C_2 (we call this move "first best-response"). Following this move, there are two options. In case C_1^{-u} remains connected, then it still has the closure property, and we can select a new first best-response (assuming we did not reach a Nash equilibrium). Otherwise, C_1^{-u} contains multiple connected components, and there is a component C_1' containing the centroid of C_1^{-u} (see Figure 3 for a graphical illustration of our notation). Let C_1'' denote $C_1^{-u} - C_1'$, *i.e.*, C_1'' contains all nodes of C_1 except for the node that made the first best-response (u) and the nodes of the connected component containing C_1^{-u}'s centroid (that is, C_1'). Next, we move all nodes of C_1'' to cluster C_2, and finish the iteration. Note that beside of the "first best-response" move, all other moves in the iteration need not be best-response moves.

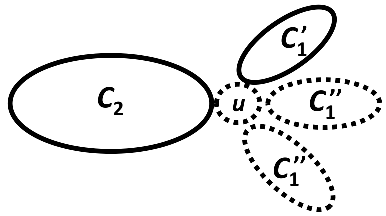

Fig. 3. Clusters C_1 and C_2. Everything displayed in the figure except for C_2 is part of C_1. Node u is adjacent to cluster C_2 and deviates to this cluster. Following the deviation, cluster C_1 is partitioned into a multiple unconnected parts. C_1' is the part of C_1 containing the centroid of C_1 after the deviation of u. C_1'' represents the rest of C_1 (without u and C_1'). The iteration of our schedule completes when C_1'' deviates as a whole to C_2. Notice that following the last deviation both C_1 and C_2 regain the closure property.

Both C_1' and $C_2 \cup \{u\} \cup C_1''$ are closed clusters. No other cluster is affected by the iteration, hence, by the end of the iteration, we get back to a closed configuration. We are now ready to use our potential function, and show that it strictly decreases with each iteration of the above schedule. We conceptually divide the iteration into three steps, and show that none of them can increase the potential function, and at least one of them strictly decreases it. The three steps are as follows.

(1) The first best-response move made by u, including the possible relocation of the clusters' centroids.
(2) The move of the nodes of C_1'' to C_2, assuming the clusters' centroids are not allowed to move.
(3) A possible move of the clusters' centroids to new locations (following the move of the nodes of C_1'').

LEMMA 3.15. *In step 1 of the iteration, the potential function P strictly decreases.*

PROOF. We consider a node u moving from C_1 to C_2, and show that the sum of distances from all nodes in C_1^{-u} and C_2^{+u} to their respective centers decreases. This is clearly true for u (the distance strictly decreases since u had an incentive to move) as well as for C_1^{-u} (as the centroid of C_1^{-u} is relocated in order to minimize its total distance from the cluster's nodes). As for C_2, observe that it has the closure property by the above schedule, hence Lemma 3.14 applies. We consider two cases. Let $Q_1, \dots, Q_m$ denote the connected components in C_2 formed after the removal of $c(C_2)$.

(1) Node u joins a connected component Q_j of size at most $\left\lceil \frac{|C_2|}{2} \right\rceil - 1$ in C_2. Then, after it joins C_2, $c(C_2)$ remains a median node of C_2^{+u} and thus the centroid remains at the same location.
(2) Node u joins a connected component Q_j of size $\left\lceil \frac{|C_2|}{2} \right\rceil$ in C_2. This can occur if $|C_2|$ is even, and then $\left\lceil \frac{|C_2|}{2} \right\rceil = \left\lceil \frac{|C_2|-1}{2} \right\rceil$ and Q_j is a maximal connected component. In this case, there are two optional centroids (median nodes) in C_2, that is $c(C_2)$ and $v_{in}(Q_j)$ (the sole neighbor of $c(C_2)$ in the connected component Q_j). Clearly, $D(c(C_2), C_2) = D(v_{in}(Q_j), C_2)$. Now, after v joins Q_j, the centroid is forced to move to $v_{in}(Q_j)$, but the total distance of the vertices in C_2 from their centroid remains unchanged. □

LEMMA 3.16. *In steps 2 and 3 of the iteration, the potential function P cannot increase.*

PROOF. In step 3, the centroids of $C_2 \cup \{u\} \cup C_1''$ and C_1' are relocated to their optimal locations. Whenever a centroid is relocated, it minimizes its total distance from the nodes of the cluster. Hence, relocation of centroids cannot increase P.

The rest of the proof is devoted for step 2. Since u had an incentive to make the first best-response move, $d(u, c(C_1)) > d(u, c(C_2^{+u}))$. We also observe that $d(u, c(C_1)) \leq d(u, c(C_1^{-u}))$, otherwise, the centroid of C_1 could decrease the total distance to the cluster's nodes by moving to $c(C_1^{-u})$. Combining, we get $d(u, c(C_1^{-u})) > d(u, c(C_2^{+u}))$. The shortest path from every node $w \in C_1''$ to either $c(C_2^{+u})$ or $c(C_1^{-u}) \in C_1'$ pass through u, hence,

$$d(w, c(C_1^{-u})) = d(w, u) + d(u, c(C_1^{-u})) > d(w, u) + d(u, c(C_2^{+u})) = d(w, c(C_2^{+u})) \ .$$

Therefore, moving the nodes of C_1'' from C_1^{-u} to C_2^{+u}, without relocating the clusters' centroids, can only decrease P. □

COROLLARY 3.17. *The potential function P strictly decreases in each iteration of the above schedule.*

As the strategy space of all players is finite, the potential function decreases following each iteration performed by the schedule, until reaching a local minimum (or global). By definition, the above schedule continues as long as there are nodes that can perform best-response moves. Therefore, the above local minimum corresponds to a Nash equilibrium.

COROLLARY 3.18. *A Nash equilibrium always exists for* median|tree.

We consider only dynamics of closed configurations. The next lemma states that it is not restrictive.

LEMMA 3.19. *For any configuration of* median|tree*, there always exists a closed configuration which is at least as good with respect to social cost.*

PROOF. Consider a two steps schedule. In the first step, fix the locations of the centroids, and let every node move to the cluster whose centroid is closest to the node. In the second step, allow the centroids to relocate themselves.

During the first step, the social cost can only decrease, as the cost of each player can only decrease. Moreover, after the step completes, the configuration is closed, assuming a consistent tie breaking rule is used. By definition, relocating a centroid decreases the social cost. Hence, the cost cannot increase during the second step as well. □

As the potential function of the game is equal to the objective function, an optimal k-clustering closed configuration for the k-median model is also a Nash equilibrium (no move can further decrease the global minimum point of the objective function). Moreover, following Lemma 3.19, the cost of such a configuration is equal to the cost of an optimal configuration. We thus get the following corollary.

COROLLARY 3.20. *The price of stability of* median|tree *is 1.*

Note that since an optimal solution for the k-median problem on a tree metric can be computed in polynomial time, and the transformation into a closed configuration given by Lemma 3.19 can be efficiently implemented, one can find, in polynomial time,a Nash equilibrium which as good as the optimal solution.

4. THE K-CENTER MODEL

In the k-center model, the cost of a cluster C is defined by its radius, which is the maximal distance between its centroid and a node of the cluster. The centroid $c(C)$ is defined as the node minimizing the cost of the cluster, that is,

$$c(C) = \min_{u \in C} \left\{ \max_{v \in C} d(u, v) \right\} \ .$$

The cost of a node v in the k-center clustering game is defined as its distance from the centroid of its cluster, $d(v, c(C_v))$. A clustering configuration of the k-center clustering game is a Nash equilibrium if no user can reduce its cost by choosing a different cluster, assuming the other users stay in their individual clusters. As mentioned, we assume history dependent tie-breaking rules. In other words, following a move performed by a player, the centroid of a cluster will not change its location unless forced.

We first consider the case of a line metric, and prove that a Nash equilibrium always exists. Then, we turn to tree metrics, and guarantee the existence of a Nash equilibrium in case centroids can be placed anywhere along edges. Finally, we consider general metrics, and prove that a Nash equilibrium does not necessarily exist, even if centroids are allowed to be placed at any location along edges. We note that by imposing a high enough penalty on a node whose move affects the location of the target cluster's centroid, existence of a Nash equilibrium is guaranteed. We further show that setting the penalty to be equal to the distance traveled by the centroid is enough. As for the price of anarchy, we show it is unbounded in all settings considered.

4.1. The Line Metric Case

In this section we prove the following theorem.

THEOREM 4.1. *A Nash equilibrium always exists for* $\mathsf{center|line}$.

First observe that Theorem 4.1 holds for G in case $|V(G)| \leq k$. This is true since for such graphs one can assign a different cluster for every node, and get a trivial Nash equilibrium. Thus, we assume from now on $|V(G)| \geq k$. A *standard form* clustering configuration $\mathcal{E}$ is a configuration in which the nodes of each cluster are consecutive nodes, and cluster i appears to the left of cluster $i+1$ (assuming $i < k$) and to the right of cluster $i-1$ (assuming $i > 1$). We denote by $\mathsf{left}_i(\mathcal{E})$ and $\mathsf{right}_i(\mathcal{E})$ the leftmost and rightmost nodes of cluster i in $\mathcal{E}$. For ease of notation, given two nodes u, v, we write $u < v$ when u is located at the left of v. The following lemma gives a useful property of standard form clustering configurations.

LEMMA 4.2. *Given a standard form clustering configuration $\mathcal{E}$. If a best-response of a node u of cluster i is to deviate to cluster i', then:*

— *If $i' < i$, then a possible best-response for $\mathsf{left}_i(\mathcal{E})$ is to deviate to cluster $i-1$.*
— *If $i' > i$, then a possible best-response for $\mathsf{right}_i(\mathcal{E})$ is to deviate to cluster $i+1$.*

PROOF. Let us assume, without loss of generality, that $i' < i$. Notice that any deviation of some node from cluster i to cluster $i'' \neq i$ will result in the same new centroid of cluster i'', and let $c_{i''}$ denote this new centroid. For consistency of the notation, let c_i denote the centroid of cluster i in $\mathcal{E}$. Since deviating to i' is a best-response of u, we learn that:

$$d(c_{i'}, u) = \min_{1 \leq i'' \leq k} d(c_{i''}, u) \ .$$

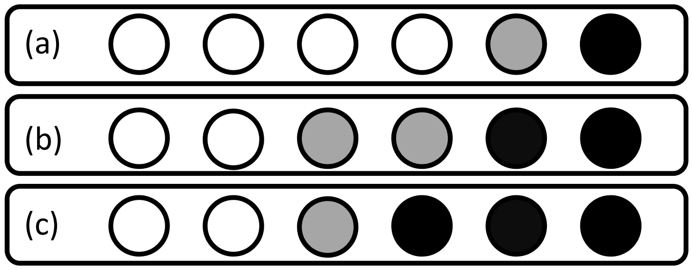

Fig. 4. Three configurations demonstrating the proof of Theorem 4.1. The circles in each case represent the nodes of a line metric and their colors represent the clusters they belong to.

On the other hand:

$$d(c_{i''}, \mathsf{left}_i(\mathcal{E})) = d(c_{i''}, u) + \begin{cases} -d(u, \mathsf{left}_i(\mathcal{E})) & \text{if } i'' < i \ , \\ d(u, \mathsf{left}_i(\mathcal{E})) & \text{if } i'' > i \ , \\ d(c_i, \mathsf{left}_i(\mathcal{E})) - d(c_i, u) & \text{otherwise} \ . \end{cases}$$

Since $i' < i$ (by our assumption), and $d(c_i, \mathsf{left}_i(\mathcal{E})) - d(c_i, u) \geq -d(u, \mathsf{left}_i(\mathcal{E}))$ (by the triangle inequality), we get that the difference $d(c_{i''}, \mathsf{left}_i(\mathcal{E})) - d(c_{i''}, u)$ is minimized for $i'' = i'$, and thus, $d(c_{i'}, u) = \min_{1 \leq i'' \leq k} d(c_{i''}, \mathsf{left}_i(\mathcal{E}))$.

To complete the proof of the lemma, we just need to observe that $c_{i'} \leq c_{i-1} \leq u$ since $c_{i''}$ must be located in i''-th cluster of $\mathcal{E}$ for every $1 \leq i'' \leq k$. Hence, $d(c_{i-1}, u) < d(c_{i''}, u)$ for every $i'' < i - 1$. □

For the next step of the proof we need an additional definition. A standard form clustering configuration $\mathcal{E}$ is a half Nash equilibrium if it contains k clusters, and for every cluster i, $\mathsf{left}_i(\mathcal{E})$ does not wish to deviate.

OBSERVATION 4.3. *Every instance of* center|line *has a half Nash equilibrium.*

PROOF. Consider a configuration $\mathcal{E}$ where the all nodes belong to cluster 1 except for the $k - 1$ rightmost nodes. The last nodes are assigned one to each cluster in a rightward increasing cluster index way (an example of $\mathcal{E}$ for the case of three clusters and 6 nodes is depicted as part a of Figure 4). Clearly, $\mathsf{left}_1(\mathcal{E})$ does not wish to deviate because any deviation will place its new centroid to the right of $\mathsf{right}_1(\mathcal{E})$. Also, for every $2 \leq i \leq k$, $\mathsf{left}_i(\mathcal{E})$ is the sole node of cluster i, and therefore, does not wish to deviate. □

Consider now a potential function $\psi(\mathcal{E})$ defined as following:

$$\psi(\mathcal{E}) = \sum_{i=2}^{k} d(\mathsf{left}_1(\mathcal{E}), \mathsf{left}_i(\mathcal{E})) \ .$$

LEMMA 4.4. *If $\mathcal{E}$ is a half Nash equilibrium which is not a Nash equilibrium, then there exists a half Nash equilibrium $\mathcal{E}'$ such that $\psi(\mathcal{E}) > \psi(\mathcal{E}')$.*

PROOF. By Lemma 4.2, there exists a cluster i such that one best-response of $\mathsf{left}_i(\mathcal{E})$ in $\mathcal{E}$ is to deviate to cluster $i - 1$ or one best-response of $\mathsf{right}_i(\mathcal{E})$ is to deviate to cluster $i + 1$. The first case cannot occur since $\mathcal{E}$ is a half Nash equilibrium. Hence, one of $\mathsf{right}_i(\mathcal{E})$'s best-responses is to deviate to cluster $i+1$. Let $\mathcal{E}'$ be the configuration resulting from this deviation.

Parts b and c of Figure 4 depict an example of a possible pair $\mathcal{E}$ and $\mathcal{E}'$ of configurations. Notice that in this case the rightmost node of the gray cluster deviates and becomes the leftmost node of the black cluster (*i.e.*, cluster i is the gray cluster in this example). Observe that the leftmost node of the black cluster "moves" left, while the leftmost nodes of the other clusters are left unchaged. Thus, ψ decreases in this example. It is easy to see that in general we have that $\mathcal{E}'$ is a standard form clustering configuration and $\psi(\mathcal{E}) > \psi(\mathcal{E}')$.

We are left to show that $\mathcal{E}'$ is a half Nash equilibrium. Consider an arbitrary cluster i', and assume, for the sake of contradiction, that $\mathsf{left}_{i'}(\mathcal{E}')$ wishes to deviate in $\mathcal{E}'$ to cluster i''. There are a two cases to consider:

— If $i'' < i'$, then by Lemma 4.2, one possible best-response of $\mathsf{left}_{i'}(\mathcal{E}')$ is to deviate to cluster $i'-1$. Since this deviation was not beneficial in $\mathcal{E}$, we must have either $i = i+1$ or $i' = i$. The first inequality implies that this deviation is exactly the reverse of the deviation which created $\mathcal{E}'$ from $\mathcal{E}$, hence, the only possibility is $i' = i$. Since cluster $i-1$ is the same in both $\mathcal{E}$ and $\mathcal{E}'$, the deviation of $\mathsf{left}_{i'}(\mathcal{E}')$ can be beneficial only if the centroid of cluster i in $\mathcal{E}'$ is to the right of its location in $\mathcal{E}$, which contradicts the way $\mathcal{E}'$ is defined.

— If $i'' > i'$, then let v be the new centroid of cluster i'' following the deviation. Clearly v is an old node of cluster i'', and thus, $\mathsf{left}_{i'}(\mathcal{E}') \leq \mathsf{right}_{i'}(\mathcal{E}') < v$, which contradicts our assumption that deviating to cluster i'' decreases the cost of $\mathsf{left}_{i'}(\mathcal{E}')$. □

We are now ready to prove Theorem 4.1.

PROOF OF THEOREM 4.1. Let $\mathcal{S}$ be the set of half Nash equilibria of a center|line game. By Observation 4.3 $\mathcal{S}$ is non-empty. On the other hand, $\mathcal{S}$ is finite since there is only a finite number of possible clustering configurations. By Lemma 4.4, any half Nash equilibrium in $\mathcal{S}$ minimizing ψ (among the equilibria of S) is a Nash equilibrium. □

Having proved that a Nash equilibrium always exists for a line metric, we consider the price of anarchy of such metrics.

PROPOSITION 4.5. *The price of anarchy of the k-center clustering game is unbounded, even for* center|line*.*

PROOF. The proof is identical to the proof of Proposition 3.7. □

4.2. The Tree Metric Case

We suspect that a Nash equilibrium always exists also in tree metrics, and best-response dynamics are guaranteed to reach it. However, we manage to prove this only in case the centroid of a cluster is allowed to be placed at any location along an edge (rather than only at a node). Formally, every edge (u, v) and $\lambda \in [0, 1]$ represent a possible location for a centroid. The distance of this centroid from a node w is $\min\{d(u, w) + \lambda \cdot d(u, v), d(v, w) + (1 - \lambda) \cdot d(u, v)\}$. The centroid of the cluster is placed at the location minimizing the maximum distance from a node of the cluster. Note that the centroid could be located at a node which does not belong to its cluster, however, such a configuration may be stable only if this node is also the centroid of its own cluster.

LEMMA 4.6. *In* center|tree|edges*, the centroid of a cluster is always located in the middle of its diameter (since we are dealing with a tree metric, all diameters share their middle point).*

PROOF. Let u_1 and u_2 be the end points of a diameter of cluster C, and let R be the radius of cluster C, *i.e.*, half the diameter. We denote by p the location of the centroid of C. First consider the case where p is located in the middle of the diameter of C, and consider any node $w \in C$. If the distance of w from p is larger than R, then the distance of w from at least one of the nodes u_1 or u_2 must be larger than $2R$, contradicting the definition of R as radius. Thus, no node is at distance of more than R from p. Therefore, it follows that placing p in the middle of the diameter gives a maximum distance of R between a node of C and the centroid.

Consider any other location p for the centroid. If p is not on the diameter, let p' be the closest point to p on the diameter. Since we are working with a tree metric, the distances of u_1 and u_2 from p are $d(u_1, p') + d(p, p')$ and $d(u_2, p') + d(p, p')$, respectively. Note that

$$\max\{d(u_1,p') + d(p,p'), d(u_2,p') + d(p,p')\} \geq \frac{(d(u_1,p') + d(p,p')) + (d(u_2,p') + d(p,p'))}{2} = R + d(p,p') \ .$$

Thus, there is a node whose distance from p is larger than R. If p is located on the diameter, but not in its middle, then again one of the nodes u_1 or u_2 must be at a distance of more than R from p (since $d(u_1, u_2) = 2R$).

Hence, when p is located at the middle of the diameter, no node has a distance of more than R from it, whereas placing p anywhere else guarantees the existence of a node in C with a larger distance from p. Thus, the centroid p must be in the middle of the diameter. □

The following theorem is proved using the characterization of the centroid given by Lemma 4.6.

THEOREM 4.7. *In* center|tree|edges *a Nash equilibrium always exists, and it is guaranteed to be reached by best-response dynamics. Moreover, the price of stability of* center|tree|edges *is* 1.

PROOF. Let us begin by proving the first part of the theorem. Let $\mathcal{E}$ be a clustering configuration, and let $d(v, \mathcal{E})$ be the cost of node v in this configuration (*i.e.*, the distance of v from the centroid of its cluster). Let $v_1^{\mathcal{E}}, \ldots, v_n^{\mathcal{E}}$ be the nodes of the graph sorted by non-increasing order according to their distance from the centroids of their cluster under the configuration $\mathcal{E}$, (that is a non-increasing order of their cost). We use as a potential function the distance vector $D(\mathcal{E})$, which is a vector of the costs of all nodes under this profile, sorted in non-increasing order. Formally, the distance vector is defined as:

$$D(\mathcal{E}) = (d(v_1^{\mathcal{E}}, \mathcal{E}), d(v_2^{\mathcal{E}}, \mathcal{E}), \ldots, d(v_n^{\mathcal{E}}, \mathcal{E})) \ .$$

We show that each best-response move must reduce the lexicographic value of the distance vector. Consider the clustering configuration $\mathcal{E}'$ resulting after a best-response move made by node $v = v_i^{\mathcal{E}}$ that deviated from cluster C_1 to cluster C_2. Since v deviated, its cost was reduced, and thus $d(v, \mathcal{E}) > d(v, \mathcal{E}')$. Let u_1 be the node with the maximum cost in $C_1 - \{v\}$ (assuming the centroid was not yet changed). Following the deviation of v, the centroid of C_1 will only move if it can find a location in which all nodes of $C_1 - \{v\}$ will have a distance lower than $d(u_1, \mathcal{E})$ (otherwise, the cost of the cluster will not be reduced).

Consider now a node $u_2 \in C_2$. We show that either $d(u_2, \mathcal{E}) = d(u_2, \mathcal{E}')$ or $d(u_2, \mathcal{E}') \leq d(v, \mathcal{E}')$. If $d(u_2, \mathcal{E}) \neq d(u_2, \mathcal{E}')$, then we know that the centroid of C_2 moved due to the addition of v, hence, by Lemma 4.6, $C_2 \cup \{v\}$ has a diameter which is not a diameter in C_2. This can only be the case if v is an end point of the new diameter. The new location

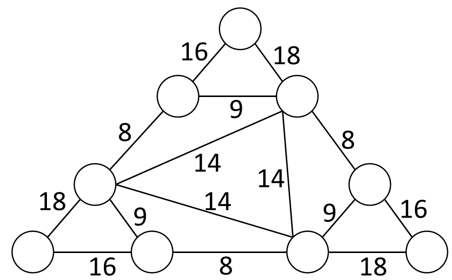

Fig. 5. Graph with no Nash equilibria (assuming two clusters).

of the centroid is the middle of the new diameter of $C_2 \cup \{v\}$, and thus v is one of the nodes with the largest distance from the centroid, *i.e.*, $d(u_2, \mathcal{E}') \leq d(v, \mathcal{E}')$.

We conclude the first part of the theorem as follows. If $d(u_1, \mathcal{E}) = d(u_1, \mathcal{E}')$, then every node whose cost changed in $\mathcal{E}'$ has now a cost of at most $d(v, \mathcal{E}')$, and we know that at least one node whose cost has changed (v itself) had a larger cost in $\mathcal{E}$. If $d(u_1, \mathcal{E}) > d(u_1, \mathcal{E}')$, then we know that every node whose cost changed in $\mathcal{E}'$ has cost of less than $\max\{d(u_1, \mathcal{E}), d(v, \mathcal{E})\}$, and at least one node (either u_1 or v) had this cost before. In both cases, the vector $D(\mathcal{E}')$ is lexicographically smaller than $D(\mathcal{E})$. Thus, each best-response move reduces the lexicographic value of the distance vector. As the number of possible configurations is finite, best-response dynamics are guaranteed to converge to a Nash equilibrium.

The second part of the theorem follows immediately from the first part since the cost of an instance can only decrease when our potential function decreases (since the cost of the instance is simply the first entry of the distance vector). □

PROPOSITION 4.8. *When centroids are allowed to be on edges, the price of anarchy of the k-center clustering game remains unbounded, even for* center|line|edges.

PROOF. Consider a line with six nodes: four nodes at locations 0, M, $M+\epsilon$, $2M$, and two nodes at locations L, $L+\epsilon$ for some large $L \gg M$. The clustering configuration with 4 clusters $\{0, M\}$, $\{M+\epsilon, 2M\}$, $\{L\}$, $\{L+\epsilon\}$ is a Nash equilibrium, as no node wishes to deviate. The cost of this configuration is $M/2$. On the other hand, the optimal solution has the following clusters: $\{0\}$, $\{M, M+\epsilon\}$, $\{2M\}$, $\{L, L+\epsilon\}$, yielding a cost of $\epsilon/2$. The price of anarchy of this instance is thus unbounded. □

4.3. The General Metric Case

Similarly to the k-median model, a Nash equilibrium might not exist for the case of a general metric.

THEOREM 4.9. *There exists an instace of* center|general *with 9 nodes having no Nash equilibria for two clusters.*

Figure 5 displays a graph representing an instance having the properties guaranteed by Theorem 4.9. The numbers on the edges of the graph represent the distances. The proof that no Nash equilibrium exists for this graph is established using a case analysis using the same method explained for the graph in Figure 1. The proof itself is long and technical, and is, therefore, deferred to Appendix A.

One can try, again, the trick that worked for tree metrics, *i.e.*, allow the centroid of a cluster to be located at any location along an edge of the graph. Unfortunately, this generalization fails to guarantee the existence of a Nash equilibrium in this case. The graph in Figure 5 represents a counterexample for this case as well. The proof consists, again, of a case analysis. Before doing the case analysis, we need the next lemma. Lemma 4.10 shows that it is sufficient to consider only a finite set of cases because we can assume the centroids of every configuration are located at half integral locations along edges.

LEMMA 4.10. *In any configuration of an instance of* center|general *having integral edge distances, centroids must be placed in half integral locations on edges.*

PROOF. Given an instance of the k-center clustering game, consider a cluster C with centroid $c(C)$ located on some edge (u, v). Let α be the fractional part of the distance between $c(C)$ and u, and assume w.l.o.g that $\alpha \in [0, 1/2]$. It can easily be noted that for every $w \in C$, its distance from $c(C)$ is given by one of two formulas: $\delta_w + \alpha$ or $\delta_w + (1-\alpha)$, where δ_w is w-dependent integral constant. Assume by way of contradiction that α is not half integral. Let $\mathcal{D}$ be the set of nodes in C with the maximum distance from $c(C)$. We consider two possible cases:

— The distances of all nodes in $\mathcal{D}$ from the centroid are given by formulas of the same type from the two options mentioned above. In this case, all nodes in $\mathcal{D}$ will benefit from either increasing or decreasing α. This will decrease the maximum distance of a node in C from $c(C)$, contradicting the definition of a centroid. Notice that since we assumed α is not half-integral, it is possible to both increase and decrease it without getting outside of the range $[0, 1/2]$.
— There are two nodes v_1 and v_2 in $\mathcal{D}$ whose distances from $c(C)$ are given by $\delta_{v_1} + \alpha$ and $\delta_{v_2} + (1 - \alpha)$ respectively. Since both nodes belong to $\mathcal{D}$ the two formulas yield the same distance, implying:

$$\delta_{v_1} + \alpha = \delta_{v_2} + (1 - \alpha) \Rightarrow \alpha = \frac{\delta_{v_2} + 1 - \delta_{v_1}}{2} .$$

Hence, α is half integral (since δ_{v_1} and δ_{v_2} are both integral). □

To verify that Figure 5 allows no weak Nash equilibria with two clusters, we do the following:

(1) Enumerate every partition of the nodes into two clusters.
(2) Enumerate for each partition all the optimal centroid locations which are half integral locations on edges.
(3) Show that no choice of partition and and two centroids induces a Nash equilibrium.

Notice that the last step can be performed due to the second part of Lemma 4.10. Due to the huge number of cases to check, we do the check using a computer program.[5] The next theorem follows by checking all the cases.

THEOREM 4.11. *There exists an instace of* center|general|edges *with* 9 *nodes having no Nash equilibria for two clusters.*

In order to guarantee the existence of a Nash equilibrium, we add to the game a rule penalizing a node whose move to a cluster C changes the location of the centroid of C. That is, the cost of a node u performing an improvement move will consist of two values:

[5]This check can be done by the program whose source code is given in http://theory.epfl.ch/moranfe/Resources/Journals/ACM%20Transactions%20on%20Parallel%20Computing%202014/NoNashValidation.vb. Notice that this source code checks, in fact, whether there exists a configuration having a somewhat weaker property: namely, can the clusters' centroids be positioned in such a way that no node will wish to deviate, assuming that following a deviation of a node, the centroid of the cluster to which the node deviated will move to the half integral position furthest away from the deviating node (among the positions minimizing the cluster's cost).

— The distance from the updated centroid of the new cluster, $d(u, c(C^{+u}))$.
— A cost equal to the distance between the original and new centroids of C, *i.e.*, $d(c(C), c(C^{+u}))$.

THEOREM 4.12. *An instance of* center|general|penalties *always has a Nash equilibrium solution, and such an equilibrium is guaranteed to be reached by best-response dynamics. Moreover, the price of stability of any such instance is* 1.

PROOF. Consider the distance vector defined as in the proof of Theorem 4.7. We prove the theorem by using the distance vector as a potential function, showing that every best-response move yields a distance vector with lower lexicographical order. Let $\mathcal{E}$ be a strategy configuration, and let $\mathcal{E}'$ be the strategy configuration resulting from $\mathcal{E}$ following a node v making a best-response move from cluster C_1 to cluster C_2. Let r be the distance by which the center of C_2 moved following the addition of v to C_2. Due to the choice of the penalties, it holds that:

$$d(v, \mathcal{E}) > r + d(v, \mathcal{E}') \ ,$$

(where $d(v, \mathcal{E})$ is the cost of v under configuration $\mathcal{E}$). We bound the maximum cost paid by a node in C_2. If the centroid of C_2 does not move when v joins it, then the costs of C_2's other nodes do not change following the deviation. Otherwise, consider the configuration after v joins C_2, but before the center $c(C_2)$ moves. Clearly, changing the center of C_2 cannot reduce the maximal distance of the centroid from any node in C_2 other than v, otherwise, it would have contradict the definition of $c(C_2)$ as centroid. Since the centroid did move after v joins C_2, it follows that v is the node furthest away from $c(C_2)$. Before the centroid moved, the distance of v from $c(C_2)$ was at most $r + d(v, \mathcal{E}')$, hence, this is an upper bound on the distance from the centroid of any node in C_2. The centroid moves to a location that minimizes the maximum distance to a node in C_2, thus, it also holds that after the move of the centroid there is no node in C_2^{+v} with a cost of more than $r + d(v, \mathcal{E}') < d(v, \mathcal{E})$. Combining the two cases (when C_2's centroid moves and when it does not), we get that no node of C_2 whose distance changed following v's deviation has a new cost of $d(v, \mathcal{E})$ or more.

If the center of C_1 does not move following v's best response, then the distance vector lexicographically decreases, as all nodes whose cost changed following v's move have a new cost of less than $d(v, \mathcal{E})$. If the center of C_1 moves following v's best-response move, then let u be the node of maximum cost in C_1. Since the centroid of C_1 moved, it holds that after the move there is no node in C_1 with distance of $d(u, \mathcal{E})$ or more. In addition, we know that all nodes in C_2^{+v} whose cost changed have a distance of less than $d(v, \mathcal{E}) \leq d(u, \mathcal{E})$. Therefore, every node whose distance changed has now a distance of less than $d(u, \mathcal{E})$, including u itself, and the distance vector lexicographically decreases. □

5. CORRELATION CLUSTERING

In this section we consider the clustering game in the correlation clustering model. We investigate both the minimization and maximization variants. In the minimization variant, a clustering configuration of the correlation clustering game is a Nash equilibrium if no user can unilaterally reduce its cost $\sum_{u \in C_v} w_u \cdot d(u, v) + \sum_{u \notin C_v} w_u \cdot (1 - d(u, v))$ by choosing a different cluster (respectively, in the maximization variant, a user cannot increase its profit $\sum_{u \in C_v} w_u \cdot (1 - d(u, v)) + \sum_{u \notin C_v} w_u \cdot d(u, v)$). For ease of notation, given a cluster C, we denote the total weight of its nodes by $w(C)$. We denote by V the set of elements, and by E the set of all pairs of (different) elements (E is the set of edges in a complete graph having V as the set of nodes).

The following lemma shows that the two variants are closely related.

LEMMA 5.1. *A configuration of the game is a Nash equilibrium for the minimization variant if and only if it is a Nash equilibrium for the maximization variant.*

PROOF. Given a clustering configuration, for each node v, the cost paid by v in the minimization variant plus the value gained by v in the maximization variant, for a given configuration, is $\sum_{u \in V} w_v$. Hence, a change of strategy strictly decreases the cost paid by v in the minimization variant if and only if it strictly increases the value gained by v in the maximization variant. □

Correlation clustering is a hedonic game and it is closely related to the class of additively-separable hedonic games [Bogomolnaia and Jackson 2002]. An additively-separable hedonic game consists of a digraph with arc weights. Each node is a player, and her strategy is which cluster to join. The utility of a player is the total weight of the arcs going from her to the other nodes in her cluster. The next observation shows that for every correlation clustering game, there exists an additively-separable hedonic game with the same set of Nash equilibria. However, the reverse is not true, *i.e.*, additively-separable hedonic games strictly generalize correlation clustering games. The proof of the observation is deferred to Appendix B.

OBSERVATION 5.2. *Every correlation clustering game can be converted into an additively-separable hedonic game with the same set of players and the same set of Nash equilibria. The reverse is not true.*

Another interesting subclass of additively-separable hedonic games are symmetric additively-separable hedonic games [Bogomolnaia and Jackson 2002]. A potential function argument shows that every symmetric additively-separable hedonic game has a Nash equilibrium, but finding it is PLS-hard [Gairing and Savani 2010]. Although these results do not extend immediately to correlation clustering games, their techniques can be used for this type of games as well. The next two theorems follow from the use of these techniques, and their proofs are also deferred to Appendix B.

THEOREM 5.3. *There always exists a Nash equilibrium for the correlation clustering game. Moreover, best-response dynamics of this game always converge to a Nash equilibrium.*

THEOREM 5.4. *The problem of computing a Nash equilibrium in the correlation clustering game is PLS-Complete.*

In the rest of this section we focus on the prices of stability and anarchy of correlation clustering games. Unlike the set of Nash equilibria, these values are different for each of the game classes mentioned above, and therefore, we cannot use the relations between our model and these classes to derive results in this context. We begin with several observations on the price of stability.

LEMMA 5.5. *In the special case where all elements have an identical weight w, the price of stability is 1.*

PROOF. In the case of identical weights, the potential function Φ from the proof of Theorem 5.3 can be rewritten as

$$\Phi(\mathcal{F}) = w^2 \sum_{v \in V} \left(\sum_{u \in C_v} d(u, v) + \sum_{u \notin C_v} (1 - d(v, u)) \right) .$$

Note that given a clustering configuration $\mathcal{F}$, the value of the potential function is identical to the value of the social objective function up to a multiplicative factor of w. This implies that any best-response move performed by a player will strictly decrease

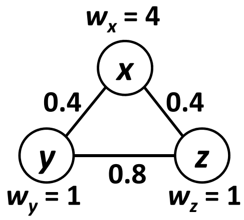

Variant	Nash Equilibrium	Optimal Configuration	Price of Stability
Minimum	5.6	5.4	≈ 1.037
Maximum	6.4	6.6	1.03125

Fig. 6. Graph demonstrating a price of stability greater than 1. The single Nash equilibrium puts all nodes in a single cluster, while the optimal solutions cluster x with either y or z (but not with both).

the social objective value. Thus, an optimal clustering configuration is also a Nash equilibrium, as no player can further reduce its cost. The proof for the maximization variant follows directly from Lemma 5.1. □

Remark: When all the distances are either 0 or 1, Lemma 5.5 is equivalent to Lemma 8.3 of Hoefer [2007].

In the case of arbitrary weights of elements, the price of stability differs between the minimization and maximization variants and can be strictly larger than 1, as shown in the following example. Consider the graph depicted in Figure 6. The similarity between any two nodes appears on the edge between them, and the weight of a node appears next to it. It is easy to verify that a Nash equilibrium in this graph only occurs when all nodes share a single cluster. However, in the optimal configurations x shares a cluster with either y or z, while the third node is in a different cluster. The table in Figure 6 presents the price of stability demonstrated by this example for both variants.

We now turn to analyze the price of anarchy. Despite the close relationship between the minimization and maximization variants of the game, the results obtained for their price of anarchy are quite different, as shown in the next sections. Before turning to analyzing each variant, we present general properties of Nash equilibria which are later used for establishing the bounds on the price of anarchy for both game variants.

Properties of Nash Equilibria. We prove two lemmata bounding distances in Nash equilibria.

LEMMA 5.6. *Consider two nodes u and v. If there exists a Nash equilibrium where u and v share a common cluster C, then:*

$$d(u,v) \leq 1 - \frac{w_u + w_v}{2w(C)} \quad .$$

PROOF. In a Nash equilibrium, u has no incentive to leave its cluster C and initiate a new one, thus:

$$\sum_{\substack{z \in C \\ z \neq u}} w_z d(u,z) \leq \sum_{\substack{z \in C \\ z \neq u}} w_z (1 - d(u,z)) \quad . \tag{2}$$

Rearranging terms, we get $2 \cdot \sum_{z \in C, z \neq u} w_z d(u,z) \leq w(C) - w_u$. A similar inequality holds for v. Adding up the inequalities for u and v and dividing by 2, we get:

$$\sum_{\substack{z \in C \\ z \neq u}} w_z d(u,z) + \sum_{\substack{z \in C \\ z \neq v}} w_z d(v,z) \leq w(C) - \frac{w_u + w_v}{2} \quad . \tag{3}$$

Combining the triangle inequality $d(u,z)+d(v,z) \geq d(u,v)$ with Inequality (3), we get:

$$w_v d(u,v) + w_u d(u,v) + \sum_{\substack{z \in C \\ z \neq u,v}} w_z d(u,v) \leq w(C) - \frac{w_u + w_v}{2} \quad . \tag{4}$$

Note that the left hand side of the last inequality is in fact $w(C) \cdot d(u,v)$. Dividing by $w(C)$ gives:

$$d(u,v) \leq \frac{w(C) - (w_u + w_v)/2}{w(C)} = 1 - \frac{w_u + w_v}{2w(C)} \quad . \quad \square$$

LEMMA 5.7. *Consider two nodes u and v. If there exists a Nash equilibrium where u and v belong to two different clusters C_u and C_v, then*

$$d(u,v) \geq \frac{w_u + w_v}{2(w(C_u) + w(C_v))} \quad .$$

PROOF. In a Nash equilibrium, u has no incentive to move to C_v. Thus,

$$\sum_{\substack{z \in C_u \\ z \neq u}} w_z d(u,z) + \sum_{z \in C_v} w_z (1 - d(u,z)) \leq \sum_{\substack{z \in C_u \\ z \neq u}} w_z (1 - d(u,z)) + \sum_{z \in C_v} w_z d(u,z) \quad .$$

Rearranging terms, we get:

$$2 \sum_{\substack{z \in C_u \\ z \neq u}} w_z d(u,z) + w(C_v) \leq w(C_u) - w_u + 2 \sum_{z \in C_v} w_z d(u,z) \quad .$$

A similar inequality holds for v. Adding up the inequalities for u and v and dividing by 2, we get:

$$\frac{w_u + w_v}{2} + \sum_{\substack{z \in C_u \\ z \neq u}} w_z d(u,z) + \sum_{\substack{z \in C_v \\ z \neq v}} w_z d(v,z) \leq \sum_{z \in C_v} w_z d(u,z) + \sum_{z \in C_u} w_z d(v,z) \quad .$$

Rearranging and applying triangle inequalities $d(v,z) - d(u,z) \leq d(u,v)$ and $d(u,z) - d(v,z) \leq d(u,v)$:

$$\frac{w_u + w_v}{2} \leq w_u d(u,v) + w_v d(u,v) + \sum_{\substack{z \in C_v \\ z \neq v}} w_z d(u,v) + \sum_{\substack{z \in C_u \\ z \neq u}} w_z d(u,v) \quad . \tag{5}$$

Note that the right hand side of Inequality (5) is exactly $d(u,v)$ times the total weight of $C_u \cup C_v$. Thus, Inequality (5) is equivalent to

$$d(u,v) \geq \frac{w_u + w_v}{2(w(C_u) + w(C_v))} \quad . \quad \square$$

5.1. Price of Anarchy – Minimization Variant

We present an upper bound of $O(n^2)$ and a lower bound of $n-1$ on the price of anarchy of the minimization variant of the correlation clustering game. We begin with the special case of identical node weights, for which we prove tight bounds, and then proceed to arbitrary node weights. The next definition is used in the sequel.

Definition 5.8. Given a clustering configuration, an edge e is of one of two types: either it is an *internal* edge within a single cluster, or it is an *external* edge between two different clusters.

THEOREM 5.9. *If all nodes have the same weight w, the price of anarchy is at most $n-1$.*

PROOF. Without loss of generality, assume $w = 1$. Let $\mathcal{E}$ be a Nash equilibrium of the game, and let $\mathcal{O}$ be an optimal solution. In order to evaluate the contribution of an edge $e = (u, v)$ (of distance $d(u, v)$) to the total cost of $\mathcal{E}$ and $\mathcal{O}$, there are four cases to be considered.

— e is an internal edge in $\mathcal{E}$ and $\mathcal{O}$. Then, e contributes a cost of $2d(u, v)$ to both $\mathcal{E}$ and $\mathcal{O}$.
— e is an external edge in $\mathcal{E}$ and $\mathcal{O}$. Then, e contributes a cost of $2(1 - d(u, v))$ to both $\mathcal{E}$ and $\mathcal{O}$.
— e is an internal edge within cluster C in $\mathcal{E}$, but is an external edge in $\mathcal{O}$. Since e is an edge within C in a Nash equilibrium, Lemma 5.6 implies that $d(u, v) \leq 1-2/2w(C) \leq 1-1/n$. Therefore, the cost contribution of e to $\mathcal{E}$ is $2d(u, v) \leq 2(1-1/n)$. On the other hand, the cost contribution of e to $\mathcal{O}$ is $2(1 - d(u, v)) \geq 2/n$. Thus, the ratio between the cost contribution of e to $\mathcal{E}$ and its contribution to $\mathcal{O}$ is at most $\frac{2(1-1/n)}{2/n} = n - 1$.
— e is an external edge between two clusters C_1 and C_2 in $\mathcal{E}$, but is an internal edge in $\mathcal{O}$. Since e is an edge between clusters in a Nash equilibrium, Lemma 5.7 implies that $d(u, v) \geq 2/(2w(C_1) + 2w(C_2)) \geq 1/n$. Therefore the cost contribution of e to $\mathcal{E}$ is $2(1 - d(u, v)) \leq 2(1 - 1/n)$. On the other hand, the cost contribution of e to $\mathcal{O}$ is $2d(u, v) \geq 2/n$. Thus, the ratio between the cost contribution of e to $\mathcal{E}$ and its contribution to $\mathcal{O}$ is again at most $n - 1$.

It follows that the cost contribution of all edges to the cost of $\mathcal{E}$ is at most $n - 1$ times their contribution to the cost of $\mathcal{O}$, completing the proof. □

For general weights we are only able to prove a weaker result.

THEOREM 5.10. *The price of anarchy of the minimization variant of the correlation clustering game is $O(n^2)$.*

PROOF. Let $\mathcal{E}$ be a Nash equilibrium of the game, and let $\mathcal{O}$ be an optimal solution. Let E' be the set of edges that have different types in $\mathcal{E}$ and $\mathcal{O}$, *i.e.*, an edge e is in E' if it is an internal edge in $\mathcal{E}$ and an external edge in $\mathcal{O}$, or vice versa. An edge $e \in E - E'$ has an equal cost contribution to both the cost of $\mathcal{E}$ and $\mathcal{O}$, and thus, we can ignore its contributions to the costs of $\mathcal{E}$ and $\mathcal{O}$ when bounding from above the price of anarchy.

A cluster of $\mathcal{E}$ is said to be "correct" if there is an identical cluster in $\mathcal{O}$ (that is, a cluster including the same nodes), and is "incorrect" otherwise. Let V' be the set of nodes belonging to incorrect clusters of $\mathcal{E}$. Note that if an edge of E' is adjacent to a node u, then u is part of an incorrect cluster ($u \in V'$). For each edge $e = (u, v) \in E'$, there are two cases.

— If e is an internal edge within cluster C in $\mathcal{E}$, then its cost contribution to $\mathcal{E}$ is $(w_u + w_v) \cdot d(u, v)$, which following Lemma 5.6 is at most:

$$(w_u + w_v) \cdot \left(1 - \frac{w_u + w_v}{2w(C)}\right) \leq (w_u + w_v) - \frac{(w_u + w_v)^2}{2w(V')} \quad ,$$

where $w(V')$ denotes the total weight of nodes in V'.
— If e is an external edge between two clusters C_1 and C_2 in $\mathcal{E}$, then its cost contribution to $\mathcal{E}$ is $(w_u + w_v) \cdot (1 - d(u, v))$, which following Lemma 5.7 is at most:

$$(w_u + w_v) \cdot \left(1 - \frac{w_u + w_v}{2w(C_1) + 2w(C_2)}\right) \leq (w_u + w_v) - \frac{(w_u + w_v)^2}{2w(V')}.$$

In both cases, the contribution of e to the cost of $\mathcal{O}$ is exactly $w_u + w_v$ minus its contribution to the cost of $\mathcal{E}$. On the other hand an edge $e \in E \setminus E'$ contributes the same to the cost of $\mathcal{E}$ and $\mathcal{O}$. Let D be the total contribution of these edges to the costs of $\mathcal{E}$ and $\mathcal{O}$. Then, the price of anarchy is at most:

$$\frac{D + \sum_{uv \in E'} (w_u + w_v) - \frac{(w_u+w_v)^2}{2w(V')}}{D + \sum_{uv \in E'} \frac{(w_u+w_v)^2}{2w(V')}} \leq \frac{\sum_{uv \in E'} (w_u + w_v) - \frac{(w_u+w_v)^2}{2w(V')}}{\sum_{uv \in E'} \frac{(w_u+w_v)^2}{2w(V')}}$$
$$= 2w(V') \cdot \frac{\sum_{uv \in E'} (w_u + w_v)}{\sum_{uv \in E'} (w_u + w_v)^2} - 1$$
$$= 2w(V') \cdot |E'| \cdot \frac{\sum_{uv \in E'} (w_u + w_v)}{\sum_{uv \in E'} (w_u + w_v)^2 \cdot \sum_{uv \in E'} 1} - 1$$
$$\leq 2w(V') \cdot |E'| \cdot \frac{\sum_{uv \in E'} (w_u + w_v)}{\left(\sum_{uv \in E'} (w_u + w_v)\right)^2} - 1$$
$$= \frac{2w(V') \cdot |E'|}{\sum_{uv \in E'} (w_u + w_v)} - 1 \ ,$$

where the first inequality follows since $w_u + w_v \geq (w_u + w_v)^2 / w(V)$ for every edge e, and the second inequality follows from the Cauchy-Schwarz inequality. Note that $\sum_{uv \in E'} (w_u + w_v)$ (in the denominator) contains a contribution from each node in V', and hence is at least $w(V')$. Thus, we obtain a bound of $2|E'| - 1 = O(n^2)$ on the price of anarchy. □

Theorem 5.12 establishes a lower bound of $n - 1$ (for even n) on the price of anarchy of the minimization variant, which holds even for he special case of identical node weights. Note that it matches the upper bound proved above for this special case.

LEMMA 5.11. *Consider an instance of the correlation clustering game where all nodes have a weight of* 1 *and each node belongs to exactly one of* r *equal size sets* $A_1, A_2, \ldots, A_r$*, and the distance function between nodes is defined as following:*

$$d(u, v) = \begin{cases} 0 & \text{if } u \text{ and } v \text{ belong to a single subset } A_i \ , \\ \alpha & \text{otherwise} \ . \end{cases}$$

Then, the clustering where the nodes of each set form their own cluster is a Nash equilibrium if $2n\alpha \geq r$.

PROOF. By Lemma 5.1, we can assume, without loss of generality, that we are dealing with the minimization variant. Due to symmetry, it suffices to show that a node $u \in A_i$ does not have an incentive to deviate. The cost of u in the current clustering is

$$(n/r - 1) \cdot 0 + (n - n/r) \cdot (1 - \alpha) = n(1 - 1/r)(1 - \alpha) \ .$$

Observe that u has two deviation options. The first one is to move to a different cluster A_j for some $i \neq j$, which will change the cost of u to

$$(n/r - 1) \cdot 1 + (n - 2n/r) \cdot (1 - \alpha) + n/r \cdot \alpha = n(1 - 1/r)(1 - \alpha) + 2n\alpha/r - 1 \ ,$$

which is no less than the current profit given our assumption that $2n\alpha \geq r$. The other deviation option of u is to form a new cluster, which will change its cost to

$$(n/r - 1) \cdot 1 + (n - n/r) \cdot (1 - \alpha) = n(1 - 1/r)(1 - \alpha) + n/r - 1 \ ,$$

which is no less than the current profit since $r \leq n$. □

THEOREM 5.12. *For even n, the price of anarchy of the minimization variant is at least $n-1$. This holds even when all nodes have the same weight and the metric is a line metric.*

PROOF. We show an instance in which the price of anarchy is at least $n-1$. Assume the weight of all nodes is 1. Let A and B be two disjoint sets of $n/2$ nodes. The points of A are located at location 0 on a line, while the points of B are located at location $1/n$ of the same line. Notice that these locations imply the following distances between the nodes:

$$d(u,v) = \begin{cases} 0 & u,v \in A \text{ or } u,v \in B \ , \\ 1/n & \text{otherwise} \ . \end{cases}$$

Hence, the above instance obeys the conditions of Lemma 5.11 with $r = 2$ and $\alpha = 1/n$. Since $2n\alpha = 2 = r$, this implies that the clustering in which the nodes of A form one cluster and the nodes of B form another one is a Nash equilibrium. The social cost of this clustering is $n(n-1)/2$. On the other hand, the cost of the configuration in which all nodes belong to a single cluster is:

$$n \cdot [(n/2 - 1) \cdot 0 + n/2 \cdot 1/n] = n/2 \ .$$

Thus, the price of anarchy of this game instance is at least $\frac{n(n-1)/2}{n/2} = n - 1$. □

5.2. Price of Anarchy – Maximization Variant

In this section we prove a bound of $\Theta(\sqrt{n})$ on the price of anarchy of the maximization variant of correlation clustering. For line metrics we provide an improved bound of $\Theta(n^{1/3})$. We also show how to get a significantly better price of anarchy via a slight modification of the "rules".

THEOREM 5.13. *The price of anarchy of the maximization variant is $O(\sqrt{n})$.*

The proof of Theorem 5.13 proceeds as follows. We first note that an upper bound on the maximum social welfare (total profit) is $(n-1) \cdot w(V)$. Then, we establish a lower bound of $\Omega(\sqrt{n}) \cdot w(V)$ on the total profit of any Nash equilibrium solution. The proof of the lower bound is split between Lemmas 5.14 and 5.15, which, together, complete the proof of the theorem.

LEMMA 5.14. *Let $\mathcal{E}$ be a Nash equilibrium in which all clusters have weight at most $w(V)/\sqrt{n}$. Then, the social welfare of $\mathcal{E}$ is $\Omega(\sqrt{n}) \cdot w(V)$.*

PROOF. First we need the following technical result. By the Cauchy-Schwarz inequality:

$$\begin{aligned} w^2(V) &= \left(\sum_{v\in V} w_v\right)^2 = \sum_{v\in V} w_v^2 + 2\sum_{uv\in E} w_u w_v \\ &= \sum_{v\in V} w_v^2 + \frac{n-2}{n-1} \cdot \sum_{uv\in E} w_u w_v + \frac{n}{n-1} \cdot \sum_{uv\in E} w_u w_v \\ &\leq \sum_{v\in V} w_v^2 + \frac{n-2}{2(n-1)} \cdot \left(\sqrt{(n-1) \cdot \sum_{v\in V} w_v^2}\right)^2 + \frac{n}{n-1} \cdot \sum_{uv\in E} w_u w_v \\ &= \frac{n}{2}\sum_{v\in V} w_v^2 + \frac{n}{n-1}\sum_{uv\in E} w_u w_v = \frac{n}{2(n-1)}\left((n-1)\sum_{v\in V} w_v^2 + 2\sum_{uv\in E} w_u w_v\right) \end{aligned}$$

$$= \frac{n}{2(n-1)} \sum_{uv \in E} (w_u + w_v)^2 \ .$$

Therefore, for any choice of weights for the nodes it holds that:

$$w^2(V) \leq \frac{n}{2(n-1)} \cdot \sum_{uv \in E} (w_u + w_v)^2 \ . \tag{6}$$

Consider an edge (u, v), and assume it is an external edge between two clusters C_u and C_v of $\mathcal{E}$. Following Lemma 5.7, it holds that:

$$d(u, v) \geq \frac{w_u + w_v}{2(w(C_u) + w(C_v))} \geq \frac{w_u + w_v}{4w(V)/\sqrt{n}} \ ,$$

where the last inequality follows from our assumed upper bound on the weight of the clusters. On the other hand, if (u, v) is an internal edge within cluster C, by Lemma 5.6 it holds that:

$$d(u, v) \leq 1 - \frac{w_u + w_v}{2w(C)} \leq 1 - \frac{w_u + w_v}{2w(V)/\sqrt{n}} \ ,$$

where the last inequality follows, again, from our assumed upper bound on the weight of the clusters. In both cases, the total profit gained by u and v from edge (u, v) is at least:

$$(w_u + w_v) \cdot \frac{w_u + w_v}{4w(V)/\sqrt{n}} = \frac{(w_u + w_v)^2}{4w(V)/\sqrt{n}} \ .$$

Summing up over all edges and using Inequality (6), we get a total profit of at least:

$$\sum_{uv \in E} \frac{(w_u + w_v)^2}{4w(V)/\sqrt{n}} \geq \frac{\sqrt{n}}{4w(V)} \cdot \frac{2(n-1)}{n} \cdot w^2(V) = \Omega(\sqrt{n}) \cdot w(V) \ . \quad \square$$

LEMMA 5.15. *Let $\mathcal{E}$ be a Nash equilibrium containing a cluster C having weight at least $\alpha \cdot w(V)$. Then, the social welfare of $\mathcal{E}$ is $\Omega(\alpha \cdot n) \cdot w(V)$. Hence, for $\alpha = 1/\sqrt{n}$, $\mathcal{E}$ must have a social welfare of $\Omega(\sqrt{n}) \cdot w(V)$.*

PROOF. Let u be a node outside cluster C. We consider two of the options that u has: either staying in the cluster it currently belongs to or joining C. Consider an edge $e = (u, v)$ for some $v \in C$. As u and v belong to different clusters in $\mathcal{E}$, the profit u gains from e is $w_v \cdot d(u, v)$. In case u deviates and joins C, it gains at least $w_v \cdot (1 - d(u, v))$ from e. Since the same holds for every edge between u and a node $v \in C$, we get that the total profit of u in $\mathcal{E}$ is at least $\sum_{v \in C} w_v \cdot d(u, v)$, and by deviating it could achieve a profit of at least $\sum_{v \in C} w_v \cdot (1 - d(u, v))$. Since $\mathcal{E}$ is in equilibrium and u has no incentive to deviate, it must be that its total profit in $\mathcal{E}$ is at least the maximum between the two expressions, *i.e.*,

$$\begin{aligned} & \max \left\{ \sum_{v \in C} w_v \cdot d(u, v), \sum_{v \in C} w_v \cdot (1 - d(u, v)) \right\} \\ \geq & \frac{\sum_{v \in C} w_v \cdot d(u, v) + \sum_{v \in C} w_v \cdot (1 - d(u, v))}{2} \\ = & \frac{\sum_{v \in C} w_v}{2} = \frac{w(C)}{2} \geq \frac{\alpha \cdot w(V)}{2} \ , \end{aligned}$$

where the last inequality follows from our assumed lower bound on the weight of C. Consider now a node u in cluster C. Two of the options that u has are staying in C, or deviating to another cluster (either an existing cluster or a new one). Similarly to the previous case, we get that the profit of u must be at least $(w(C) - w_u)/2$. Summing up the profit of all nodes in C, we get a value of at least:

$$\sum_{u \in C} \frac{w(C) - w_u}{2} = \frac{(|C|-1)w(C)}{2} \geq \frac{\alpha \cdot (|C|-1)w(V)}{2} \enspace ,$$

where the inequality follows, again, from our assumed lower bound on $w(C)$. Summing up the profits of nodes both outside of C and in C, we get that the social welfare of $\mathcal{E}$ is at least:

$$\frac{\alpha \cdot (|C|-1)}{2} + \frac{\alpha \cdot (n - |C|)w(V)}{2} = \Omega(\alpha \cdot n) \cdot w(V) \enspace . \quad \square$$

Next, we prove that Theorem 5.13 is tight even in a very restricted setting.

THEOREM 5.16. *The price of anarchy of the maximization variant is $\Omega(\sqrt{n})$. This bound holds even if all nodes have the same weight and the metric is a tree metric.*

PROOF. We show an instance of the correlation clustering game in which the price of anarchy is at least $\Omega(\sqrt{n})$. We assume the weight of all nodes is 1. Let $A_1, \ldots, A_{\sqrt{n}}$ be $\sqrt{n}$ subsets, each of size $\sqrt{n}$ (if $\sqrt{n}$ is not an integer, then we replace n by the largest perfect square n' smaller than n. To have exactly n agents in our instance, we introduce $n - n'$ additional nodes which are at distance 1 from all other nodes and form singleton clusters in all the configurations we consider). The distances between the nodes are defined as follows:

$$d(u, v) = \begin{cases} 0 & \text{if } u \text{ and } v \text{ belong to a single subset } A_i \enspace , \\ \sqrt{1/n} & \text{otherwise} \enspace . \end{cases}$$

Note that this distance function is a special case of a tree metric. Observe that the above instance obeys the conditions of Lemma 5.11 with $r = \alpha^{-1} = \sqrt{n}$. Since $2n\alpha = 2\sqrt{n} > r$, the configuration where the nodes of A_i form a distinct cluster for each $i \in \{1, \ldots, \sqrt{n}\}$ is a Nash equilibrium. Due to symmetry, the social welfare of this configuration is n times the profit of a single node, *i.e.*, $2n(\sqrt{n} - 1)$. On the other hand, the cost of the configuration in which all nodes belong to a single cluster is:

$$n \cdot \left[(\sqrt{n} - 1) \cdot 1 + (n - \sqrt{n}) \cdot (1 - \sqrt{1/n})\right] = n \cdot \left[n - \sqrt{n}\right] = n^{3/2} \cdot (\sqrt{n} - 1) \enspace .$$

Thus, the price of anarchy of this game instance is at least

$$\frac{n^{3/2} \cdot (\sqrt{n} - 1)}{2n(\sqrt{n} - 1)} = \sqrt{n}/2 = \Omega(\sqrt{n}) \enspace . \quad \square$$

Remark: The construction used by Theorem 5.16 is very close to a construction used in the proof of Theorem 8.1 of Hoefer [2007].

For the line metric it is possible to prove a somewhat better bound.

THEOREM 5.17. *The price of anarchy of the maximization variant in the case of a line metric is $O(n^{1/3})$.*

We note, again, that an upper bound on the maximum social welfare (total profit) is $(n - 1) \cdot w(V)$. Then, we establish a lower bound of $\Omega(n^{2/3}) \cdot w(V)$ on the total profit of any Nash equilibrium solution. For Nash equilibria containing a cluster of weight $n^{-1/3} \cdot w(V)$ or more this follows immediately from Lemma 5.15 by plugging $\alpha = n^{-1/3}$. The following lemma completes the proof for the case when no cluster has such weight.

LEMMA 5.18. *Let $\mathcal{E}$ be a Nash equilibrium in which all clusters have weight at most $n^{-1/3} \cdot w(V)$. Then, the social welfare of $\mathcal{E}$ is $\Omega(n^{2/3}) \cdot w(V)$.*

PROOF. Let V' be the set of nodes of weight at least $w(V)/(2n)$. Clearly, $w(V') \geq w(V)/2$. Let $\mathcal{C}$ be the set of clusters which include nodes of V'. Since no cluster has weight of more than $n^{-1/3} \cdot w(V)$, the number of clusters in $\mathcal{C}$ is at least:

$$\frac{w(V)/2}{n^{-1/3} \cdot w(V)} = \frac{n^{1/3}}{2} \ .$$

By Lemma 5.7, the distance between nodes $u, v \in V'$ not belonging to the same cluster is at least:

$$\frac{w_u + w_v}{2(w(C_u) + w(C_v))} \geq \frac{w(V)/(2n) + w(V)/(2n)}{2(n^{-1/3}w(V) + n^{-1/3}w(V))} = \frac{1/n}{4n^{-1/3}} = \frac{1}{4n^{2/3}} \ ,$$

where the inequality follows from our assumed upper bound on the weights of the clusters. Let C be a cluster and let $u \notin C$. We define the distance between u and C to be the minimum distance between u and a node of C. Consider now any node $u \in V$. Since the metric is a line metric, for every $k > 0$, at most $2k$ clusters of $\mathcal{C}$ (except for C_u) can be at a distance of less than $k \cdot n^{-2/3}/4$ from u, otherwise at least two clusters of $\mathcal{C}$ will have to be too close. Observe that the profit of u is minimized when the weight of all other clusters is as close to u as possible. Hence, the value gained by u is at least:

$$\begin{aligned}\sum_{C \in \mathcal{C} - C_u} w(C) \cdot d(u, C) &\geq \sum_{i=1}^{\frac{n^{1/3}/2-3}{2}} 2(n^{-1/3} \cdot w(V)) \cdot (i \cdot n^{-2/3}/4) \\ &= \frac{w(V)}{2n} \cdot \sum_{i=1}^{n^{1/3}/4-3/2} i \ = \Omega(n^{-1/3}) \cdot w(V) \ .\end{aligned}$$

The lemma now follows since there are n nodes in V (and each one of them has a profit of at least $\Omega(n^{-1/3}) \cdot w(V)$ in $\mathcal{E}$). $\square$

Next, we prove that Theorem 5.17 is tight even for unit node weights.

THEOREM 5.19. *The price of anarchy of the maximization variant in the case of a line metric is $\Omega(n^{1/3})$. Moreover, this bound is tight even if all nodes have unit weight.*

PROOF. Assume there are $n^{1/3}$ subsets $A_1, \ldots, A_{n^{1/3}}$ of nodes, each of size $n^{2/3}$. The nodes of A_i are located on the line in location $i \cdot n^{-2/3}$. Clearly, the diameter of the metric is at most 1.

Consider the following clustering. For each $i \in \{1, \ldots, n^{1/3}\}$, the nodes of A_i form a distinct cluster. We would like to show that this configuration is a Nash equilibrium. We show that for any node $u \in A_i$, u does not have an incentive to deviate under this configuration. Node u has two possible deviations available. The first option is to move to a cluster of the nodes of A_j for some $i \neq j$. This will change u's profit by:

$$\begin{aligned}&\left[(n^{2/3} - 1) \cdot 0 + n^{2/3} \cdot \left(1 - \frac{|j-i|}{n^{2/3}}\right)\right] - \left[(n^{2/3} - 1) \cdot 1 + n^{2/3} \cdot \frac{|j-i|}{n^{2/3}}\right] \\ &= n^{2/3} - |j-i| - n^{2/3} + 1 - |j-i| = 1 - 2|j-i| \leq -1 \ .\end{aligned}$$

The other option is to form a new cluster, which will change u's profit by $(n^{2/3} - 1) \cdot (0 - 1) = 1 - n^{2/3} \leq 0$. Thus, the above configuration is a Nash equilibrium. Let us upper bound the social welfare of this equilibrium. The nodes of one of the extreme subsets (*i.e.*, A_1 or $A_{n^{1/3}}$) are farthest away from all other clusters, hence, they have the largest profit. The profit of such a node is:

$$(n^{2/3}-1)\cdot 1+\sum_{i=1}^{n^{1/3}-1} n^{2/3}\cdot\frac{i}{n^{2/3}} = n^{2/3}-1+\sum_{i=1}^{n^{1/3}-1} i = n^{2/3}-1+\frac{n^{1/3}(n^{1/3}-1)}{2}$$
$$= \frac{3n^{2/3}-n^{1/3}-2}{2} \le \frac{3n^{2/3}}{2} \ .$$

Hence, the social welfare of the above clustering is at most $n\cdot 3n^{2/3}/2 = 3n^{5/3}/2$. On the other hand, consider the social welfare of a solution in which all nodes belong to a single cluster. Again a node of an extreme subset is farthest away from all other nodes, and therefore has the least profit. The profit of such a node is:

$$(n^{2/3}-1)\cdot 1+\sum_{i=1}^{n^{1/3}-1} n^{2/3}\cdot\left(1-\frac{i}{n^{2/3}}\right) = n^{2/3}-1+\sum_{i=1}^{n^{1/3}-1}(n^{2/3}-i)$$
$$= n^{2/3}-1+n^{2/3}(n^{1/3}-1)-\frac{n^{1/3}(n^{1/3}-1)}{2}$$
$$= \frac{2n-n^{2/3}+n^{1/3}-2}{2} \ge \frac{2n-n^{2/3}}{2} \ .$$

Therefore, the social welfare of the single cluster solution is at least $n(2n-n^{2/3})/2 = n^2-n^{5/3}/2$, and the price of anarchy is at least:

$$\frac{n^2-n^{5/3}/2}{3n^{5/3}/2} = 2n^{1/3}/3-1/3 = \Omega(n^{1/3}) \ . \quad \square$$

5.2.1. Bounding the Price of Anarchy. We suggest a method for bounding the price of anarchy (at the cost of making a slight modification to the "rules").

LEMMA 5.20. *If only k clusters are allowed (for $k \ge 2$), then the price of anarchy is at most k.*

PROOF. Let $\mathcal{E}$ be a Nash equilibrium satisfying the above restriction, and let v be a node. Let S_i denote the set of nodes from $V-\{v\}$ which belong to cluster i. Assume, w.l.o.g., that v belongs to cluster 1, and let $g(v)$ denote its profit. The profit v gets after deviating to cluster $j>1$ is

$$\sum_{u\in S_j} w_u d(u,v) + \sum_{\substack{1\le i\le k\\ i\ne j}}\sum_{u\in S_i} w_u(1-d(u,v)) \ .$$

Since v has no incentive to deviate from $\mathcal{E}$, the last expression must be no larger than $g(v)$.

$$g(v) \ge \sum_{u\in S_j} w_u d(u,v) + \sum_{\substack{1\le i\le k\\ i\ne j}}\sum_{u\in S_i} w_u(1-d(u,v)) \ .$$

Summing over all clusters $j > 1$, we get

$$(k-1) \cdot g(v) \geq \sum_{j=2}^{k} \sum_{u \in S_j} w_u d(u,v) + \sum_{j=2}^{k} \sum_{\substack{1 \leq i \leq k \\ i \neq j}} \sum_{u \in S_i} w_u (1 - d(u,v))$$
$$\geq \sum_{j=2}^{k} \sum_{u \in S_j} w_u d(u,v) + \sum_{u \in S_1} w_u (1 - d(u,v)) \ .$$

It can easily be seen that the right hand side of the last expression is $\sum_{u \in V - \{v\}} w_u - g(v)$. Using this observation, the previous inequality becomes

$$g(v) \geq \frac{\sum_{u \in V - \{v\}} w_u}{k} \ .$$

On the other hand, v cannot get a profit of more than $\sum_{u \in V - \{v\}} w_u$ under any configuration, implying that $g(v)$ is at least k^{-1} of the maximal profit v can have. Since this is true for every node v under $\mathcal{E}$, the social welfare of $\mathcal{E}$ is at least k^{-1} times the optimal social welfare. □

COROLLARY 5.21. *Consider the case where all nodes have identical weights and a best-response dynamics is executed in the following way. First the dynamics is allowed to converge to a Nash equilibrium under the restriction that only $k \geq 2$ clusters are allowed. Then, this restriction is lifted, and the dynamics is allowed to continue till it finds a true Nash equilibrium. Then, the resulting Nash equilibrium has a value of at least k^{-1} times the optimal social welfare.*

PROOF. By Lemma 5.20, the Nash equilibrium reached while the number of clusters is restricted to k has a social welfare of at least OPT/k, where OPT is the optimal social welfare. Moreover, the proof of Lemma 5.20 actually shows that the social welfare of this Nash equilibrium is also at least W/k, where W is the maximal social welfare possible for any configuration (with any number of clusters). Using ideas similar to the proof of Lemma 5.5, we get that the social welfare cannot decrease by best-response dynamics. Hence, the social welfare of the final Nash equilibrium is at least as large as the social welfare of the initial Nash equilibrium. □

6. CONCLUSION AND OPEN PROBLEMS

We considered three game theoretic models of clustering corresponding to well-known combinatorial clustering problems. Our models share the common property that the nodes to be clustered are the players of the game. We studied the existence of Nash equilibria under each model as well as the price of anarchy and price of stability whenever a Nash equilibrium is guaranteed to exist. Somewhat surprisingly, our models turn out to have very different answers for the above questions.

For the k-median model, we show that in general an instance might have no Nash equilibria, unless penalties are imposed on deviations. On the other hand, there always exists a Nash equilibrium in instances corresponding to tree metrics, and best-response dynamics are guaranteed to reach such an equilibrium for line metrics. Moreover, the price of stability of the game is 1 for tree metrics, and the price of anarchy is unbounded even for line metrics. It is an open question whether best-response dynamics converge into a Nash equilibrium for trees.

For the k-center model, we show that in general an instance might have no Nash equilibria, unless penalties are imposed on deviations. On the other hand, there always exists a Nash equilibrium in instances corresponding to line metrics. When the

Fig. 7. Parition of the nodes of Figure 1 into types.

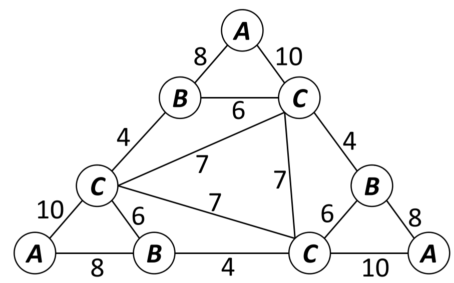

centroids are allowed to be located in arbitrary positions along edges, tree metrics are also guaranteed to have a Nash equilibrium, and best-response dynamics are guaranteed to reach such an equilibrium. The price of anarchy of the game is unbounded even for line metrics. It is an open question whether tree-metrics always have a Nash equilibrium also when centroids are required to be located in nodes of their own cluster, and whether best-response dynamics converge into a Nash equilibrium for trees/line metrics. Another interesting question is to determine the price of stability, even for line metrics.

For the correlation clustering model, we show that a Nash equilibrium always exists, and best-response dynamics are guaranteed to reach such an equilibrium. Despite of this positive result, the problem of finding a Nash equilibrium is PLS-complete. The price of stability of this model is 1 when all nodes have identical weights, and strictly larger than 1 otherwise. However, our understanding of the price of stability is very lacking – we do not have any non-trivial upper bounds, and our lower bounds are only small constants. The price of anarchy is much better understood. For the maximization variant it is $\Theta(\sqrt{n})$ for general and tree metrics and $\Theta(\sqrt[3]{n})$ for line metrics. For the minimization variant, the price of anarchy is about n for identical node weights and at most $O(n^2)$ for general node weights. Finding better bounds on the price of anarchy of the minimization variant for general and more speific metrics is another open problem.

APPENDIX

A. EXAMPLES OF NO NASH EQUILIBRIA

In this appendix we show that the metrics given by Figures 1 and 5 do not have a Nash equilibrium (under their respective model) with two clusters. We begin with the graph of Figure 1. It is possible to split the nodes of this graph into three types, such that the nodes of each type are symmetric. Notice that a Nash equilibrium must have two non-empty clusters, otherwise any non-centroid node has an incentive to deviate. We would like to check every possible pair of nodes $u \neq v$, and prove that no Nash equilirium has u and v as the centroids of its clusters. For a fixed pair, this can be done as following:

- Assume $\mathcal{E}$ is a Nash equilibrium whose centroids are u and v
- Each node $w \notin \{u, v\}$ must belong in $\mathcal{E}$ to the cluster of the centroid closer to it. This observation allows an easy reconstruction of the clusters of $\mathcal{E}$.
- Once $\mathcal{E}$ is reconstructed, we find a some node which can benefit by deviating or show that one of the centroid is wrongly placed.

The following table checks pairs of nodes where at least one node is of type A. By symmetry, we may assume the top A node of Figure 1 is in all the pairs. The left column of the table presents in bold the pair of nodes checked, and by shading the resulting potential Nash equilibrium $\mathcal{E}$ constructed using the above logic: white and gray nodes represent the two clusters, while black nodes can belong to either cluster (they have the same distance from both centoids). The right column of the table explains why $\mathcal{E}$ is not a Nash equilibrium.

Checked Pair	**Not a Nash equilibrium because...**
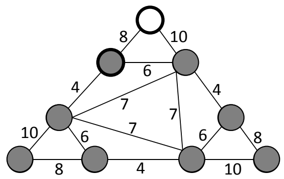	The current cost of the gray cluster is 73. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 57.
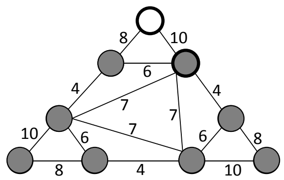	The current cost of the gray cluster is 64. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 57.
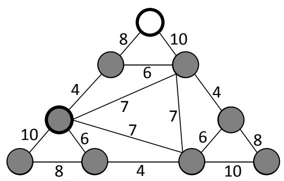	The current cost of the gray cluster is 62. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 57.
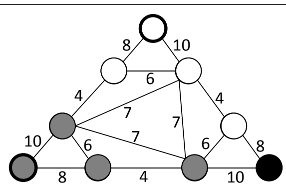	The current cost of the gray cluster is 52 or 30 (depending on which cluster the black node belongs to). If the centroid of this cluster moved to the bottom left B node, then the cost of this cluster would decrease to 32 or 18, respectively.
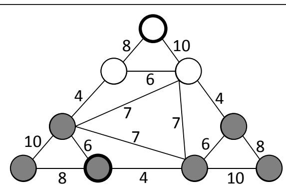	The current cost of the white cluster is 18. If the centroid of this cluster moved to the top B node, then the cost of this cluster would decrease to 14.
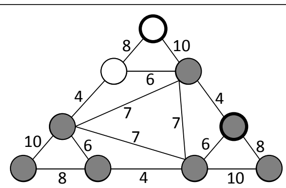	The current cost of the gray cluster is 57. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 46.

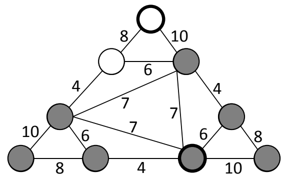	The current cost of the top C node is 7. If this node switched to the white cluster, the centroid of the white cluster would move to the top B node, which would decrease the cost of the top C node to 6.
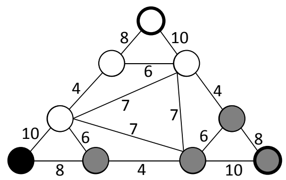	This configuration is symmetric to the one given in the forth line of this table, with the roles of the white and gray clusters reversed.

The following table checks pairs of nodes where at least one node is of type B and no node is of type A. By symmetry, we may assume the top B node of Figure 1 is in all the pairs. The table is structured in the same way as the previous one.

Checked Pair	**Not a Nash equilibrium because...**
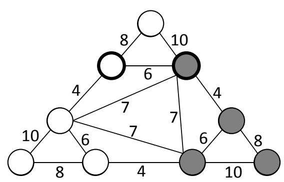	The current cost of the gray cluster is 23. If the centroid of this cluster moved to the bottom right B node, then the cost of this cluster would decrease to 18.
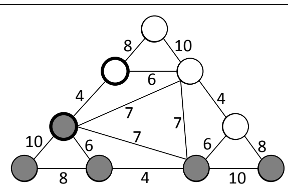	The current cost of the gray cluster is 40. If the centroid of this cluster moved to the bottom left B node, then the cost of this cluster would decrease to 32.
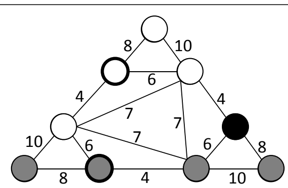	Assume first the black node belongs to the white cluster. In this case, the current cost of the white cluster is 28. If the centroid of this cluster moved to the top C node, then the cost of this cluster would decrease to 27. Assume now that the black node belongs to the gray cluster. In this case, the current cost of the gray cluster is 36. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 32.
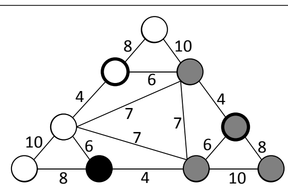	This configuration is symmetric to the previous one, with the roles of the white and gray clusters reversed.

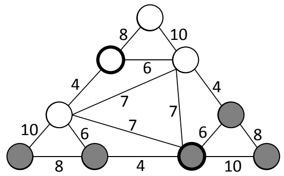	The current cost of the bottom right B node is 6. If this node switched to the white cluster, the centroid of the white cluster would move to the top C node (see the first case of the configuration before the previous one), which would decrease the cost of the bottom right B node to 4.

The following table checks pairs of C nodes. By symmetry, we may assume the top C node of Figure 1 is in all the pairs. The table is structured in the same way as the previous ones.

Checked Pair	**Not a Nash equilibrium because...**
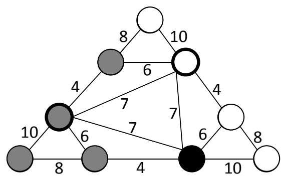	Assume first the black node belongs to the white cluster. In this case, the current cost of the white cluster is 33. If the centroid of this cluster moved to the bottom right B node, then the cost of this cluster would decrease to 32. Assume now that the black node belongs to the gray cluster. In this case, its the current cost is 7. If it switched to the white cluster, the centroid of the first cluster would have moved to the bottom right B node, which would decrease the cost of the black node to 6.
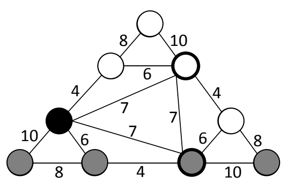	This configuration is symmetric to the previous one, with the roles of the white and gray clusters reversed.

We switch our attention to the graph of Figure 5, and show that it represents an instance with no Nash equilibrium under the k-center model. We use the same partition of the nodes into three types used for the graph of Figure 5. The proof itself is done in exactly the same way, except that the cost of the cluster is calculated according to the k-center model.

The following table checks pairs of nodes where at least one node is of type A. By symmetry, we may assume the top A node of Figure 1 is in all the pairs. The table is structured in the same way as the previous ones.

Checked Pair	**Not a Nash equilibrium because...**
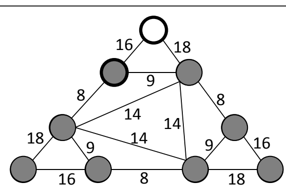	The current cost of the gray cluster is 33. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 24.

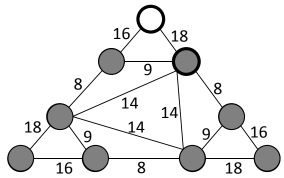	The current cost of the gray cluster is 32. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 24.
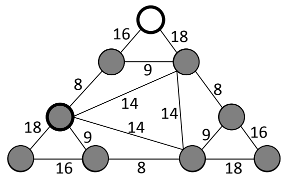	The current cost of the gray cluster is 32. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 24.
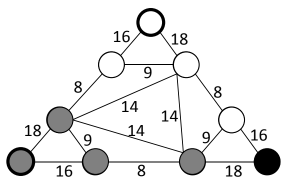	The current cost of the gray cluster is 42 or 24 (depending on which cluster the black node belongs to). If the centroid of this cluster moved to the bottom left B node, then the cost of this cluster would decrease to 26 or 16, respectively.
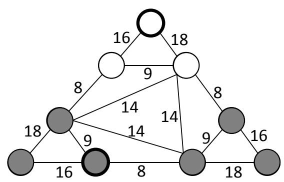	The current cost of the white cluster is 18. If the centroid of this cluster moved to the top B node, then the cost of this cluster would decrease to 16.
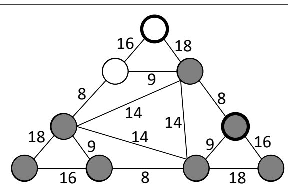	The current cost of the gray cluster is 33. If the centroid of this cluster moved to the bottom right C node, then the cost of this cluster would decrease to 24.
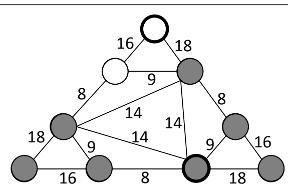	The current cost of the top C node is 14. If this node switched to the white cluster, the centroid of the white cluster would move to the top B node, which would decrease the cost of the top C node to 9.
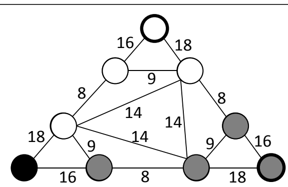	This configuration is symmetric to the one given in the forth line of this table, with the roles of the white and gray clusters reversed.

The following table checks pairs of nodes where at least one node is of type B and no node is of type A. By symmetry, we may assume the top B node of Figure 1 is in all the pairs. The table is structured in the same way as the previous one.

Checked Pair	**Not a Nash equilibrium because...**
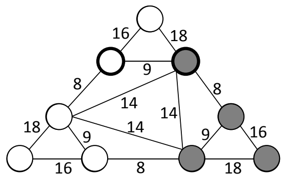	The current cost of the gray cluster is 24. If the centroid of this cluster moved to the bottom right B node, then the cost of this cluster would decrease to 16.
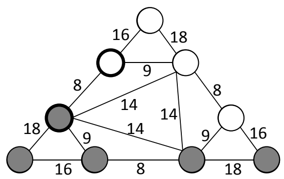	The current cost of the gray cluster is 32. If the centroid of this cluster moved to the bottom left B node, then the cost of this cluster would decrease to 26.
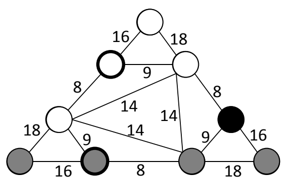	The current cost of the gray cluster is 26 (regardless of which cluster does the black node belong to). If the centroid of this cluster moved to the bottom left C node, then the cost of this cluster would decrease to 24.
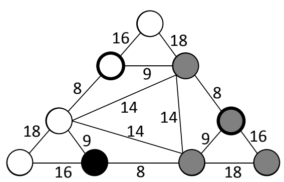	This configuration is symmetric to the previous one, with the roles of the white and gray clusters reversed.
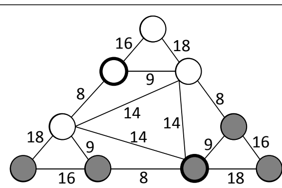	The current cost of the bottom left A node is 24. If this node switched to the white cluster, the centroid of the white cluster would move to the bottom left C node (to decrease the cost of the cluster from 26 to 24), which would decrease the cost of the bottom left A node to 18.

The following table checks pairs of C nodes. By symmetry, we may assume the top C node of Figure 1 is in all the pairs. The table is structured in the same way as the previous ones.

<table>
<tr><th>Checked Pair</th><th>Not a Nash equilibrium because...</th></tr>
<tr><td>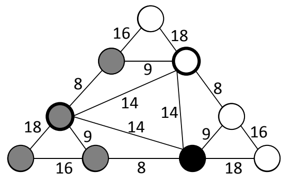
</td><td>The current cost of the gray cluster is 18 (regardless of which cluster does the black node belong to). If the centroid of this cluster moved to the bottom left B node, then the cost of this cluster would decrease to 17.</td></tr>
<tr><td>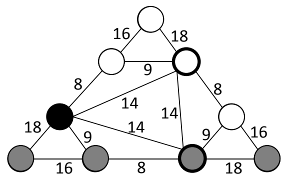
</td><td>This configuration is symmetric to the previous one, with the roles of the white and gray clusters reversed.</td></tr>
</table>

B. OMITTED PROOFS

LEMMA 3.8. *The centroid of a cluster on a line is a median of that cluster.*

PROOF. Consider a median node v^* of a cluster C of size m. Then, by definition, v^* divides the nodes in C into two sets, one set consists of the nodes at the left of v^* and the other consists of the nodes at its right, where none of the sets contains more than $\lceil \frac{m-1}{2} \rceil$ nodes. We prove the lemma by way of contradiction. Assume the centroid of C is a node v that is not the median. Then, v divides the nodes into two sets, right and left, denoted respectively by S_L^v and S_R^v, such that there is a set with at least $\lceil \frac{m-1}{2} \rceil + 1$ nodes. Assume without loss of generality that this is the left set, thus $|S_R^v| + 1 < |S_L^v|$. Denote the node at the left of v by v_L. We show that $D(v, C) > D(v_L, C)$, contradicting the assumption that v is the centroid. By moving the center to v_L, we increase the distance of all nodes in $|S_R^v|$ as well as the distance of v itself by $d(v, v_L)$, while decreasing the distance of all nodes in $|S_L^v|$ (including v_L) by $d(v, v_L)$. Thus,

$$D(v_L, C) = D(v, C) + d(v, v_L)(|S_R^v| + 1 - |S_L^v|) \ . \tag{7}$$

Since we assumed $|S_R^v| + 1 < |S_L^v|$, we get $D(v, C) > D(v_L, C)$. □

LEMMA 3.12. *There are at most 2 median nodes in a tree.*

PROOF. Consider a tree T. Let v^* be a median node of T, and let $Q_1, \ldots, Q_m$ denote the connected components formed after v^*'s removal. We consider two cases, according to the number of nodes $|V|$.

(1) In case $|V|$ is odd, the maximal size of such a component can be at most $\frac{|V|-1}{2}$. We show that a median node cannot be located in any of the m connected components, and thus, there is a single median node v^*. Assume by way of contradiction that there exists another median node w in a connected component Q_j. In case we remove w from T, a connected component containing the set of vertices $S = \{v^* \cup \{w \in \bigcup_{i \neq j} Q_i\}\}$ is formed (note that this connected component can also contain additional vertices from Q_j). As the total size of the connected components $\sum_{i \neq j} |Q_i|$ is at least $\frac{|V|-1}{2}$ (as $|Q_j| \leq \frac{|V|-1}{2}$), it follows that the cardinality of the set S is at least $\frac{|V|-1}{2} + 1$, contradicting the assumption that w is a median node.

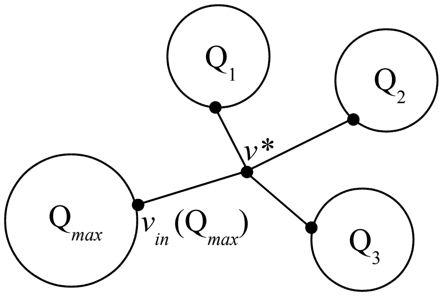

Fig. 8. Decomposition of a tree to connected components after the removal of v^*.

(2) In case $|V|$ is even, the maximal size of a connected component following v^*'s removal can be at most $q = \left\lceil \frac{|V|-1}{2} \right\rceil$, and there is at most one component Q_{max} of this size exactly. In that case, we show that there are at most two median nodes.
We first show that a median node cannot be located in any component of size less than q. Assume by way of contradiction that there exists another median node w in a connected component Q_j of size less than q, then in case w is removed from T, a connected component containing the set of vertices $S = \{v^* \cup \{w \in \bigcup_{i \neq j} Q_i\}\}$ is formed (note that this connected component can also contain additional vertices from Q_j). As the total size of the connected components $\sum_{i \neq j} |Q_i|$ is at least q (as $|Q_j| \leq \left\lfloor \frac{|V|-1}{2} \right\rfloor$), it follows that the cardinality of the set S is at least $\left\lceil \frac{|V|-1}{2} \right\rceil + 1$, contradicting the assumption that w is a median node.
Now, we show that another median node in T is the node $v_{in}(Q_{max})$, which is the (sole) neighbor of v^* in the connected component Q_{max} (see Figure 8). In case we remove this node from the tree, we get a connected component containing the vertices $\{v^* \cup \{w \in \bigcup_{Q_i \neq Q_{max}} Q_i\}\}$ of size $\left\lfloor \frac{|V|-1}{2} \right\rfloor + 1 = \left\lceil \frac{|V|-1}{2} \right\rceil$, and the connected components of $Q_{max} - v^*$, whose total size is at most $\left\lfloor \frac{|V|-1}{2} \right\rfloor$. Thus, $v_{in}(Q_{max})$ is indeed a median node in T. We show that no other node in Q_{max} can be a median node. Assume by way of contradiction that there exists another median node $w \neq v_{in}(Q_{max})$ in Q_{max}, then, if we remove it from T, a connected component is formed containing the vertices $\{\{v^*, v_{in}(Q_{max})\} \cup \{w \in \bigcup_{Q_i \neq Q_{max}} Q_i\}\}$ (note that this connected component can also contain additional vertices from Q_{max}). The size of this components is thus at least $\left\lceil \frac{|V|-1}{2} \right\rceil + 1$, contradicting the assumption that w is a median node. $\square$

OBSERVATION 5.2. *Every correlation clustering game can be converted into an additively-separable hedonic game with the same set of players and the same set of Nash equilibria. The reverse is not true.*

PROOF. Given a maximization correlation clustering game, one can convert it into an additively-separable hedonic game with the same set of players. For every two players u, v, the weight of the arc uv is set to be $w_v \cdot (1 - 2d(u,v))$. Consider a configuration $\mathcal{E}$, and let C be a cluster of this configuration and u be a node of this cluster. The utility of u under $\mathcal{E}$ in the correlation clustering game is:

$$\sum_{\substack{v \in C \\ v \neq u}} w_v(1 - d(u,v)) + \sum_{v \notin C} w_v d(u,v) = \sum_{\substack{v \in C \\ v \neq u}} w_v(1 - 2d(u,v)) + \sum_{\substack{v \in V \\ v \neq u}} w_v d(u,v) \ .$$

On the other hand, the utility of u under $\mathcal{E}$ in the additively-separable hedonic game is:

$$\sum_{\substack{v \in C \\ v \neq u}} w_v(1 - 2d(u,v)) \ ,$$

which is different from the previous expression by a constant (that depends on u, but not on $\mathcal{E}$). Since this is true for every configuration $\mathcal{E}$, u has an incentive to deviate from $\mathcal{E}$ in the correlation clustering game if and only if she has such an incentive in the additively-separable hedonic game. Thus, the two games have the same set of Nash equilibria. The same result extends to the minimization variant via Lemma 5.1.

The second part of the observation follows since an additively-separable hedonic game might not have any Nash equilibrium [Gairing and Savani 2010], while correction clustering games always have a Nash equilibrium by Theorem 5.3 (notice that although Theorem 5.3 appears later than this observation, this is not a circular reference since Theorem 5.3 does not depend on this observation). □

THEOREM 5.3. *There always exists a Nash equilibrium for the correlation clustering game. Moreover, best-response dynamics of this game always converge to a Nash equilibrium.*

PROOF. It suffices to prove the theorem for the minimization variant (the proof for the maximization variant follows directly from Lemma 5.1.) Given a clustering configuration $\mathcal{F}$, consider the potential function $\Phi(\mathcal{F})$ whose value is given by:

$$\sum_{v \in V} w_v \left(\sum_{u \in C_v} w_u d(u,v) + \sum_{u \notin C_v} w_u (1 - d(v,u)) \right) .$$

A best-response move performed by element v changes the value of Φ only through edges adjacent to v. The overall contribution of such edges is

$$2w_v \left(\sum_{u \in C_v} w_u d(u,v) + \sum_{u \notin C_v} w_u (1 - d(v,u)) \right) .$$

This expression is $2w_v$ times the cost paid by v, hence, it strictly decreases if v makes a best-response move. Hence, Φ decreases with every move performed by a player decreasing its cost. Since the correlation clustering game is a finite game, best-response dynamics are guaranteed to converge to a Nash equilibrium. □

THEOREM 5.4. *The problem of computing a Nash equilibrium in the correlation clustering game is PLS-Complete.*

PROOF. The hardness result for the problem of computing a Nash equilibrium of a correlation clustering game is achieved by a reduction from the problem called POS-NAE-3SAT. In POS-NAE-3SAT, we are given an expression ϕ of m clauses, each containing up to 3 variables (with no negative literals). Each clause c_i is associated with a weight w_i. A truth assignment to the variables satisfies a clause if it does not assign the same truth value to all the variables of the clause. The weight of the assignment is the sum of the weights of the clauses it satisfies. The objective of the problem is to find an assignment which is a local maximum, *i.e.*, its weight cannot be increased by flipping the value of a single variable. POS-NAE-3SAT is known to be a PLS-complete problem [Schäffer and Yannakakis 1991].

The theorem is proved by a reduction from POS-NAE-3SAT. We present the reduction for the maximization variant of the game, but the hardness result also applies to the minimization variant following Lemma 5.1. The reduction is established by showing the following three steps:

— Any instance ϕ of POS-NAE-3SAT can be converted into an instance I of the correlation clustering game.
— Any clustering configuration of I can be converted into a truth assignment of ϕ.
— A Nash equilibrium of I converts into a local maximum of ϕ.

We begin by proving the first step. Let us first give an intuitive explanation of our construction. The instance we construct has $n + 2$ nodes, where n is the number of variables of ϕ. Of these nodes, n represent the variables of ϕ and 2 nodes (called c_1 and c_2) represent the two truth values (TRUE and FALSE). The distances are chosen in such a way that c_1 and c_2 must belong to different clusters in any Nash equilibrium. For (a node representing) a variable, we interpret its cluster as a truth value assignment – if it share a cluster with c_1 (respectively, c_2) it is assigned the truth value TRUE (respectively, FALSE). Each variable is initially located at identical distances from c_1 and c_2, and at distance $1/2$ from every other nodes.

For every clause, we increase the distances between the variables of the clause by an amount proportional to the weight of the clause. Notice that given a clustering, every clause can be in one of two states: state 1 – all the nodes of the clause belong to a single cluster (which does not satisfy the clause), or state 2 – the nodes of the cluster are partitioned into a set of 2 nodes and a singleton (which satisfy the clause). Observe that increasing the distances between the variables of a clause favors clusterings in which the clause is in state 2.

Consider some variable u. When u deviates from one cluster to the other, it switches the state of every clause whose satisfaction is altered by the deviation (if the two other variables of a clause belong to different clusters, then the state of the clause does not change and the clause is satisfied both before and after the deviation). It is easy to see that the gain of u increases when it changes the state of clauses from state 1 to state 2 and decreases when it changes clauses of state 2 to state 1. Hence, the interest of u is correlated with the global objective.

We are now ready to describe our construction formally. Given an instance ϕ of POS-NAE-3SAT, we construct an instance I of the correlation clustering game. We denote the set of variables in ϕ by $\mathcal{V}$ (recall that $|\mathcal{V}| = n$). The set of nodes in I is defined as $V = \{c_1, c_2\} \cup \{v_x \mid x \text{ is a variable of } \phi\}$. All the nodes of V have identical weights. The nodes c_1 and c_2 represent the two truth values. The distance between c_1 or c_2 and any node v_x is set to $1/2 - 1/(4(n + 2))$, and the distance between c_1 and c_2 is set to $1-1/(2(n+2))$. Due to the large distance between c_1 and c_2, it follows from Lemma 5.6 that they cannot share a common cluster in a Nash equilibrium. Given two variables x, y, we denote by $w(x, y)$ the total weight of the clauses that include both x and y. The total weight of all clauses is denoted by W. The distance between any two nodes v_x and v_y is set to $\frac{1}{2} + \frac{w(x,y)}{5n^2W}$. An example of our construction can be found in Figure 9.

We proceed with the second step, showing how to translate a clustering configuration of I into a truth assignment of ϕ. Given such a clustering configuration, a variable x gets the truth assignment TRUE if and only if v_x shares a cluster with c_1. We now get to the third (and last) step, showing that every Nash equilibrium configuration $\mathcal{E}$ of I converts into a local optimum of ϕ. We first show that every node v_x in $\mathcal{E}$ either shares a cluster with c_1 or with c_2. Assume by way of contradiction that this is not the case. The value v_x gains is upper bounded by:

$$2\left[\frac{1}{2} - \frac{1}{4(n+2)}\right] + \sum_{y \in V - \{x\}} \left[1 - \left(\frac{1}{2} + \frac{w(x,y)}{5n^2W}\right)\right] \leq \frac{n+1}{2} - \frac{n-1}{5n^2} - \frac{1}{2(n+2)} \quad . \tag{8}$$

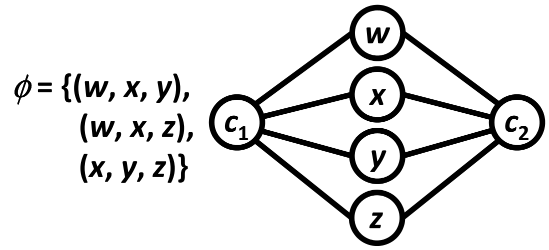

	w	x	y	z
w	0	1/2+2/240	1/2+1/240	1/2+1/240
x	1/2+2/240	0	1/2+2/240	1/2+2/240
y	1/2+1/240	1/2+2/240	0	1/2+1/240
z	1/2+1/240	1/2+2/240	1/2+1/240	0

Fig. 9. An example of the construction used by the proof of Theorem 5.4. ϕ is an instance of POS-NAE-3SAT having 3 clauses of weight 1. The construction produces the metric on the right. All the edges appearing in the image correspond to distances of $1/2 - 1/24$. The metric also assigns $d(c_1, c_2) = 1 - 1/12$ and the distances given by the above table.

In case v_x deviates to the cluster of c_1, its profit is at least:

$$\left[\frac{1}{2} - \frac{1}{4(n+2)}\right] + \left[1 - \left(\frac{1}{2} - \frac{1}{4(n+2)}\right)\right] + \sum_{y \in V - \{x\}} \left[\frac{1}{2} + \frac{w(x,y)}{5n^2W}\right] \geq \frac{n+1}{2} + \frac{n-1}{5n^2} \ . \tag{9}$$

The value of expression (9) is larger than the value of expression (8), hence, v_x has an incentive to deviate, and we get a contradiction to the choice of $\mathcal{E}$ as a Nash equilibrium. Hence, there are exactly two clusters in $\mathcal{E}$, the cluster of c_1 corresponding to the truth value TRUE and the cluster of c_2 corresponding to the truth value FALSE.

Assume by way of contradiction that $\mathcal{E}$ corresponds to an assignment ψ that is not a local maximum, *i.e.*, ψ can be improved by flipping the truth value of a single variable x. Let ψ' denote the assignment ψ with the value of x flipped, and let $\mathcal{E}'$ denote the clustering configuration corresponding to ψ' (*i.e.*, all variables assigned the truth value of TRUE by ψ' are in the cluster of c_1, and all other variables are in the cluster of c_2). Clearly, $\mathcal{E}'$ can result from $\mathcal{E}$ following a deviation of v_x, and thus, as $\mathcal{E}$ is an equilibrium, the profit of v_x in $\mathcal{E}$ is at least as large as in $\mathcal{E}'$. We use a few definitions:

— Let W_1 be the total weight of the clauses *satisfied* by ψ in which x appears, and in which the two other variables are assigned equal values by ψ.
— Let W_2 be the total weight of clauses *not satisfied* by ψ in which x appears, and in which the two other variables are assigned equal values by ψ.
— Let W_3 be the total weight of clauses in which x appear and in which the two other variables are assigned different values by ψ. Observe that these clauses are satisfied regardless of the value assigned to x.

The difference between the weights of ψ and ψ' is $W_1 - W_2$, and following our assumption, it is a negative value. We now consider the value gained by v_x in $\mathcal{E}$ and $\mathcal{E}'$. We know that exactly one of the nodes c_1 and c_2 is in the cluster of v_x, hence the profit from the edges between v_x and these nodes is:

$$\left[\frac{1}{2} - \frac{1}{4(n+2)}\right] + \left[1 - \left(\frac{1}{2} - \frac{1}{4(n+2)}\right)\right] = 1 \ .$$

Consider an edge connecting v_x to some other variable v_y sharing a cluster with v_x. The contribution of this edge to the profit of v_x is:

$$1 - \left(\frac{1}{2} + \frac{w(x,y)}{5n^2W}\right) = \frac{1}{2} - \frac{w(x,y)}{5n^2W} \ .$$

Denoting by ℓ the number of nodes sharing a cluster with v_x in $\mathcal{E}$, we get that the total contribution of the above edges is:

$$\frac{\ell}{2} - \frac{2W_2 + W_3}{5n^2W} \ .$$

Similarly, we consider the edges connecting v_x to other variables of the form v_z that do not share a cluster with v_x. The total contribution of these edges to the profit of v_x is:

$$\frac{n-1-\ell}{2} + \frac{2W_1 + W_3}{5n^2W} \ .$$

The total profit of v_x from all its adjacent edges in $\mathcal{E}$ is thus:

$$1 + \left[\frac{\ell}{2} - \frac{2W_2 + W_3}{5n^2W}\right] + \left[\frac{n-1-\ell}{2} + \frac{2W_1 + W_3}{5n^2W}\right] = \frac{n+1}{2} + \frac{2}{5n^2W} \cdot (W_1 - W_2) < \frac{n+1}{2} \ ,$$

where the inequality holds since $W_1 - W_2 < 0$. A similar analysis shows that the values gained by v_x from all its adjacent edges in $\mathcal{E}'$ is:

$$1 + \left[\frac{\ell}{2} + \frac{2W_2 + W_3}{5n^2W}\right] + \left[\frac{n-1-\ell}{2} - \frac{2W_1 + W_3}{5n^2W}\right] = \frac{n+1}{2} + \frac{2}{5n^2W} \cdot (W_2 - W_1) > \frac{n+1}{2} \ .$$

This is a contradiction, since $\mathcal{E}$ is a Nash equilibrium and the profit of v_x in $\mathcal{E}'$ cannot be larger than its profit in $\mathcal{E}$. The contradiction results from the assumption that ψ is not a local maximum; hence ψ is a local maximum, which completes the proof of step 3 and the whole theorem. □

ACKNOWLEDGMENT

We thank Ari Freund, Yossi Richter, and Elad Yom-Tov for helpful discussions. Special thanks go to Adam Kalai, Ronny Lempel, and Nir Ailon for fruitful comments on our work.

REFERENCES

A. Abbasi and M. Younis. 2007. A survey on clustering algorithms for wireless sensor networks. *Computer Communications* 30, 14–15 (June 2007), 2826–2841.

S. Bandyopadhyay and E.J. Coyle. 2003. An energy efficient hierarchical clustering algorithm for wireless sensor networks. In *Proceedings of INFOCOM*. 1713–1723.

S. Banerjee, H. Konishi, and T. Sonmez. 2001. Core in a Simple Coalition Formation Game. *Social Choice and Welfare* 18, 1 (2001), 135–153.

N. Bansal, A. Blum, and S. Chawla. 2004. Correlation clustering. *Machine Learning J.* 56, 1–3 (2004), 86–113.

F. Bloch and E. Diamantoudi. 2010. Noncooperative formation of coalitions in hedonic games. *International Journal of Game Theory* 40, 2 (2010), 263–280.

A. Blum. 2009. Thoughts on clustering. In *NIPS Workshop on Clustering Theory*.

A. Bogomolnaia and M.O. Jackson. 2002. The stability of hedonic coalition structures. *Games and Economic Behavior* 38, 2 (2002), 201–230.

S. R. Bulò. 2009. *A game-theoretic framework for similarity-based data clustering*. Ph.D. Dissertation. Università Ca' Foscari di Venezia.

M. Charikar, S. Guha, E. Tardos, and D. Shmoys. 1999. A constant factor approximation algorithm for the k-median problem. In *Proceedings of STOC*. ACM, New York, NY, USA, 1–10.

M. Charikar, V. Guruswami, and A. Wirth. 2005. Clustering with qualitative information. *J. Comput. Syst. Sci.* 71, 3 (2005), 360–383.

E. Demaine and N. Immorlica. 2003. Correlation clustering with partial information. In *Proceedings of APPROX*. Springer, Berlin Heidelberg, 71–80.

J.H. Dreze and J. Greenberg. 1980. Hedonic coalitions: Optimality and stability. *Econometrica* 48, 4 (1980), 987–1003.

Z. Drezner. 1995. *Facility Location: A survey of Applications and Methods*. Springer Verlag, New York, NY, USA.

D. Emanuel and A. Fiat. 2003. Correlation clustering – Minimizing disagreements on arbitrary weighted graphs. In *Proceedings of ESA*. Springer, Berlin, Heidelberg, 208–220.

M. Gairing and R. Savani. 2010. Computing Stable Outcomes in Hedonic Games. *Lecture Notes in Computer Science* 10 (2010), 174–185.

T. F. Gonzalez. 1985. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science* 38 (1985), 293–306.

L. Gourvès and J. Monnot. 2009. On Strong Equilibria in the Max Cut Game. In *Proceedings of WINE*. Springer-Verlag, Berlin, Heidelberg, 608–615.

T. Groves. 1973. Incentives in teams. *Econometrica* 41, 4 (1973), 617–631.

D. S. Hochbaum and D. B. Shmoys. 1985. A best possible heuristic for the k-center problem. *Mathematics of Operations Research* 10 (1985), 180–184.

M. Hoefer. 2007. *Cost Sharing and Clustering under Distributed Competition*. Ph.D. Dissertation. Universität Konstanz.

H. Hotelling. 1929. Stability in competition. *Economic Journal* 39, 153 (1929), 41–57.

G. Koltsidas and F. N. Pavlidou. 2011. A game theoretic approach to clustering of ad-hoc and sensor networks. *Telecommunication Systems* 47, 1–2 (June 2011), 81–93.

P. B. Mirchandani and R. L. Francis. 1990. *Discrete Location Theory*. Wiley Interscience, New York, NY, USA.

M. Pelillo. 2009. What is a cluster? perspectives from game theory. In *NIPS Workshop on Clustering Theory*.

A. A. Schäffer and M. Yannakakis. 1991. Simple local search problems that are hard to solve. *Siam Journal of Computing* 20, 1 (February 1991), 56–87.

A. Tamir. 1996. An $O(pn^2)$ algorithm for the p-median and related problems on tree graphs. *Operations Research Letters* 19, 2 (August 1996), 59–64.

V. Vazirani. 2001. *Approximation Algorithms*. Springer Verlag, New York, NY, USA.

E. B. Yanovskaya. 1968. Equilibrium points in polymatrix games (in Russion). *Litovskii Matematicheskii Sbornik* 8 (1968), 381–384.

Received Month Year; revised Month Year; accepted Month Year