Electronic Colloquium on Computational Complexity, Report No. 91 (2006)

Symmetries and the Complexity of Pure Nash Equilibrium

Felix Brandt* Felix Fischer* Markus Holzer†

*Institut für Informatik, Universität München,
80538 München, Germany
{brandtf,fischerf}@tcs.ifi.lmu.de

†Institut für Informatik, Technische Universität München,
85748 Garching, Germany
holzer@in.tum.de

Abstract

Strategic games may exhibit symmetries in a variety of ways. A common aspect of symmetry, enabling the compact representation of games even when the number of players is unbounded, is that players cannot (or need not) distinguish between the other players. We define four classes of symmetric games by considering two additional properties: *identical payoff functions* for all players and the ability to *distinguish oneself* from the other players. Based on these varying notions of symmetry, we investigate the computational complexity of pure Nash equilibria. It turns out that in all four classes of games equilibria can be found efficiently when only a constant number of actions is available to each player, a problem that has been shown intractable for other succinct representations of multi-player games. We further show that identical payoff functions simplify the search for equilibria, while a growing number of actions renders it intractable. Finally, we show that our results extend to wider classes of *threshold symmetric* games where players are unable to determine the exact number of players playing a certain action.

1 Introduction

In recent years, the computational complexity of game-theoretic solution concepts, both in cooperative and non-cooperative game theory, has come under increasing scrutiny. A major obstacle when considering non-cooperative normal-form games with an unbounded number of players is the exponential size of the naive representation of payoffs. More precisely, a general game in normal-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 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 the other players (von Neumann and Morgenstern, 1947; Luce and Raiffa, 1957). When explicitly looking at *multi-player* games,

ISSN 1433-8092

there are other conceivable notions 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. It turns out that in all four classes equilibria can be found efficiently if only a constant number of actions is available to each player. Moreover, identical payoff functions for all players further reduce the computational complexity associated with pure Nash equilibria, an effect that is nullified as soon as there are two different payoff functions. Anonymity, *i.e.*, the fact that a player cannot (or does not) distinguish himself from the other players, does not seem to offer any computational advantage. Finally, computing equilibria becomes intractable in all four classes of symmetric games when the number of actions grows linearly in the number of players.

Unlike Nash equilibria in *mixed* strategies, *i.e.*, probabilistic combinations of actions, pure Nash equilibria are not guaranteed to exist. They nevertheless form an interesting subset of equilibria for three reasons. First, requiring randomization in order to reach a stable outcome has been criticized on various grounds. In multi-player games, where action probabilities in equilibrium can be irrational, randomization is particularly questionable. Secondly, computation of pure equilibria may be tractable in cases where that of mixed ones is not. Finally, pure equilibria as computational objects are usually much smaller in size than mixed ones.

We assume the reader to be familiar with the well-known chain of complexity classes $\mathrm{AC}^0 \subset \mathrm{TC}^0 \subseteq \mathrm{L} \subseteq \mathrm{P} \subseteq \mathrm{NP}$, and the notions of constant-depth and polynomial-time reducibility (see, *e.g.*, Papadimitriou, 1994; Johnson, 1990). AC^0 is the class of problems solvable by uniform constant-depth Boolean circuits with unbounded fan-in. TC^0 adds so-called threshold gates which output *true* if and only if the number of *true* inputs exceeds a certain threshold. L is the class of problems solvable by deterministic Turing machines using only logarithmic space. P and NP are the classes of problems that can be solved in polynomial time by deterministic and nondeterministic Turing machines, respectively. Furthermore, #P is the class of counting problems associated with polynomially balanced polynomial-time decidable relations. The class PLS of polynomial local search problems and an appropriate notion of reduction (Johnson et al., 1988) will be introduced as needed.

The remainder of this paper is organized as follows: In the following section, we survey relevant work on symmetric games, succinct representations, and the computational complexity of pure Nash equilibrium. In Section 3, we then formally introduce four different notions of symmetry in strategic games and the solution concept of Nash equilibrium. The main results of this paper, including efficient algorithms as well as hardness results for all four symmetry classes, are given in Section 4. In Section 5, we provide additional results for a more general notion of symmetry. Section 6 concludes the paper and points to some open problems.

2 Related Work

Symmetries in games have been investigated since the earliest days of game theory. Von Neumann (1928) and von Neumann and Morgenstern (1947) were the first to consider symmetries of *cooperative* games, calling a game in characteristic form symmetric if the value of a coalition depends only on its size. In the context of two-player (non-cooperative) normal-form games, the term symmetric is used to refer to games with a skew-symmetric payoff matrix (see, *e.g.*, Borel, 1921; Gale et al., 1950), corresponding to strong symmetry in the vocabulary of this paper. Gale et al. (1950) provided a (polynomial-time) reduction from arbitrary games to symmetric games which preserves equilibria. Since finding a (possibly mixed) equilibrium in general games has recently been shown PPAD-complete even for just two players (Chen and Deng, 2005; Daskalakis et al., 2006), the same holds for symmetric games as well.

To date, most research on symmetries in games has concentrated on strongly symmetric games, which require identical payoff functions for all players. One of the reasons for this may have been the strong focus of the early research in non-cooperative game theory on two-player games, where weak symmetry as defined in this paper does not impose any restrictions. An early result by Nash (1951) implies the existence of a symmetric equilibrium in (again, strongly) symmetric games. Papadimitriou and Roughgarden (2005) capitalize on this existence result and show that a Nash equilibrium of a strongly symmetric game with n players and k actions can be computed in P if $k = O(\log n / \log\log n)$. While their related results about the tractability of correlated equilibrium (Aumann, 1974) do not rely on identical payoff functions and hence apply to weakly symmetric games as well, this is not the case for their results about Nash equilibria. The aforementioned existence of *symmetric* Nash equilibria does neither extend to pure equilibria, nor does it hold for the classes of weakly symmetric and weakly anonymous games. For example, Figure 3 on page 7 shows a weakly symmetric game without a symmetric equilibrium.

Obviously, deciding the existence of a pure equilibrium is easy if the number of candidates for such an equilibrium, *i.e.*, the number of action profiles, is polynomial in the size of the game. This is certainly the case for the "naive" representation of a game as a multi-dimensional table of payoffs, but no longer holds if the game is represented succinctly. For example, deciding the existence of a pure Nash equilibrium has been shown to be NP-complete for games in graphical normal-form (Fischer et al., 2006; Gottlob et al., 2005) or circuit form (Schoenebeck and Vadhan, 2006). Apart from these generic types, many succinct representations are related to symmetries in that they exploit similarities between players. In congestion games (Rosenthal, 1973), the available actions consist of sets of resources, and the payoff depends on the number of other players that have selected the same resources (*i.e.*, played the same action). Congestion games always have a pure Nash equilibrium (Rosenthal, 1973), and finding one is PLS-complete in the general case and in P in the symmetric network case (Fabrikant et al., 2004). For singleton (or simple) congestion games, where only a single resource can be selected, there is a polynomial-time algorithm for finding a social-welfare-maximizing Nash equilibrium (Ieong et al., 2005). In local-effect games (Leyton-Brown and Tennenholtz, 2003), the payoff from an action may also depend on (a function of) the number of agents playing "neighboring" actions. Unlike congestion games and local-effect games, action-graph games (Bhat and Leyton-Brown, 2004; Jiang and Leyton-Brown, 2006) can encode arbitrary payoffs. For action-graph games of bounded degree, expected payoffs and the Jacobian of the payoff function can be computed in polynomial time. In practice, the latter forms the practical bottleneck step of the algorithm of Govindan and Wilson (2003) for finding Nash equilibria, but the algorithm may still take exponentially many steps to converge even for bounded degree. While closely related to the general notion of symmetry as studied in this paper, the main idea behind all the above representations is to exploit some form of independency among certain actions, or among players playing these actions. We do not make such assumptions in this paper.

3 Preliminaries

In this section, we formally define essential game-theoretic concepts, introduce four notions of symmetry in games, and state several facts concerning these notions.

3.1 Strategic Games

An accepted way to model situations of strategic interaction is by means of a *normal-form game* (see, *e.g.*, Luce and Raiffa, 1957).

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

Table 1: Four classes of symmetric games

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*.*

A combination of actions $s \in \times_{i\in N} A_i$ is also called a profile of *pure strategies*. This concept can be generalized to *mixed strategy profiles* $s \in S = \times_{i\in N} S_i$, by letting players randomize over their actions. We have S_i denote the set of probability distributions over player i's actions, or *mixed strategies* available to player i. We further write $n = |N|$ for the number of players in a game, s_i for the ith strategy in profile s, and s_{-i} for the vector of all strategies in s but s_i.

3.2 Symmetries in Multi-Player Games

A central aspect of our view on symmetry is the inability to distinguish between other players. We will therefore mainly talk about games where the set of actions is the same for all players and write $A = A_1 = \cdots = A_n$ and $k = |A|$, respectively, to denote this set and its cardinality. In the following definition, we formally introduce four classes of symmetric games by considering two additional characteristics: *identical payoff functions* for all players and the ability to *distinguish oneself* from the other players. An overview of the different classes and their properties is given in Table 1.

Definition 2 (symmetries) *Let* $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ *be a normal-form game,* A *a set of actions such that* $A_i = A$ *for all* $i \in N$*. 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* $\pi'((a_1, \ldots, a_n)) = (a_{\pi(1)}, \ldots, a_{\pi(n)})$*.* Γ *is called*

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

Inclusions and separations of the different classes follow directly from the definition and are illustrated in Figure 1. Figure 2 details the relationship for $n = 3$ and $k = 2$.

π' is an automorphism on the set of action profiles that preserves the number of players that play a particular action. Thus, an intuitive and convenient way to describe a symmetric game is in terms of the equivalence classes induced by π', or by the number of players playing the different actions in each of these classes. We use a notion introduced by Parikh (1966) in the context of context-free languages.

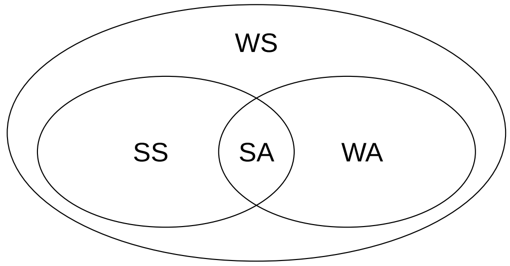

Figure 1: Inclusion relationships between weakly symmetric (WS), weakly anonymous (WA), strongly symmetric (SS), and strongly anonymous (SA) games

Γ_1:

$(\cdot,\cdot,\cdot)$	(a,g,b)	$(a,c,\cdot)$	$(\cdot,e,f)$
$(\cdot,c,b)$	$(d,e,\cdot)$	$(d,\cdot,f)$	$(\cdot,\cdot,\cdot)$

Γ_2:

(a,a,a)	(b,c,b)	(b,b,c)	(e,d,d)
(c,b,b)	(d,d,e)	(d,e,d)	(f,f,f)

Γ_3:

$(\cdot,\cdot,\cdot)$	(a,b,c)	(a,b,c)	(d,e,f)
(a,b,c)	(d,e,f)	(d,e,f)	$(\cdot,\cdot,\cdot)$

Γ_4:

(a,a,a)	(b,b,b)	(b,b,b)	(c,c,c)
(b,b,b)	(c,c,c)	(c,c,c)	(d,d,d)

Figure 2: Relationships between the payoffs of weakly symmetric (Γ_1), strongly symmetric (Γ_2), weakly anonymous (Γ_3), and strongly anonymous (Γ_4) games for $n = 3$ and $k = 2$. Players 1, 2, and 3 choose rows, columns, and tables, respectively. As an example for the separation of the different classes, Γ_1 is not strongly symmetric if $a \neq b$ and not weakly anonymous if $c \neq g$. Γ_2 is not anonymous if $b \neq c$. Γ_3 is not strongly symmetric if $a \neq c$.

Definition 3 (commutative image) *Let $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ be a normal-form game, A a set of actions such that $A_i = A$ for all $i \in N$. Then, the* commutative image *of an action profile $s \in A^N$ is defined as*

$$\begin{aligned} \#(s) &= (\#(a,s))_{a\in A} \quad \textit{where} \\ \#(a,s) &= |\{\, i \in N \mid s_i = a \,\}|. \end{aligned}$$

That is, $\#(a,s)$ denotes the number of players playing action a in action profile s, and $\#(s)$ is the vector of these numbers for all the different actions. This definition naturally extends to action profiles for subsets of the players. Since for every permutation π' induced by a permutation π of the set of players and for every action profile s, $\#(s) = \#(\pi'(s))$, the following is easily verified.

Fact 4 *Let A be a set of actions. A normal-form game $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ with $A_i = A$ for each $i \in N$ is*

- *weakly symmetric iff $p_i(s) = p_i(t)$ for all $i \in N$ and all $s,t \in A^N$ with $s_i = t_i$ and $\#(s_{-i}) = \#(t_{-i})$,*
- *strongly symmetric iff $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s,t \in A^N$ with $s_i = t_j$ and $\#(s_{-i}) = \#(t_{-j})$,*
- *weakly anonymous iff $p_i(s) = p_i(t)$ for all $i \in N$ and all $s,t \in A^N$ with $\#(s) = \#(t)$, and*

- *strongly anonymous iff* $p_i(s) = p_j(t)$ *for all* $i, j \in N$ *and all* $s,t \in A^N$ *with* $\#(s) = \#(t)$.

When talking about symmetric games, we will write $p_i(s_i, x_{-i})$ to denote the payoff of player i under any action profile s with $\#(s_{-i}) = x_{-i}$. For anonymous games, we will write $p_i(x)$ for the payoff of player i under any profile s with $\#(s) = x$. In terms of this characterization, a game is *weakly symmetric* if the payoff $p_i(s)$ of player $i \in N$ in action profile s depends, besides his own action s_i, only on the number $\#(a, s_{-i})$ of other players playing each of the actions $a \in A$, but not on who plays them. If two players exchange actions, all other players' payoffs remain the same. For two-player games, weak symmetry does not impose any restrictions (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 in the past. A game is *strongly symmetric* if it is weakly symmetric and if the payoff function is the same for all players. Hence, if two players exchange actions, their payoffs are also exchanged while all other players' payoffs remain the same. Many well-known games like the Prisoner's Dilemma, Rock-Paper-Scissors, or Chicken are examples of (two-player) strongly symmetric games. Multi-player simple congestion games (Ieong et al., 2005) are also strongly symmetric. In a *weakly anonymous* game the payoff of each player depends only on the number $\#(a, s)$ of players playing each of the actions $a \in A$, *including the player himself*. If two players exchange actions, the payoffs of all players remain the same. Matching Pennies is a weakly anonymous two-player game, voting with identical weights can be seen as an example for the multi-player case. Finally, in a *strongly anonymous* game the payoff is always the same for all players and stays the same if two players exchange actions. Strongly anonymous games are a special case of common payoff (or pure coordination) games, in which every action profile with maximum payoff is a Nash equilibrium (no player can gain by deviating). Obviously, every common payoff game, and hence every strongly anonymous game, is guaranteed to possess a pure Nash equilibrium. Other games with this property, and the complexity of finding an equilibrium in this case, have recently been investigated by Fabrikant et al. (2004).

The most basic way to specify a normal-form game is by means of a multi-dimensional table of payoffs for every single action profile. Certain games are succinctly representable simply because the payoff is the same for action profiles that are equivalent according to some equivalence relation, and needs only be specified once. For symmetric games, this equivalence relation is given by the number of players playing each action. The representation that lists the payoffs for every equivalence class will henceforth be referred to as the *naive representation* of a symmetric game. There are $\binom{n+k-1}{k-1}$ distributions of n players among k actions. Since these are exactly the equivalence classes of the set of action profiles for $n-1$ players under the commutative image, we have the following.

Fact 5 *A weakly symmetric game can be represented using at most* $n \cdot k \cdot \binom{n+k-2}{k-1}$ *numbers, and is succinctly representable* in general *if and only if* k *is bounded by a constant.*

In turn, the size of the game may become super-polynomial in n even for the slightest growth of k. Nevertheless, a succinct representation may exist for *certain classes* of games with a larger number of actions.

3.3 Nash Equilibrium

One of the best-known solution concepts for strategic games is Nash equilibrium (Nash, 1951). In a Nash equilibrium, no player is able to increase his payoff by *unilaterally* changing his strategy.

Definition 6 (Nash equilibrium) *A strategy profile* $s \in S$ *is called a* Nash equilibrium *if for each player* $i \in N$ *and each strategy* $s'_i \in S_i$,

$$p_i(s) \geq p_i((s_{-i}, s'_i)).$$

A Nash equilibrium is called pure *if it is a pure strategy profile.*

$(0,1,1)$	(0,0,1)		(0,1,0)	$(0,0,0)$
(1,1,1)	$(0,0,0)$		$(0,1,0)$	$(1,0,1)$

Figure 3: A weakly symmetric game with a unique, non-symmetric Nash equilibrium at the action profile with payoff $(1,1,1)$. Players 1, 2, and 3 choose rows, columns, and tables, respectively. Action profiles with the same commutative image as the equilibrium are shaded.

For general games, simply checking the equilibrium condition for each action profile takes time polynomial in the size of their natural representation (*i.e.*, a table of payoffs for the different action profiles). Using a succinct representation for games where the size of the natural representation grows exponentially in the number of players, which is the case for $k \geq 2$ already, quickly renders the problem NP-complete (see, *e.g.*, Fischer et al., 2006; Schoenebeck and Vadhan, 2006). In turn, the polynomial size of the naive representation for symmetric games with a constant number of actions might suggest that finding pure Nash equilibria is easy by a similar argument as above. This reasoning is flawed, however, since a single entry in the payoff table corresponds to an exponential number of action profiles, and it is very well possible that only a single one of them is a Nash equilibrium while all others are not. The weakly symmetric game given in Figure 3 illustrates this fact.

Interestingly, the ability to distinguish oneself from the other players does not extend the expressive power of anonymous games when players only have two actions.

Fact 7 *When there are only two actions available to each player, there is an AC^0-reduction from symmetric games to anonymous games that preserves pure Nash equilibria and strong symmetry.*

Let $\Gamma = (N, \{a_1, a_2\}^n, (p_i)_{i \in N})$ be a weakly symmetric game. This game induces a weakly anonymous game $\Gamma' = (N, \{a_1, a_2\}^n, (p'_i)_{i \in N})$ by defining p' so that for all $i \in N$ and for all $x \in \{0, \ldots, n-1\}$ the following statements hold:

1. $p'_i(x) > p'_i(x+1)$ if $p_i(a_1, x) > p_i(a_2, x)$
2. $p'_i(x) < p'_i(x+1)$ if $p_i(a_1, x) < p_i(a_2, x)$
3. $p'_i(x) = p'_i(x+1)$ if $p_i(a_1, x) = p_i(a_2, x)$

Depending on the original game Γ, it may be necessary to use up to n different payoffs in Γ', even when Γ contains only two. Moreover, the procedure cannot in general be extended to games where players have more than two actions, because it can lead to cyclic preference relations. For example, the (strongly) symmetric two-player game Rock-Paper-Scissors cannot be mapped to a corresponding anonymous game using the above technique.

4 Solving Symmetric Games

In this section, we analyze the computational complexity associated with pure Nash equilibrium in symmetric games with a constant number of actions and a growing number of actions, respectively.

4.1 Games with a Constant Number of Actions

As we have noted earlier, the potential hardness of finding pure Nash equilibria in games with succinct representations stems from the fact that the number of action profiles that are candidates for being an equilibrium is exponential in the size of the representation of the game. While weakly symmetric games certainly satisfy this property, the following lemma shows that the problem of deciding whether such a game possesses a pure Nash equilibrium is nevertheless tractable.

Lemma 8 *The problem of deciding whether a weakly symmetric or weakly anonymous game with a constant number of actions has a pure Nash equilibrium is in L.*

Proof: We propose an algorithm to decide whether there exists a pure Nash equilibrium s with $\#(s) = (w_{a_1},\dots,w_{a_k})$. Fixing a particular $x = (w_{a_1},\dots,w_{a_k})$, this algorithm can be divided into two phases:

1. For each $C \subseteq A$, compute the number w_C of players for which C is the set of *potential* pure best responses in x. We say that an action $a_\ell \in A$ is a potential best response for player i in the commutative image $x = (w_{a_1},\dots,w_{a_k})$ of an action profile for all players including i if $w_{a_\ell} > 0$ and

$$p_i(a_\ell, x_{-\ell}) \geq p_i(a_m, x_{-\ell}) \qquad \text{for all } a_m \in A, \tag{1}$$

   where $x_{-\ell} = (w_{a_1},\dots,w_{a_{\ell-1}}, w_{a_\ell}-1, w_{a_{\ell+1}},\dots,w_{a_k})$.

2. Check if the numbers computed in the first step are consistent with x, *i.e.*, if for each $C \subseteq A$ and each $c \in C$ there exists a non-negative integer $w_{(C,c)}$ such that

$$\sum_{c\in C} w_{(C,c)} = w_C \tag{2}$$

   and

$$\sum_{C\subseteq A} w_{(C,a)} = w_a \quad \text{for all } a \in A. \tag{3}$$

   In other words, $w_{(C,c)}$ denotes the number of players that have $C \subseteq A$ as their set of potential best responses in x and actually play $c \in A$, and Equations 2 and 3 ensure this number is consistent with the number w_C of players having C as possible best responses and the number w_c of players playing c.

Weights $(w_{(C,c)})_{C\subseteq A, c\in C}$ exist if and only if there is an action profile in which every player plays a best response in x, *i.e.*, a Nash equilibrium. Furthermore, if k is a constant, x, a constant number of values not larger than n, can be stored using only logarithmic space. It hence suffices to show that both of the above steps require only logarithmic space. The number of different coefficients w_C equals the cardinality of 2^A and is a constant if k is. Since each coefficient is an integer not larger than n, all of them can be stored using logarithmic space. Their computation from the input involves checking Inequality 1 for each player $i \in N$ and can be done using logarithmic space as well.

The problem faced in the second phase of the algorithm can alternatively be written as a homologous flow problem in a directed (almost bipartite) graph $G = (V,E)$ with nodes $V = A \cup 2^A \cup \{s,t\}$ and edges $E = \{\,(C,a) \in 2^A \times A \mid a \in C\,\} \cup (\{s\} \times 2^A) \cup (A \times \{t\})$. In a homologous flow problem, both a lower bound $\ell(e)$ and an upper bound $u(e)$ are given for the capacity of every edge $e \in E$ (see, *e.g.*, Greenlaw et al., 1995). If we let $\ell(s,C) = u(s,C) = w_C$ and $\ell(a,t) = u(a,t) = w_a$ for all $C \subseteq A$ and $a \in A$, and $\ell(e) = 0$, $u(e) = n$ otherwise, then a feasible flow for G exists if and only if $(w_C)_{C\subseteq A}$ and $(w_a)_{a\in A}$ are compatible. To see this, observe that Equations 2 and 3 constitute flow consistency conditions for all nodes but s and t, and that the size of every feasible flow through G must equal n. While this problem can be solved in polynomial time in the general case if G has a polynomial number of nodes and is in fact P-complete (Greenlaw et al.,

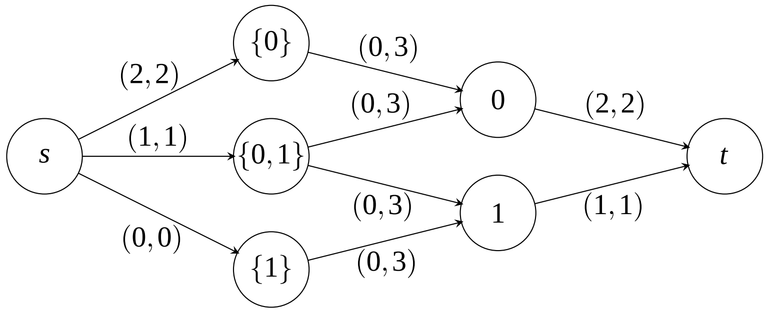

Figure 4: Integer flow network used in the proof of Lemma 8, example for the game of Figure 3. Edge e is labeled $(u(e), \ell(e))$.

1995), we will give an algorithm for our special case of the problem that requires only logarithmic space. As an example, the flow network for the game in Figure 3 is given in Figure 4. Edge capacities have been computed by checking for each player if his action in the respective (shaded) action profile of Figure 3 is a best response. This particular instance can easily be solved by assigning $w_{(\{0\},0)} = 2$, $w_{(\{0,1\},0)} = 0$, $w_{(\{0,1\},1)} = 1$, and $w_{(\{1\},1)} = 0$. In general, however, there need not be a unique solution, and the graph may contain (undirected) cycles, preventing a direct assignment of the weights. We therefore claim that the existence of a feasible flow means that there are particular weights $(w_{C,c})_{C\subseteq A, c\in C}$ satisfying an additional property, namely that there is a sequence $\langle e_1, \dots, e_m\rangle$ of all edges $e_j = (C_j, c_j)$ and an integer $i \leq m$ such that $w_{(C_j,c_j)} = 0$ if $j \leq i$ and

$$w_{(C_j,c_j)} = \min((w_{C_j} - \sum_{\substack{m<j \\ C_m=C_j}} w_{(C_m,c_m)}), (w_{c_j} - \sum_{\substack{m<j \\ c_m=c_j}} w_{(C_m,c_m)})) \tag{4}$$

otherwise. To see that this is indeed the case, consider weights satisfying Equations 2 and 3. Further assume w.l.o.g. that *every* (undirected) cycle in the flow network contains an edge e with weight zero. Otherwise, while the graph contains a cycle $\langle e_1, \dots, e_m\rangle$ (which must have even length, since the graph is bipartite) such that the weight w_{e_1} of e_1 is positive and minimal among all edges in the cycle, we modify the weights w_{e_i}, $1 \leq i \leq m$ according to

$$w_{e_i} := \begin{cases} w_{e_i} + w_{e_1} & \text{if } i \text{ is even} \\ w_{e_i} - w_{e_1} & \text{if } i \text{ is odd.} \end{cases}$$

Observe that after this modification, (i) Equations 2 and 3 are still satisfied and (ii) $w_{e_1} = 0$. If now we remove all edges with zero weight from the graph, we obtain an acyclic graph, which must have a node v with degree 1 if it contains at least one edge. If the latter was not the case, we could construct a sequence $\langle v_1, \dots, v_m\rangle$ for arbitrary m with $(v_i, v_{i+1}) \in E$ for $1 \leq i \leq m-1$ and $v_i \neq v_{i+2}$ for $1 \leq i \leq m-2$. For $m > |V|$, we would necessarily have $v_i = v_j$ for some $i \neq j$, and hence a cycle. Returning to the node v of degree one, we can greedily assign the weight to the sole edge (v, v') incident to v, remove (v, v') from the graph, and update the weights of v and v' accordingly. Repeating this process until no more edges remain, we obtain all weights $w_{C,c}$, $C \subseteq A$, $c \in C$. These weights satisfy Equations 2 and 3 if and only if all vertices in the remaining graph (with $E = \emptyset$) have weight zero.

Based on this observation, we can design a simple algorithm that enumerates all possible pairs $(\overline{e}, i)$ of a sequence $\overline{e} = \langle e_1, \dots, e_m\rangle$ and an index $i \leq m$ and tries to assign weights in a particular order. All e_j with $j \leq i$ receive weight 0. If one of the nodes incident to e_j, $j > i$ has degree 1 in the graph $(V, \{e_j, \dots, e_m\})$, w_{e_j} is set according to Equation 4. Otherwise the sequence is rejected. If k and hence the number of different pairs $(\bar{e}, i)$ is a constant, the algorithm requires only logarithmic space. The inclusion relationship between the different classes of symmetric games implies that the pure Nash equilibrium problem is in L for all kinds of symmetric games with a constant number of actions, and for weakly anonymous games in particular. $\square$

The flow network used to prove the above lemma has some rather unusual properties. On the one hand, its structure only depends on the number of actions in the game, and is predetermined if this number is a constant. On the other hand, edge capacities greatly depend on the number of players and on the payoff structure of the game. Hence, while showing that the sequential execution of the above algorithm requires only logarithmic space has been quite illustrative, the fixed structure of the flow network for a fixed number of actions raises the question if we can do better than that. As we will see, this is indeed the case. The following theorem states that the problem under consideration can be solved in TC^0, and is in fact TC^0-complete.

Theorem 9 *Deciding whether a weakly symmetric or weakly anonymous game with a constant number of actions has a pure Nash equilibrium is TC^0-complete. Hardness holds even if there is only a constant number of payoffs and only two different payoff functions.*

Proof: To show *membership*, we will return to the algorithm used in the proof of Lemma 8 to decide whether a weakly symmetric game Γ has a pure Nash equilibrium, and show that it can be realized as a threshold circuit with unbounded fan-in, constant depth, and a polynomial number of gates.

We start with the edge capacities computed in the first phase of the algorithm. For a fixed commutative image x, a particular player $i \in N$, and a particular action $a \in A$, we can easily construct a circuit of constant depth that checks whether Equation 1 is satisfied. To compute whether $C \subseteq A$ is the set of best responses for player i under x, we simply combine the outputs of the above circuits for all actions $a \in A$ into a single *AND* gate, negating the outputs of that for actions $a \notin C$. w_C is then obtained by adding up the outputs of these gates for all players $i \in N$. Clearly, the number of gates in this circuit is polynomial if the number of actions is a constant.

As for the the second phase of the algorithm, we construct a circuit that computes whether a feasible flow can be found by assigning weights to edges using the algorithm of Lemma 8 and according to a particular pair $(\langle e_1, \ldots, e_m \rangle, i)$ of a sequence of all edges and an index in this sequence. When a weight is assigned to edge e_j, $j > i$, the weights of the nodes incident to e_j have to be updated by subtracting the weight just assigned. Clearly, the new weights can be computed using a constant-depth circuit, and since m is a constant, these circuits can be layered. The jth layer receives as inputs the weights before a weight has been assigned to e_j, and outputs the updated weights after this has been done. Furthermore, it outputs an additional bit that is *true* if and only if the assignments up to and including the jth step have been consistent. If the latter is *true* at the final layer, and the updated weights are zero, the circuit outputs *true*. Finally, since there is only a constant number of pairs $(\bar{e}, i)$, the outputs of the above circuits can be combined into a single *OR* gate to obtain a circuit with constant depth and a polynomial number of gates that decides whether Γ has a pure Nash equilibrium.

For *hardness*, we reduce the problem of deciding whether exactly ℓ bits of a string of m bits are 1 to deciding the existence of a pure Nash equilibrium in a weakly anonymous game. Hardness of the former problem is immediate from that of MAJORITY (see, *e.g.*, Chandra et al., 1984). For a particular m-bit string b, we define a game Γ with $m+2$ players of two different types 0 and 1 and actions $A = \{0, 1\}$. The ith player of Γ is of type 0 or 1 if the ith bit of b is 0 or 1, respectively. Player $m+1$ is of type 0, player $m+2$ is of type 1. The payoffs p_0 and p_1 for the two types are given in Figure 5, the column labeled j specifies the payoff when exactly j players, including the player himself, play action 1. It is easily verified that this is an AC^0 reduction. We claim that Γ possesses a pure Nash equilibrium if and only if exactly ℓ bits of b are 1. We observe the following:

p_0	0		$\dots$			$\ell+1$			$\dots$		$m+2$
	$\dots$	0	1	0	2	1	0	1	0	1	$\dots$

p_1	0		$\dots$			$\ell+1$			$\dots$		$m+2$
	$\dots$	1	0	1	0	1	2	0	1	0	$\dots$

Figure 5: Payoffs of the game Γ used in the proof of Theorem 9

- An action profile s cannot be a Nash equilibrium of Γ if $\#(1,s) \neq \ell+1$. In this case, the players of one of the two types get a higher payoff at both $\#(1,s)-1$ and $\#(1,s)+1$ (or at one of these in case $\#(1,s)=0$ and $\#(1,s)=m+2$). Since by construction we have at least one player of each type, there always exists a player who can change his action to get a higher payoff.
- If there are $\ell+1$ players of type 1, the action profile where all players of type 0 play action 0 and all players of type 1 play action 1 is a Nash equilibrium. None of the players of type 0 can gain by changing his action to 1, and none of them can change his action to 0 (because all of them already play 0). A symmetric condition holds for players of type 1.
- In turn, if the number of players of type 1 does not equal $\ell+1$, an action profile s with $\#(s,1)=\ell+1$ cannot be a Nash equilibrium. In this case, there must be (i) a player of type 0 playing action 1 in s, or (ii) a player of type 1 playing 0. This player can change his action to get a higher payoff.

Hence, a pure Nash equilibrium exists if and only if there are $\ell+1$ players of type 1, *i.e.*, if and only if b has ℓ 1-bits. This completes the reduction. □

In contrast to weakly symmetric games, if s is a Nash equilibrium of a *strongly* symmetric game, so are all t satisfying $\#(t)=\#(s)$. This is due to the fact that the payoff functions of all players, and thus the situation of all players playing the same action $a \in A$, is identical, as would be the situation of any other player exchanging actions with someone playing a. We exploit this property to show that deciding the existence of a Nash equilibrium in strongly symmetric games with a constant number of actions is strictly easier than for weakly symmetric or weakly anonymous games.

Theorem 10 *The problem of deciding whether a strongly symmetric game with a constant number of actions has a pure Nash equilibrium is in AC^0.*

Proof: Like with weakly symmetric games, an action profile s is a Nash equilibrium of a strongly symmetric game if and only if, for all $i \in N$, s_i is a best response to $\#(s_{-i})$, *i.e.*, if

$$p_i(s_i,\#(s_{-i})) \geq p_i(a,\#(s_{-i})) \qquad \text{for all } a \in A. \tag{5}$$

For a particular player $i \in N$ and for constant k, checking this inequality requires only a constant number of comparisons and can be done using a circuit of constant depth and polynomial size (see, *e.g.*, Chandra et al., 1984). When it comes to checking Equation 5 for the different players, the observation about action profiles with identical commutative images affords us a considerable computational advantage as compared to, say, weakly symmetric or weakly anonymous games. More precisely, we only have to check if Equation 5 is satisfied for a player *playing a certain action*, of which there are at most k. Again, this can be done using a circuit of constant depth and polynomial size if k is a constant.

Finally, to decide whether game Γ has a pure Nash equilibrium, we have to check Equation 5 for the different values of $\#(s)$ for $s \in A^N$. If k is constant, there are only polynomially many of these, so the complete check requires only polynomial size and constant depth. □

This proof provides a nice illustration of the fact that every strongly symmetric game with two actions possesses a pure Nash equilibrium, as recently shown by Cheng et al. (2004). In this case, p_i depends only on player i's action (0 or 1) and on the number of other players playing action 1. A pure Nash equilibrium exists if for some m neither the players playing 0 (who see m players playing 1) nor the players playing 1 (who see $m-1$ *other* players playing 1) have an incentive to deviate, *i.e.*, $p_i(0,m) \geq p_i(1,m)$ and $p_i(1,m-1) \geq p_i(0,m-1)$. For $m=0$ and $m=n$, one of the conditions is trivially satisfied, because there are no players playing 1 or 0, respectively. It is easily verified that at least one such m must exist. Alternatively, the existence of pure Nash equilibria in strongly symmetric games with two actions can be obtained as an immediate consequence of Fact 7. We can transform every strongly symmetric game with two actions into a strongly anonymous game with the same set of equilibria, and every strongly anonymous is guaranteed to have at least one pure equilibrium.

As we have already said, strongly *anonymous* games always possess a pure Nash equilibrium. We proceed to show that we can find one that maximizes the sum of payoffs of all players in AC^0.

Theorem 11 *The problem of finding a social-welfare-maximizing pure Nash equilibrium of a strongly anonymous game with a constant number of actions is in AC^0.*

Proof: Since strongly anonymous games belong to the class of common payoff games, any action profile with maximum payoff (for all players) is a social-welfare-maximizing Nash equilibrium (and Pareto-dominates any other strategy profile). Finding such an equilibrium is thus equivalent to finding the maximum of $\binom{n+k-2}{k-1}$ integers. The exact number is irrelevant as long as it is polynomial in the size of the input which, according to Fact 5, is certainly the case if k is bounded by a constant. Chandra et al. (1984) have shown that the maximum of m m-bit binary numbers can be computed by an unbounded fan-in, constant-depth Boolean circuit of size polynomial in m. Since m is of course polynomial in the size of the input, the size of this circuit is as well. $\square$

4.2 Games with a Growing Number of Actions

The proofs we have seen in the previous section rely on the fact that for constant k the naive representation of a symmetric game (*i.e.*, in terms of payoff tables) is computationally equivalent to any kind of polynomially computable payoff function because we can transform the latter representation into the former by means of a log-space reduction. This is no longer the case for unbounded k, because the size of the naive representation grows exponentially in n. However, a succinct representation of the payoff function (*e.g.*, a Boolean circuit) might exist for certain classes of games.

We will now show that deciding the existence of a pure Nash equilibrium in weakly and strongly symmetric and weakly anonymous games becomes NP-hard if the number of actions grows in n. For strongly anonymous games, which always have a Nash equilibrium, the associated search problem will be shown to be PLS-hard. In the following, we will only consider games where (i) the payoff *to all players* can be computed in polynomial time and (ii) a single player can check in polynomial time whether a particular action is a best response to a given action profile for the other players. Under this assumption, which is quite reasonable for "natural" games, we will be able to obtain membership in NP or PLS, respectively. All *hardness* results hold irrespective of this assumption. While there certainly are meaningful games with an exponential number of players or actions, the complexity in this case mainly stems from the sheer size of the game rather than the actual problem of finding a Nash equilibrium.

For the following proofs, recall that circuit satisfiability (CSAT), i.e., deciding whether a Boolean circuit has a satisfying assignment, is NP-complete (see, *e.g.*, Papadimitriou, 1994). We provide a reduction from CSAT to the problem of deciding the existence of a pure Nash equilibrium in a special class of games. For a particular circuit $\mathscr{C}$ with inputs $M = \{1, \ldots, m\}$, we define a game Γ with players $N = M$ and actions

$A = \{a_i^0, a_i^1 \mid i \in M\}$. An action profile s of Γ where $\#(a_i^0, s) + \#(a_i^1, s) = 1$ for all $i \in M$, *i.e.*, one where exactly one action of each pair a_i^0, a_i^1 is played, directly corresponds to an assignment c of $\mathscr{C}$, the ith bit c_i of this assignment being $j \in \{0,1\}$ if a_i^j is played. We can thus distinguish between the action profiles of Γ corresponding to a satisfying assignment of $\mathscr{C}$, those corresponding to a non-satisfying assignment, and those not corresponding to an assignment at all.

Theorem 12 *Deciding whether a weakly anonymous game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Proof: If the number of players and actions is polynomial in the input size, and if the payoff function is computable in polynomial time, *membership* in NP is immediate. We can simply guess an action profile and verify that it satisfies the equilibrium condition.

To show *hardness*, we reduce CSAT to the problem at hand by mapping a particular circuit $\mathscr{C}$ with inputs $M = \{1, \dots, m\}$ to a game Γ with players $N = M$, actions $A = \{a_i^0, a_i^1 \mid i \in M\}$, and payoff functions p_i as follows:

- If s corresponds to a *satisfying* assignment of $\mathscr{C}$, we let $p_i(s) = 2$ for all $i \in N$.
- If s corresponds to an assignment that does not satisfy $\mathscr{C}$, we let
  - $p_1(s) = 2,\ p_2(s) = 1$ if $|\{i \in M \mid \#(a_i^0, s) > 0\}|$ is even, *i.e.*, an even number of 0-actions is played by at least one player, and
  - $p_1(s) = 1,\ p_2(s) = 2$ if this number is odd.
  - For all $i \in N \setminus \{1,2\}$, we let $p_i(s) = 2$.
- If s does not correspond to an assignment of $\mathscr{C}$, we let $p_i(s) = 1$ if $\#(a_i^0, s) + \#(a_i^1, s) > 0$, and $p_i(s) = 0$ otherwise.

We observe the following:

- Γ (*e.g.*, Boolean circuits that compute p_i) can be constructed from $\mathscr{C}$ in polynomial time.
- For all of the above cases, the payoff of player i only depends on the number of players playing certain actions. If two players exchange actions, the payoff to all other players remains the same. Hence, Γ is weakly anonymous.
- Clearly, every action profile s corresponding to a satisfying assignment of $\mathscr{C}$ is a Nash equilibrium, because in this case all players receive the maximum payoff of 2.
- In any other case, s cannot be a Nash equilibrium. If s corresponds to a non-satisfying assignment of $\mathscr{C}$, either player 1 or player 2 can change his action to get a higher payoff, depending on whether the number of actions a_i^0 played by at least one player is even or odd. If s does not correspond to an assignment of $\mathscr{C}$, there exists $i \in M$ such that $\#(a_i^0, s) + \#(a_i^1, s) = 0$, so player i can change to either a_i^0 or a_i^1 to get a higher payoff.

Hence, there is a direct correspondence between satisfying assignments of $\mathscr{C}$ and Nash equilibria of Γ. This completes the reduction. $\square$

Theorem 13 *Deciding whether a strongly symmetric game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Proof: Under the same conditions as in the previous theorem, *membership* in NP is immediate.

For *hardness*, we again provide a reduction from CSAT, mapping a circuit $\mathscr{C}$ with inputs $M = \{1,\dots,m\}$ to a game Γ with players $N = M$, actions $A = \{\, a_i^0, a_i^1 \mid i \in M \,\}$, and payoff functions p_i as follows:

- If s corresponds to a satisfying assignment of $\mathscr{C}$, we let $p_i(s) = 3$ for all $i \in N$.
- If s does not correspond to a satisfying assignment of $\mathscr{C}$, we let
  - $p_i(s) = 2$ if $s_i = a_j^1$ for some $j \in M$, $\#(a_j^0, s) > 0$, and $\#(a_j^1, s) > 0$,
  - $p_i(s) = 1$ if $s_i = a_j^0$ for some $j \in M$, $\#(a_j^0, s) > 0$, and $\#(a_j^1, s) = 0$, and
  - $p_i(s) = 0$ otherwise.

We observe the following:

- Γ (*e.g.*, Boolean circuits that compute p_i) can be constructed from $\mathscr{C}$ in polynomial time.
- For all of the above cases, the payoff of player i only depends on his own action and on the number of players playing certain other actions. If two players exchange actions, their payoffs are also exchanged. Hence, Γ is strongly symmetric.
- Clearly, any action profile corresponding to a satisfying assignment of $\mathscr{C}$ is a Nash equilibrium, because in this case all players receive the maximum payoff of 3. In turn, if s does not correspond to a satisfying assignment, we have one of two cases, in both of which s is not a Nash equilibrium:
  - If $\#(a_j^0, s) = 1$ for all $j \in M$, player $i \in N$ can change to some a_m^1 such that $s_i \neq a_m^0$ to get a higher payoff.
  - Otherwise, there has to be some player $i \in N$ who gets payoff 0, and, by the pigeonhole principle, some $j \in M$ such that $\#(a_j^0, s_{-i}) = \#(a_j^1, s_{-i}) = 0$. Then, player i can change to a_j^0 to get a higher payoff.

Again, there is a direct correspondence between Nash equilibria of Γ and satisfying assignments of $\mathscr{C}$. This completes the reduction. □

By each of the previous two theorems and by the inclusion relationships between the different classes of symmetric games, we also have the following.

Corollary 14 *Deciding whether a weakly symmetric game has a pure Nash equilibrium is NP-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.* □

In the proofs of Theorems 12 and 13, every satisfying assignment of circuit $\mathscr{C}$ corresponds to a certain number of pure Nash equilibria of game Γ. This allows us to show that counting the number of Nash equilibria in these games is hard.

Corollary 15 *For weakly symmetric, weakly anonymous, and strongly symmetric games, counting the number of pure Nash equilibria is #P-complete, even if the number of actions is linear in the number of players and there is only a constant number of different payoffs.*

Proof: Recall that in the proof of Theorem 12, actions of the game Γ are identified with inputs of the Boolean circuit $\mathscr{C}$. As a direct consequence of anonymity or symmetry, it does not matter which player plays a particular action to assigns a value to the corresponding gate. Every satisfying assignment of $\mathscr{C}$ thus corresponds to $n!$ equilibria of Γ, so the number of satisfying assignments can be determined by counting the number of Nash equilibria, of which there are at most $2^n n!$, and dividing this number by $n!$. Division of two m-bit binary numbers can be done using a circuit with bounded fan-in and depth $O(\log m)$ (Beame et al., 1986). For $m = \log(2^n n!) = O(n^2)$, we have depth $O(\log n^2) = O(\log n)$, so the above division can be carried out in NC^1. We have thus found a reduction of the problem #SAT of counting the number of satisfying assignments of $\mathscr{C}$, which is #P-complete (see, *e.g.*, Papadimitriou, 1994), to the problem of counting the Nash equilibria of Γ. The same line of reasoning applies to the proof of Theorem 13. By Corollary 14, #P-completeness extends to weakly symmetric games. □

As we have already outlined above, every strongly anonymous game possesses a pure Nash equilibrium. Theorem 11 states that finding even a social-welfare-maximizing one is very easy as long as the number of actions is bounded by a constant. If now the number of actions is growing but polynomial in the size of the input, an assumption we have made throughout the paper, we can start at an arbitrary action profile and check in polynomial time whether some player can change his action to increase the (common) payoff. If this is not the case, we have found an equilibrium. Otherwise, we can repeat the process for the new profile, resulting in a procedure called *best-response dynamics* in game theory. Since the payoff strictly increases in every step, we are guaranteed to find a Nash equilibrium in polynomial time if the number of different payoffs is polynomial. In turn, we will show that, given a strongly anonymous game with a growing number of actions and an exponential number of different payoffs, finding a Nash equilibrium is at least as hard as finding a *locally optimal* solution to an NP-hard optimization problem. For this, we formally introduce the class of search problems for which a solution is guaranteed to exist by a local optimality argument.

Definition 16 (local search, PLS) *A* local search problem *is given by (i) a set $\mathscr{I}$ of instances, (ii) a set $\mathscr{F}(x)$ of feasible solutions for each $x \in \mathscr{I}$, (iii) an integer measure $\mu(S,x)$ for each $S \in \mathscr{F}(x)$, and (iv) a set $\mathscr{N}(S,x)$ of neighboring solutions for each $S \in \mathscr{F}(x)$. A solution is* locally optimal *if it does not have a strictly better neighbor,* i.e.*, one with a higher or lower measure depending on the kind of optimization problem.*

A local search problem is in the class PLS of polynomial local search *problems (Johnson et al., 1988) if for every $x \in \mathscr{I}$ there exist polynomial time algorithms for (i) computing an initial feasible solution in $\mathscr{F}(x)$, (ii) computing the measure $\mu(S,x)$ of a solution $S \in \mathscr{F}$, and (iii) determining that S is locally optimal or finding a better solution in $\mathscr{N}(S,x)$.*

*A problem P in PLS is PLS-*reducible *to another problem Q in PLS if there exist polynomial time computable functions Φ and Ψ mapping (i) instances x of P to instances $\Phi(x)$ of Q and (ii) solutions S of an instance $\Phi(x)$ of Q to solutions $\Psi(S,x)$ of the corresponding instance x of P such that locally optimal solutions are mapped to locally optimal solutions. A PLS reduction from P to Q is called* tight *(Schäffer and Yannakakis, 1991) if for any instance x of P there exists a set $\mathscr{R} \subseteq \mathscr{F}(\Phi(x))$ with the following properties:*

1. *$\mathscr{R}$ contains all local optima of $\Phi(x)$.*

2. *For every $p \in \mathscr{F}(x)$, a solution $q \in \mathscr{R}$ satisfying $\Psi(q,x) = p$ can be computed in polynomial time.*

3. *Consider $q_0, \dots, q_\ell \in \mathscr{F}(\Phi(x))$ such that $q_0, q_\ell \in \mathscr{R}$, $q_i \notin \mathscr{R}$ for all $0 < i < \ell$, $q_{i+1} \in \mathscr{N}(q_i, \Phi(x))$ for all $i < \ell$, and $\mu(q_i) > \mu(q_j)$ if $i > j$. Let $p = \Psi(q_0, x)$, $p' = \Psi(q_\ell)$. Then, either $p = p'$ or $p' \in \mathscr{N}(p,x)$.*

Theorem 17 *The problem of finding a pure Nash equilibrium in a strongly anonymous game is PLS-complete, even if the number of actions is linear in the number of players.*

Proof: Neighborhood among action profiles is given by a single player changing his action. If the number of players and actions is polynomial in the input size, and if the payoff function is computable in polynomial time, *membership* in PLS is immediate.

For *hardness*, consider a Boolean circuit $\mathscr{C}$ with inputs $M = \{1,\dots,m\}$ and ℓ outputs. Finding an assignment such that the output interpreted as an ℓ-bit binary number is a local maximum under the FLIP neighborhood (*i.e.*, changing a single input bit) is known to be PLS-complete (see Johnson et al., 1988; Schäffer and Yannakakis, 1991). We provide a PLS reduction to the problem of finding a Nash equilibrium in a strongly anonymous game by mapping a particular circuit $\mathscr{C}$ as described above to a game Γ with players $N = M$, actions $A = \{\, a_i^0, a_i^1 \mid i \in M \,\}$, and a (common) payoff function p as follows:

- If s corresponds to an assignment c of $\mathscr{C}$, we let $p(s) = n + \mathscr{C}(c)$, where $\mathscr{C}(c)$ denotes the output of $\mathscr{C}$ for input c, interpreted as a binary number.
- Otherwise, we let $p(s) = |\{\, i \in M \mid \#(a_i^0, s) + \#(a_i^1, s) > 0 \,\}|$. That is, the payoff is at most $n-1$ and decreases in the minimum number of players that would have to change their action in order to make s correspond to an assignment of $\mathscr{C}$.

We observe the following:

- Obviously, Γ is a common payoff game. Since p is invariant under any permutation of the players in both of the above cases, Γ is strongly anonymous.
- Γ (*e.g.*, a Boolean circuit that computes p) can be constructed from $\mathscr{C}$ in polynomial time. Hence, there exists a polynomial time computable function that maps instances of FLIP to instances of the problem under consideration.
- An action profile a that does not correspond to a valid assignment of $\mathscr{C}$ cannot be a Nash equilibrium of Γ. In this case there always exist $i, j \in M$ such that a_i^0 and a_i^1 are played by more than one player while no one plays a_j^0 or a_j^1. If one of the players playing the former changes to the latter, he gets a higher payoff (actually, all players do).
- There is a direct correspondence between the FLIP neighborhood of $\mathscr{C}$ and a single player changing between a_i^0 and a_i^1 for some $i \in M$. Furthermore, changing to an action profile that does not correspond to an assignment of $\mathscr{C}$ will get the player strictly less payoff. Thus, there is a direct correspondence between Nash equilibria of Γ and local maxima of $\mathscr{C}$ under the FLIP neighborhood. Obviously, the assignment corresponding to an action profile can be computed in polynomial time (if such an assignment exists). The conditions of Definition 16 do not require that we map solutions of Γ that are not locally optimal to solutions of $\mathscr{C}$ that are not locally optimal. This means that action profiles not corresponding to an assignment can simply be mapped to an arbitrary assignment.

We observe that this satisfies the properties of a PLS reduction. $\square$

Implicit in the definition of PLS is a *standard algorithm* for finding a locally optimal solution for a given input $x \in \mathscr{I}$: start with an arbitrary feasible solution $S \in \mathscr{F}(x)$ and repeatedly find a strictly better neighbor until a locally optimal solution $T \in \mathscr{F}(x)$ has been found. The *standard algorithm problem* can be phrased as follows: given x, find the locally optimal solution T output by the standard algorithm on input x. By the above proof, we can draw some additional conclusions about the worst-case running time of the standard algorithm and about the hardness of the standard algorithm problem.

Corollary 18 *The standard algorithm for finding Nash equilibria in strongly anonymous games has an exponential worst-case running time. The standard algorithm problem is NP-hard.*

Proof: To show tightness of the reduction used in the proof of the previous theorem, choose $\mathscr{R}$ to be the set of actions profiles of Γ that correspond to an assignment of $\mathscr{C}$. Obviously, $\mathscr{R}$ contains all optimal solutions, and a payoff profile corresponding to a particular assignment can be computed in polynomial time. The third condition is trivially satisfied because the measure of any solution inside $\mathscr{R}$ is strictly greater than that of any solution outside of $\mathscr{R}$. The corollary then follows directly from Lemma 3.3 in (Schäffer and Yannakakis, 1991). $\square$

By a slight modification of the proof of Theorem 17, PLS-completeness, exponential worst-case running time of the standard algorithm, and NP-hardness of the standard algorithm problem can also be shown for general (*i.e.*, not necessarily symmetric) common payoff games with $k = 2$. This fact nicely illustrates the influence of symmetry on the hardness of finding (or deciding the existence of) a Nash equilibrium.

5 Threshold Symmetries

In order to extend the basic concept of symmetry as the indistinguishability of players, we will now consider games where the players cannot even observe the exact number of players playing a certain action, but only whether this number reaches certain *thresholds*. Let $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ be a normal-form game and A a set of actions such that $A_i = A$ for all $i \in N$. For $T \subseteq \{1, \ldots, n\}$, let $\sim_T \subseteq A^N \times A^N$ be defined as follows: $s \sim_T t$ if for all $a \in A$ and all $x \in T$, $\#(a,s) < x$ if and only if $\#(a,t) < x$. $\sim_T$ naturally extends to action profiles for subsets of N. The following is easily verified.

Fact 19 *For any $T \subseteq \{1, \ldots, n\}$, $\sim_T$ is an equivalence relation on the set A^M of action profiles for players $M \subseteq N$.*

Based on $\sim_T$, we can give a more general version of Definition 2.

Definition 20 (threshold symmetry) *Let $\Gamma = (N, (A_i)_{i\in N}, (p_i)_{i\in N})$ be a normal-form game, A a set of actions such that $A_i = A$ for all $i \in N$. Let $T \subseteq \{1, \ldots, n\}$. Γ is called*

- weakly T-symmetric *if $p_i(s) = p_i(t)$ for all $i \in N$ and all $s, t \in A^N$ with $s_i = t_i$ and $s_{-i} \sim_T t_{-i}$,*
- strongly T-symmetric *if $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s, t \in A^N$ with $s_i = t_j$ and $s_{-i} \sim_T t_{-j}$,*
- weakly T-anonymous *if $p_i(s) = p_i(t)$ for all $i \in N$ and all $s, t \in A^N$ with $s \sim_T t$, and*
- strongly T-anonymous *if $p_i(s) = p_j(t)$ for all $i, j \in N$ and all $s, t \in A^N$ with $s \sim_T t$.*

For $T = \{1, \ldots, n\}$, these classes are equivalent to those of Definition 2. This is immediate from Fact 4. Moreover, we obtain *Boolean symmetry*, where payoffs only depends on the *support* of an action profile (*i.e.*, the actions that are played by at least one player), for $T = \{1\}$. In general, we call a game *threshold symmetric* (for one of the above classes) if it is T-symmetric for some T (and the corresponding class).

Obviously, the number of payoffs that need to be written down for each player to specify a general weakly T-symmetric game is exactly the number of equivalence classes of $\sim_T$ for action profiles of the other players.

Fact 21 *A weakly T-symmetric game can be represented using at most $n \cdot k \cdot |A^{n-1}/\sim_T|$ numbers, where $X/\sim$ denotes the quotient set of set X by equivalence relation $\sim$. For Boolean weak symmetry, the number of equivalence classes equals the number of k-bit binary numbers where at least one bit is* 1, i.e., $2^k - 1$. *More generally, there cannot be more than $(|T|+1)^k$ equivalence classes if $|T|$ is bounded by a constant (since for every action, the number of players playing this action must be between two thresholds), while for $T = \{n\}$ there are as few as $k+1$. Hence, any T-symmetric game with constant $|T|$ is succinctly representable if $k = O(\log n)$.*

Theorem 22 *For threshold symmetric games with $k = O(\log n)$ and a constant number of thresholds, deciding the existence of a Nash equilibrium is in P.*

Proof: Like in the proof of Lemma 8, we provide an algorithm that checks whether there is a Nash equilibrium in a particular quotient set $A^N/\sim_T$ of the set of payoff profiles. Since for $k = O(\log n)$ and $|T| = O(1)$, the cardinality of $A^N/\sim_T$ is polynomial in n, it suffices to show that the algorithm requires only polynomial time for every such set.

For a particular element $X \in A^N/\sim_T$, the algorithm is again divided into two phases: (i) computing the set of best responses for each player under X, and (ii) checking whether there is a particular action profile $s \in X$ where each player plays a best response.

In the first phase, and unlike the case $T = \{1, \dots, n\}$ covered by Lemma 8, the action a played by player $i \in N$ may or may not yield a different element of $A^{N \setminus \{i\}}/\sim_T$ against which a should be a best response. Instead of just looking for best responses under elements of T^N, we thus look for best responses under those of U^N, where $U = \{\, u \leq n \mid u \in T \text{ or } (u-1) \in T \,\}$. Since the cardinalities of both U^N and of the set possible best responses is polynomial if $|T| = O(1)$ and $k = O(\log n)$, the first phase requires only polynomial time.

As for the second phase, we recall that it can be reduced to deciding the existence of a feasible flow in a homologous flow network with $O(2^k)$ nodes. Since this problem is in P if the number of nodes is polynomial (see, *e.g.*, Greenlaw et al., 1995), observing that 2^k is polynomial if $k = O(\log n)$ completes the proof. □

In turn, it is easily verified that all the games defined in the proofs of Theorems 12, 13, and 17 are Boolean. Action profiles corresponding to an assignment of a circuit trivially satisfy the conditions of Definition 20, since each action is played by either zero or one players. For all other action profiles, the conditions have to be checked individually. In the proof of Theorem 12, for example, the payoff of player i only depends on whether a_i^0 or a_i^1 is played by at least one player. We thus have the following corollary.

Corollary 23 *Deciding the existence of a pure Nash equilibrium is NP-complete for threshold weakly symmetric, threshold weakly anonymous, and threshold strongly symmetric games, even if thresholds are Boolean, the number of actions is linear in the number of players, and there is only a constant number of different payoffs. For the same classes, counting the number of pure Nash equilibria is #P-complete.*

For threshold strongly anonymous games, finding a pure Nash equilibrium is PLS-complete, even if thresholds are Boolean. □

6 Conclusion and Future Work

In this paper, we have introduced four notions of symmetry in strategic multi-player games and investigated the computational complexity of finding pure Nash equilibria. This problem has been shown tractable for games with a constant number of actions, but intractable if the number of actions is linear in the number of players. It is worth noting that, for games with a constant number of actions, the Nash equilibrium problem

<table>
<tr><th></th><th>$k = O(1)$</th><th>$k = O(n)$</th></tr>
<tr><td>weakly symmetric</td><td rowspan="2">TC^0-complete</td><td rowspan="3">NP-complete</td></tr>
<tr><td>weakly anonymous</td></tr>
<tr><td>strongly symmetric</td><td rowspan="2">in AC^0</td></tr>
<tr><td>strongly anonymous</td><td>PLS-complete</td></tr>
</table>

Table 2: Complexity of Nash equilibrium in symmetric games

happens to lie in NC^1 for all types of symmetry and is thus open to parallel computation. For games in which the number of actions grows slowly (*e.g.*, logarithmically) in the number of players, the complexity remains open. The main results are summarized in Table 2.

In future work, it would further be interesting to investigate the notion of a *player type* to obtain efficient algorithms for more general classes of games. For example, games where indistinguishability holds only for players of the same type can be obtained by restricting Definition 2 to permutations that map players from a certain subset to players of the same set. We conjecture that using the algorithm of Lemma 8, pure Nash equilibria can still be found in polynomial time if the number of player types is constant. A different notion, such that players of the same type have identical payoff functions, does not seem to provide additional structure. As we have already shown, only two different payoff functions suffice to make the Nash equilibrium problem TC^0-hard for a constant number of actions and NP-hard for a growing number of actions. More generally, one might investigate games where payoffs are invariant under particular sets of permutations. Von Neumann and Morgenstern (1947) regard the number of permutations under which the payoffs of a game are invariant as a measure for the degree of symmetry. The question is in how far the computational complexity of solving a game depends on the degree of symmetry.

Acknowledgements

We thank Jan Johannsen for enlightening discussions on circuit complexity and local search, Rob Powers for introducing the first author to the ambiguity of symmetry in games, and Christian Schallhart for pointing us to homologous flows.

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

References

R. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1: 67–96, 1974.

P. Beame, S. A. Cook, and H. J. Hoover. Log depth circuits for division and related problems. *SIAM Journal of Computing*, 15(4):994–1003, 1986.

N. Bhat and K. Leyton-Brown. Computing Nash equilibria of action-graph games. In *Proceedings of the 20th International Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 35–42. AUAI Press, 2004.

E. Borel. La théorie du jeu et les équations intégrales à noyau symétrique. *Comptes Rendus de l'Académie des Sciences*, 173:1304–1308, 1921.

A. K. Chandra, L. Stockmeyer, and U. Vishkin. Constant depth reducibility. *SIAM Journal on Computing*, 13(2):423–439, 1984.

X. Chen and X. Deng. Settling the complexity of 2-player Nash-equilibrium. Technical Report TR05-140, Electronic Colloquium on Computational Complexity (ECCC), 2005.

Shih-Fen Cheng, Daniel M. Reeves, Yevgeniy Vorobeychik, and Michael P. Wellman. Notes on equilibria in symmetric games. In *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents (GTDT)*, 2004.

C. Daskalakis, P. Goldberg, and C. Papadimitriou. The complexity of computing a Nash equilibrium. In *Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC)*. ACM Press, 2006.

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.

G. Gottlob, G. Greco, and F. Scarcello. Pure Nash equilibria: Hard and easy games. *Journal of Artificial Intelligence Research*, 24:195–220, 2005.

S. Govindan and R. Wilson. A global Newton method to compute Nash equilibria. *Journal of Economic Theory*, 110(1):65–86, 2003.

R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. *Limits to Parallel Computation*. Oxford University Press, 1995.

S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI)*, pages 489–494. AAAI Press, 2005.

A. X. Jiang and K. Leyton-Brown. A polynomial-time algorithm for action-graph games. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*. AAAI Press, 2006. To Appear.

D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2, pages 67–161. Elsevier, 1990.

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

K. Leyton-Brown and M. Tennenholtz. Local-effect games. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 772–777, 2003.

R. D. Luce and H. Raiffa. *Games and Decisions: Introduction and Critical Survey*. Wiley, 1957.

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

C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

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

R. Parikh. On context-free languages. *Journal of the ACM*, 13(4):570–581, 1966.

R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM Journal on Computing*, 20(1):56–87, 1991.

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.

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

J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, 2nd edition, 1947.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de/