

Computing Desirable Partitions in Coalition Formation Games

Doctoral Consortium

Martin Bullinger
Technische Universität München
München, Germany
bullinge@in.tum.de

ABSTRACT

Coalition formation games aim at predicting the cooperative behavior of agents when forming alliances. Agents entertain preferences over coalition structures, and the goal is to find a coalition structure that is good for both individual agents and the society as an entity. We measure the quality of partitions in terms of Pareto optimality and popularity. We give both efficient algorithms and hardness results for computing partitions that satisfy these properties for various classes of coalition formation games, including roommate games, flatmate games, and cardinal hedonic games.

KEYWORDS

Coalition Formation; Social Choice Theory; Hedonic Games; Pareto optimality; Popularity

ACM Reference Format:

Martin Bullinger. 2020. Computing Desirable Partitions in Coalition Formation Games. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 3 pages.

1 INTRODUCTION

Social choice theory deals with the question of how to aggregate some voters’ preferences over a set of candidates to enable a collective choice among these candidates. On the other hand, game theory investigates the outcome of strategic actions of individual agents. An important subfield is cooperative game theory, where one considers coalition formation games, which have been a central aspect of game theory ever since the publication of von Neumann and Morgenstern’s *Theory of Games and Economic Behavior* in 1944. The traditional models involve a formal specification of a value that each group of agents can achieve on their own. This value can, for example, be interpreted as their bargaining power for their treatment in a larger coalition of which they are part. Drèze and Greenberg [9] noted that in many situations, assigning such a value is not feasible, possible, or even relevant to the coalition formation process, for example in the formation of social clubs, teams, or societies. Instead, in coalition formation games, the agents’ preferences are defined directly in terms of the coalition structures, i.e., partitions of the agents in disjoint coalitions. In the special case of *hedonic games*, these preferences only depend on the agent’s own coalition. Formally, coalition formation can therefore be considered as a special case within the voting setting, where the agents

entertain preferences over coalition structures. This places coalition formation at the intersection of social choice theory and game theory, and a combination of ideas from both fields is a fruitful approach to tackle problems of coalition formation. The goal of the PhD project is therefore to apply methods from social choice theory to game theory and vice versa.

A central concern of coalition formation is to measure the quality of a partition. Various such measures—also called solution concepts—have been proposed in the literature. Most of these measures aim at guaranteeing a certain degree of stability—preventing single agents or groups of agents to break apart from their coalitions—or optimality—guaranteeing a globally measured outcome that is good for society as a whole. A good overview of solution concepts has been made by Aziz and Savani [3]. Many such concepts are inspired by analogs in non-cooperative game theory, such as Nash equilibria and the core. *Pareto optimality* is defined by requiring that there exists no partition of the agents that is weakly preferred by all agents and strictly preferred by some agent. This assumption seems quite mild, and thus a reasonable solution concept that violates Pareto optimality is hard to imagine. While Pareto optimal partitions do not allow the grand coalition, i.e., the whole group of agents, to deviate, stronger notions of stability might also consider small groups. On the other hand, while Pareto optimality guarantees a certain efficiency for the whole society, individual agents can be very poorly off.

2 PARETO OPTIMALITY IN CARDINAL HEDONIC GAMES

First results of the thesis for Pareto optimality are for classes of hedonic games that can be succinctly represented by cardinal valuation functions that every agent entertains over the other agents. Individual values can be aggregated to utilities over coalitions (and partitions) by computing the sum or average of values of agents in their own coalition. Taking the sum of values defines the class of additively separable hedonic games [4], and taking the average of values including or excluding oneself results in the classes of fractional and modified fractional hedonic games [2, 18].

In addition to Pareto optimality, one often also requires *individual rationality*, requiring that every agent receives as much utility as she would receive in a singleton coalition. Aumann [1] says about the conjunction of these two properties that “The requirement that a feasible outcome be undominated via one-person coalitions (individual rationality) and via the all-person coalition (efficiency or Pareto optimality) is thus quite compelling.” Indeed, it is always possible to fulfill both properties at once via local search heuristics, starting with the partition that places every agent in a singleton

coalition and subsequent application of Pareto improvements. However, the intractability of the computability of outcomes satisfying both demands in certain classes of hedonic games suggests that this procedure need not run in polynomial time. A stronger notion of optimality is *welfare optimality*, which seeks to maximize the sum of utilities of the agents. We summarize the main results obtained for Pareto optimality [7].

- Pareto optimal outcomes can be computed in polynomial time for symmetric additively separable and for simple fractional hedonic games by variations of serial dictatorship that apply dictatorships to subcoalitions.
- Pareto optimal and individually rational outcomes can be computed in polynomial time for symmetric modified fractional hedonic games by using a strong connection to the problem of covering as many vertices of a graph as possible with disjoint cliques, a problem known to be tractable [15].
- In contrast, computing Pareto optimal and individually rational outcomes for symmetric additively separable and symmetric fractional hedonic games is NP-hard.
- On simple, symmetric modified fractional hedonic games, Pareto optimal and welfare optimal outcomes coincide. This solves an open problem by Elkind et al. [10] who proved the result for underlying bipartite graphs.
- Welfare optimality for general modified fractional hedonic games can be attained by partitions of size 2 and 3 only, and can be efficiently approximated within a factor of 2.

3 POPULAR COALITION STRUCTURES

In the second part of the thesis, we focus on the notion of popularity [13], which involves aspects of both stability and optimality. A partition is *popular* if there is no other partition that is preferred by a majority of the agents. Moreover, a partition is *strongly popular* if it is preferred to every other partition by some majority of agents. Popularity thus corresponds to the notion of weak and strong Condorcet winners in voting theory, i.e., candidates that are at least as good as any other candidate in pairwise majority comparisons. A recent survey of popular matchings is provided by Cseh [8].

In contrast to Pareto optimal partitions, popular partitions are not guaranteed to exist. Allowing for randomization over partitions, we therefore also consider *mixed* popular partitions, as proposed by Kavitha et al. [16] and whose existence follows from the minimax theorem. Mixed popular partitions are a special case of *maximal lotteries*, a randomized voting rule that has recently gathered increased attention in social choice theory [5, 12].

We study the computational complexity of popular, strongly popular, and mixed popular partitions in a variety of hedonic coalition formation settings, including additively separable hedonic games, fractional hedonic games, and hedonic games where the coalition size is bounded. The latter includes flatmate games (which only allow coalitions of up to three agents) and roommate games (which only allow coalitions of up to two agents).

Our main findings are as follows [6]:

- Generalizing earlier results by Kavitha et al. [16], we show how mixed popular partitions in roommate games can be computed in polynomial time via linear programming and

a separation oracle on a subpolytope of the matching polytope for non-bipartite graphs. This stands in contrast to a recent result showing that computing popular partitions in roommate games is NP-hard [11, 14].

- As a corollary, we obtain that finding strongly popular partitions can be done in polynomial time in roommate games, even when preferences allow for ties. This resolves an acknowledged open problem.¹
- We provide the first negative computational results for mixed popular partitions and strongly popular partitions by showing that finding these partitions in flatmate games is NP-hard. Moreover, it turns out that verifying whether a given partition is popular, strongly popular, or mixed popular in flatmate games is coNP-complete. All of these results hold for strict and globally ranked preferences, i.e., coalitions appear in the same order in each individual preference ranking. This is interesting because finding popular partitions in roommate games becomes tractable under the same restrictions.
- We prove that computing popular, strongly popular, and mixed popular partitions is NP-hard in symmetric additively separable hedonic games, and that computing popular partitions is NP-hard in symmetric fractional hedonic games.

4 CONCLUSION AND OUTLOOK

We have mainly seen results on two measures for the quality of a partition in coalition formation. Pareto optimality, a measure considered in many game-theoretic settings, was investigated for hedonic games with cardinal utilities, and popularity, a measure that implements a general idea from voting, was considered in broad classes of games of both ordinal and cardinal preferences.

In the further process of the PhD thesis, we will tackle further problems around social choice, cooperative game theory, and fair division by investigation of axioms that model demands on multi-agent systems and algorithms that find outcomes satisfying these axioms.

ACKNOWLEDGMENTS

This material is based on work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/12-1. The author would like to thank Felix Brandt for his constant guidance and support as supervisor.

REFERENCES

- [1] R. J. Aumann. 1987. Game Theory. In *The New Palgrave: A Dictionary of Economics*, J. Eatwell, M. Milgate, and P. Newman (Eds.). Vol. 2. MacMillan, 460–482.
- [2] H. Aziz, F. Brandl, F. Brandt, P. Harrenstein, M. Olsen, and D. Peters. 2019. Fractional Hedonic Games. *ACM Transactions on Economics and Computation* 7, 2 (2019).
- [3] H. Aziz and R. Savani. 2016. Hedonic Games. In *Handbook of Computational Social Choice*, F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. D. Procaccia (Eds.). Cambridge University Press, Chapter 15.
- [4] A. Bogomolnaia and M. O. Jackson. 2002. The Stability of Hedonic Coalition Structures. *Games and Economic Behavior* 38, 2 (2002), 201–230.
- [5] F. Brandl, F. Brandt, and H. G. Seedig. 2016. Consistent Probabilistic Social Choice. *Econometrica* 84, 5 (2016), 1839–1880.

¹See, for example, Manlove [17, p. 380]: “Our last open problem concerns the complexity of the problem of finding a strongly popular matching, or reporting that none exists, given an instance of SRTI [Stable Roommates with Ties and Incomplete lists], which is unknown at the time of writing”

- [6] F. Brandt and M. Bullinger. 2020. Finding and Recognizing Popular Coalition Structures. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS. Forthcoming.
- [7] M. Bullinger. 2020. Pareto-Optimality in Graph-Based Hedonic Games. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. IFAAMAS. Forthcoming.
- [8] Á. Cseh. 2017. Popular Matchings. In *Trends in Computational Social Choice*, U. Endriss (Ed.). AI Access, Chapter 6.
- [9] J. H. Drèze and J. Greenberg. 1980. Hedonic Coalitions: Optimality and Stability. *Econometrica* 48, 4 (1980), 987–1003.
- [10] E. Elkind, A. Fanelli, and M. Flammini. 2016. Price of Pareto optimality in hedonic games. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*. AAAI Conference on Artificial Intelligence (AAAI), 475–481.
- [11] Y. Faenza, T. Kavitha, V. Power, and X. Zhang. 2019. Popular Matchings and Limits to Tractability. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2790–2809.
- [12] P. C. Fishburn. 1984. Probabilistic Social Choice Based on Simple Voting Comparisons. *Review of Economic Studies* 51, 4 (1984), 683–692.
- [13] P. Gärdenfors. 1975. Match Making: Assignments based on bilateral preferences. *Behavioral Science* 20, 3 (1975), 166–173.
- [14] S. Gupta, P. Misra, S. Saurabh, and M. Zehavi. 2019. Popular Matching In Roommates Setting is NP-hard. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2810–2822.
- [15] P. Hell and D. G. Kirkpatrick. 1984. Packings by cliques and by finite families of graphs. *Discrete Mathematics* 49, 1 (1984), 45–59.
- [16] T. Kavitha, J. Mestre, and M. Nasre. 2011. Popular mixed matchings. *Theoretical Computer Science* 412, 24 (2011), 2679–2690.
- [17] D. F. Manlove. 2013. *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company.
- [18] M. Olsen. 2012. On defining and computing communities. In *Proceedings of the 18th Computing: Australasian Theory Symposium (CATS) (Conferences in Research and Practice in Information Technology (CRPIT))*, Vol. 128. 97–102.
- [19] J. von Neumann and O. Morgenstern. 1947. *Theory of Games and Economic Behavior* (2nd ed.). Princeton University Press.