Symmetries and Efficient Solvability in Multi-Player Games

(Extended Abstract)

Felix Brandt* Felix Fischer* Markus Holzer†

Abstract

There are various ways in which strategic games may exhibit forms of symmetry. A common aspect of symmetry, which enables the compact representation of games even when the number of players is unbounded, is that players are incapable of distinguishing between the other players. We define four classes of symmetric games by additionally considering the following two characteristics: the availability of *identical payoff functions* and the ability to *distinguish oneself* from the other players. Based on these varying notions of symmetry, we investigate the computational complexity of finding pure Nash equilibria. It turns out that in all four classes of games equilibria can be found efficiently when the number of actions available to each player is held constant. For most succinct representations of multi-player games, the same computational problem has been shown to be intractable. Furthermore, we show that the availability of identical payoff functions greatly simplifies the search for equilibria.

1 Introduction

In recent years, the computational complexity of game-theoretic solution concepts (both in cooperative as well as in non-cooperative game theory) has come under increasing scrutiny. A major obstacle when considering non-cooperative strategic-form games with an unbounded number of players is the exponential size of the naive representation of the payoffs. More precisely, a general game in strategic form with n players and k actions per player comprises $n \cdot k^n$ numbers. Computational statements over such large objects are somewhat questionable for two reasons [*cf.* Papadimitriou and Roughgarden, 2005]. First, the value of efficient, *i.e.*, polynomial-time, algorithms for problems whose input size is already exponential in a natural parameter (the number of players) is doubtful. Secondly, most, if not all, "natural" multi-player games will hardly be given as multi-dimensional payoff matrices but rather in terms of some more intuitive (and compact) representation. A natural and straightforward way to simplify the representation of multi-player games is to somehow formalize the similarities between players. As a matter of fact, *symmetric games* have been studied since the early days of game theory [see, *e.g.*, von Neumann, 1928, Gale et al., 1950, Nash, 1951]. The established definition states that a game is symmetric if the payoff functions of all players are identical and symmetric in the other players' actions, *i.e.*, it is impossible to distinguish between other players. When explicitly looking at *multi-player* games, there are other conceivable concepts of symmetry. For instance, dropping the requirement of identical payoff functions yields a more general class of multi-player games that still admits a compact representation. In this paper, we define four classes of succinctly representable symmetric multi-player games and study the computational complexity of finding pure Nash equilibria in games belonging to these classes. The complexity classes appearing in the paper

*Computer Science Department, University of Munich, Germany, {brandtf,fischerf}@tcs.ifi.lmu.de
†Computer Science Department, Technical University of Munich, Germany, holzer@in.tum.de

	Indistinguishability of oneself and other players	Identical payoff functions	Indistinguishability of other players
weakly symmetric	–	–	✓
strongly symmetric	–	✓	✓
weakly anonymous	✓	–	✓
strongly anonymous	✓	✓	✓

Table 1: Four classes of symmetric games

are the following: the class AC^0 of problems that can be solved using *constant-depth Boolean circuits* with unbounded fan-in[1], the class L of problems solvable using only *logarithmic space*,[2] the class NP of problems whose *solutions can be verified in polynomial time*. The following relations between these complexity classes are currently known (P is the class of problems that can be solved in polynomial time, which is commonly regarded as efficient solvability): $AC^0 \subset L \subseteq P \subseteq NP$. The hardest problems in a given class are called "complete" for that class. For example, NP-complete problems are problems for which no efficient (polymonial-time) algorithm is known.

2 Symmetries in Multi-Player Games

We start by introducing strategic games and the different notions of symmetry studied in this paper.

Definition 1 (Normal-form game) *A* game in normal-form *is a tuple* $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ *where* N *is a set of* players *and for each player* $i \in N$, A_i *is a nonempty set of* actions *available to player* i, *and* $p_i : (\times_{i\in N} A_i) \to \mathbb{R}$ *is a function mapping each action profile of the game (*i.e., *combination of actions) to a real-valued* payoff *for player* i.

The unifying feature of the different notions of symmetry we consider is that the available actions are identical for all players, denoted $A = A_1 = \cdots = A_n$, and that players are incapable of distinguishing between the other players. The four different classes of symmetric games are then defined by two additional characteristics, namely (i) *identical payoff functions* for all players and (ii) the ability to *distinguish oneself* from the other players (see also Table 1). We formally define these classes as follows.

Definition 2 *Let* $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ *a normal-form game with* $A_i = A_j = A$ *for all* $i, j \in N$. *For any permutation* $\pi : N \to N$ *of the set of players, let* $\pi' : A^N \to A^N$ *be the permutation of the set of action profiles given by* $\pi'((a_1, \ldots, a_n)) = (a_{\pi(1)}, \ldots, a_{\pi(n)})$.[3] Γ *is called*

- weakly symmetric *if* $p_i(a) = p_i(\pi'(a))$ *for all* $i \in N$ *and all* π *with* $\pi(i) = i$,
- strongly symmetric *if* $p_i(a) = p_j(\pi'(a))$ *for all* $i, j \in N$ *and all* π *with* $\pi(j) = i$
- weakly anonymous *if* $p_i(a) = p_i(\pi'(a))$ *for all* $i \in N$, *and*
- strongly anonymous *if* $p_i(a) = p_j(\pi'(a))$ *for all* $i, j \in N$.

[1]The fan-in is the number of input wires of a logical gate.
[2]Both space and time are measured relative to the size of the input.
[3]That is, π' is an automorphism on the set of action profiles that preserves the number of players that play a particular action.

If two players of a *weakly symmetric* game exchange their actions, all other players' payoffs remain the same. For two-player games, weak symmetry is not a restriction (action sets of equal size can simply be achieved by adding dummy actions for one of the players). This may be one of the reasons why weak symmetry has not received much attention so far.

The additional restriction of identical payoff functions in *strongly symmetric* games means that two players exchanging actions exchange their payoffs as well (while all other players' payoffs again remain the same). Numerous well known games like the Prisoner's Dilemma, Matching Pennies, or Chicken are examples of (two-player) strongly symmetric games. Nash [1951] has shown that strongly symmetric games always possess a *symmetric* equilibrium in *mixed* strategies. However, this theorem affords us no computational advantage when considering pure Nash equilibria.

In a *weakly anonymous* game, the payoffs of all players remain the same if two players exchange their actions. Voting with identical weights can be seen as an example of a weakly anonymous game.

Finally, in a *strongly anonymous* game, all players receive the same payoff in each outcome and the payoff remains the same if two players exchange their actions. This means that strongly anonymous games are a special case of common payoff (or pure coordination) games, in which every action profile with maximum payoff trivially is a Nash equilibrium (no player can gain by deviating). Common payoff games thus belong to the class of games guaranteed to possess a pure Nash equilibrium, as recently discussed by Fabrikant et al. [2004].

3 Results

We investigated the computational complexity of finding pure Nash equilibria in symmetric games with an *unbounded* number of players n. The number of actions available to each player (k) was either *held constant* ($k = O(1)$) or *growing linearly* in n ($k = O(n)$). Table 2 summarizes our results.

Two main conclusions that can be drawn from Table 2 are that, for a constant number of actions, symmetric games with identical payoff functions can be solved very easily (in AC^0), and that symmetry significantly facilitates the computation of Nash equilibria: For most succinct representations of multi-player games [Fischer et al., 2006, Schoenebeck and Vadhan, 2006], deciding the existence of a pure equilibrium has been shown to be NP-complete, even when players just have a constant number of actions; in common-payoff games (a superclass of strongly anonymous games) finding a pure equilibrium is PLS-complete.[4] This emphasizes the computational benefit of taking into account symmetries when looking for Nash equilibria.

Future work includes studying the computational complexity of other game-theoretic solution concepts in symmetric games. Preliminary results indicate that solvability via iterated weak dominance [Moulin, 1979] is a computationally more expensive concept than Nash equilibrium.

	constant k	growing k
weakly symmetric	L	NP-complete
strongly symmetric	AC^0	NP-complete
weakly anonymous	L	NP-complete
strongly anonymous	AC^0	PLS-complete

Table 2: Computational complexity of finding pure Nash equilibria in symmetric games

[4]PLS is the class of polynomial local search problems [Johnson et al., 1988].

Acknowledgements

We thank Jan Johannsen for enlightening discussions on circuit complexity and local search and Rob Powers for introducing the first author to the ambiguity of symmetry in games.

This material is based upon work supported by the Deutsche Forschungsgemeinschaft under grant BR 2312/3-1.

References

A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 604–612. ACM Press, 2004.

F. Fischer, M. Holzer, and S. Katzenbeisser. The influence of neighbourhood and choice on the complexity of finding pure Nash equilibria. *Information Processing Letters*, 2006. To Appear.

D. Gale, H. W. Kuhn, and A. W. Tucker. On symmetric games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games*, volume 1, pages 81–87. Princeton University Press, 1950.

D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37:79–100, 1988.

H. Moulin. Dominance solvable voting schemes. *Econometrica*, 47:1337–1351, 1979.

J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

J. von Neumann. Zur Theorie der Gesellschaftspiele. *Mathematische Annalen*, 100:295–320, 1928.

C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *Proceedings of the 16thAnnual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 82–91. SIAM, 2005.

G. Schoenebeck and S. Vadhan. The computational complexity of Nash equilibria in concisely represented games. In *Proceedings of the 7th ACM Conference on Electronic Commerce (ACM-EC)*. ACM Press, 2006.