Algorithmic Game Theory	**Winter 2022/23**

Lecture 26: Hedonic Games

Lecturers: Prof. Dr. Martin Hoefer, Dr. Giovanna Varricchio *Scribes: Dr. Giovanna Varricchio*

Summary: In this lecture we discuss Hedonic Games. Here some selfish agents have to gather to form coalitions and they have preferences among coalitions the might be involved to. We will have a short overview on classical solutions concepts like core and Nash stability. We then focus on specific classes of Hedonic Games where every agent classify any other as a friend or an enemy and we discuss about the existence of core stable outcomes.

Resources:

- *Handbook of computational social choice.* F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A.D. Procaccia, 2016, Cambridge University Press (Chapther 15, Hedonic Games).
- *Simple priorities and core stability in hedonic games.* D. Dimitrov, P. Borm, R. Hendrickx, S.C. Sung.
- *Enemies and friends in hedonic games: individual deviations, stability and manipulation.* D. Dimitrov, S.C. Sung.

1 Setting

Hedonic Games (HGs) represent a subclass of Coalition Formation Games where agents have *hedonic preferences*, that is, their preferences depend only on the coalition they are involved in and not on how the other agents aggregate. Formally, we are given a set $N = \{1, \dots, n\}$ of selfish agents, and the goal is to partition them into a collection of disjoint coalitions $\mathcal{C} = \{C_1, \dots, C_m\}$ such that $\cup_{i=1}^m C_i = N$. Such a partition is also called an *outcome* or a *coalition structure*.
The *grand coalition* GC is a coalition structure that consists of only one coalition containing all the agents and a *singleton coalition* is any coalition of size 1. We denote by $\mathcal{O}$ the set of all possible outcomes and by $\mathcal{C}(i)$ the coalition that agent i belongs to in outcome $\mathcal{C} \in \mathcal{O}$. Each agent i has a *preference relation*, or simply called *preference*, $\succeq_i$ over $N_i = \{S \subseteq N \,|\, i \in S\}$ the family of subsets of N containing i. According to $\succeq_i$, for every $X, Y \in N_i$, we say that agent i prefers or equally prefers X to Y whenever $X \succeq_i Y$, and we say that agent i equally prefers X to Y whenever $X \sim Y$. We are now ready to define an HG instance.

Definition 1 (Hedonic Games)**.** *A* Hedonic Game instance *$\mathcal{I}$ is given by a pair (N, P), where N is the set of agents and $P = \{\succeq_i\}_{i\in N}$ is the* preference profile*, that is, the collection of agents' preferences.*

According to the properties satisfied by P different types of HGs can be defined.
When preferences are in a cardinal form they are described through a *utility function*, instead of a preference relation, that is, in general, a map $u_i : N_i \to \mathbb{R}$ in HGs, for each $i \in N$. Without loss of generality, we assume the preference profile P to be either the collection of agents preferences $\{\succeq_i\}_{i\in N}$ or the collection of agents utilities $\{u_i\}_{i\in N}$, depending on the game instance.
A well-known class of HGs with cardinal valuations is the one of *Additively Separable Hedonic Games.*

Definition 2 (Additively Separable Hedonic Games)**.** *An HG (N, P) is additively separable (ASHG) if for each $i \in N$ there exists a mapping $w_i : N \setminus \{i\} \to \mathbb{R}$, where $w_i(j)$, or simply $w_{i,j}$, is the value that agent i gives to agent j. Given a coalition structure $\mathcal{C}$ the utility agent i achieves in the coalition $\mathcal{C}(i)$ is given by $u_i(\mathcal{C}(i)) = \sum_{j\in\mathcal{C}(i)\setminus\{i\}} w_{i,j})$. Clearly, $u_i(\{i\}) = 0$.*

2 Solution Concepts

Traditionally, in HGs the focus has been put onto the existence of several solution concepts, based either on individual or group deviations.

Stability Concepts based on Single-Agent Deviations An outcome $\mathcal{C} \in \mathcal{O}$ for an instance $\mathcal{I} = (N, P)$ is said to be:

- *Individually rational*, if no agent benefits by moving from her current coalition in $\mathcal{C}$ and forming a new coalition by herself.
- *Nash stable*, if no agent benefits by either moving from her current coalition in $\mathcal{C}$ to an already formed one or forming a new coalition by herself.

Notice that Nash stability implies individually rational.

Stability Concepts based on Group Deviations An outcome $\mathcal{C} \in \mathcal{O}$ for an instance $\mathcal{I}$ is said to be:

- *Core stable*, if no subset of agents can (strictly) improve the utility of each member by forming together a new coalition, i.e., there is no strong core-blocking coalition.
- *Strictly core stable*, if no subset of agents can improve the gain of at least one member and not decrease for the others by forming together a new coalition, i.e., there is no core-blocking coalition.

Clearly, any strictly core stable outcome is also core stable and any core stable is individually rational.
Let us start by noticing that (strictly) core stable outcomes may not exist for ASHGs.

Proposition 1. *The (strict) core of an ASHG might be empty.*

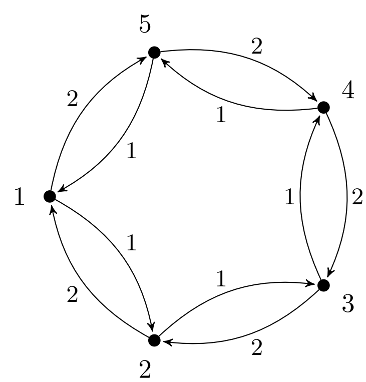

Figure 1: The core of an ASHG may be empty. All hidden edges have value -10.

Proof. Consider the instance depicted in Figure 1 where edge weights represent individual values. First, observe that no coalition of size ≥ 3 is core stable, since the utility of at least one agent, say i, is negative and therefore $\{i\}$ is a blocking coalition.
Second, the coalition structure where two consecutive nodes are in a singleton, namely $\{i\}$, $\{j\}$, is not stable as $\{i, j\}$ is a blocking coalition.
Hence, the only possible stable outcome is made by a singleton and two pairs of consecutive nodes.
Assume without loss of generality (by symmetry, the same arguments apply to any coalition structure with the very same properties) that the possible stable outcome is $\{1\}, \{2, 3\}, \{4, 5\}$. However, $\{1, 2\}$ is a blocking coalition.
In conclusion, no core stable outcome exists. $\square$

In turn, if all the weights are positive then a strict core stable outcome always exists.

Proposition 2. *For any ASHG with positive weights, that is, $w_{i,j} \geq 0$ for each $i \neq j$, there always exists a strictly core stable coalition structure.*

Proof. Consider $\mathcal{C} = \{\{1, \ldots, n\}\}$, that is, the agents form the grand coalition. Such a partition cannot be core blocked by any coalition $S \subseteq N$ as every agent has maximum possible utility. Therefore, the grand coalition is strictly core stable. □

In what follows, we focus on the existence of (strictly) core stable outcomes in special subclasses of ASHGs.

3 Friends and Enemies games

In *Friends and Enemies games* every agent $i \in N$ partitions the other agents into a *set of friends* F_i and a *set of enemies* E_i, with $F_i \cup E_i = N \setminus \{i\}$ and $F_i \cap E_i = \emptyset$.

Example 1. *Let $N = \{1, 2, 3\}$ be the set of agents, and let $F_1 = \{2\}$, $F_2 = \{3\}$, $F_3 = \{2\}$ and $E_1 = \{3\}$, $E_2 = \{1\}$, $E_3 = \{1\}$ be the agents' sets of friends and enemies, respectively. The sets N, F_i and E_i, $i = 1, 2, 3$, define an instance of friends and enemies game. The just described instance is depicted in Figure 2a, where a directed edge from agent i to agent j represents i's opinion of j. Moreover, solid edges and dashed edges represent friend and enemy relations, respectively.*

Given such a friend and enemy partition, different settings can be defined.
A preference profile P is based on *Friends Appreciation* (FA) when $X \succeq_i Y$, for $X, Y \in N_i$, iff

$$|X \cap F_i| > |Y \cap F_i| \qquad \text{or} \qquad |X \cap F_i| = |Y \cap F_i| \text{ and } |X \cap E_i| \leq |Y \cap E_i|,$$

and it is based on *Enemies Aversion* (EA) when $X \succeq_i Y$, for $X, Y \in N_i$, iff

$$|X \cap E_i| < |Y \cap E_i| \qquad \text{or} \qquad |X \cap E_i| = |Y \cap E_i| \text{ and } |X \cap F_i| \geq |Y \cap F_i|.$$

In other words, under FA, a coalition is preferred over another one if it contains a higher number of friends; if the number of friends is the same, the coalition with less enemies is preferred. On the other hand, under EA, a coalition is preferred if it contains a smaller number of enemies; if the number of enemies is the same, the coalition with more friends is preferred.
Friends and Enemies Games are a proper subclass of ASHGs, because each agent i can be seen as having a suitable valuation $w_i(j)$ for every other agent j (according to the fact that she is "friend" or an "enemy"), and her utility for being in a given coalition C is $u_i(C) = \sum_{j \in C \setminus \{i\}} w_i(j)$. In the FA case, such valuation functions for each agent $i \in N$ can be defined as

$$w_i(j) = \begin{cases} n & \text{, if } j \in F_i, \\ -1 & \text{, if } j \in E_i, \end{cases}$$

which, as it can be easily checked, correctly encodes the setting. In other words, the positive effect of one friend is greater than the overall possible negative effect due to enemies. Similarly, in the EA case, the valuations can be set in such a way that, for every agent $i \in N$,

$$w_i(j) = \begin{cases} 1 & \text{, if } j \in F_i, \\ -n & \text{, if } j \in E_i. \end{cases}$$

Example 2. *Let us consider again the instance described in Example 1 and the grand coalition $GC = \{\{1, 2, 3\}\}$ as a possible outcome. Then, under FA preferences profiles, $u_1(GC) = w_1(2) + w_1(3) = 3 - 1 = 2$. Similarly, $u_2(GC) = u_3(GC) = 3$. Under EA preferences profiles, instead, $u_1(GC) = w_1(2) + w_1(3) = 1 - 3 = -2$, and $u_2(GC) = u_3(GC) = -2$.*

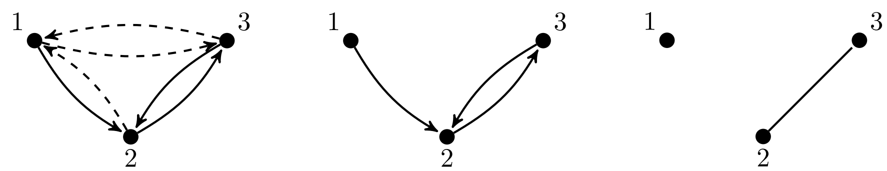

(a) Instance (b) $G^{\rightarrow}$ (c) $G^{\leftrightarrows}$

Figure 2: Example of a Friends and Enemies Game instance and the corresponding graphs $G^{\rightarrow}_{\mathbf{d}}$ and $G^{\leftrightarrows}_{\mathbf{d}}$. Solid (resp. dashed) edges represent friend (resp. enemy) relations.

Direct revelation of friends and enemies. In what follows, we assume that friend and enemy relationships are declared by the agents themselves. We denote by $\mathbf{d}$ the declaration of the agents, this will define the profile P and therefore the instance. In other words, we are assuming that friends and enemies are private information of the agents, and hence they could misreport this information if profitable for them. We will be back to this while computing stable outcomes. Furthermore, since for any i and any $j \neq i$, j is for i either a friend or an enemy, it is sufficient to ask the agents about the set of friends which will be denoted by F_{d_i}.

3.1 Graph Representation

As already mentioned, Friends and Enemies Games, being a proper subclass of ASGHs, can be suitably represented by means of graphs. More specifically, the following representations will be useful for our purposes. For a given profile $\mathbf{d}$, $G^{\rightarrow}_{\mathbf{d}} = (N, F_{\mathbf{d}})$ is a directed graph with the edge set

$$F_{\mathbf{d}} = \{(i,j) \mid i,j \in N,\ j \in F_{d_i}\} \quad ,$$

i.e., $G^{\rightarrow}_{\mathbf{d}}$ contains only directed edges corresponding to friendship relations;
$G^{\leftrightarrows}_{\mathbf{d}} = (N, F^{\leftrightarrows}_{\mathbf{d}})$ is an undirected graph with the edge set

$$F^{\leftrightarrows}_{\mathbf{d}} = \big\{\{i,j\} \,|\, i,j \in N,\ j \in F_{d_i} \wedge i \in F_{d_j}\big\} \quad ,$$

i.e., each edge corresponds to a mutual friendship relation. Graphs $G^{\rightarrow}_{\mathbf{d}}$ and $G^{\leftrightarrows}_{\mathbf{d}}$ for the instance given in Example 1 can be seen in Figure 2b and 2c, respectively.
Given a directed graph G and a pair of nodes x, y in G, x and y are *weakly connected* if they are connected in the undirected version of G. Moreover, a *weakly connected component* in G is a maximal subset of pairwise weakly connected nodes in G. Similarly, given a directed graph G and a pair of nodes x, y in G, x and y are *strongly connected* if there exists a directed path from x to y and a directed path from y to x in G. A *strongly connected component* in G is a maximal subset of strongly connected nodes in G.

4 Friends oriented preferences

Given an instance of FA with agents declarations $\mathbf{d}$, let us compute a partition of agents as follows.

Strongly connected components partition (SCCP)

- Build $G^{\rightarrow}_{\mathbf{d}}$
- Compute the maximal strongly connected components of $G^{\rightarrow}_{\mathbf{d}}$
- Create a coalition for each strongly connected component

Theorem 3. *The SCCP algorithm is IC. Moreover, the outcome is in the strict core.*

Proof. (IC) Consider the strongly connected components of $G_{\mathbf{d}}^{\rightarrow}$ without i, let's call them $C_1, \ldots, C_k$. Notice that these strongly connected components do not depend on the declaration of i
We will show the mechanism is monotone w.r.t. F_i, that is, declaring a larger set of friends cannot decrease the utility of i.
Notice that whatever is the declaration of i in the outcome $\mathcal{C}$ of the mechanisms, $\mathcal{C}(i)$ either contains C_h or they are disjoint. In particular, since C_h and $\mathcal{C}(i)$ are strongly connected, $C_h \subset \mathcal{C}(i)$ if and only $F_i \cap C_h \neq \emptyset$ and exists $j \in C_h$ such that $i \in F_j$. Therefore, if i increases the number of friends she has in C_h she can only increase the number of friends in her coalition. Since it is always preferable to have one more friend, no matter how many enemies we consequently include, increasing the number of friends in C_h does not decrease the utility of i. Applying the same argument on each $C_1, \ldots, C_k$ shows the SCCP algorithm is IC.
(Strict core) Let $\mathcal{C}$ be the outcome of the algorithm. We first notice that every agent has utility ≥ 0. Therefore, any core-blocking coalition has no sink node in the graph induced by the coalition otherwise the sink node (agent) has negative utility in the core-blocking coalition – a contradiction. Hence, our core-blocking coalition is made by strongly connected groups of agents and directed paths ending into a group of strongly connected agents. Such a coalition S actually cannot core-block the outcome of SCCP. If any of the groups of connected agents is not maximal it means there exists an agent in S who is losing a friend by deviating. Therefore, we can assume that each group of strongly connected agents is maximal. Hence, each agent in S has the same number of friends as in $\mathcal{C}$. Furthermore, if there is an agent in S, say j who is not strongly connected to another agent in S, say i, then in S agent i has a higher number of enemies than in $\mathcal{C}(i)$ and hence lower utility, which contradicts the fact that S is core-blocking. The only possibility is to have S is a strongly connected component wich means that there is no agent in S who is strictly increasing her utility. ☐

5 Enemies oriented preferences

Given an instance of EA with agents declarations $\mathbf{d}$, let us compute a partition of agents as follows.

(Greedy) maximum clique covering partition (MCCP)

- Build $G_{\mathbf{d}}^{\leftrightarrows}$
- While $G_{\mathbf{d}}^{\leftrightarrows} \neq \emptyset$
  - Extract a maximum clique from $G_{\mathbf{d}}^{\leftrightarrows}$ and create a coalition
  - Update $G_{\mathbf{d}}^{\leftrightarrows}$

REMARK! Computing a maximum clique is an NP-hard problem.

Theorem 4. *The GMCP outputs a core stable partition.*

Proof. Let us denote by $C_1, \ldots, C_k$ the computed coalitions. Let us assume that C_h is the clique built during the h-th iteration of the while loop.
First notice that every agent has utility ≥ 0 in the computed coalition structure.
Let be i in C_1. The only way to create a blocking coalition containing i is to create a coalition S where i has a stricter number of friends. Moreover, to be a strictly core-blocking coalition it must hold that every agent has positive utility. This implies that S is a coalition inducing a clique in $G_{\mathbf{d}}^{\leftrightarrows}$ and since i has larger utility we have that $|S| > |C_1|$ – a contradiction to the choice of C_1. In conclusion, no agent in C_1 can be involved in any core-blocking coalition.
Let us proceed inductively.
Let us assume that no agent in $C_1 \cup \ldots \cup C_{h-1}$ is involved in a strictly core-blocking coalition, let us show that it holds true for every agent in C_h. The arguments are similar to the previous case.

Let $i \in C_h$ and S a blocking coalition containing i. By the inductive hypothesis, $S \subseteq N \setminus \{C_1 \cup \ldots \cup C_{h-1}\}$. Since S is a strictly core-blocking coalition, i has to increase the number of friends wrt C_h and S has to be a clique strictly larger than C_h, otherwise, at least one agent has negative utility or i is not increasing her utility, but this is a contradiction to the choice of C_h. In fact, C_h is a maximum clique of agents in $S \subseteq N \setminus \{C_1 \cup \ldots \cup C_{h-1}\}$. Repeating the same argument on all h the thesis follows. □

Is the algorithm IC? As for kidney exchange, it depends on the implementation. If we set up a lexicographic order on cliques of the same size it becomes IC. Specifically, a clique may be represented as a binary n-tuple, the i-th component is equal to 1 if and only if i is in the clique. According to this tuple, we can define a lexicographic order among cliques of the same size and always select the one lexicographically larger. Formally, assume we have two n-tuples $x = (x_1, \ldots, x_n)$ and $y = (y_1, \ldots, y_n)$ with the same number of 1 entries (which means they are cliques of the same size), we then have that $x \succ_{lex} y$ if in the first entry they differ, let's say in the k-th entry, we have $x_k \geq y_k$. In other words, we give priorities to maximum cliques containing agent 1, if any, in case of ties we give priority to the ones containing agent 2, if any, and so forth.