Social Distance Games

Alkida Balliu, Michele Flammini, Giovanna Melideo, **Dennis Olivetti**

Aalto University

Nash Stability in Social Distance Games, AAAI 2017
On Pareto Optimality in Social Distance Games, AAAI 2017

Slides freely inspired by “Social Distance Games” of Simina Brânzei and Kate Larson

Algorithmic Game Theory

Study how networks, not controlled by a single entity, behave:

- Individual nodes can make autonomous decisions
- Input is divided among many rational players
- Design algorithms in strategic environments
- Concerned with the computational questions that arise in game theory

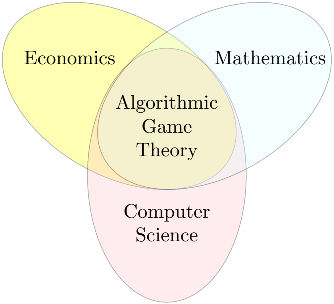

Coalition formation games

Context:

- Agents are part of a social network
- Agents must divide in disjoint groups
- Agents prefer to form groups with their friends
- Agents' interactions are constrained by an underlying network

Questions:

- What structure appears in social networks?
- How much the optimum for the society and a stable solution can differ?

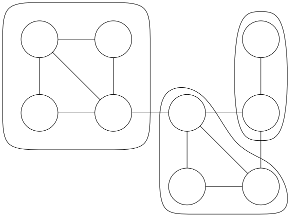

Outline

- Model
- Social welfare in Social Distance Games
- Stability in Social Distance Games
  - Core
  - Pareto
  - Nash

Notation

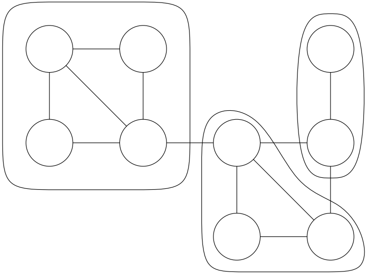

- **Social graph**: $G = (V, E)$, nodes represent agents, edges represent preferences
- **Coalition structure**: a partition $\mathcal{P} = \{P_1, \ldots, P_k\}$ of V into disjoint coalitions
- **Grand coalition**, a single coalition N containing all nodes
- **Non transferable utility**: The utility of each agent depends on the agent and the structure of its coalition

Social Distance Games

The utility of agent i in the coalition C is:

$$u(i, C) = \frac{1}{|C|} \sum_{j \in C \setminus \{i\}} \frac{1}{d_C(i, j)}$$

where $d_C(i, j)$ is the distance between agents i and j in the subgraph of G induced by the coalition C.

Social Distance Games

The utility of agent i in the coalition C is:

$$u(i, C) = \frac{1}{|C|} \sum_{j \in C \setminus \{i\}} \frac{1}{d_C(i, j)}$$

where $d_C(i, j)$ is the distance between agents i and j in the subgraph of G induced by the coalition C.

- The utility is the harmonic centrality of the agent, divided by the size of the coalition.

Social Distance Games

The utility of agent i in the coalition C is:

$$u(i, C) = \frac{1}{|C|} \sum_{j \in C \setminus \{i\}} \frac{1}{d_C(i, j)}$$

where $d_C(i, j)$ is the distance between agents i and j in the subgraph of G induced by the coalition C.

- The utility is the harmonic centrality of the agent, divided by the size of the coalition.
- Harmonic centrality has been elected as the *best* centrality measure for social networks, since it's the only one satisfying a set of desirable properties [Axioms for Centrality, Boldi and Vigna, 2014].

Example

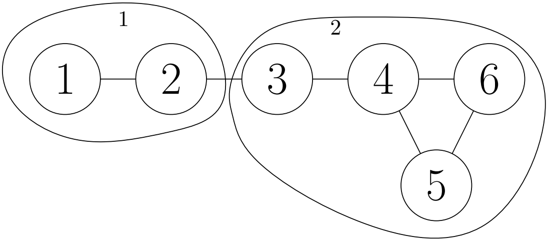

$\mathcal{P} = \{\{1, 2\}, \{3, 4, 5, 6\}\}$

- $u(1) = u(2) = \frac{1}{2}$
- $u(3) = \frac{1+2\cdot\frac{1}{2}}{4} = \frac{1}{2}$
- $u(4) = \frac{3}{4}$
- $u(5) = u(6) = \frac{2+\frac{1}{2}}{4} = \frac{5}{8}$

Some Properties: Singletons

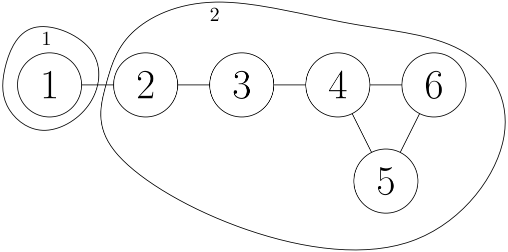

$$u(1) = 0$$

- Singletons always receive zero utility

Some Properties: Adding edges

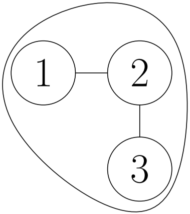

$u(1) = \frac{1}{2}$ $\qquad$ $u(1') = \frac{2}{3}$

- An agent prefers direct connections over indirect ones.

Some Properties: Adding agents

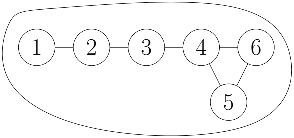

- $u(1) = \frac{1+\frac{1}{2}+\frac{1}{3}+2\cdot\frac{1}{4}}{6} \approx 0.39$

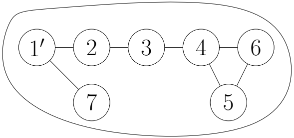

- $u(1') = \frac{2+\frac{1}{2}+\frac{1}{3}+2\cdot\frac{1}{4}}{7} \approx 0.48$
  - Adding a close connection positively affects an agent's utility.

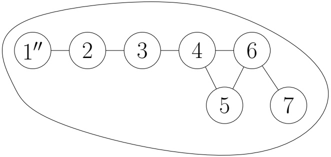

- $u(1'') = \frac{1+\frac{1}{2}+\frac{1}{3}+2\cdot\frac{1}{4}+\frac{1}{5}}{7} \approx 0.36$
  - Adding a distant connection negatively affects an agent's utility.

Social Welfare

The social welfare of a coalition structure $\mathcal{P} = \{P_1, \ldots, P_k\}$ is

$$SW(\mathcal{P}) = \sum_{i=1}^{k} \sum_{j \in P_i} u(j, C_i)$$

Example

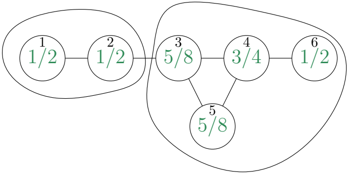

$\mathcal{P} = \{\{1, 2\}, \{3, 4, 5, 6\}\}$

$SW(\mathcal{P}) =$
$3 \cdot \frac{1}{2} + 2 \cdot \frac{5}{8} + \frac{3}{4} = 3.5$

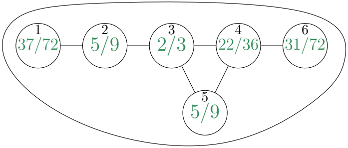

$\mathcal{P} = N = \{\{1, 2, 3, 4, 5, 6\}\}$

$SW(\mathcal{P}) =$
$\frac{37}{72} + 2 \cdot \frac{5}{9} + \frac{2}{3} + \frac{22}{36} + \frac{31}{72} \approx 3.3$

We are interested in social welfare maximizing coalitions
(the overall best result for the society)

Some Properties

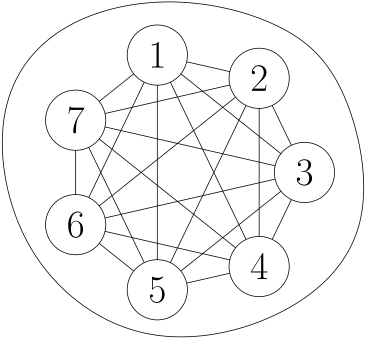

- The social welfare is at most $n - 1$

Some Properties

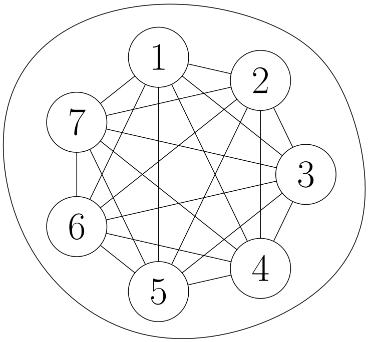

- On cliques, N is the only social welfare maximizing coalition

Some Properties

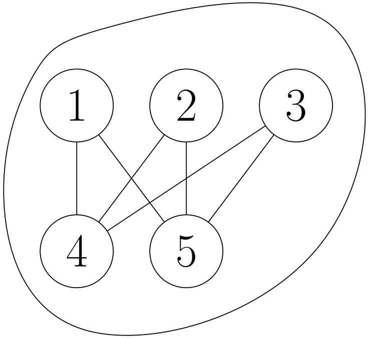

- On complete bipartite graphs, N maximizes the social welfare, and it is fair (guarantees $\frac{1}{2}$ to each agent)

Some known results [Brânzei and Larson, AAMAS'11]

- Finding the optimal social welfare is NP-hard
- We can easily find a $\frac{1}{2}$-approximation, that guarantees fairness.

$\frac{1}{2}$-approximation

Star decomposition:

- Compute a spanning tree
- Split in stars, starting from the leaves

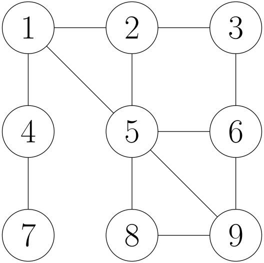

$\frac{1}{2}$-approximation

Star decomposition:

- Compute a spanning tree
- Split in stars, starting from the leaves

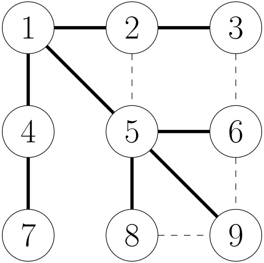

$\frac{1}{2}$-approximation

Star decomposition:

- Compute a spanning tree
- Split in stars, starting from the leaves

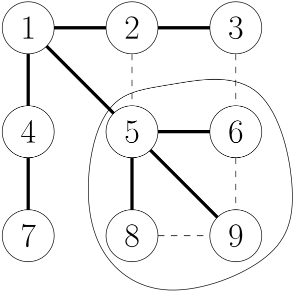

$\frac{1}{2}$-approximation

Star decomposition:

- Compute a spanning tree
- Split in stars, starting from the leaves

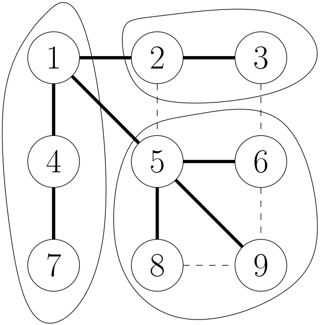

$\frac{1}{2}$-approximation

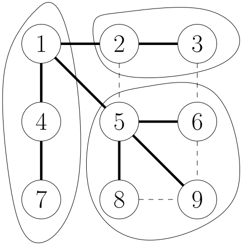

Properties:

- Each leaf has utility $\frac{1}{2}$
- Each center has utility at least $\frac{1}{2}$
- This partition is fair (guarantees half of the optimum to each agent)

Stability of coalitions

Different notions of stability:

- **Core**: a group of agents decide to form a new coalition.
- **Nash**: an agent decides to move to a different coalition.
- **Pareto**: all agents can simultaneously deviate and form new coalitions, and no agent loses utility.

Core stability

A coalition $\mathcal{P} = \{P_1, \ldots, P_k\}$ is core stable if there is no coalition $B \subseteq N$ such that $\forall x \in B$, $u(x, B) \geq u(x, \mathcal{P})$, and for some $x \in B$ the inequality is strict.

Example

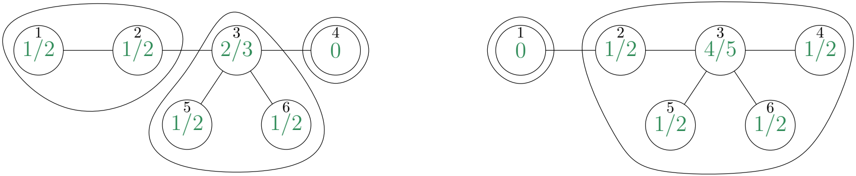

- Agents $\{2, 3, 4, 5, 6\}$ can deviate and form a new coalition, they do not lose utility, and agents $\{3, 4\}$ increase their utility.
- In the new coalition structure, $u(1)$ is zero.
- The new coalition structure is also unstable: $\{1, 2\}$ can deviate

Some known results [Brânzei and Larson, AAMAS'11]

For some games, the core is empty.

Any stable coalition has diameter at most 14

Core Stability vs Social Welfare

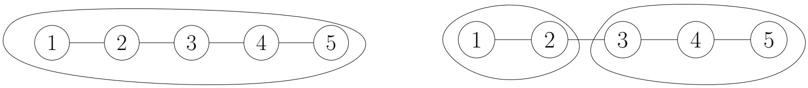

- The only stable coalition structure is N
  - $SW = 2 \cdot \frac{1+\frac{1}{2}+\frac{1}{3}+\frac{1}{4}}{5} + 2 \cdot \frac{2+\frac{1}{2}+\frac{1}{3}}{5} + \frac{2+2\cdot\frac{1}{2}}{5} \approx 2.56$
- The social welfare is maximized by $\{\{1,2\},\{3,4,5\}\}$
  - $SW = 4 \cdot \frac{1}{2} + \frac{2}{3} \approx 2.66$
- Social welfare maximizing structures are not always stable

The Core Stability Gap

Let:

- G be an arbitrary graph for Social Distance Games
- $\mathcal{P}^*$ be a social welfare maximizing coalition structure
- $\mathcal{P}$ be a core stable coalition structure

The stability gap is:

$$Gap(G) = \frac{SW(\mathcal{P}^*)}{min_{\mathcal{P} \in Core(G)} SW(\mathcal{P})}$$

The Core Stability Gap [Brânzei and Larson, AAMAS'11]

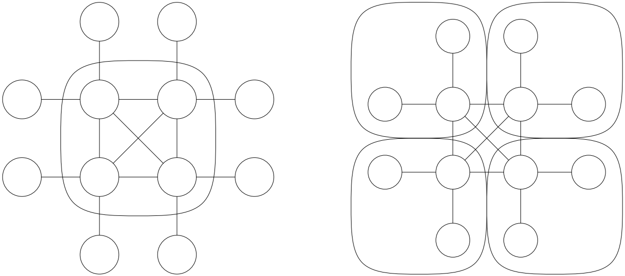

$Gap(G)$ in the worst case is $\Theta(\sqrt{n})$.

$Gap(G) < 4$ if it is not allowed to leave isolated agents while deviating (*no man left behind* policy).

Pareto Stability

A coalition $\mathcal{P} = \{P_1, \ldots, P_k\}$ is Pareto stable if it does not allow a simultaneous deviation by all the agents that makes all agents weakly better off and some agents strictly better off.

Example: Invalid deviation

- Core stability would allow this deviation
- Pareto stability does not allow this deviation.
- $\{\{1, 2\}, \{3, 4, 5\}\}$ is stable.

Example: Valid deviation

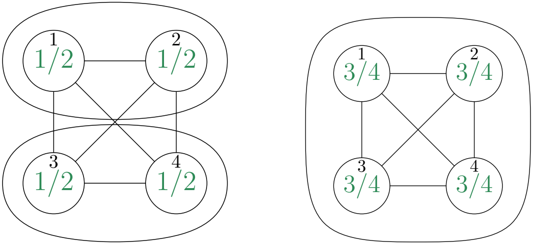

- If agents deviate, nobody loses utility
- $\{\{1,2\},\{3,4\}\}$ is not stable, $\{\{1,2,3,4\}\}$ dominates it

Pareto stability

- No agent can be abandoned
- The coalition structure maximizing the social welfare is always stable
- How far the social welfare of a Pareto stable coalition structure can be from the optimum?

Price of Pareto Optimality Definition

Let:

- G be an arbitrary graph for Social Distance Games
- $\mathcal{P}^*$ be a social welfare maximizing coalition structure
- $\mathcal{P}$ be a Pareto stable coalition structure

The Price of Pareto Optimality is:

$$PPO(G) = \frac{SW(\mathcal{P}^*)}{min_{\mathcal{P} \in PPO(G)} SW(\mathcal{P})}$$

Price of Pareto Optimality

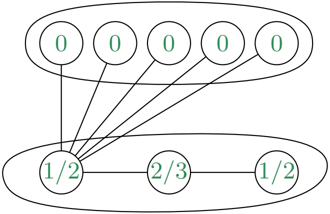

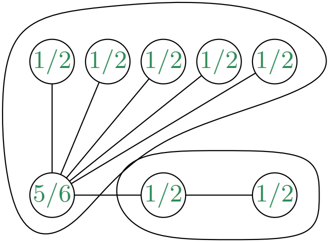

- The social welfare is $O(1)$
- This coalition structure is stable

- The social welfare is $\Theta(n)$

$$PPO = \Theta(n)$$

Price of Pareto Optimality Results

Undirected	Unweighted	Weighted
General	$\Theta(n)$	$\Theta(nW)$
Δ-bounded degree	$\Theta(\Delta)$	$\Omega(\Delta W)$, $O(\min(nW, \Delta W^2))$

Directed	Unweighted	Weighted
General	$\Theta(n)$	$\Theta(nW)$
$(1, 1)$ bounded degree	$\Theta(\frac{n}{\log n})$	$\Theta(\frac{nW}{W+\log n} + W)$
$(\Delta, 1)$ bounded degree	$\Theta(\frac{n}{\log \log_\Delta n})$	$\Theta(\frac{nW}{\log \log_\Delta n})$

Forcing fair outcomes

- Compute a star decomposition
- The coalition structure may be not stable
- If agents deviate, the social welfare increases

It is possible to compute a (possibly unstable) coalition stricture where all agents achieve at least $\frac{1}{2}$

Note that this implies the existence of a *stable* coalition structure where all agents achieve at least $\frac{1}{2}$

Forcing fair outcomes

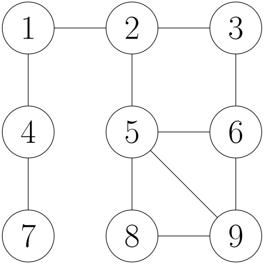

Forcing fair outcomes

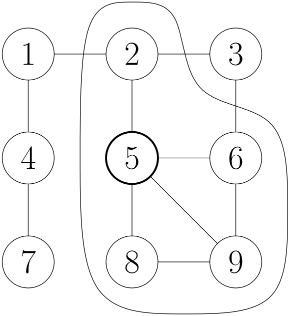

Forcing fair outcomes

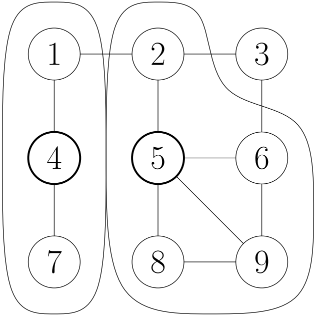

Forcing fair outcomes

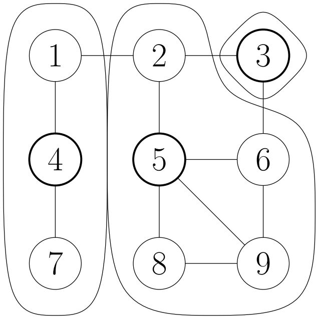

Forcing fair outcomes

It is possible to compute a stable coalition structure that $\frac{1}{2\sqrt{n}}$-approximates the optimum

Forcing fair outcomes

It is possible to compute a (possibly unstable) coalition stricture where all agents achieve at least $\frac{1}{2}$

It is possible to compute a stable coalition structure that $\frac{1}{2\sqrt{n}}$-approximates the optimum

Open problem

Is it possible to find a stable coalition structure that achieves high social welfare?

Nash Stability

A coalition $\mathcal{P} = \{P_1, \ldots, P_k\}$ is Nash stable there is no agent that can increase its utility by moving to a different coalition.

Example

Differently from Pareto stability, this deviation is allowed.

Example (2)

Agent 1 increased its utility, but agents 3, 4, 5 lost utility!

Nash Stability

- The grand coalition is always stable
- An agent can freely move to a different coalition:
  - even if agents of the old coalition lose utility
  - even if agents of the new coalition lose utility
- How far the social welfare of a Nash stable coalition structure can be from the optimum?

Price of Anarchy and Price of Stability

Let:

- G be an arbitrary graph for Social Distance Games
- $\mathcal{P}^*$ be a social welfare maximizing coalition structure
- $\mathcal{P}$ be a Nash stable coalition structure

The Price of Anarchy is:

$$PoA(G) = \frac{SW(\mathcal{P}^*)}{min_{\mathcal{P} \in PPO(G)} SW(\mathcal{P})}$$

The Price of Stability is:

$$PoS(G) = \frac{SW(\mathcal{P}^*)}{max_{\mathcal{P} \in PPO(G)} SW(\mathcal{P})}$$

Results

$PoA = \Theta(n)$

$PoS > 1$

Computing the best Nash equilibrium is NP-hard

Social Distance Games may not converge to a Nash equilibria.

Open problems

- Find an upper bound of the Price of Stability
- Find an efficient way to compute a Nash stable coalition structure that is not too far from the optimum

Conclusion

- Considering core stability, stable outcomes can be a factor $\Theta(\sqrt{n})$ from the best outcome
- If no agent can lose utility when all agents deviate, stable outcomes are even worse, a $\Theta(n)$ factor from the best outcome
- If agents can freely move without coordination, stable outcomes can still be $\Theta(n)$ factor from the best outcome
- Which policies are both permissive and fair?
- What about different centrality measures?