Computing Desirable Outcomes in Specific Multi-Agent Scenarios

Doctoral Consortium

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

ABSTRACT

Coalition formation and Schelling segregation are important scenarios in algorithmic game theory. While the former considers the strategic behavior of agents gathering in coalitions, the latter is a setting in which agents of two groups seek to surround themselves with like-minded agents. In each case, the quality of outcomes can be measured in form of axioms of optimality and stability. The thesis investigates how to compute such desirable outcomes efficiently and how to deal with computational intractability by means of approximation algorithms, randomization, or domain restrictions.

KEYWORDS

Social Choice Theory; Coalition Formation; Hedonic Games; Schelling Segregation

ACM Reference Format:
Martin Bullinger. 2021. Computing Desirable Outcomes in Specific Multi-Agent Scenarios: Doctoral Consortium. In *Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 3 pages.

1 INTRODUCTION

Game theory investigates the outcome of the strategic behavior of individual agents. A central concern is coalition formation, a scenario in cooperative game theory, which has received great attention ever since the development of game theory. 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 the bargaining power for their treatment in a larger coalition. 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, agents entertain preferences over coalition structures, i.e., partitions of the agents into disjoint coalitions. In the special case of *hedonic games*, these preferences only depend on the agent's own coalition.

Another interesting subject matter in multi-agent systems is to observe the strategic behavior of individuals which desire to surround themselves with like-minded agents. This idea traces back to a model by Schelling [15, 16] aiming to explain racial segregation in metropolitan areas.

Proc. of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

Outcomes are usually evaluated by measures of stability—concerning the likeliness of single agents or groups of agents to stick to their behavior—or optimality—guaranteeing an outcome that is good for the society as a whole. The goal of the thesis is to investigate the efficient computability of desirable outcomes according to these measures in the aforementioned settings. Often, we face intractability of such computations. We show, how these are caused by demanding requirements on the desired properties of solutions or by the generality of the domain of agents' preferences. We thus identify transitions to tractability by means of approximation algorithms, randomized versions of deterministic solution concepts, or domain restrictions.

2 INDIVIDUALLY STABLE PARTITIONS

We start by investigating stability in hedonic games. The measures we consider are based on the happiness of single agents in a coalition structure. Clearly, an agent would not join a coalition that she prefers over being on her own. Otherwise, she would just leave her coalition to form a singleton coalition. A coalition structure that places every agent in a coalition at least as good as her singleton coalition is called *individually rational.* However, an agent might also leave her coalition to improve by joining another coalition, which accepts her. In this case we say that an agent performs an *individual deviation.* If no such deviation is possible, a coalition structure is called *individually stable.* Consent of the joint coalition is a reasonable requirement and distinguishes individual stability from Nash stability, under which changing the coalition without further consent is forbidden. For instance, joining international bodies like the European Union or the NATO requires unanimous agreement of all current parties. Individual deviations give rise to an *individual dynamics*, a process of transitions amongst coalition structures by means of individual deviations.

While individually rational coalition structures always exist (e.g., the coalition structure consisting of all singleton coalitions), individually stable coalition structures need not exist and the individual dynamics is not guaranteed to converge. The corresponding decision problem is often NP-hard. We prove such hardness results in a broad range of hedonic games including symmetric fractional hedonic games (FHGs) and anonymous hedonic games (AHGs), introduced by Aziz et al. [1] and Bogomolnaia and Jackson [3], respectively. These intractabilities can be met with appropriate domain restrictions such as assuming binary utility functions or single-peakedness, a structural condition on the domain.

THEOREM 2.1 (BRANDT ET AL. [6]). *The individual dynamics converges on FHGs with binary utility functions, when starting with the singleton partition. The individual dynamics converges in AHGs for strict single-peaked preferences, regardless of the starting partition.*

3 WELFARE GUARANTEES IN SCHELLING SEGREGATION

The next section deals with measuring the optimality of an outcome for a society as a whole. If the agents entertain cardinal utilities, the *social welfare* of an outcome is defined as the sum of the agents' valuations of this outcome. A *welfare optimal* outcome maximizes the social welfare and would thus be desirable for the society of the agents. We study the problem of finding such an outcome in the case of Schelling segregation. In Schelling instances, two classes of agents are assigned to vertices of a topology graph and their utility is the fraction of agents from the same class in their neighborhood [10]. Hence, their respective utility is bounded by 1, and consequently, the social welfare is always bounded by the number of agents. Unfortunately, finding a welfare optimal outcome is intractable under quite restrictive assumptions. However, an outcome which approximates social welfare well, can be computed in polynomial time. The algorithm is a derandomization of the randomized procedure selecting an assignment of the agents uniformly at random. The tight bound on the welfare can be shown to be slightly higher and is achieved by the algorithm.

THEOREM 3.1 (BULLINGER ET AL. [8]). *For any Schelling instance with n agents, an assignment with social welfare at least $\frac{n}{2} - 1$ can be computed in polynomial time.*

4 COMBINING OPTIMALITY AND STABILITY

So far, we considered stability and optimality only separately. Finally, we want to address the problem of satisfying both notions simultaneously. The first approach is to unify the axioms in the previous sections for cardinal models of hedonic games. Secondly, we implement the idea of weak Condorcet winners from voting to the special domain of hedonic games to define a measure satisfying both aspects of stability and optimality.

Since welfare optimality is usually hard to achieve, we consider the important weakening of Pareto optimality. An outcome is *Pareto optimal* if it is not Pareto dominated, i.e., there exists no other outcome weakly preferred by all agents, and strictly preferred by some agents. Our goal is to find outcomes both Pareto optimal and individually rational in FHGs, symmetric additively separable hedonic games (ASHGs) and symmetric modified fractional hedonic games (mFHGs). The latter two classes were introduced by Bogomolnaia and Jackson [3] and Olsen [14], respectively. Note that such outcomes are guaranteed to exist and can in principle be computed by a local search algorithm. This simple algorithm starts with the coalition structure consisting of singleton coalitions and subsequently applies Pareto improvements. Then, individual rationality is maintained, and the terminal state is Pareto optimal. However, this algorithm is not guaranteed to run in polynomial time.

Still, outcomes which are only Pareto optimal can be computed efficiently by variations of serial dictatorship or a generalization of largest matchings, respectively. On the other hand, individual rationality is far more difficult to satisfy in addition and the computational complexity depends on the underlying domain [7].

THEOREM 4.1 (BULLINGER [7]). *Pareto optimal outcomes can be computed in polynomial time for FHGs with binary utility functions and ASHGs. Pareto optimal and individually rational partitions can be computed in polynomial time for mFHGs. In contrast, computing Pareto optimal and individually rational partitions is NP-hard for FHGs and ASHGs.*

A second approach is to consider an axiom that implements ideas from both optimality and stability. An outcome is *popular* if it never looses a majority vote amongst the agents against another outcome. Thus, popular outcomes are Pareto optimal and stable in the sense that they prevent a deviation towards another outcome that the agents agreed on by means of a simple vote. Moreover, an outcome is *strongly popular* if it wins every majority vote against another outcome. Therefore, (strongly) popular outcomes correspond to the notion of weak and strong Condorcet winners in social choice theory. We consider popularity in the important domain of roommate games, where the agents have ordinal preferences over possible partners to form coalitions of size 2. In general, it is hard to compute popular outcomes even under strict preferences [11, 12]. Generalizing an earlier result by Kavitha et al. [13], we were able to prove existence and efficient computability of a randomized version of popularity called *mixed popularity* by an approach using linear programming [4]. Apart from the consequences for mixed popularity, the method is very general and yields new polynomial-time algorithms for both known and new problems. In particular, it can be used to efficiently compute strongly popular outcomes, a problem whose complexity was open in the case of arbitrary (i.e., weak) preferences [2].

THEOREM 4.2 (BRANDT AND BULLINGER [4]). *Mixed popular and strongly popular outcomes can be computed in polynomial time in roommate games.*

5 FUTURE DIRECTIONS

The aim of future work of the thesis is to strengthen the results described so far and to obtain similar results in related areas. Interesting follow-up questions concern the speed of convergence of the individual dynamics and more insights on the simultaneous satisfiability of individual stability and Pareto optimality. Apart from the models described so far, intriguing further settings encompass voting and fair division. In the former, we recently obtained some interesting results concerning strategyproofness in social choice correspondences [5].

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] 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), 1–29.

[2] P. Biró, R. W. Irving, and D. F. Manlove. 2010. Popular Matchings in the Marriage and Roommates Problems. In *Proceedings of the 7th Italian Conference on Algorithms and Complexity (CIAC)*. 97–108.
[3] A. Bogomolnaia and M. O. Jackson. 2002. The Stability of Hedonic Coalition Structures. *Games and Economic Behavior* 38, 2 (2002), 201–230.
[4] 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)*. 195–203.
[5] F. Brandt, M. Bullinger, and P. Lederer. 2021. On the Indecisiveness of Kelly-Strategyproof Social Choice Functions. In *Proceedings of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. Forthcoming.
[6] F. Brandt, M. Bullinger, and A. Wilczynski. 2021. Reaching Individually Stable Coalition Structures in Hedonic Games. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
[7] M. Bullinger. 2020. Pareto-Optimality in Cardinal Hedonic Games. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 213–221.
[8] M. Bullinger, W. Suksompong, and A. Voudouris. 2021. Welfare Guarantees in Schelling Segregation. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
[9] J. H. Drèze and J. Greenberg. 1980. Hedonic Coalitions: Optimality and Stability. *Econometrica* 48, 4 (1980), 987–1003.
[10] E. Elkind, J. Gan, A. Igarashi, W. Suksompong, and A. A. Voudouris. 2019. Schelling games on graphs. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. 266–272.
[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] 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.
[13] T. Kavitha, J. Mestre, and M. Nasre. 2011. Popular mixed matchings. *Theoretical Computer Science* 412, 24 (2011), 2679–2690.
[14] 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.
[15] T. C. Schelling. 1969. Models of Segregation. *American Economic Review* 59, 2 (1969), 488–493.
[16] T. C. Schelling. 1971. Dynamic Models of Segregation. *Journal of Mathematical Sociology* 1 (1971), 143–186.