

Coalitional Affinity Games and the Stability Gap

Simina Brânzei

Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada
sbranzei@uwaterloo.ca

Kate Larson

Cheriton School of Computer Science
University of Waterloo
Waterloo, Canada
klarson@uwaterloo.ca

Abstract

We present and analyze *coalitional affinity games*, a family of hedonic games that explicitly model the value that an agent receives from being associated with other agents. We provide a characterization of the social-welfare maximizing coalition structures, and study the stability properties of affinity games, using the core solution concept. Interestingly, we observe that members of the core do not necessarily maximize social welfare. We introduce a new measure, the *stability-gap* to capture this difference. Using the stability gap, we show that for an interesting class of coalitional affinity games, the difference between the social welfare of a stable coalition structure and a social welfare maximizing coalition structure is bounded by a factor of two, and that this bound is tight.

1 Introduction

Imagine the following scenario. You are organizing a party and have to come up with a seating arrangement, but this arrangement should take into consideration the relationships between the guests. For example, Alice is a good friend of Bob and would like to sit with him. However, Alice is feuding with Chris, Bob's best friend, and refuses to remain at any table with Chris. You want to make all the guests as happy as possible, but you also do not want guests changing the seating arrangement when they arrive. How should you assign the guests to tables?

The scenario just described has some interesting features. First, the agents (or guests) have value from interacting or being associated with others, and this value may be positive or negative. Second, the agents would like to coordinate with others, but ideally only with agents with whom they have a positive relationship. Finally, we would like stability. That is, once a group is seated at a table, they have no incentive to move.

In this paper we propose a model which explicitly captures the value, which we call the *affinity*, that an agent receives from being associated with another agent. In particular we study situations where an agent is interested in belonging to groups (coalitions) that contain agents for which it has high affinity, while avoiding groups that contain agents for which

it has negative affinity. By placing our affinity model into the context of non-transferable utility coalitional games [Pegel and Sudhölter, 2003], we are able to characterize and compare cooperative structures (*i.e.* coalition structures) that maximize social welfare with those that are stable (*i.e.* in the core). We argue that our affinity model is a rich representation that can be used to model many situations, while at the same time has enough structure to provide interesting characterizations of coalitions and coalition structures.

We organize the rest of the paper as follows. In the next section we introduce our affinity model, the *affinity graph* and *coalitional affinity game*. We also introduce the important concepts from the literature on coalitions that we use throughout the paper. In Section 3 we provide a complete characterization of the social-welfare maximizing coalition structure. We then study stability properties of the coalitional affinity game, using the core as the solution concept. In Section 5 we compare the social-welfare maximizing coalition structure with stable coalition structures. We ask the question “Assuming that stable coalition structures exist, how large a sacrifice, in terms of social welfare, does the group of agents have to make in order to be stable?” Experimental work, described in Section 6, support our theoretical findings from Section 5, while also showing that in practice, the core of coalitional-affinity games is often non-empty.

2 The Model

Let there be a set of agents $N = \{x_1, \dots, x_n\}$. For any pair of agents, we denote the *affinity* that agent x_i has for x_j as $a(x_i, x_j) \in \mathbb{R}$ which represents the value that agent x_i receives from being associated with agent x_j . We represent the agents and their affinities with an *affinity graph*.

Definition 1 An affinity graph, $A = (N, E)$, is a weighted directed graph where

- N is a set of agents, $N = \{x_1, \dots, x_n\}$,
- edge $(x_i, x_j) \in E$ represents an affinity relation between agents x_i and x_j , and
- weight $a(x_i, x_j) \in \mathbb{R}$ is the value that agent x_i receives from being associated with agent x_j .¹

¹If $(x_i, x_j) \notin E$ then $a(x_i, x_j) = 0$.




Figure 1: An affinity graph. Here agent x and y have equal affinity, as do y and z . However, while x has positive affinity for agent z ($a(x, z) = 5$), agent z has negative affinity for agent x ($a(z, x) = -3$).

Figure 1 is an example of an affinity graph. While the definition of an affinity graph is general, we will sometimes be interested in *symmetric* affinity graphs.

Definition 2 An affinity graph $A = (N, E)$ is symmetric if $a(x_i, x_j) = a(x_j, x_i)$ for all $(x_i, x_j), (x_j, x_i) \in E$.

Given an affinity graph, $A = (N, E)$, we are interested in understanding how agents will choose to interact with each other by forming coalitions. We define the utility of agent x_i from belonging in coalition $S \subseteq N$ as follows:

$$u(S, x_i) = \begin{cases} 0 & \text{if } |S| = 1 \\ \sum_{(x_i, x_j) \in E | x_j \in S} a(x_i, x_j) & \text{otherwise} \end{cases}$$

Referring back to Figure 1 where $N = \{x, y, z\}$, we have $u(N, x) = 1 + 5 = 6$, $u(N, y) = 1 + 0 = 1$ and $u(N, z) = -3 + 0 = -3$. Since an agent can always decide to not belong to a coalition, we are interested in coalitions where the utility of an agent is at least zero. If for agent $x_i \in S$, $u(S, x_i) \geq 0$ then we say that S is *individually rational* for agent x_i .

It is now possible to define a *coalitional affinity game* which is modeled as a characteristic function game with non-transferable utility [Peleg and Sudhölter, 2003].

Definition 3 Given an affinity graph $A = (N, E)$, the coalitional affinity game, $G(A)$, is the pair $\langle N, v \rangle$ where

- N is the set of agents defined by A , and
- for any $S \subseteq N$, $v(S) \subset \mathbb{R}^{|S|}$, such that for $x_i \in S$, $v_i(S) = u(S, x_i)$.

If A is symmetric, then we say that $G(A)$ is a symmetric affinity game. While the value function, v , of a coalitional affinity game returns a vector given a coalition, where entry i is the value that agent x_i receives from being in the coalition, we will sometimes abuse notation and say that the value of coalition S is

$$V(S) = \sum_{x_i \in S} v_i(S) = \sum_{x_i \in S} u(S, x_i) = \sum_{x_i, x_j \in S} a(x_i, x_j).$$

This definition implicitly assumes that agents' utilities are comparable. While this assumption is quite strong, it is commonly made in much of the coalition literature (see, for example [Bachrach and Rosenschein, 2008]) including work studying coalitions in networks (see, for example [Jackson and Wolinsky, 1996]).

As is standard, we define a *coalition structure*, P , to be a partition of the set of agents into coalitions. We are interested in properties of different coalition structures, such as the *social welfare* of a coalition structure and whether it is *stable*.

Definition 4 The social welfare of coalition structure $P = (S_1, \dots, S_m)$ is $SW(P) = \sum_{i=1}^m \sum_{x_j \in S_i} u(S_i, x_j)$.

The notion of stability we use in this paper is the *core* solution concept. We emphasize that this definition does not rely on any assumptions concerning inter-agent utility comparisons.

Definition 5 A coalition structure $P = (S_1, \dots, S_m)$ is in the core if there is no coalition $B \subseteq N$ such that $\forall x \in B$, if $x \in S_i$ then $u(B, x) \geq u(S_i, x)$ and for some j , $1 \leq j \leq m$, $\exists y \in S_j$ such that $u(B, y) > u(S_j, y)$. If the inequality is strict then P is in the weak core.

If a coalition structure P is in the core, then it is resistant against group deviations. No set of agents, B , can break away by forming a new coalition and improve the utility for at least one member, while not degrading the other agents. If such a coalition exists, then it is called a *blocking* coalition.

We introduce a weaker stability concept, *inner stability*, which limits the type of group deviations that can occur.

Definition 6 A coalition structure, $P = (S_1, \dots, S_m)$ exhibits inner stability if there is no blocking coalition B such that $B \subset S_i$ for some i , $1 \leq i \leq m$.

If a coalition structure exhibits inner stability then any blocking coalition must “cross coalition boundaries” by drawing members from at least two different coalitions.

2.1 Related Models

In this section we describe two models which are closely related to coalitional affinity games. Coalitional affinity games are a special subclass of *hedonic games* [Drèze and Greenberg, 1980; Bogomolnaia and Jackson, 2002]. Hedonic games are non-transferable utility games where each agent's utility depends on the identity of the other members of its coalition. Thus our model represents hedonic games where the pair-wise relationships between coalition members are important.² While we cannot represent all hedonic games, we capture an interesting subclass and our representation allows us to leverage the underlying relationships between agents when providing characterizations of social-welfare maximization and stable coalition structures.

Another related model was proposed by Deng and Papadimitriou [Deng and Papadimitriou, 1994]. They also used a weighted (undirected) graph to model the relationship between agents, and defined the value of a coalition S , $v(S)$, as we do in this paper. However, an important distinction between their work and ours is that they assumed *transferable utility*, and thus agents were able to freely redistribute the value of a coalition amongst themselves. This key difference in the models (transferable vs. non-transferable utility) means that the results obtained (and the techniques used) do not apply in our setting, as we illustrate later.

3 Social Welfare and Affinity Games

Given an affinity graph one question we are interested in is how the agents should be assigned to coalitions so as to maximize the social welfare. In general, this *coalition-structure*

²We can also model certain situations where the utility depends on the *number* of members, and not on their identities.

generation problem is challenging since if there are n agents, then there are $O(n^n)$ possible coalition structures. While there has been much work on searching for social-welfare maximizing coalitions for general coalition problems (see, for example [Rahwan *et al.*, 2007]), we are interested in understanding whether there are any particular properties in the affinity-graph representation that can be used to characterize the social-welfare maximizing coalition structure.

In this section we note the relationship between the social-welfare maximizing coalition structure and the *minimal cut* of the affinity graph. This observation provides us with a complete characterization of the structure of social-welfare maximizing coalition structures. Recall that a k -cut of weighted graph $G = (N, E)$ is a partition of N into k disjoint sets $P = (S_1, S_2, \dots, S_k)$, where $S_i \subseteq N$ and $S_i \cap S_j = \emptyset$ if $i \neq j$. The size of the cut is equal to the sum of the weight of the edges between each $S_i, S_j \in P$. A *minimal k -cut* is a k -cut that is no larger than any other k -cut.

Lemma 1 *Given an affinity game $G(A)$, and a fixed value k , $1 \leq k \leq n$, a minimal k -cut of A has the highest social welfare amongst all coalition structures of size k .*

Proof: For any coalition structure of size k , $P_k = (S_1, \dots, S_k)$, let $w(S_i, S_j)$ be the sum of the weights of the edges in the cut between coalitions S_i and S_j . That is

$$w(S_i, S_j) = \sum_{(x,y) | x \in S_i, y \in S_j} a(x, y).$$

Then, the weight of the cut of P_k is

$$\text{Cut}(P_k) = \sum_{i=1}^k \sum_{j=1, j \neq i}^k w(S_i, S_j).$$

Since the value of a coalition, S_i is equal to $\sum_{x,y \in S_i} a(x, y)$, the social welfare of P_k is

$$SW(P_k) = \sum_{i=1}^k \sum_{x,y \in S_i} a(x, y).$$

Therefore, the social welfare of P_k is equal to the sum of all edges in A minus the edges in the cut. That is

$$SW(P_k) = \sum_{i=1}^n \sum_{j=1}^n a(i, j) - \text{Cut}(P_k).$$

Therefore, the coalition structure of size k that maximizes social welfare, P_k^* , is the one that minimizes $\text{Cut}(P_k)$. $\square$

Theorem 1 *Given affinity game $G(A)$, let P_k^* be a minimal k -cut of affinity graph $A = (N, E)$. Then a social welfare maximizing coalition is $P^* = \max_k [P_1^*, \dots, P_n^*]$.*

Proof: Proof follows immediately from Lemma 1. $\square$

Theorem 1 provides a characterization of the social-welfare maximizing coalition structure for an arbitrary affinity graph. If the affinity graph is *symmetric* (Definition 2) then we are able to describe additional properties of the social-welfare maximizing coalition structure, P^* .




Figure 2: An affinity graph for which the coalitional affinity game has an empty core.

Theorem 2 *Let $G(A)$ be a symmetric affinity game, and let $P^* = (S_1^*, \dots, S_k^*)$ be the social-welfare maximizing coalition structure. Then, for any $S_i^* \in P^*$, any cut of S_i^* is non-negative.*

Due to space limitations we are unable to include the proof of Theorem 2. However, the proof involves assuming that a social-welfare maximizing coalition structure, P^* , has a negative cut, and then deriving a contradiction by finding another coalition structure P' such that $SW(P') > SW(P^*)$.

Theorem 2 does not imply that coalition structures that maximize social welfare in symmetric affinity games only contain coalitions whose members all have positive affinity. Instead, if agents $x, y \in S_i^*$ and $a(x, y) < 0$, then there must exist other agents in $S_i^* \setminus \{y\}$ such that

$$\sum_{z \in S_i^* \setminus \{y\}} a(x, z) \geq -a(x, y).$$

Informally, if a coalition belongs to the social-welfare maximizing coalition structure, and if two agents in the coalition dislike each other, then there must be other agents in the coalition that the two agents like.

Corollary 1 follows directly from Theorem 2.

Corollary 1 *If $G(A)$ is a symmetric affinity game and $P^* = (S_1^*, \dots, S_k^*)$ is a social-welfare maximizing coalition structure, then $\forall S_i^* \in P^*$, S_i^* is individually rational.*

Proof: Let $x \in S_i^*$. If $|S_i^*| = 1$ then $u(S_i^*, x) = 0$ and so S_i^* is individually rational. Assume that $|S_i^*| > 1$. Then $u(S_i^*, x) = \sum_{(x,y) | y \in S_i^*} a(x, y)$. From Theorem 2, we know that every cut of S_i^* is non-negative. Therefore,

$$\sum_{(x,y) | y \in S_i^*} a(x, y) \geq 0$$

since otherwise we could find a cut $(\{x\}, S_i^* \setminus \{x\})$ which is strictly negative. That is $u(S_i^*, x) \geq 0$. $\square$

4 Stability and Affinity Games

In this section we study the stability of different coalition structures for affinity games, and in particular how the existence of stable coalition structures (where we use the core as our definition of stability) depends on the underlying affinity graph. We first note that for a general affinity game, the core may be empty. Figure 2 is an example of an affinity game for which there is no coalition structure in the core.³ While

³The grand coalition, $(\{x, y, z\})$ is blocked by coalition $\{z\}$, $(\{x\}, \{y\}, \{z\})$ is blocked by $\{x, z\}$, $(\{x\}, \{y, z\})$ is blocked by

for general affinity games, the core may be empty, there are interesting affinity-graph structures for which there are positive results. Our first two results follow immediately from the definition of the utility functions of the agents.

Theorem 3 *Let $G(A)$ be an affinity game. If for all $(x_i, x_j) \in E$, $a(x_i, x_j) \geq 0$, then the grand coalition is in the core.*

Theorem 4 *Let $G(A)$ be an affinity game. If for all $(x_i, x_j) \in E$, $a(x_i, x_j) \leq 0$ then the coalition structure $P = (\{x_1\}, \{x_2\}, \dots, \{x_n\})$ is in the core.*

We contrast Theorem 4 with the transferable-utility case. In particular, if agents have transferable utilities then the core is non-empty if and only if there is no negative cut (Lemma 2) [Deng and Papadimitriou, 1994]. This discrepancy highlights the differences in our models.

If the affinity graph is symmetric then we can further expand our understanding of stable coalition structures. Our next result shows that the social-welfare maximizing coalition structure, while not necessarily belonging in the core, still satisfies an interesting stability property.

Theorem 5 *If affinity game $G(A)$ is symmetric, then the social-welfare maximizing partition, P^* exhibits inner stability.*

Proof: Let $P^* = (S_1^*, \dots, S_k^*)$ be the social-welfare maximizing coalition structure. Assume that there exists S_j^* with blocking coalition $B \subset S_j^*$. Since B blocks S_j^* then for all $x_i \in B$, $u(B, x_i) \geq u(S_j^*, x_i)$ and the inequality is strict for at least one agent. However, since $B \subset S_j^*$, $u(S_j^*, x_i) = u(B, x_i) + \sum_{x_k \in S_j^* \setminus B} a(x_i, x_k)$, and so

$$\begin{aligned} V(S_j^*) &= \sum_{x_i \in S_j^*} u(S_j^*, x_i) \\ &= \sum_{x_i \in B} u(B, x_i) + \sum_{x_i \in B} \sum_{x_k \in S_j^* \setminus B} a(x_i, x_k) \\ &= V(B) + \text{Cut}(B, S_j^* \setminus B) \end{aligned}$$

where $\text{Cut}(B, S_j^* \setminus B)$ is the weight of the cut between B and S_j^* . Therefore, $\text{Cut}(B, S_j^* \setminus B) < 0$. But this is not possible (Theorem 2). Therefore B can not exist, and thus P^* exhibits inner stability. $\square$

5 Stability and Social Welfare

In this section we study the relationship between the core and social welfare. Our first observation is that for affinity games, if a coalition structure, P , is in the core, then this does not imply that it must also be a social-welfare maximizing coalition structure.⁴ Figure 3 illustrates this. The social-welfare maximizing coalition structure is $P^* = (\{x, t\}, \{y, z\})$. However P^* is not in the core when $W > 3$ since $\{x, y\}$ is a blocking coalition. In fact, the coalition structure $P^C =$

$\{y\}$, and coalition structure $(\{x, z\}, \{y\})$ is blocked by $\{y, z\}$ while $(\{y, z\}, \{x\})$ is blocked by $\{x, z\}$.

⁴Researchers on coalitions in other network-based games have also made similar observations [Jackson and Wolinsky, 1996].




Figure 3: An affinity graph where if $W > 3$ then the (non-empty) core does not contain the social-welfare maximizing coalition structure.

$(\{x, y\}, \{t, z\})$ is in the core but $SW(P^C) = 2W + 2 < 4W - 4 = SW(P^*)$. While the core members may not maximize social welfare, and the social welfare maximizing coalition structure may not be in the core, we are still interested in understanding the relationship between the two concepts. In particular, we are interested in understanding the potential loss of social welfare that comes from being in the core. We call this loss the *stability gap*.

Definition 7 *Let $G(A)$ be an affinity game with a non-empty core. Let P^* be the social-welfare maximizing coalition structure, and let P^C be a member of the core. The stability gap of P^C is*

$$\text{Gap}(P^C) = \frac{SW(P^*)}{SW(P^C)}.$$

If $\text{Gap}(P^C) = 1$ then P^C is a social-welfare maximizing coalition structure. If $\text{Gap}(P^C) > 1$ then P^C sacrifices social welfare in exchange for stability. For a given affinity game $G(A)$ we are particularly interested in measuring the stability gap of the member of the core with the *lowest* social welfare,

$$\text{Gap}_{\min}(A) = \frac{SW(P^*)}{\min_{P \in \text{Core}(A)} SW(P)}$$

and the stability gap of the member of the core with the *highest* social welfare

$$\text{Gap}_{\max}(A) = \frac{SW(P^*)}{\max_{P \in \text{Core}(A)} SW(P)}.$$

Clearly, $\text{Gap}_{\min}(A) \geq \text{Gap}_{\max}(A)$ for any affinity graph with a non-empty core. We also note that $\text{Gap}_{\min}(A)$ has parallels with the *price of anarchy* while $\text{Gap}_{\max}(A)$ has parallels with the *price of stability* [Nisan et al., 2007].

We start by looking at general affinity graphs. Unfortunately, our first result is negative.

Theorem 6 *Let $G(A)$ be an affinity game with a non-empty core. Then, $\text{Gap}_{\max}(A)$ can be unbounded.*

Proof: Consider the graph $A = (N, E)$ where $N = \{x_0, x_1, \dots, x_{n-1}\}$. Let

$$E = \{(x_0, x_i) | 1 \leq i < n\} \cup \{(x_i, x_0) | 1 \leq i < n\}.$$

That is, A is a star with x_0 as the center. Now, let $a(x_0, x_1) = a(x_1, x_0) = 1$ and for all x_i such that $1 < i < n$ let $a(x_0, x_i) = -1$ and $a(x_i, x_0) = W$ for some $W \geq 1$. The social-welfare maximizing coalition structure P^* is the

grand coalition, and $SW(P^*) = (n-2)(W-1) + 2$. The only coalition structure in the core, however, is $P^C = (\{x_0, x_1\}, \{x_2\}, \dots, \{x_{n-1}\})$, and $SW(P^C) = 2$. Therefore, for arbitrary $W > 1$,

$$\text{Gap}_{\max}(A) = \frac{SW(P^*)}{SW(P^C)} = \frac{(n-2)(W-1) + 2}{2}.$$

□

While Theorem 6 is distressing since it states that even core members with the highest social welfare can still be arbitrarily worse than the maximum social welfare, if we place *some* restrictions on the affinity graph, then the sacrifice in terms of social welfare is significantly reduced.

Theorem 7 *Let $G(A)$ be a symmetric affinity game with a non-empty core. Then $\text{Gap}_{\min}(A)$ is bounded by 2, and this bound is tight.*

Proof: Let P^C be a member of the core of affinity graph A . We will show that $SW(P^C) \geq \frac{1}{2}SW(P)$ for any coalition structure $P = (S_1, \dots, S_m)$. Before starting the proof, we need to introduce the notation used in the rest of the proof. First, for agents $i, j \in N$ define

$$e^*(x_i, x_j) = \begin{cases} a(x_i, x_j) & \text{if } x_i, x_j \in S_k \text{ for some } S_k \in P \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$SW(P) = 2 \sum_{1 \leq i < j \leq n} e^*(x_i, x_j) \quad (1)$$

since it is the sum of all weights of edges inside each coalition of P (and does not count the cut edges between coalitions).

Second, for agents $\{x_1, x_2, \dots, x_t\}$ let $P - x_1 - \dots - x_t$ denote a new coalition structure

$$P' = (S_1 \setminus \{x_1, \dots, x_t\}, \dots, S_m \setminus \{x_1, \dots, x_t\}, \{x_1\}, \dots, \{x_t\}).$$

We say that a coalition structure $P_1 = (C_1, \dots, C_k)$ *blocks* another coalition structure $P_2 = (S_1, \dots, S_m)$ if there exists some $C_i \in P_1$ such that $\forall x \in C_i, u(C, x) \geq u(S_j, x)$ where $x \in S_j$.

Finally, we will abuse notation and for any partition $P = (S_1, \dots, S_m)$ and agent $x \in S_i$, we say $u(P, x) = u(S_i, x)$.

We can now start the proof. Since P^C is in the core, by definition of the core it follows that it is not blocked by P . Therefore, there exists at least one agent, say x_1 , such that $u(P, x_1) \leq u(P^C, x_1)$. Now, consider coalition structure $P - x_1$. $P - x_1$ does not block P^C and so there must exist some agent x_2 such that $u(P - x_1, x_2) \leq u(P^C, x_2)$ by definition of the core. We can continue this process and iteratively remove agents in the order $x_1, x_2, \dots, x_{n-1}$. Each new coalition structure does not block P^C since P^C is in the core. Therefore, we get the following inequalities;

$$\begin{aligned} u(P, x_1) &\leq u(P^C, x_1) \\ u(P - x_1, x_2) &\leq u(P^C, x_2) \\ &\vdots \\ u(P - x_1 - \dots - x_{n-1}, x_n) &\leq u(P^C, x_n) \end{aligned}$$

Summing, we get

$$\begin{aligned} u(P, x_1) + \sum_{i=2}^n u(P - x_1 - \dots - x_{i-1}, x_i) &\leq \sum_{i=1}^n u(P^C, x_i) \\ &= SW(P^C). \end{aligned}$$

Since

$$u(P - x_1 - \dots - x_{i-1}, x_i) = u(P, x_i) - \sum_{j=1}^{i-1} e^*(x_i, x_j)$$

we have

$$\sum_{i=1}^n u(P, x_i) - \sum_{1 \leq i < j \leq n} e^*(x_i, x_j) \leq SW(P^C)$$

and so

$$SW(P) - \sum_{1 \leq i < j \leq n} e^*(x_i, x_j) \leq SW(P^C).$$

Substituting in Equation 1 we get

$$\frac{1}{2}SW(P) \leq SW(P^C).$$

Since this holds for all coalition structures, it must also hold for the social-welfare maximizing coalition structure, P^* .

To show that the bound is tight, consider the example in Figure 3. Recall that the coalition structure $P^* = (\{x, t\}, \{y, z\})$ maximizes welfare with $SW(P^*) = 4W - 4$. The coalition structure $P^C = (\{x, y\}, \{t, z\})$ is in the core, and $SW(P^C) = 2W + 2$. The stability gap of P^C is $\frac{4W-4}{2W+2}$, which converges to 2 as $W \rightarrow \infty$. □

6 Experiments

Section 5 showed that for symmetric affinity games, the sacrifice in social welfare in order to achieve stability was bounded by a constant (2), while in general affinity games it could be unbounded. These results relied on the assumption that the core was non-empty. In this section we describe a series of experiments that we conducted on randomly generated symmetric affinity games, to answer the following questions: *i)* How often is the core non-empty? and *ii)* What are $\text{Gap}_{\max}(A)$ and $\text{Gap}_{\min}(A)$ in practice?

To generate random symmetric affinity games we generated graphs with $|N|$ ranging from 4 to 10. For each edge in the graph we assigned an integer weight from interval $[-W, W]$, chosen uniformly at random. We generated 1000 graphs for each W . In the rest of this section we show results for $W \in \{1, 5, 25, 125, 700\}$ since these were illustrative of general trends, irrespective of the number of agents. For each graph we exhaustively searched the space of coalition structures to find the social-welfare maximizing coalition structure, and if the core was non-empty, we found the member of the core with the highest social welfare, and with the lowest social welfare.

In our first set of experiments we studied the frequency with which the core and weak core actually exists. Table 1 presents our findings. We note that the weak core was never empty, and that core was *non-empty* for a significant number of the random affinity games.

	Core	Weak Core
W	% non-empty	% non-empty
1	73.3%	100%
5	93.3%	100%
25	99.1%	100%
125	99.7%	100%
700	100%	100%

Table 1: The percentage of graphs which had a non-empty core and weak core. Note the weak core was always non-empty.

	Core		Weak Core	
W	$\text{Gap}_{\max}$	$\text{Gap}_{\min}$	$\text{Gap}_{\max}$	$\text{Gap}_{\min}$
1	1.0017	1.0159	1.0000	1.6886
5	1.0089	1.1210	1.0019	1.3070
25	1.0104	1.1841	1.0063	1.2161
125	1.0090	1.1894	1.0091	1.2008
700	1.0100	1.1936	1.0065	1.2005

Table 2: Average values for $\text{Gap}_{\max}(A)$ and $\text{Gap}_{\min}(A)$ for games where the core and weak core were non-empty.

In our next set of experiments we restricted ourselves to affinity games where the core is non-empty. Table 2 contains our findings. We note that the stability gap is always very close to one, indicating that the sacrifice with respect to social welfare in order to gain stability is actually very low, and for many instances we found that the core and weak core did contain the social-welfare maximizing coalition structure.⁵ In the rest of this section we propose a hypothesis as to why this was so.

In Theorem 5 we proved that for any symmetric affinity game, the social-welfare maximizing coalition structure, $P^* = (S_1^*, \dots, S_k^*)$, satisfies inner stability. This means that in order for P^* to not belong in the core, there must exist a blocking coalition, B , that draws its members from different coalitions in P^* . Thus, these agents, once inside B , derive their utility from edges that belong to the minimum cut. Since we allow negative affinities, then it is likely that many of the edges in the minimal cut will also be negative while edges inside coalitions in P^* will likely be positive. Therefore, it is unlikely that agents in B will actually be able to improve their utility compared to if they stay in S_i^* . Thus, we hypothesize that on many randomly generated graphs, the existence of such a blocking coalition, B , is rare.

7 Conclusion and Future Work

In this paper we introduced *coalitional affinity games*, a family of non-transferable utility games that explicitly model the values that agents receive from being associated with others. In these games, an agent is interested in joining coalitions with agents for which it has high affinity, while avoiding coalitions that contain agents for which it has low affinity.

⁵For both the core and weak core, the standard deviation for $\text{Gap}_{\max}$ was always less than 0.05 and for $\text{Gap}_{\min}$ it was always less than 0.2.

Given our model, we provided a characterization of the social-welfare maximizing coalition structure, and showed that it corresponds to the minimal k -cut of the underlying affinity graph. We then studied stability properties of affinity games using the core as the solution concept. We investigated the relationship between the core and social welfare and proposed a new measure, the *stability gap*, that represents the sacrifice in social welfare that is made in exchange for stability. We showed that for general affinity games, this sacrifice is unbounded. However, for a subclass of affinity games (symmetric affinity games), we showed that the bound is two.

There are several research directions we would like to pursue. First, our experimental results illustrated that for symmetric affinity games, the core was rarely empty. We are interested in general properties of random graphs and would like, if possible, to prove that with high probability, the core is non-empty. Second, our current model only captures affinities between pairs of agents. We would like to extend this so that affinities among *groups* of agents could be captured, thus modelling all hedonic games. Preliminary work using *hypergraphs* looks promising.

References

- [Bachrach and Rosenschein, 2008] Yoram Bachrach and Jeffrey S. Rosenschein. Coalitional skill games. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 1023–1030, Estoril, Portugal, 2008.
- [Bogomolnaia and Jackson, 2002] Anna Bogomolnaia and Matthew O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38:201–230, 2002.
- [Deng and Papadimitriou, 1994] Xiaotie Deng and Christos H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.
- [Drèze and Greenberg, 1980] J. H. Drèze and J. Greenberg. Hedonic coalitions: Optimality and stability. *Econometrica*, 48(4):987–1003, 1980.
- [Jackson and Wolinsky, 1996] Matthew O. Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71:44–74, 1996.
- [Nisan et al., 2007] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [Peleg and Sudhölter, 2003] Bezalel Peleg and Peter Sudhölter. *Introduction to the Theory of Cooperative Games*. Springer, 2003.
- [Rahwan et al., 2007] Talal Rahwan, Sarvapali D. Ramchurn, Viet Dung Dang, and Nicholas R. Jennings. Near-optimal anytime coalition structure generation. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 2365–2371, Hyderabad, India, 2007.