

Symmetric Nash equilibria

Steen Vester

December 2012

Research report LSV-12-23



Laboratoire Spécification & Vérification

École Normale Supérieure de Cachan
61, avenue du Président Wilson
94235 Cachan Cedex France

Symmetric Nash Equilibria

Steen Vester

Supervisors: Patricia Bouyer-Decitre & Nicolas Markey,
Laboratoire Spécification et Vérification, ENS de Cachan

August 27, 2012

Summary

In this report we study the computational problem of determining existence of pure, symmetric Nash equilibria in concurrent games as well as in symmetric concurrent game structures. Games have traditionally been used mostly in economics, but have recently received an increasing amount of attention in computer science with applications in several areas including logic, verification and multi-agent systems. Our motivation for studying symmetric Nash equilibria is to let the players in a game represent programmable devices and let strategies represent programs. Nash equilibria then correspond to a system where all devices are programmed optimally in some sense, whereas symmetric Nash equilibria correspond to systems where all devices are programmed in the same way while still preserving the optimality. The idea of modelling the interaction in distributed systems by means of games is in many cases more realistic than traditional analysis where opponents are assumed to act in the worst possible way instead of acting rationally and in their own interest.

Symmetry in games has been studied to some extent in classical normal-form games and our goal is to extend the notions of symmetry to concurrent games and investigate the computational complexity of finding symmetric Nash equilibria in these games. A number of different settings and types of symmetry are introduced and analyzed. Since infinite concurrent games have not been studied thoroughly for a lot of years yet there are still many unexplored branches in the area. Some of the settings studied in the report are completely new, whereas others have a lot of resemblance with problems that have already been analyzed quite extensively. In this case we can reuse some of the same proof techniques and obtain similar results.

Initially, we will study the problem of finding pure Nash equilibria in concurrent games where every player uses the same strategy. This

work closely resembles the work done in [2, 3, 4, 5] where the problem of finding pure Nash equilibria in concurrent games are studied. We provide a number of hardness proofs as well as algorithms altered to handle the requirement that every player must use the same strategy for a number of different objectives. In these problems we obtain the same upper and lower bounds as in the non-symmetric case.

We then proceed to define concurrent games where the symmetry is built into the game, which means that the games should be the same from the point of view of every player and that players should in some sense be interchangeable. Here we still keep perfect information and only let players act based on the (sequences of) global states of the game. In this section we question a definition of symmetry in normal-form games provided in [8] when trying to extend the notion of symmetry to concurrent games. We also prove that in our definition of a symmetric concurrent game there does not necessarily exist a symmetric Nash equilibrium in mixed strategies as is the case for normal-form games.

In the last part of the report we present a family of games which we call *symmetric concurrent game structures* in which each player controls his own personal transition system. Each player also has some information about the state of the other players and has objectives that concern both the state of his own system as well as the other players. The setting is quite general with possibility of incomplete information and even of making teams of players with the same objectives, while still retaining symmetry between the players. In the general case we prove undecidability of the existence problem of (symmetric) pure Nash equilibria even for two players and perfect information. It follows that finding pure Nash equilibria in concurrent games is undecidable, which was already known but with a different proof. We also show that in some sense the existence problem for symmetric pure Nash equilibria is at least as hard as the existence problem for pure Nash equilibria. Finally, we study the problem of finding (symmetric) Nash equilibria in m -bounded suffix strategies which is a generalization of memoryless strategies. In this case we provide a connection with model-checking path logics over interpreted, rooted transition systems in the sense that if it is decidable to model-check such a logic \mathcal{L} , then the (symmetric) existence problem is decidable when the objectives of the players are given by formulas in \mathcal{L} .

This report opens a large class of new and interesting problems regarding the complexity of computing (symmetric) Nash equilibria. In addition, some of these problems are solved while many still remain unexplored.

The report is written in English, since the author does not speak or write French.

1 Concurrent Games

1.1 Basic Definitions

The first type of game we study is concurrent games which has been studied by a number of other authors, including [1, 2, 3, 4, 5, 16]. A concurrent game in our context is played on a finite graph where the nodes represent the different states of the game and edges represent transitions between game states. The game is played an infinite number of rounds and in each round the players must concurrently choose an action. The choices of the actions of all the players then determines the successor of the current game state. An example of a concurrent game can be seen in Figure 1. This is a mobile phone game (also used in [2]), where two mobile phones fight for bandwidth in a network. They each have three transmitting power levels (0,1 and 2) and each turn a phone can choose to increase, decrease or keep its current power level. State t_{ij} corresponds to phone 1 using power level i and phone 2 using power level j . The goal of the phones is to obtain as much bandwidth as possible, but at the same time reduce energy consumption. More formally, we define a concurrent game as follows

Definition 1. *A Concurrent Game is a tuple $\mathcal{G} = (\text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab})$ where*

- *States is a finite set of states*
- *Agt is a finite set of agents*
- *Act is a finite set of actions*
- *Mov: $\text{States} \times \text{Agt} \rightarrow 2^{\text{Act}} \setminus \{\emptyset\}$ is the set of actions available to a given player in a given state*
- *Tab: $\text{States} \times \text{Act}^{\text{Agt}} \rightarrow \text{States}$ is a transition function that specifies the next state, given a state and an action of each player.*

A *move* $(m_A)_{A \in \text{Agt}}$ consists of an action for each player. It is said to be *legal* in the state s if $m_A \in \text{Mov}(s, A)$ for all $A \in \text{Agt}$. The legal plays $\text{Play}_{\mathcal{G}}$ of \mathcal{G} is the infinite sequences $\rho = s_0 s_1 \dots \in \text{States}^{\omega}$ such that for all i there is a legal move m_i such that $\text{Tab}(s_i, m_i) = s_{i+1}$. In the same way the set of *histories* $\text{Hist}_{\mathcal{G}}$ is the set of finite sequences of states that respects the Mov and Tab functions. We define the paths of \mathcal{G} as $\text{Path}_{\mathcal{G}} = \text{Hist}_{\mathcal{G}} \cup \text{Play}_{\mathcal{G}}$. For a given path ρ we write ρ_i for the i th state of the path (the first state has index 0), $\rho_{\leq i}$ for the prefix containing the first i states of ρ and $|\rho|$ for the

number of transitions in ρ when ρ is finite. When ρ is finite we also write $\text{last}(\rho)$ to denote the last state in ρ . For a path ρ the set of states occurring at least once and the set of states occurring infinitely many times is denoted $\text{Occ}(\rho)$ and $\text{Inf}(\rho)$ respectively.

We define a *strategy* $\sigma_A : \text{Hist}_{\mathcal{G}} \rightarrow \text{Act}$ of player A as a function specifying a legal move for each finite history. We call a subset $P \subseteq \text{Agt}$ a *coalition* of players and define a strategy $\sigma_P = (\sigma_A)_{A \in P}$ of a coalition P as a tuple of strategies, one for each player in the coalition. A strategy for the coalition Agt is called a *strategy profile*. We denote the set of strategy profiles $\text{Prof}_{\mathcal{G}}$. Given a strategy $(\sigma_A)_{P \in A}$ for a coalition P we say that a (finite or infinite) path $\pi = s_0 s_1 \dots$ is *compatible* with the strategy if for all consecutive states s_i and s_{i+1} there exists a move $(m_A)_{A \in \text{Agt}}$ such that $\text{Tab}(s_i, (m_A)_{A \in \text{Agt}}) = s_{i+1}$ and $m_A = \sigma_A(s_0 s_1 \dots s_i)$ for all $A \in P$. The set of infinite plays compatible with $(\sigma_A)_{A \in P}$ from a state s is called the *outcomes* from s and is denoted $\text{Out}_{\mathcal{G}}((\sigma_A)_{A \in P}, s)$. The set of finite histories compatible with $(\sigma_A)_{A \in P}$ from a state s is denoted $\text{Out}_{\mathcal{G}}^f((\sigma_A)_{A \in P}, s)$. In particular, note that for a strategy profile, the set of outcomes from a given state is a singleton.

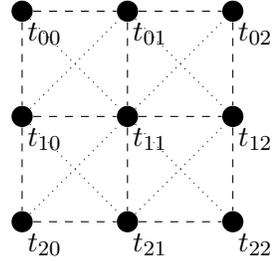


Figure 1: Mobile game. Self-loops and actions are omitted.

1.2 Objectives and Preferences

In our setting there are a number of *objectives* that can be accomplished in a concurrent game. An objective is simply a subset of all possible plays. Examples of objectives which have been analyzed in the literature include reachability and Büchi conditions defined on sets T of states:

$$\begin{aligned} \Omega_{reach}(T) &= \{\rho \in \text{Play}_{\mathcal{G}} \mid \text{Occ}(\rho) \cap T \neq \emptyset\} \\ \Omega_{\text{Büchi}}(T) &= \{\rho \in \text{Play}_{\mathcal{G}} \mid \text{Inf}(\rho) \cap T \neq \emptyset\} \end{aligned}$$

which corresponds to reaching a state in T and visiting a state in T infinitely many times respectively. We use the same framework as in [2] where players can have a number of different objectives and the payoff of a player given a play can depend on the combination of objectives that are accomplished in a quite general way. In this framework there are given a number of objectives

$(\Omega_i)_{1 \leq i \leq n}$. The *payoff vector* of a given play ρ is defined as $v_\rho \in \{0, 1\}^n$ such that $v_i = 1$ iff $\rho \in \Omega_i$. Thus, the payoff vector specifies which objectives are accomplished in a given play. We write $v = \mathbb{1}_T$ where $T \subseteq \{1, \dots, n\}$ to denote $v_i = 1 \Leftrightarrow i \in T$. We simply denote $\mathbb{1}_{\{1, \dots, n\}}$ by $\mathbb{1}$. Each player A in a concurrent game is given a total preorder $\lesssim_A \subseteq \{0, 1\}^n \times \{0, 1\}^n$ which intuitively means that $v \lesssim_A w$ if player A prefers the objectives accomplished in w over the objectives accomplished in v . This preorder induces a preference relation $\lesssim \subseteq \text{Play}_{\mathcal{G}} \times \text{Play}_{\mathcal{G}}$ over the possible plays defined by $\rho \lesssim \rho' \Leftrightarrow v_\rho \lesssim v_{\rho'}$. Additionally we say that A strictly prefers $v_{\rho'}$ to v_ρ if $v_\rho \lesssim_A v_{\rho'}$ and $v_{\rho'} \not\lesssim_A v_\rho$. In this case we write $v_\rho <_A v_{\rho'}$ and $\rho \prec_A \rho'$.

1.3 Nash Equilibria, Symmetry and Decision Problems

We have now defined the rules of the game and are ready to look at the solution concept of a Nash equilibrium, which is a strategy profile in which no player can improve by changing his strategy, given that all the other players keep their strategies fixed. The idea is that a Nash equilibrium corresponds to a stable state of the game since none of the players has interest in deviating unilaterally from their current strategy and therefore is in some sense rational. Formally, we define it as follows

Definition 2. *Let \mathcal{G} be a concurrent game with preference relations $(\lesssim_A)_{A \in \text{Agt}}$ and let s be a state. Let $\sigma = (\sigma_A)_{A \in \text{Agt}}$ be a strategy profile with $\text{Out}(\sigma, s) = \{\pi\}$. Then σ is a Nash equilibrium from s if for every $B \in \text{Agt}$ and every $\sigma'_B \in \text{Strat}^B$ with $\text{Out}(\sigma[\sigma_B \mapsto \sigma'_B], s) = \{\pi'\}$ we have $\pi \lesssim_B \pi'$.*

The concept was first introduced in [12] in normal-form games where it was proven that a Nash equilibrium always exists in mixed strategies (where players can choose actions probabilistically), however a Nash equilibrium in pure strategies does not always exist. We will only focus on pure Nash equilibria which will be called Nash equilibria in the sequel. In this report we are not merely interested in Nash equilibria, but in symmetric Nash equilibria in which all players use the same strategy. This is motivated by the idea of modelling distributed systems as games, where players correspond to programmable devices (possibly with different goals) and strategies correspond to programs. Then a Nash equilibrium is in some sense an optimal configuration of a system, since no device has any interest in deviating from the chosen program, whereas symmetric Nash equilibria are configurations where all devices use the same program while preserving optimality. This way of modelling distributed systems is in many cases more realistic since we assume that other devices act rationally and not as worst-case opponents

which is often done in theoretical models of distributed systems. In the setting of a concurrent game where players act solely based on the sequence of global states of the game, we define a symmetric strategy profile as follows

Definition 3. *A symmetric strategy profile is a strategy profile σ such that $\sigma_A(\pi) = \sigma_B(\pi)$ for all histories π and all $A, B \in \text{Agt}$.*

We denote the set of symmetric strategy profiles $\text{Prof}^{\text{Sym}} \subseteq \text{Prof}$. This naturally leads to the following definition.

Definition 4. *A strategy profile σ is a symmetric Nash equilibrium if it is a symmetric strategy profile and a Nash equilibrium.*

The computational problem of deciding existence of pure strategy Nash equilibria in concurrent games are analyzed for different types of objectives in [2, 3, 4, 5] with the same setting used here. In this chapter we will use some of the same techniques for analyzing the following two computational problems restricted to particular types of objectives and preference relations. The same problems will be analyzed for different games in later chapters.

Definition 5 (Symmetric Existence Problem). *Given a game \mathcal{G} and a state s does there exist a symmetric Nash equilibrium in \mathcal{G} from s ?*

Definition 6 (Symmetric Constrained Existence Problem). *Given a game \mathcal{G} , a state s and two vectors u^A and w^A for each player A , does there exist a symmetric Nash equilibrium in \mathcal{G} from s with some payoff $(v^A)_{A \in \text{Agt}}$ such that $u^A \lesssim v^A \lesssim w^A$ for all $A \in \text{Agt}$?*

1.4 General Reachability Objectives

In this section we focus on the case where the objectives of all players are reachability objectives and each player has a preorder on the payoff vectors specified by a boolean circuit, which is quite general. We first present an algorithm solving the problem using polynomial space and afterwards show that the problem is \mathcal{PSPACE} -hard by a reduction from QSAT. In [2] the same complexity is obtained for regular Nash equilibria. In this section as well as Section 1.5 we reuse techniques from [2, 4] and adjust them to deal with the symmetry constraint.

1.4.1 \mathcal{PSPACE} -algorithm

In [2] it is shown that if the preferences of all players only depend on the states visited and the states visited infinitely often, then there is a Nash

equilibrium with payoff v from state s if and only if there is a Nash equilibrium with payoff v from s that has an outcome of the form $\pi \cdot \tau^\omega$ where $|\pi|, |\tau| \leq |\text{States}|^2$. We obtain a similar result for symmetric Nash equilibria by using the same proof idea, but adjusting it to symmetric strategy profiles. The proof is in Appendix A.1.

Lemma 7. *Assume that every player has a preference relation which only depends on the set of states that are visited and the states that are visited infinitely often. If there is a symmetric Nash equilibrium with payoff v from state s then there is a symmetric Nash equilibrium with payoff v from s that has an outcome of the form $\pi \cdot \tau^\omega$ where $|\pi|, |\tau| \leq |\text{States}|^2$.*

This lemma helps us, since we only need to look for symmetric Nash equilibria with outcomes of the desired shape. The idea for the algorithm is the same as in [2], where a given payoff vector with reachability objectives can be encoded as a 1-weak deterministic Büchi automaton for each player. A 1-weak deterministic Büchi automaton is a Büchi automaton $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$ where all strongly connected components of the transition graph contains exactly one state. A player with payoff given by a Büchi automaton $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$ with $\Sigma = \text{States}$ gets payoff 1 for a run ρ if $\rho \in L(\mathcal{A})$ and 0 otherwise. For a concurrent game $\mathcal{G} = (\text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab})$ and a given payoff vector $v = (v_A)_{A \in \text{Agt}}$ we generate the game $\mathcal{G}(v)$ with the same structure as \mathcal{G} but with objectives given by 1-weak deterministic Büchi automata $\mathcal{A}_A = (Q, \Sigma, \delta, q_0, F_A)$ for each player A with the same structure, but different acceptance conditions. It is designed so $Q = 2^{\text{States}}$, $\Sigma = \text{States}$, $q_0 = \emptyset$ and $\delta(q, s) = q \cup \{s\}$ for all $q \in Q$ and $s \in \text{States}$. For every player A we further define $F_A = \{q \in Q \mid \mathbb{1}_{\{i \mid q \cap \Omega_i^A \neq \emptyset\}} \lesssim_A v_A\}$. Intuitively, the automata will loop infinitely in some state q , which is exactly the subset of States which is visited. Then player A will win in $\mathcal{G}(v)$ if and only if he gets a strictly better payoff than v_A . From this one can see that there is a Nash equilibrium in $\mathcal{G}(v)$ with payoff $(0, \dots, 0)$ if and only if there is a Nash equilibrium in \mathcal{G} with payoff v . Our algorithm works by non-deterministically guessing a payoff vector v and an outcome ρ of the form from Lemma 7 with the guessed payoff and then using an altered version of the following lemma from [2, 3] to check if the guessed outcome is the outcome of a symmetric Nash equilibrium with payoff $(0, \dots, 0)$ in $\mathcal{G}(v)$

Lemma 8. *Let \mathcal{G} be a concurrent game with a single objective per player given by 1-weak Büchi automata $(\mathcal{A}_A)_{A \in \text{Agt}}$. Let l_A be the length of the longest acyclic path in \mathcal{A}_A , let m be the space needed for deciding whether a state of \mathcal{A}_A is final and whether a state q' of \mathcal{A}_A is the successor of a*

state q for an input symbol s . Then whether a path of the form $\pi \cdot \tau^\omega$ is the outcome of a Nash equilibrium can be decided in space $O((|\mathcal{G}| \cdot (\sum_{A \in \text{Agt}} l_A) + \sum_{A \in \text{Agt}} \log |\mathcal{A}_A| + m)^2)$.

However, this lemma needs to be altered somewhat to be applicable since it is possible to have a play ρ which is the outcome of a Nash equilibrium and the outcome of a symmetric strategy profile, but not the outcome of a symmetric Nash equilibrium (see Figure 2). The technicalities behind the lemma is proven in [3] and the algorithm can be altered to check if a path is the outcome of a symmetric Nash equilibrium by using essentially the same technique for changing an iterated repellor computation to deal with symmetry as we use in Section 1.5.1 on repellor sets to deal with symmetry in games with a single reachability objective. This means that we have an algorithm running in \mathcal{PSPACE} for solving both the symmetric existence as well as the constrained symmetric existence since the only change is the non-deterministic guess of the payoff.

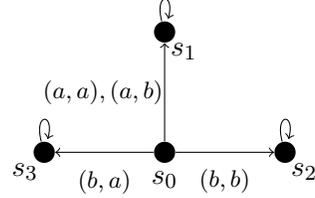


Figure 2: A prefers reaching s_3 and B prefers reaching s_2 . $s_0 s_1^\omega$ is the outcome of a Nash equilibrium and a symmetric strategy profile, but not a symmetric Nash equilibrium

Theorem 9. *The symmetric existence problem and constrained symmetric existence problem are decidable in \mathcal{PSPACE} for reachability objectives and preference relations represented by boolean circuits.*

1.4.2 \mathcal{PSPACE} -hardness

The following is proven by a reduction from QSAT in Appendix A.2.

Theorem 10. *The (constrained) symmetric existence problem is \mathcal{PSPACE} -hard for games with 2 players with reachability objectives and preference relations represented by boolean circuits.*

1.5 Single Reachability Objective

In this section we focus on the case where each player A has one reachability objective, denoted Ω_A . First we will provide algorithms solving the symmetric existence problem and the constrained symmetric existence problem in polynomial time for a bounded number of players and in non-deterministic

polynomial time for an unbounded number of players. Then we will proceed to prove that the bounds are tight.

1.5.1 Suspects and repellor sets

In the following we fix a game $\mathcal{G} = (\text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab})$. We begin by defining the set of suspect players for a move m_{agt} and an edge (s, s') :

$$\text{Susp}((s, s'), m_{\text{agt}}) = \{B \in \text{Agt} \mid \exists m' \in \text{Mov}(s, B). \text{Tab}(s, m_{\text{agt}}[B \mapsto m']) = s'\}$$

Intuitively, the set of suspect players for move m_{agt} and edge (s, s') is the set of players that can choose an action so that the play moves from s to s' given that all other players play according to m_{agt} . In the same spirit we define the suspect players of a path $\pi = (s_p)_{p \leq |\pi|}$ given strategy profile $\sigma = (\sigma_A)_{A \in \text{Agt}}$ as the players that can choose actions to enforce the path, given that all other players play according to σ :

$$\text{Susp}(\pi, \sigma) = \bigcap_{i < |\pi|} \text{Susp}((s_i, s_{i+1}), (\sigma_A(\pi_{\leq i}))_{A \in \text{Agt}})$$

As a basis for the algorithms we introduce *repellor sets* which are used in [4], but we define them in a different way so that they can be used to solve the symmetric strategy problems on which we focus here. We call the altered version *symmetric repellor sets*. The symmetric repellor sets $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ in a game \mathcal{G} is defined inductively on subsets $P \subseteq \text{Agt}$ such that

- $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(\emptyset) = \text{States}$
- If $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P')$ is calculated for all $P' \subsetneq P$ then $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ is the largest set such that
 - $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P) \cap \Omega_B = \emptyset$ for all $B \in P$.
 - $\forall s \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P). \exists a \in \text{Act}. \forall s' \in \text{States}. s' \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}((s, s'), m_a^{\text{Sym}}))$

where m_a^{Sym} is the symmetric move where all players choose action a . The symmetric repellor set $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ is then the largest set of states that does not contain any target states of players in P and such that for each state in the symmetric repellor set there exists a symmetric move m so any player in P capable of moving the play somewhere else than m dictates can only move the play to a state contained in a new repellor set for this player.

The idea is roughly that if all players play according to these symmetric moves in all states, then a player in P from a state in $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ will not be able to unilaterally deviate from his strategy and reach a target state. These sets of symmetric moves, called secure moves, are defined more precisely as

$$\text{Secure}_{\mathcal{G}}^{\text{Sym}}(s, P) = \{m_a^{\text{Sym}} \mid a \in \text{Act} \wedge \forall s' \in \text{States}. s' \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}((s, s'), m_a^{\text{Sym}}))\}$$

Next, we define a transition system $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P) = (\text{States}, \text{Edg})$ with the same set of states as in \mathcal{G} and with an edge $(s, s') \in \text{Edg}$ if and only if for all $A \in P$ we have $s \notin \Omega_A$ and there exists $m \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(s, P)$ such that $\text{Tab}(s, m) = s'$. These notions will be useful in the search for symmetric Nash equilibria because of the following theorem

Theorem 11. *Let \mathcal{G} be a concurrent game with one reachability objective Ω_A for each $A \in \text{Agt}$. Let $P \subseteq \text{Agt}$ be a subset of players and let v be the payoff where Ω_B is accomplished iff $B \notin P$. Let $s \in \text{States}$. Then there is a symmetric Nash equilibrium from s with payoff v iff there is an infinite path π in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ starting in s which visits Ω_A for every $A \notin P$.*

Some preliminary lemmas are needed to prove this result. Proofs of these, as well as of the theorem, can be found in Appendix A.3. The result leads to an algorithm for both a bounded and an unbounded number of players, which works by searching for infinite paths in the transition systems $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$. The details are in Appendix A.4.

Theorem 12. *The symmetric constrained existence problem is in \mathcal{P} for a bounded number of players and in \mathcal{NP} for an unbounded number of players when each player has a single reachability objective.*

1.5.2 Hardness Results

The following results show that the bounds found in the previous section are tight. Proofs by reduction from CIRCUITVALUE and 3SAT are in Appendix A.5 and A.6.

Theorem 13. *The symmetric existence problem is \mathcal{P} -hard for a bounded number of players and \mathcal{NP} -hard for an unbounded number of players when each player has a single reachability objective*

1.6 Symmetric vs Regular NE in reachability games

Our results in this chapter show that for general reachability objectives the symmetric existence problem as well as the constrained symmetric existence problem are \mathcal{PSPACE} -complete for preferences given by boolean circuits. For single reachability objectives we have \mathcal{NP} -completeness for an unbounded number of players and \mathcal{P} -completeness for a bounded number of players. These complexity results are the same as for Nash equilibria [2, 4] and we have not currently any reason to believe that reachability objectives are special in this regard. It remains to investigate if the symmetric and regular problems will have the same complexity for other interesting objectives such as safety, Büchi, parity, etc.

2 Symmetric Concurrent Games

The concurrent games studied in the previous section are interesting, but there are some phenomena that cannot really be captured with the notion of symmetry introduced there. We have until now required that a symmetric strategy profile was simply a strategy profile where every player takes the same action at every decision point in the game, which is quite restrictive in some sense. For instance in the mobile phone game, it would make more sense in a symmetric strategy profile to require player 1 to play in state t_{12} as player 2 does in state t_{21} since t_{12} has the same meaning to player 1 as t_{21} has to player 2. In this section we propose a definition of a symmetric concurrent game and another definition of symmetric strategy profiles than in the previous section. In some sense we would like the game to be the same from the point of view of all players, which means there should be some constraints on the structure of the arena. We would like the game to be the same when switching positions between the different players. Further, the payoff of a player should only depend on the number of other players who choose a particular strategy and not on which of the other players who choose a particular strategy.

2.1 Symmetric Normal-Form Games

We begin by looking at symmetric normal-form games which have been defined in the literature. We would like our definition of symmetric concurrent games to agree with the corresponding definition for normal-form games, since concurrent games generalize normal-form games. Therefore we use this as a stepping stone to get some intuition for a definition of sym-

metric concurrent games. We use a definition of a symmetric normal-form game resembling the definition used in [8]¹ but equivalent to the definition used in [6, 7, 10, 14], since we believe [8] is erroneous. This is discussed in Appendix B. A symmetric normal-form game should be a game where all players have the same set of actions and the same utility functions where the utility of a player should depend exactly on which action he chooses himself and on the number of other players choosing any particular action.

Definition 14. A symmetric normal-form game $\mathcal{G} = (\{1, \dots, n\}, S, (u_i)_{1 \leq i \leq n})$ is a 3-tuple where $\{1, \dots, n\}$ is the set of players, S is a finite set of strategies and $u_i : S^n \rightarrow \mathbb{R}$ are utility functions such that for all strategy vectors $(a_1, \dots, a_n) \in S^n$, all permutations π of $(1, \dots, n)$ and all i it holds that

$$u_i(a_1, \dots, a_n) = u_{\pi^{-1}(i)}(a_{\pi(1)}, \dots, a_{\pi(n)})$$

The intuition behind this definition, which is reused later, is as follows. Suppose we have a strategy profile $\sigma = (a_1, \dots, a_n)$ and a strategy profile where the actions of the players have been rearranged by permutation π , $\sigma_\pi = (a_{\pi(1)}, \dots, a_{\pi(n)})$. We would prefer that the player j using the same action in σ_π as player i does in σ gets the same utility. Since j uses $a_{\pi(j)}$ this means that $\pi(j) = i \Rightarrow j = \pi^{-1}(i)$. Now, from this intuition we have that $u_i(a_1, \dots, a_n) = u_j(a_{\pi(1)}, \dots, a_{\pi(n)}) = u_{\pi^{-1}(i)}(a_{\pi(1)}, \dots, a_{\pi(n)})$. Apart from this intuition, the new definition can be shown to be equivalent to the one from [6, 7, 10, 14].

2.2 Symmetric Concurrent Games

Following the intuition of the previous sections, we are ready to propose a definition of a symmetric concurrent game generalising the definition of symmetry in normal-form games. In the following, the set of agents will be $\text{Agt} = \{1, \dots, n\}$. In addition, if $\alpha = ((a_1(i), \dots, a_n(i)))_{0 \leq i \leq k}$ is a sequence of moves, we write $\alpha_\pi = ((a_{\pi(1)}(i), \dots, a_{\pi(n)}(i)))_{0 \leq i \leq k}$ for the sequence of moves where the actions of every move are reordered according to the permutation π of $(1, \dots, n)$. Finally, we write $s \xrightarrow{\alpha} s'$ to denote that if the players play according to the sequence of moves α then the play moves from s to s' .

Definition 15. A concurrent game $\mathcal{G} = (\text{States}, \text{Agt}, \text{Act}, \text{Mov}, \text{Tab}, s_0)$ is symmetric if the following points are satisfied

- If $s_0 \xrightarrow{\alpha} s$ and $s_0 \xrightarrow{\beta} s$ for some s, α, β then $s_0 \xrightarrow{\alpha_\pi} s'$ and $s_0 \xrightarrow{\beta_\pi} s'$ for some s' for all permutations π of $(1, \dots, n)$.

¹The same definition is used on Wikipedia citing this source at the time of writing

- For every infinite sequence α of moves, every player i and every permutation π of $\{1, \dots, n\}$ we have
 - $u_i(\text{Out}(\alpha)) = u_{\pi^{-1}(i)}(\text{Out}(\alpha_\pi))$ for preferences given by utility functions.
 - $\text{Out}(\alpha) \in \Omega_i^j \Leftrightarrow \text{Out}(\alpha_\pi) \in \Omega_{\pi^{-1}(i)}^j$ for preferences given by ordered objectives.

When $s_0 \xrightarrow{\alpha} s$ and $s_0 \xrightarrow{\alpha_\pi} s'$ we denote $h_\pi(s) = s'$. Because of point 1 in the definition, this is unique. In addition, we have that h_π is bijective for all permutations π since $(\alpha_\pi)_{\pi^{-1}} = \alpha$. Further, $h_\pi(s_0) = s_0$ for all permutations π since all permutations of the empty sequence of moves equals the empty sequence of moves.

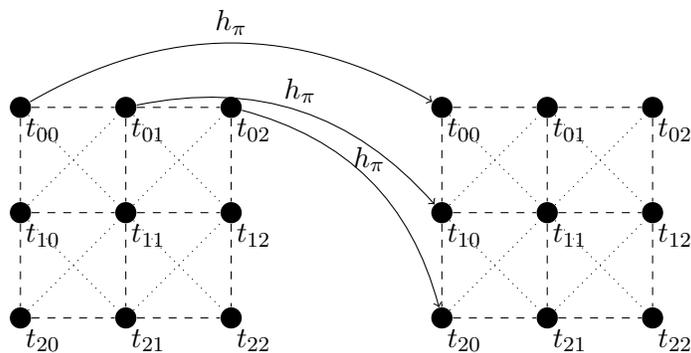


Figure 3: Mobile game and part of h_π illustrated for $\pi = (2, 1)$

The intuitive understanding of the mappings h_π is that they reorder the states corresponding to the reordering of the players. For instance in the mobile game in Figure 3 we have for $\pi = (2, 1)$ that $h_\pi(t_{01}) = t_{10}$ since the state t_{10} means the same after players change positions as t_{01} did before. But on the other hand $h_\pi(t_{00}) = t_{00}$ since this state means the same thing to the two players. We can also show that all the mappings h_π are in some sense automorphisms of the arena:

Lemma 16. For all permutations π of $(1, \dots, n)$, all states s and legal moves m from s we have

$$\text{Tab}(s, m) = s' \Leftrightarrow \text{Tab}(h_\pi(s), m_\pi) = h_\pi(s')$$

These mappings are also used to define symmetric strategy profiles

Definition 17. *A symmetric strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ in a symmetric concurrent game \mathcal{G} is a strategy profile such that*

$$\sigma_i(\rho) = \sigma_{\pi^{-1}(i)}(h_\pi(\rho))$$

for all $i \in \{1, \dots, n\}$, all $\rho \in \text{Hist}_{\mathcal{G}}$ and all permutations π of $(1, \dots, n)$.

Intuitively, this means that when we switch the players they will keep using the same strategy, but adjusted to their new position defined by the permutation π and the corresponding mapping h_π between states. In addition, all players in this sense uses the same strategy and in this strategy a player will not distinguish between which of the other players choose particular moves, but only how many of the other players choose particular moves. A symmetric Nash equilibrium is again defined as a symmetric strategy profile which is a Nash equilibrium. A natural question is now whether some of the results holding for symmetric normal-form games also hold for this generalization. For instance, every symmetric normal-form game has a symmetric Nash Equilibrium in mixed strategies [7]. We start by defining a mixed strategy for player i in game \mathcal{G} as a mapping from any finite history $h \in \text{Hist}_{\mathcal{G}}$ to $\Delta(\text{Mov}_{\mathcal{G}}(h_{|h|}, i))$ where $\Delta(\cdot)$ is the set of discrete probability distributions over a finite set. In addition, a Nash equilibrium is now defined as a mixed strategy profile, so no player can unilaterally change to improve his expected utility. We show that the result for normal-form symmetric games does not hold for our generalization in Appendix A.8

Theorem 18. *There exists symmetric concurrent games with no symmetric Nash equilibrium in mixed strategies*

This result seems to be a consequence of infinite runs rather than symmetry. As described in [16] we can have two-player zero-sum games with no winning strategies (where one player can be sure to win), no surely winning strategies (where one player can win with probability 1) but limit surely winning strategies (where one player can win with probability $1 - \epsilon$ for ϵ arbitrarily close to 0, but not 0). In these cases the limit surely winning player can always ensure a higher utility by choosing smaller values of ϵ , which entails that no Nash equilibrium will exist even in mixed strategies.

It would be interesting to investigate computational problems in these games, both with respect to Nash equilibria, symmetric Nash equilibria and other solution concepts to see if the symmetry defined will have an impact

on the computational complexity. Especially the investigation of symmetric Nash equilibria might be different since symmetry is not defined locally as in the previous section. The choice of action for a player in one game state affects the action of another player in another game state in a symmetric strategy profile.

3 Symmetric Concurrent Game Structures

In this section we investigate a different type of game which is not completely symmetric in all aspects like the games investigated in the previous section. In these games the rules should be the same from the point of view of every player, but we allow the possibility of players to distinguish between the actions of some of the other players, which we call the neighbours of a player. A player will not be able to distinguish between his non-neighbours. The games include scenarios where teams of agents can work together and have shared information about their personal state. This was not possible in the games defined in the previous section. However, the game is still the same from the point of view of every player in some sense. For instance they all have the same number of neighbours. In addition to this "loosening" of symmetry, we also allow the play of a game to start in asymmetric states which should give rise to more interesting behaviors. In our setting each player has a local transition system which he controls. The whole game structure will be a product of these transition systems in some sense.

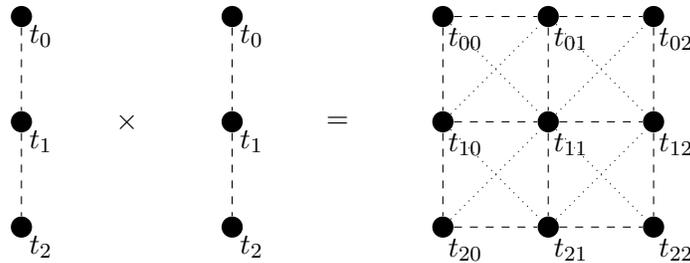


Figure 4: Mobile game, as a symmetric concurrent game structure

Each player will get some information about the state of the other players, but this information can be partial in a number of different ways. Formally, we define a symmetric game structure with n players as follows:

Definition 19. *A symmetric game structure is a tuple*

$\mathcal{G} = (G, (\text{Nei}_i)_{1 \leq i \leq n}, (\pi_{i,j})_{1 \leq i,j \leq n}, (\equiv_i)_{1 \leq i \leq n}, (\Omega_i^j)_{1 \leq i \leq n, 1 \leq j \leq c}, k, c)$ where

- $G = (\text{States}, \text{Act}, \text{Mov}, \text{Tab})$ is a one-player arena where
 - States is a finite set of states
 - Act is a finite set of actions
 - $\text{Mov} : \text{States} \rightarrow 2^{\text{Act}}$ is the legal moves from a particular state
 - $\text{Tab} : \text{States} \times \text{Act} \rightarrow \text{States}$ is the transition function.
- $\text{Nei}_i : \text{Agt} \rightarrow \text{Agt}^k$ is a neighbour function for each player i . If $\text{Nei}_i = (i_1, \dots, i_k)$ denote $N(i) = \{i, i_1, \dots, i_k\}$.
- $\pi_{i,j} : \text{Agt} \rightarrow \text{Agt}$ is a partially defined permutation of $(1, \dots, n)$ for each pair of players i, j restricted to $N(i)$ such that $\pi_{i,j}(i) = j$ and $\text{Nei}(j) = (\pi_{i,j}(i_1), \dots, \pi_{i,j}(i_k))$ where $\text{Nei}(i) = (i_1, \dots, i_k)$.
- \equiv_i is an equivalence relation between state configurations for each player i such that for every π with $\pi(j) = j$ for $j \in N(i)$ we have $(s_1, \dots, s_n) \equiv_i (s_{\pi(1)}, \dots, s_{\pi(n)})$. For symmetry between players we also require $(s_1, \dots, s_n) \equiv_i (s'_1, \dots, s'_n) \Leftrightarrow (s_{\pi^{-1}(1)}, \dots, s_{\pi^{-1}(n)}) \equiv_j (s'_{\pi^{-1}(1)}, \dots, s'_{\pi^{-1}(n)})$ for every i, j and every π extending $\pi_{i,j}$.
- $\Omega_i^l \subseteq \text{Play}_{\mathcal{G}}$ is objective l of player i and we require $\rho \in \Omega_i^l \Leftrightarrow \pi^{-1}(\rho) \in \Omega_j^l$ for all paths ρ , all l , all players i, j and all π which extend $\pi_{i,j}$.

Note that plays are infinite sequences of state configurations, where a state configuration is an n -tuple of states. As in concurrent games, in each round each player chooses an action which gives the next state configuration.

For a state configuration $t = (s_1, \dots, s_n)$ we define $\pi(t) = (s_{\pi(1)}, \dots, s_{\pi(n)})$ for all permutations π of $\{1, \dots, n\}$ and for a path $\rho = t_1 t_2 \dots$ of state configurations we define $\pi(\rho) = \pi(t_1) \pi(t_2) \dots$. The intuition behind the mapping $\pi_{i,j}$ is that it can be used to map the neighbour relationships of player i to the neighbour relationships of player j while keeping the game symmetric between all the players, thus making players interchangeable. For instance in Figure 5 there is a card game tournament with 6 players, 3 on each table. Here each player has a left neighbour, a right neighbour and 3 opponents at a different table. One could use $\text{Nei}_1 = (2, 3), \text{Nei}_2 = (3, 1)$ and $\text{Nei}_3 = (1, 2)$ to model

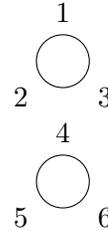


Figure 5: A card game tournament

this. Then $\pi_{1,2}(1) = 2, \pi_{1,2}(2) = 3$ and $\pi_{1,2}(3) = 1$ would map the neighbours of 1 to neighbours of 2 since indeed $\text{Nei}_2 = (3, 1) = (\pi_{1,2}(2), \pi_{1,2}(3))$. The requirements for the equivalence relations \equiv_i for each player i then makes sure that all players have the same information partitions, adjusted to their specific position and that non-neighbours can be interchanged without affecting this. In addition \equiv_i induce equivalence classes of state configurations that player i cannot distinguish. We call these equivalence classes information sets and denote by \mathcal{I}_i the set of information sets for player i . Then we define a strategy σ for player i to be a map from \mathcal{I}_i^* to Act where the action must be legal in the states of the final information set in the sequence. For a state configuration t we denote $I_i(t)$ the information set of player i that t is contained in. For a sequence $\rho = t_1 t_2 \dots$ of state configurations we define $I_i(\rho) = I_i(t_1) I_i(t_2) \dots$. We say that a strategy profile is symmetric if for every pair of players i and j and every sequence of state configurations ρ we have

$$\sigma_i(I_i(\rho)) = \sigma_j(I_j(\pi^{-1}(\rho))) = \sigma_{\pi(i)}(I_{\pi(i)}(\pi^{-1}(\rho)))$$

for every π that extends $\pi_{i,j}$. We define a Nash equilibrium in the usual way and a symmetric Nash equilibrium as a symmetric strategy profile which is a Nash equilibrium. Our first result is that even though symmetric Nash equilibria are Nash equilibria with special properties they are in some sense at least as hard to find as Nash equilibria. This is due to the following result

Theorem 20. *From a symmetric game structure \mathcal{G} we can in polynomial time construct a symmetric game structure \mathcal{H} , which is polynomial in the size of \mathcal{G} with the same type of objectives and such that there exists a symmetric Nash equilibrium in \mathcal{H} if and only if there exists a Nash equilibrium in \mathcal{G} .*

This means that we cannot in general hope to have an algorithm with better complexity for the symmetric problem by using properties of symmetry. This is of course unfortunate, but is a good thing to know. Next, we can show that the problem of finding Nash equilibria as well as symmetric Nash equilibria is undecidable even in the case of 2 players and perfect information. The proof of undecidability is in Appendix A.10 and is a reduction of the halting problem for deterministic 2-counter machines.

Theorem 21. *The existence problem is undecidable for symmetric game structures*

Corollary 22. *The symmetric existence problem is undecidable for symmetric game structures*

Corollary 23. *The existence problem is undecidable for concurrent games.*

Corollary 23 was known already [5], but with a different proof than the one used in this report. We now define a generalization of a memoryless strategy called an m -bounded suffix strategy as a strategy which only depends on the last m infosets seen. For $m = 1$ this coincides with a memoryless strategy.

Definition 24. *An m -bounded suffix strategy σ is a strategy such that $\sigma(\rho) = \sigma(\rho')$ whenever $\rho_{\geq|\rho|-m+1} = \rho'_{\geq|\rho'|-m+1}$*

Next, suppose every information set is labeled with propositions from some finite set Prop. This is done by the labelling function $L_i : \mathcal{I}_i \rightarrow 2^{\text{Prop}}$ for each player i . We say that the objective of player i is given by a formula φ_i from a logic interpreted over finite words, if i wins in a play ρ when $L_i(\rho) \models \varphi_i$ and loses otherwise. For instance, a single reachability objective can be represented by the LTL formula $\mathbf{F} p$ where the information sets of the reachability objective are labelled with p . In the same way, a single Büchi objective can be represented by the LTL formula $\mathbf{GF} p$. We now get the following result

Theorem 25. *Suppose \mathcal{L} is a logic interpreted over infinite words. If model checking a formula $\varphi \in \mathcal{L}$ in a rooted, interpreted transition system (S, T, L, s_0) is decidable in time $f(|S|, |T|, |L|, |\varphi|)$ then the (symmetric) existence problem for m -bounded suffix strategies in a symmetric game structure $\mathcal{G} = (G, (\text{Nei}_i), (\pi_{i,j}), (\equiv_i), (\varphi_i), k)$ is decidable in time*

$$O(n \cdot |\text{States}|^{n^2 \cdot m \cdot |\text{Act}|} \cdot f((|\text{States}|^n + 1)^{n \cdot m}, |\text{Act}| \cdot (|\text{States}|^n + 1)^{n \cdot m}, |\text{Prop}|, |\varphi|))$$

when every player i has an objective given by a formula $\varphi_i \in \mathcal{L}$ where the proposition symbols occurring are Prop and $|\varphi| = \max |\varphi_i|$.

This result means that we can decide the existence of Nash equilibria which are m -bounded suffix strategy profiles with many different objectives. Indeed, since model-checking LTL is decidable, the problem is decidable for all objectives which can be described by an LTL-formula. Since model-checking an LTL formula φ over an interpreted, rooted transition system (S, T, L) can be done in time $2^{O(|\varphi|)} \cdot O(|S|)$ (see [15]) it follows that

Corollary 26. *The (symmetric) existence problem for m -bounded suffix strategies in a symmetric game structure $\mathcal{G} = (G, (\text{Nei}_i), (\pi_{i,j}), (\equiv_i), (\varphi_i), k)$ is decidable in time*

$$O(n \cdot (|\text{States}|^n + 1)^{2mn(1+|\text{Act}|)} \cdot 2^{O(|\varphi|)})$$

when every player i has an objective given by an LTL formula φ_i and $|\varphi| = \max|\varphi_i|$.

4 Conclusion

We have presented a number of different versions of symmetric games and symmetric Nash equilibria as well as discussed the intuition behind them. A number of computational problems have been analyzed in this area where we have seen problems ranging from being solvable in polynomial time to being undecidable. It should be clear that in all the games we have looked at in this report there are still many unanswered problems depending on the exact subclass of games considered and that these games are quite expressive. When analysing symmetric Nash equilibria in concurrent games we obtained the same complexity as for regular Nash equilibria in a number of cases, but it would be interesting to see if this is also the case for the non-local versions of symmetry defined in the later chapters. Another obvious extension in symmetric game structures is to investigate the effect of incomplete information on computational complexity which we have not done thoroughly in this report.

References

- [1] *R. Alur, T. A. Henzinger & O. Kupferman* Alternating-Time Temporal Logic, *Journal of the ACM*, Vol. 49, No. 5, September 2002, pp. 672-713
- [2] *P. Bouyer, R. Brenquier, N. Markey & M. Ummels*, Concurrent Games with Ordered Objectives, Research report LSV-11-22, Version 2 - January 11, 2012
- [3] *P. Bouyer, R. Brenquier, N. Markey & M. Ummels*, Nash Equilibria in Concurrent Games with Büchi Objectives, in *FSTTCS'11, LIPIcs*, Mumbai, India, 2011. Leibniz-Zentrum für Informatik
- [4] *P. Bouyer, R. Brenquier & N. Markey*, Nash Equilibria for Reachability Objectives in Multi-Player Timed Games, Research report LSV-10-12 - June 2010

- [5] *P. Bouyer, R. Brenquier & N. Markey* Nash equilibria in concurrent games, Part 1: Qualitative Objectives, Preliminary Version
- [6] *F. Brandt, F. Fischer & M. Holzer* Symmetries and the complexity of pure Nash equilibrium, *Journal of Computer and System Sciences* 75 (2009) 163-177
- [7] *S.-F. Cheng, D. M. Reeves, Y. Vorobeychik & M. P. Wellman*, Notes on Equilibria in Symmetric Games, *Proceedings of the 6th International Workshop On Game Theoretic And Decision Theoretic Agents*, 71-78, 2004.
- [8] *P. Dasgupta & E. Maskin* The Existence of Equilibrium in Discontinuous Economic Games, 1: Theory, *The Review of Economic Studies*, 53(1): 1-26, 1986
- [9] *J. E. Hopcroft, R. Motwani & J. D. Ullman* Automata Theory, Languages, and Computation, 3rd Edition, Addison Wesley, 2007
- [10] *A. X. Jiang, C. T. Ryan & K. Leyton-Brown* Symmetric Games with Piecewise Linear Utilities
- [11] *R. Mazala*, Infinite Games, E. Grädel et al. (Eds.): *Automata, Logics, and Infinite Games*, LNCS 2500, pp. 23-38. Springer 2002
- [12] *J. F. Nash Jr.* Equilibrium points in n-person games. *Proc. National Academy of Sciences of the USA*, 36(1):48-49, 1950.
- [13] *M. J. Osborne & A. Rubinstein* A Course in Game Theory, MIT Press, 1994
- [14] *C. Papadimitriou* The complexity of finding Nash equilibria, Chapter 2 in *Algorithmic Game Theory*, edited by N. Nisam et al., Cambridge University Press, 2007
- [15] *P. Schnoebelen* The Complexity of Temporal Logic Model Checking, 2003
- [16] *O. Serre* Game Theory Techniques in Computer Science - Lecture Notes, MPRI 2011-2012, January 4, 2012
- [17] *Y. Shoham & K. Leyton-Brown* Multiagent Systems - Algorithmic, Game-Theoretic, and Logical Foundations, Cambridge University Press, 2009

A Proofs

A.1 Proof of Lemma 7

Lemma 7 *Assume that every player has a preference relation which only depends on the set of states that are visited and the states that are visited infinitely often. If there is a symmetric Nash equilibrium with payoff v then there is a symmetric Nash equilibrium with payoff v that has an outcome of the form $\pi \cdot \tau^\omega$ where $|\pi|, |\tau| \leq |\text{States}|^2$.*

Proof. If there is no symmetric Nash equilibrium, then there is clearly no symmetric Nash equilibrium with the desired outcome. On other hand, suppose there is a symmetric Nash equilibrium with payoff v and let it have outcome ρ . From this symmetric Nash equilibrium we will generate a symmetric Nash equilibrium σ' with payoff v and outcome $\pi \cdot \tau^\omega$ where $|\pi|, |\tau| \leq |\text{States}|^2$. The idea is to define π and τ so π contains exactly the states that appear in ρ and τ contains exactly the states that appear infinitely often in ρ . In this way σ and σ' will have the same payoff. We divide π into subpaths π^0 and π^1 , where $\text{Occ}(\pi^0) = \text{Occ}(\rho)$. Thus, π^0 will make sure the proper states are included in π whereas π^1 is responsible for connecting the last state of π^0 with the first state of τ . In the same way we divide τ into subpaths τ^0 and τ^1 where τ^0 will make sure the proper states are included in τ and τ^1 is responsible for connecting the last state of τ^0 with its first state. The situation is illustrated in Figure 6.

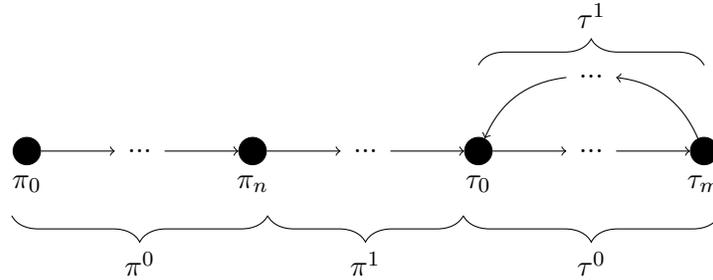


Figure 6: Illustration of $\pi \cdot \tau^\omega$

We define π^0 inductively. Start by setting $\pi_0^0 = \rho_0$. Then assume that we have created $\pi_{\leq k}^0$ for some k which only contains states occurring in $\rho_{\leq k'}$ for some k' . If $\text{Occ}(\pi_{\leq k}^0) = \text{Occ}(\rho)$ then the construction of π^0 is finished. Otherwise let i be the least index such that ρ_i does not occur in $\pi_{\leq k}^0$. Let $j <$

i be the largest index such that $\pi_k^0 = \rho_j$. Now we continue the construction by setting $\pi_{k+1}^0 = \rho_{j+1}$. In addition, we define a correspondence function c_1 , by setting $c_1(k) = j$ and initially, $c_1(0) = 0$. In this way, the correspondence function maps the index of every state in $\pi \cdot \tau^\omega$ to the index of a corresponding state in ρ .

By continuing in this way ρ_i will be "reached" after at most $|Occ(\pi_{\leq k}^0)|$ steps since we will not add the same state twice on the way from π_k^0 to ρ_i and we will not add a state that is not already in $\pi_{\leq k}^0$. All the states of ρ will be added to π^0 in this way and when the procedure terminates we have $|\pi^0| \leq \sum_{i=1}^{|\text{States}|-1} i = \frac{|\text{States}|(|\text{States}|-1)}{2}$.

We choose l as the smallest number so all states appearing finitely often in ρ appears before ρ_l and let $\tau_0^0 = \rho_l$. We need to create π^1 so it connects the last state π_n^0 of π^0 with τ_0^0 . This is done in the same way as when we created π^0 by inductively choosing the successor state of π_k^1 (initially, the successor of π_n^0) as the successor of ρ_j in ρ where $j \leq l$ is the largest number so $\rho_j = \pi_k^1$ (initially, $\rho_j = \pi_n^0$) until $j = k$ in which case the state is not added to π^1 . This will connect π^0 and τ^0 and $|\pi^1| \leq |\text{States}| - 1$. Thus $|\pi| = |\pi^0| + |\pi^1| + 1 = \frac{(|\text{States}|+2)(|\text{States}|-1)}{2} + 1 \leq |\text{States}|^2$. The correspondence function c_1 is defined in the same way for the final part of π .

We use the same technique to create τ^0 and τ^1 , which gives us the same bounds as for π^0 and π^1 . In addition, since the first state $\tau_0^0 = \rho_l$ appears after all states that appear finitely often in ρ , the construction is done so all states in τ^0 and τ^1 appear infinitely often in ρ . A new correspondence function c_2 is defined in the same way for mapping indexes of τ to indexes of the corresponding states in ρ . Note that we have $\pi_{k+1} = \rho_{c_1(k)+1}$ and $\tau_{k+1} = \rho_{c_2(k)+1}$.

We now define a symmetric strategy profile σ' with the outcome $\pi \cdot \tau^\omega$ as follows

- $\sigma'(\pi_{\leq k}) = \sigma(\rho_{\leq c_1(k)})$ for prefixes $\pi_{\leq k}$ of π
- $\sigma'(\pi \cdot \tau^{k_1} \cdot \tau_{\leq k_2}) = \sigma(\rho_{\leq c_2(k_2)})$ for all other prefixes of $\pi \cdot \tau^\omega$
- $\sigma'(\pi_{\leq k} \cdot h) = \sigma(\rho_{\leq c_1(k)} \cdot h)$ for finite paths deviating from $\pi \cdot \tau^\omega$ after $\pi_{\leq k}$
- $\sigma'(\pi \cdot \tau^{k_1} \cdot \tau_{\leq k_2} \cdot h) = \sigma(\rho_{\leq c_2(k_2)} \cdot h)$ for finite paths deviating from $\pi \cdot \tau^\omega$ after $\pi \cdot \tau^{k_1} \cdot \tau_{\leq k_2}$

Now for prefixes of $\pi \cdot \tau^\omega$ we have

$$\sigma'(\pi_{\leq k}) = \sigma(\rho_{\leq c_1(k)}) = \rho_{c_1(k)+1} = \pi_{k+1} \text{ and}$$

$$\sigma'(\pi \cdot \tau^{k_1} \cdot \tau_{\leq k_2}) = \sigma(\rho_{\leq c_2(k_2)}) = \rho_{c_2(k_2)+1} = \tau_{k_2+1}$$

which means that the outcome of σ' is indeed $\pi \cdot \tau^\omega$. Since σ' only uses moves which are also used by σ and σ is symmetric, then σ' is also symmetric. Now we just need to show that σ' is also a Nash equilibrium. Suppose a player A deviates from σ' at some point. If the outcome is unchanged then he does not improve. If the outcome deviates from $\pi \cdot \tau^\omega$ then from the point of deviation $\pi_{\leq k}$ or $\pi \cdot \tau^{k_1} \cdot \tau_{\leq k_2}$ the players in σ' will play as they would in σ from $\rho_{\leq c_1(k)}$ or $\rho_{\leq c_2(k_2)}$ from that point on respectively. And since the states occurring in both these cases are the same, this means that if A can deviate and improve from σ' then he can also deviate and improve from σ . But since σ is a Nash equilibrium he cannot deviate and improve from σ . Thus he can neither deviate and improve from σ' which makes σ' a Nash equilibrium. \square

A.2 Proof of Theorem 10

Theorem 10. *The (constrained) symmetric existence problem is \mathcal{PSPACE} -hard for games with 2 players with reachability objectives and preference relations represented by boolean circuits.*

Proof. We do the proof as a reduction from the QSAT problem which is \mathcal{PSPACE} -complete. In this problem one tries to determine whether a quantified boolean formula is satisfiable or not. We can without loss of generality assume that an input formula φ is of the form $\varphi = Q_1 Q_2 \dots Q_m \varphi'(x_1, \dots, x_m)$ where $Q_i \in \{\exists, \forall\}$ and $\varphi'(x_1, \dots, x_m)$ is a boolean formula in conjunctive normal form with free variables x_1, \dots, x_m .

For a given input formula φ of the form above we construct a two-player game \mathcal{G} with players A and B which has a symmetric Nash equilibrium if and only if φ is satisfiable. Let $\varphi' = \bigwedge_{i=1}^n C_i$ be in CNF with n clauses of the form $C_i = \bigvee_{j=1}^{s(i)} l_{i,j}$ where $l_{i,j}$ is a literal (a variable or its negation) and $s(i)$ is the number of literals in clause C_i . We construct the game \mathcal{G} where the set of states is $\text{States} = \bigcup_{i=1}^m \{Q_i\} \cup \bigcup_{i=1}^n \bigcup_{j=1}^{s(i)} \{l_{i,j}\} \cup \{z, u\}$, thus we have one state for each quantifier, one state for each literal and two extra states z and u . The set of legal actions for each player is $\{T, F\}$ in each state and from each state $Q_i = \forall$ the possible transitions are defined by $\text{Mov}(Q_i, \langle T, T \rangle) =$

$\text{Mov}(Q_i, \langle F, T \rangle) = x_i$ and $\text{Mov}(Q_i, \langle T, F \rangle) = \text{Mov}(Q_i, \langle F, F \rangle) = \neg x_i$. In other words, these states are controlled entirely by player B . In the same way the states $Q_i = \exists$ are controlled by player A such that if he chooses T then the play goes to x_i and otherwise the play goes to $\neg x_i$. Further, from states x_i and $\neg x_i$ any move makes the play go to Q_{i+1} except when $i = m$ in which case the play goes to z . From z the play stays in z if the players choose the same action, otherwise it goes to u where it will stay infinitely. An example of the game for the formula $\varphi = \forall x_1 \exists x_2 \forall x_3 (x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2)$ can be seen in Figure 7.

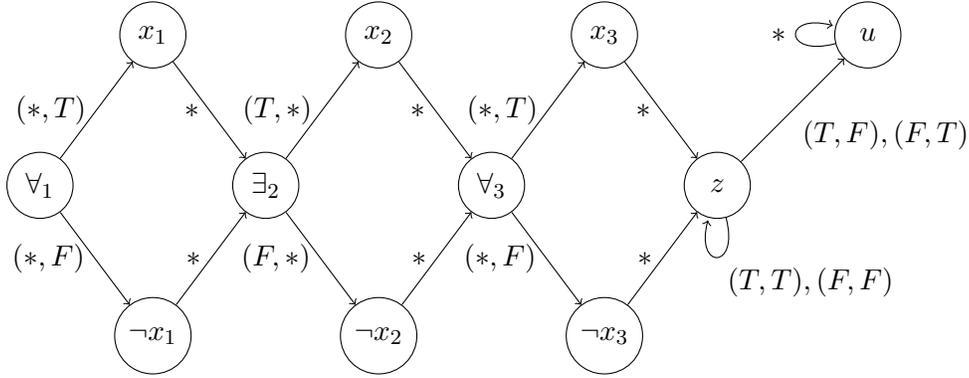


Figure 7: Game constructed from $\varphi = \forall x_1 \exists x_2 \forall x_3 (x_1 \vee x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2)$ where $*$ means any move or any action depending on the context.

We create one reachability objective Ω_i for each clause C_i such that $\Omega_i = \text{Reach}(\{l_{i,1}, \dots, l_{i,s(i)}\})$ and let $\Omega_{m+1} = \text{Reach}(\{u\})$. We let player A prefer a payoff where $\Omega_1, \dots, \Omega_m$ is accomplished over all other payoffs. However, if this is not obtainable he will strictly prefer obtaining Ω_{m+1} over any other payoff. Player B strictly prefers any payoff where Ω_{m+1} is not accomplished over any other payoff. In the case where Ω_{m+1} is not accomplished he will strictly prefer that not all of $\Omega_1, \dots, \Omega_m$ is accomplished.

We now wish to prove that φ is satisfiable if and only if there is a symmetric Nash equilibrium in the game defined above. First, consider the case where φ is satisfiable. Recall that player A controls the existential quantifiers and player B controls the universal quantifiers. Since φ is satisfiable then player A can choose a strategy such that no matter which strategy player B chooses $\Omega_1, \dots, \Omega_m$ is accomplished since each objective Ω_i corresponds to satisfying the clause C_i of φ . We then define a strategy profile $\sigma = (\sigma_A, \sigma_B)$ such that σ_A makes sure that $\Omega_1, \dots, \Omega_m$ are accomplished no

matter which strategy player B uses. Define in addition $\sigma_B(h) = \sigma_A(h)$ for every history h . This strategy profile is clearly symmetric and we will also show that it is a Nash equilibrium. Since the outcome of σ gives the maximum payoff for A he cannot deviate and improve his payoff. Player B can only improve his payoff if he can deviate in such a way that one of $\Omega_1, \dots, \Omega_m$ is not accomplished. However, σ_A is defined in a way such that this is not possible. Thus, B cannot improve by deviating and σ is a symmetric Nash equilibrium.

Now we consider the case where φ is not satisfiable and assume for contradiction that there exists a symmetric Nash equilibrium $\sigma = (\sigma_A, \sigma_B)$. Using this profile, u is never reached since it is symmetric. The objectives $\Omega_1, \dots, \Omega_m$ cannot all be accomplished, because if they were then B would be able to deviate from his strategy (possible at more than node) and improve given that φ is unsatisfiable in which case σ would not be a symmetric Nash equilibrium. Since $\Omega_1, \dots, \Omega_m$ are not all accomplished player A can improve his strategy by changing his action in z , because then going to u will give him a larger payoff. This contradicts that σ is a symmetric Nash equilibrium and therefore there exists no symmetric Nash equilibrium when φ is not satisfiable. \square

A.3 Proof of Theorem 11

To prove Theorem 11 we first need a few lemmas.

Lemma 27. *Let $P \subseteq \text{Agt}$ and $s_0 \in \text{States}$. Then $s_0 \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ if and only if there exists an infinite path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ starting from s_0 .*

Proof. Suppose first that there exists an infinite path $\pi = s_0 s_1 \dots$ in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ starting from s_0 . Since $(s_0, s_1) \in \text{Edg}$ then by definition there exists $m \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(s_0, P)$ such that $\text{Tab}(s_0, m) = s_1$. In addition $s_0 \notin \Omega_A$ for all $A \in P$. This implies that $s_0 \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$.

On the other hand suppose that $s_0 \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$. We now construct an infinite path using induction. When generating s_{i+1} we assume that $s_i \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ and then at each step choose (s_i, s_{i+1}) so $\text{Tab}(s_i, m) = s_{i+1}$ and $m \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(s_i, P)$ which is possible since $s_i \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$. It then holds that (s_i, s_{i+1}) is an edge in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}$. Now $s_{i+1} \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}((s_i, s_{i+1}), m))$. But since m moves the play from s_i to s_{i+1} we have $P \cap \text{Susp}((s_i, s_{i+1}), m) = P$ which means that $s_{i+1} \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$. Since initially $s_0 \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ the construction can be continued like this inductively and thus there exists an infinite path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ starting from s_0 . \square

From [4] we have the following lemma

Lemma 28. *Let $P \subseteq \text{Agt}$ be a subset of agents and $(\sigma_A)_{A \in P}$ be a strategy of coalition P . Let $s \in \text{States}$ and $\pi \in \text{Out}^f(s, (\sigma_A)_{A \in P})$ ending in some state s' . For any history π' starting in s' define $\sigma_A^{-\pi}(\pi') = \sigma_A(\pi \cdot \pi')$. Then*

$$\pi \cdot \text{Out}(s', (\sigma_A^{-\pi})_{A \in P}) \subseteq \text{Out}(s, (\sigma_A)_{A \in P})$$

Next we wish to show the following

Lemma 29. *Let $P \subseteq \text{Agt}$ and $\pi \in \text{Play}_{\mathcal{G}}(s)$ for some state s . Then π is a path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ if and only if there exists $\sigma = (\sigma_A)_{A \in \text{Agt}} \in \text{Prof}^{\text{Sym}}$ so $\text{Out}(s, \sigma) = \{\pi\}$ and for all $B \in P$ and $\sigma'_B \in \text{Strat}^B$ it holds that $\text{Occ}(\pi') \cap \Omega_B = \emptyset$ where $\text{Out}(s, \sigma[\sigma_B \mapsto \sigma'_B]) = \{\pi'\}$*

Proof. First assume that π is a path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$. We construct a symmetric strategy profile σ with the desired properties as follows:

- If π' is a prefix $\pi_{\leq k}$ of π then let $\sigma_A(\pi') = m_A$ for all $A \in \text{Agt}$ where $m = (m_A)_{A \in \text{Agt}} \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(\pi_k, P)$ and $\text{Tab}(\pi_k, m) = \pi_{k+1}$. Such a move exists since π is a path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$.
- If π' is not a prefix of π then let $\sigma_A(\pi') = m_A$ for all $A \in \text{Agt}$ where $m = (m_A)_{A \in \text{Agt}} \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(\text{last}(\pi'), P \cap \text{Susp}(\pi', \sigma))$. The fact that such a move exists can be seen by induction on the length of π' :
 - If $|\pi'| = 1$ then π'_1 is the first state where the play deviates from π . Since $\pi'_0 = s \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P)$ we must have $\pi'_1 \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}((\pi'_0, \pi'_1), m))$ where m is a secure symmetric move as defined in the previous point. From this follows that there is a secure symmetric move $m' \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(\pi'_1, P \cap \text{Susp}((\pi'_0, \pi'_1), m)) = \text{Secure}_{\mathcal{G}}^{\text{Sym}}(\text{last}(\pi'), P \cap \text{Susp}(\pi', \sigma))$.
 - If $|\pi'| > 1$ the induction hypothesis says that there is an $m \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(\pi'_{|\pi'|-1}, P \cap \text{Susp}(\pi'_{\leq |\pi'|-1}, \sigma))$ and therefore $\pi'_{|\pi'|-1} \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}(\pi'_{\leq |\pi'|-1}, \sigma))$. From this follows that $\pi'_{|\pi'|} \in \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}(\pi'_{\leq |\pi'|-1}, \sigma) \cap \text{Susp}((\pi'_{|\pi'|-1}, \pi'_{|\pi'|}), m)) = \text{Rep}_{\mathcal{G}}^{\text{Sym}}(P \cap \text{Susp}(\pi', \sigma))$ which means that there is an $m' \in \text{Secure}_{\mathcal{G}}^{\text{Sym}}(\text{last}(\pi'), P \cap \text{Susp}(\pi', \sigma))$.

According to the first point above, we have $\text{Out}(s, \sigma) = \{\pi\}$. In addition, σ is a symmetric strategy profile since it only uses secure symmetric moves. Now we let B be a player in P and let him deviate and use strategy σ'_B . We wish to prove that when $\text{Out}(s, \sigma[\sigma_B \mapsto \sigma'_B]) = \{\pi'\}$ with $\pi' = (s_i)_{i \geq 0}$ it holds that $s_i \in \text{Rep}_G^{\text{Sym}}(\{B\})$ for all $i \geq 0$ which means that $\text{Occ}(\pi') \cap \Omega_B \subseteq \text{Rep}_G^{\text{Sym}}(\{B\}) \cap \Omega_B = \emptyset$. This will be done by induction:

Since π is an infinite path from s_0 in $\mathcal{S}_G^{\text{Sym}}(P)$ by Lemma 27 we have $s_0 \in \text{Rep}_G^{\text{Sym}}(P)$. And since $\{B\} \subseteq P$ we have $s_0 \in \text{Rep}_G^{\text{Sym}}(P) \subseteq \text{Rep}_G^{\text{Sym}}(\{B\})$. As the induction hypothesis we assume that $s_i \in \text{Rep}_G^{\text{Sym}}(\{B\})$ for some $i \geq 0$. By definition of σ , we have $\sigma(\pi'_{\leq i}) = m$ where $m \in \text{Secure}_G^{\text{Sym}}(\pi'_i, P \cap \text{Susp}(\pi'_{\leq i}, \sigma))$. From this follows that $\pi'_{i+1} \in \text{Rep}_G^{\text{Sym}}(P \cap \text{Susp}(\pi'_{\leq i}, \sigma) \cap \text{Susp}((\pi'_i, \pi'_{i+1}), m))$. Since B is suspect for both the finite history $\pi'_{\leq i}$ and the transition (π'_i, π'_{i+1}) when all other players play according to σ we have that $\{B\} \subseteq P \cap \text{Susp}(\pi'_{\leq i}, \sigma) \cap \text{Susp}((\pi'_i, \pi'_{i+1}), m)$ and therefore

$$\pi'_{i+1} \in \text{Rep}_G^{\text{Sym}}(P \cap \text{Susp}(\pi'_{\leq i}, \sigma) \cap \text{Susp}((\pi'_i, \pi'_{i+1}), m)) \subseteq \text{Rep}_G^{\text{Sym}}(\{B\})$$

which concludes the induction and the first direction of the proof.

For the other direction assume that there exists $\sigma = (\sigma_A)_{A \in \text{Agt}} \in \text{Prof}^{\text{Sym}}$ so $\text{Out}(s, \sigma) = \{\pi\}$ and for all $B \in P$ and $\sigma'_B \in \text{Strat}^B$ it holds that $\text{Occ}(\pi') \cap \Omega_B = \emptyset$ where $\text{Out}(s, \sigma[\sigma_B \mapsto \sigma'_B]) = \{\pi'\}$. We need to show that π is a path in $\mathcal{S}_G^{\text{Sym}}(P)$. This will be done by induction on P .

In the case where $P = \emptyset$ we have $\text{Rep}_G^{\text{Sym}}(P) = \text{Rep}_G^{\text{Sym}}(\emptyset) = \text{States}$ in which case the set $\text{Secure}_G^{\text{Sym}}(s, P)$ equals the set of all symmetric moves for every state s . This means that $\mathcal{S}(P)$ contains exactly the edges corresponding to symmetric moves and in particular it contains all edges in π .

As induction hypothesis we assume that the theorem holds for all proper subsets of P . Now, consider the set

$$S = \{s \in \text{States} \mid \exists \rho \in \text{Hist}_G. \rho|_{|\rho|} = s \wedge P \subseteq \text{Susp}(\rho, \sigma)\}$$

Consider some state $s' \in S$ where $\rho' \in \text{Hist}_G$, $\rho'|_{|\rho'|} = s'$ and $P \subseteq \text{Susp}(\rho', \sigma)$. Since every player in P can deviate to reach s' given that anyone else plays according to σ we have $s' \notin \Omega_B$ for all $B \in P$. Since s' is an arbitrary state in S , we have

$$S \cap \Omega_B = \emptyset \text{ for all } B \in P$$

Next, consider an arbitrary $s'' \in \text{States}$. Let $P' = P \cap \text{Susp}((s', s''), (\sigma_A(\rho'))_{A \in \text{Agt}})$. We look at the two possibilities

- $P' = P$
- $P' \subset P$

In the first case, $s'' \in S$ by definition. In the second case for all $B \in P'$ and all $\sigma'_B \in \text{Strat}^B$ we have

$$\rho' \cdot \text{Out}(s'', \sigma^{-\rho'}[\sigma_B \mapsto \sigma'_B]) \subseteq \text{Out}(s', \sigma[\sigma_B \mapsto \sigma'_B]) \subseteq (\text{States} \cap \Omega_B)^\omega$$

according to Lemma 28. In other words, the premises of the induction hypothesis are fulfilled for s'' , which means that $s'' \in \mathcal{S}_G^{\text{Sym}}(P')$ and therefore by Lemma 27 $s'' \in \text{Rep}_G^{\text{Sym}}(P') = \text{Rep}_G^{\text{Sym}}(P \cap \text{Susp}((s', s''), (\sigma_A(\rho'))_{A \in \text{Agt}}))$. If we look at the definition of $\text{Rep}_G^{\text{Sym}}(P)$ we have that S satisfies the same requirements using the results obtained above. Since $\text{Rep}_G^{\text{Sym}}(P)$ is the largest set satisfying the requirements and it is the union of all sets satisfying them, we have that $S \subseteq \text{Rep}_G^{\text{Sym}}(P)$. \square

Using the lemmas introduced in this section, we can now prove Theorem 11, which will provide us with an important connection between the transition systems $\mathcal{S}_G^{\text{Sym}}(P)$ and the existence of symmetric Nash equilibria in \mathcal{G} when each player has one reachability objective.

Theorem 11. *Let \mathcal{G} be a concurrent game with one reachability objective Ω_A for each $A \in \text{Agt}$. Let v be the payoff where Ω_B is accomplished iff $B \notin P$ for some $P \subseteq \text{Agt}$. Let $s \in \text{States}$. Then there is a symmetric Nash equilibrium from s with payoff v iff there is an infinite path π in $\mathcal{S}_G^{\text{Sym}}(P)$ starting in s which visits Ω_A for every $A \notin P$.*

Proof. If there is a symmetric Nash equilibrium $\sigma = (\sigma_A)_{A \in \text{Agt}}$ from s with payoff v and outcome π then by the definition of Nash equilibrium we have that for all $B \subseteq P$ and all $\sigma'_B \in \text{Strat}^B$ it holds that $\text{Occ}(\pi') \cap \Omega_B = \emptyset$ where $\text{Out}(s, \sigma[\sigma_B \mapsto \sigma'_B]) = \{\pi'\}$. According to Lemma 29 π is an infinite path in $\mathcal{S}_G^{\text{Sym}}(P)$ starting in s . In addition, π visits Ω_A for all $A \notin P$ by definition.

On the other hand, if there is an infinite path $\pi \in \mathcal{S}_G^{\text{Sym}}(P)$ starting in s which visits Ω_A for all $A \notin P$ then according to 29 there exists a symmetric

strategy profile σ with outcome π and such that no $B \in P$ can unilaterally deviate from σ to reach Ω_B . Thus, since no player in P can improve his payoff and all players in $\text{Agt} \setminus P$ gets their maximum payoff in π no agent can improve by deviating from σ and it must be a symmetric Nash equilibrium from s with payoff v . \square

A.4 Proof of Theorem 12

Theorem 12. *The symmetric constrained existence problem is in \mathcal{P} for a bounded number of players and in \mathcal{NP} for an unbounded number of players when each player has a single reachability objective.*

Proof. We start by providing an algorithm solving the problem in polynomial time for a bounded number of players. Since the number of players is bounded, so is the number of possible payoffs. Our algorithm works by checking for every payoff satisfying the constraints whether there is a symmetric Nash equilibrium with the given payoff. This check needs to be done in polynomial time. The checking algorithm works by using Theorem 11 which tells us that there is a symmetric Nash equilibrium from s with payoff $v = \mathbb{1}_{\text{Agt} \setminus P}$ if and only if there is an infinite path π in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ starting from s which visits Ω_A for every $A \notin P$. To do the check we generate $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$. To do this we first calculate $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P')$ for all $P' \subseteq P$ from which we will obtain $\text{Secure}_{\mathcal{G}}^{\text{Sym}}(s', P)$ for all $s' \in \text{States}$. Then we have $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ from which we can check if there is an infinite path from s that visits Ω_A for all $A \notin P$. The way to calculate $\text{Rep}_{\mathcal{G}}^{\text{Sym}}(P')$ for $P' \subseteq P$ is done by using an alternative (but equivalent) definition of the symmetric repeller sets:

Definition 30. *Assume that $\text{Rep}^{\text{Sym}}(P')$ have been defined for all $P' \subsetneq P$. Then let*

- $\text{Rep}_0^{\text{Sym}}(P) = \text{States} \setminus \bigcup_{B \in P} \Omega_B$
 - $\text{Rep}_{i+1}^{\text{Sym}}(P) = \text{Rep}_i^{\text{Sym}}(P) \setminus \{s \in \text{Rep}_i^{\text{Sym}}(P) \mid \neg \exists a \in \text{Act} \forall t \in \text{States}. t \in \text{Rep}_i^{\text{Sym}}(P \cap \text{Susp}((s, t), m_a^{\text{Sym}}))\}$
- where we let $\text{Rep}_i^{\text{Sym}}(P') = \text{Rep}^{\text{Sym}}(P')$ for all $P' \subsetneq P$.

Define $\text{Rep}^{\text{Sym}}(P)$ to be the limit of the decreasing sequence, i.e. $\text{Rep}^{\text{Sym}}(P) = \text{Rep}_{\infty}^{\text{Sym}}(P)$

This inductive definition is used to create the repellor sets "bottom up" for every $P' \subseteq P$. For each P' the fixpoint of the sequence is calculated in at most $|\text{States}|$ steps. Since there are at most $2^{|\text{Agt}|}$ subsets of P there are at most $\mathcal{O}(|\text{States}| \cdot 2^{|\text{Agt}|})$ iterations in total. In each of these iterations we need to calculate for each action (symmetric move) and each state a set of suspect players and then make a lookup in the previously calculated repellor sets. Thus, the calculation takes $\mathcal{O}(|\text{States}| \cdot 2^{|\text{Agt}|} \cdot |\text{Act}| \cdot |\text{States}| \cdot |\text{Tab}|) = \mathcal{O}(|\text{States}|^2 \cdot 2^{|\text{Agt}|} \cdot |\text{Act}| \cdot |\text{Tab}|)$ which is polynomial since the number of agents is bounded. Thus the transition system $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ can be calculated in polynomial time. In this transition system we need to find an infinite path that contains a state from each of the sets Ω_A for $A \in \text{Agt} \setminus P$. Choosing a state from each of these sets can be done in $\prod_{A \in \text{Agt} \setminus P} |\Omega_A| \leq |\text{States}|^{|\text{Agt}|}$ ways, which is polynomial when the number of players is bounded. We can easily check paths for all these possibilities in polynomial time by checking all the different possible orders of occurrence of states, of which there are $|\text{Agt}|!$ in each case. In total, this means we can solve the symmetric constrained existence problem in polynomial time.

We now turn our attention to the problem with an unbounded number of players where we wish to create an \mathcal{NP} algorithm. The basic idea is the same as above. However, we cannot calculate all of $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ because of the exponential number of subsets of P and the exponential number of possible payoffs. What we do is to non-deterministically guess a payoff v and then guess a path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ and then try to check if the path is indeed a path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ which visits Ω_A for all $A \notin P$. According to Lemma 7 we only need to guess paths of the form $\rho = \pi \cdot \tau^\omega$ where $|\pi|, |\omega| \leq |\text{States}|^2$. This means that checking that it visits Ω_A for all $A \notin P$ can be done in polynomial time. However, we also need to be able to check that the path is a path in $\mathcal{S}_{\mathcal{G}}^{\text{Sym}}(P)$ in polynomial time which requires some more analysis. The idea is that we do not have to calculate repellor sets for all subsets of P , in fact we only need to calculate it for polynomially many subsets which is shown in [4]. The analysis just needs to be adjusted to the case of symmetric repellor sets. Then we can non-deterministically guess which subsets to use and then the checking algorithm will run in polynomial time, which means the problem is solvable in \mathcal{NP} . \square

A.5 Proof of Theorem 13 (Bounded number of players)

Theorem 13a. *The symmetric existence problem is \mathcal{P} -hard for a bounded number of players when each player has a single reachability objective*

Proof. The proof is done as a reduction from the CIRCUITVALUE problem which is known to be \mathcal{P} -complete. Suppose we have an instance C of the CIRCUITVALUE-problem where without loss of generality we can assume that there are only AND-gates and OR-gates. From a given instance to this problem we create a 2-player turn based game \mathcal{G} where the states are the gates, player A controls the OR-gates and player B controls the AND-gates. This means that the players can move the play either to the left subcircuit or the right subcircuit using actions L and R respectively. The other player can choose the same actions, but since the game is turn-based only the player controlling a gate has influence on the transition taken. The initial state is the output-gate of the circuit and the final states are the inputs. The goal of player A is to reach a positive input, whereas the goal of player B is to reach a negative input. The situation is illustrated above the dashed line in Figure 8. In addition, at each negative input we add an extra module (below the dashed line), such that if both players choose the same action at a negative output state then the play will move to a state in Ω_A , otherwise it will move to a state that is not and stay there.

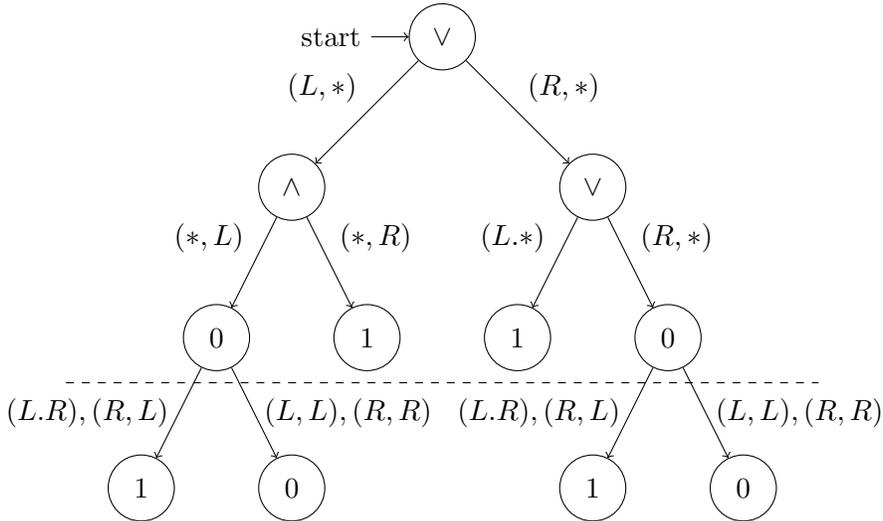


Figure 8: Game \mathcal{G} where $*$ means any action.

The idea is now that if the original circuit C evaluates to 1, then player A has a strategy to make sure he reaches one of the nodes in Ω_A corresponding to an input gate with label 1. If B does the same as player A after all histories, this gives us a symmetric strategy profile in which none of the player

can improve by deviating. Thus, there is a symmetric Nash equilibrium in \mathcal{G} if C evaluates to 1.

If C evaluates to 0 assume for contradiction that there is a symmetric Nash equilibrium σ in \mathcal{G} . Since B can make sure that a negative input gate is reached, the run of σ cannot reach one of the positive input gates because then B would be able to deviate and improve. Since σ is symmetric then after reaching a state corresponding to a negative input gate, the next state will be a state which is not in Ω_A . But then A can deviate by changing his action. Thus, there cannot be a symmetric Nash equilibrium in this case. In total, C evaluates to 1 if and only if there is a symmetric Nash equilibrium in \mathcal{G} which means that the symmetric existence problem is \mathcal{P} -hard for a bounded number of players when each player has a single reachability objective. \square

A.6 Proof of Theorem 13 (Unbounded number of players)

Theorem 13b. *The symmetric existence problem is \mathcal{NP} -hard for an unbounded number of players when each player has a single reachability objective.*

Proof. The proof is done as a reduction from the 3SAT problem which is known to be \mathcal{NP} -complete. An instance of 3SAT is a formula $\varphi = \bigwedge_{i=1}^n C_i$ in conjunctive normal form over the set of variables $V = \{x_1, \dots, x_m\}$ where each clause C_i is of the form $C_i = l_{i,1} \vee l_{i,2} \vee l_{i,3}$ where each of the literals $l_{i,j} \in \{x|x \in V\} \cup \{\neg x|x \in V\}$. From this formula we generate a game \mathcal{G}_φ such that φ is satisfiable if and only if there is a symmetric Nash equilibrium in \mathcal{G}_φ . There are $n + m + 2$ players, one for each clause called C_i , one for each variable called l_i and two additional players A_1 and A_2 . The game is designed as shown in Figure 9.

In every state every player can choose the actions 0 and 1. In state C if any of the players C_i chooses 0 then the play goes to A otherwise it goes to L_1 . In state L_i if player l_i chooses 1 the play goes to x_i otherwise it goes to $\neg x_i$. Finally, in state A if A_1 and A_2 chooses the same action the play goes to T_1 otherwise it goes to T_2 . Every player has one reachability objective defined as

$$\Omega_{l_i} = \Omega_{\text{Reach}}(\{x_i, \neg x_i\})$$

$$\Omega_{C_i} = \Omega_{\text{Reach}}(\{A, l_{i,1}, l_{i,2}, l_{i,3}\})$$

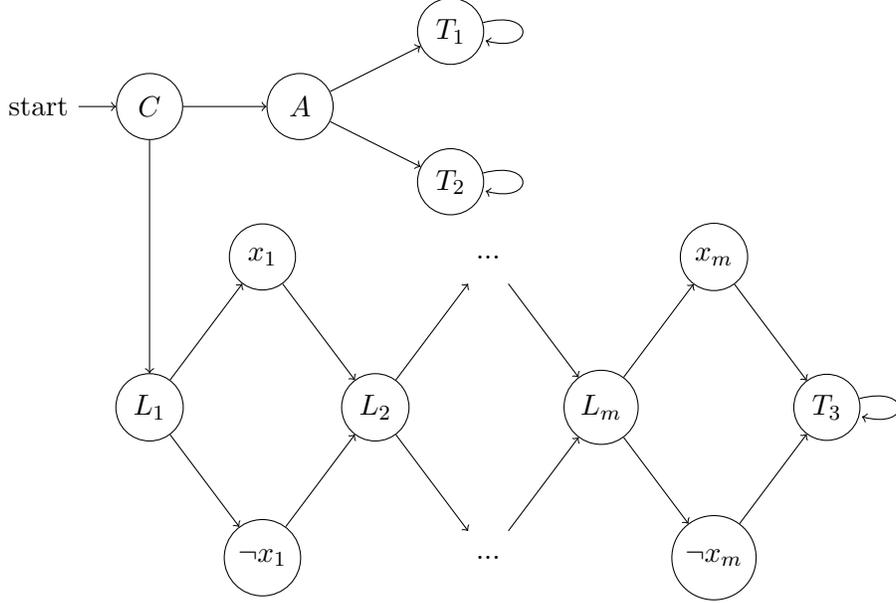


Figure 9: Game \mathcal{G}_φ where $*$ means any move or any action depending on the context.

$$\Omega_{A_1} = \Omega_{\text{Reach}}(\{T_i\})$$

Now assume φ is satisfiable by the values $x_1^*, x_2^*, \dots, x_m^* \in \{0, 1\}$. Design a strategy σ so all players choose 1 in state C , so all players choose the same action in histories ending in A, T_1, T_2 , so all players choose x_i^* in state L_i and all players choose the same action after histories ending in T_3 . Then all players except A_1 and A_2 reach their target. Thus all players except A_1 and A_2 cannot improve their payoff. Since neither of the players A_1 and A_2 can deviate to make the play reach their target σ is a Nash equilibrium. In addition, since σ is symmetric it is a symmetric Nash equilibrium.

Next assume φ is unsatisfiable. Then it is not possible to choose a path from L_1 to T_3 that corresponds to a satisfying assignment of φ . This means that if L_1 is reached, then there exists j so player C_j does not reach his target since the targets corresponds to clauses being satisfied. This means that there cannot be a Nash equilibrium where L_1 is reached in this case since C_j can deviate to move the play to A and improve his payoff. When A is reached only one of the players A_1 and A_2 can reach their target. Since the player not reaching his target in the outcome of a strategy profile can

always improve by deviating there cannot be a Nash equilibrium when A is reached which means that there exists no Nash equilibrium when φ is unsatisfiable and thus no symmetric Nash equilibrium. \square

A.7 Proof of Lemma 16

Lemma 16. *For all permutations π of $(1, \dots, n)$, all states s and legal moves m from s we have*

$$\text{Tab}(s, m) = s' \Leftrightarrow \text{Tab}(h_\pi(s), m_\pi) = h_\pi(s')$$

Proof. Let α be a sequence of moves such that $s_0 \xrightarrow{\alpha} s$ and $m = (a_1, \dots, a_n)$. Denote by $\alpha' = \alpha \cdot m$ the sequence of moves beginning with α and adding the move m at the end. We now have

$$\begin{aligned} \text{Tab}(s, m) &= s' \\ \Leftrightarrow h_\pi(\text{Tab}(s, m)) &= h_\pi(s') \end{aligned}$$

since h_π is bijective. By rewriting the left-hand side we obtain the result

$$\begin{aligned} &h_\pi(\text{Tab}(s, m)) \\ &= h_\pi(\text{last}(\text{Out}(\alpha'))) \\ &= \text{last}(\text{Out}(\alpha'_\pi)) \\ &= \text{Tab}(\text{last}(\text{Out}(\alpha_\pi)), m_\pi) \\ &= \text{Tab}(h_\pi(\text{last}(\text{Out}(\alpha))), m_\pi) \\ &= \text{Tab}(h_\pi(s), m_\pi) \end{aligned} \quad \square$$

A.8 Proof of Theorem 18

Theorem 18. *There exists symmetric concurrent games with no symmetric Nash equilibrium in mixed strategies*

Proof. We will present such a game: Two agents must pass through a door, but if they try to pass through it at the same time they will break. The game can run for an infinite number of time steps and at each time step

each agents has the actions s (stay) and g (go). If they choose go at the same time step they will both break and lose, if one agent chooses go at some time step and the other chooses stay, they will both win. If no agent ever goes, they will both lose.

Define \mathcal{G} as the game in Figure 10 representing the situation where the initial state is s_0 and where both players has the single goal of reaching s_2 , giving utility 1. All other paths gives utility 0. We will show that there is no symmetric Nash equilibrium in mixed strategies in this game.

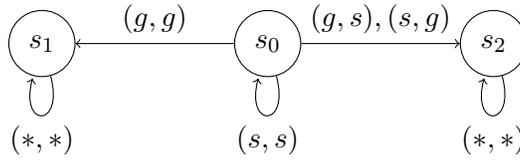


Figure 10: Concurrent game \mathcal{G} modelling the symmetric normal-form game \mathcal{N}

We have the mappings $h_{\pi_1}(s_i) = h_{\pi_2}(s_i) = s_i$ for $0 \leq i \leq 2$, $\pi_1 = (1, 2)$ and $\pi_2 = (2, 1)$. This means that in a symmetric Nash equilibrium, the players should choose the same action distribution after each history in \mathcal{G} .

Let p_i be the probability that each agent will go at time step i given that the game is not finished at that time step, where $i = 0$ in the first time step. The p_i defines a symmetric strategy profile in mixed strategies. Suppose for contradiction that this strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ is a Nash equilibrium. If $p_i = 0$ for some i , then the strategy profile is not a Nash equilibrium, since a player can deviate to improve his expected utility by waiting until time step i and then going. Thus, $p_i > 0$ for all i . Denote the pure strategy for player 1 of waiting until time step i and then going by σ^i . If player 1 deviates and chooses a pure strategy, he should get the same expected utility when p_2 plays according to σ^* no matter which choice of pure strategy. If this was not the case, then there would exist pure strategies σ^i and σ^j such that $u_1(\sigma^i, \sigma_2^*) > u_1(\sigma^j, \sigma_2^*)$ in which case player 1 could improve from σ_1^* by changing to $p'_i = p_i + p_j$ and $p'_j = 0$ in the Nash equilibrium. Thus, we have $u_1(\sigma^i, \sigma_2^*) = u_1(\sigma^j, \sigma_2^*)$ for all i, j . In particular, this is the case for $i = n$ and $i = n + 1$. The different ways player 1 can win when using σ^i is if player 2 either goes in one of the first $i - 1$ steps or if he does not go in the first i steps. This gives us

$$u(\sigma^n, \sigma_2^*) = u(\sigma^{n+1}, \sigma_2^*)$$

$$\begin{aligned}
&\Rightarrow \sum_{i=0}^{n-1} p_i \prod_{j=0}^{i-1} (1-p_j) + \prod_{i=0}^n (1-p_i) = \sum_{i=0}^n p_i \prod_{j=0}^{i-1} (1-p_j) + \prod_{i=0}^{n+1} (1-p_i) \\
&\Rightarrow \prod_{i=0}^n (1-p_i) = p_n \prod_{j=0}^{n-1} (1-p_j) + \prod_{i=0}^{n+1} (1-p_i) \\
&\Rightarrow 0 = \left(\prod_{j=0}^{n-1} (1-p_j) \right) \cdot (p_n - (1-p_n) + (1-p_n) \cdot (1-p_{n+1})) \\
&\Rightarrow 2p_n - 1 + (1-p_n) \cdot (1-p_{n+1}) = 0 \\
&\Rightarrow p_n - p_{n+1} + p_n \cdot p_{n+1} = 0 \\
&\Rightarrow p_{n+1} = \frac{p_n}{1-p_n}
\end{aligned}$$

Now, by using induction we can prove that $p_n = \frac{p_0}{1-n \cdot p_0}$. Indeed $p_0 = \frac{p_0}{1-0 \cdot p_0}$ and for the induction step we assume the formula holds for n and get

$$p_{n+1} = \frac{p_n}{1-p_n} = \frac{\frac{p_0}{1-n \cdot p_0}}{1 - \frac{p_0}{1-n \cdot p_0}} = \frac{p_0}{1-n \cdot p_0 - p_0} = \frac{p_0}{1-(n+1) \cdot p_0}$$

completing the induction step. But this means that for $n > \frac{1}{p_0}$ we have $1-n \cdot p_0 < 0$ and $p_n < 0$ which is not possible. Thus, our assumption that a symmetric Nash equilibrium in mixed strategies exists must be false. \square

A.9 Proof of Theorem 20

Theorem 20. *From a symmetric game structure \mathcal{G} we can in polynomial time construct a game \mathcal{H} , which is polynomial in the size of \mathcal{G} with the same type of objectives and such that there exists a symmetric Nash equilibrium in \mathcal{H} if and only if there exists a Nash equilibrium in \mathcal{G} .*

Proof. Let $\mathcal{G} = (G, (\text{Nei}_i), (\pi_{i,j}), (\equiv_i), (\Omega_l^i))$ be a symmetric game structure and $(s_{1,0}, \dots, s_{n,0})$ be a configuration as an input to the existence problem. We will generate a symmetric game structure \mathcal{H} which has a symmetric Nash equilibrium from some particular configuration if and only if \mathcal{G} has a Nash equilibrium from $(s_{1,0}, \dots, s_{n,0})$. We generate \mathcal{H} as follows.

Let $\mathcal{H} = (H, (\text{Nei}_i), (\pi_{i,j}), (\sim_i), (\Theta_l^i))$ be a symmetric game structure with n players as \mathcal{G} with the same neighbour functions and the same permutations $\pi_{i,j}$. We design H as n disconnected copies of G as shown in Figure

11.

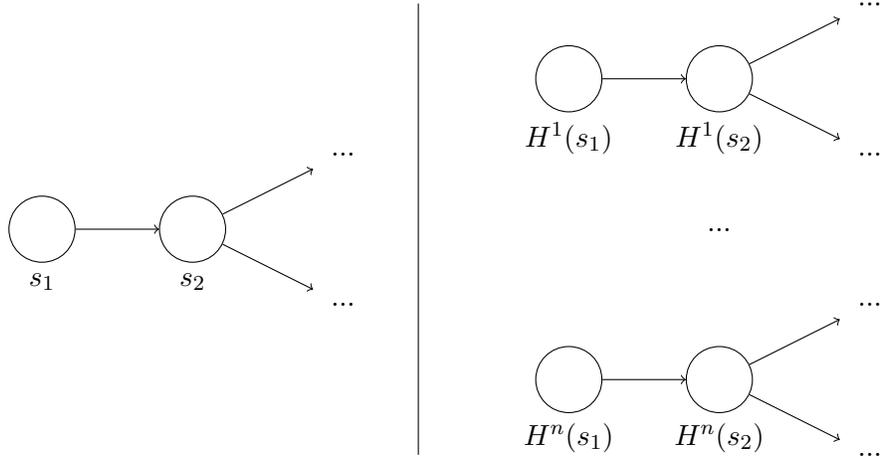


Figure 11: Arena G on the left and H on the right

These copies will be denoted H^1, \dots, H^n . To denote the state in some H^j corresponding to some state s in G we write $H^j(s)$. We introduce the mapping π_1 such that $\pi_1(H^j(s)) = s$ for all states s and all players j . Then we let the initial configuration in \mathcal{H} be $(H^1(s_{1,0}), H^2(s_{2,0}), \dots, H^n(s_{n,0}))$. In other words, in H each player starts in the same states as in \mathcal{G} , but in different copies of G . We define \sim_i such that for all i , all states $s_1, \dots, s_n, s'_1, \dots, s'_n$ in G and all $m_1, \dots, m_n, j_1, \dots, j_n \in \{1, \dots, n\}$:

$$(H^{m_1}(t_1), H^{m_2}(t_2), \dots, H^{m_n}(t_n)) \sim_i (H^{j_1}(v_1), H^{j_2}(v_2), \dots, H^{j_n}(v_n))$$

if and only if

$$(t_1, \dots, t_n) \equiv_i (v_1, \dots, v_n) \wedge m_i = j_i$$

Finally, for every player i , every objective Ω_i^l and every infinite play $\rho = (s_1^1, \dots, s_n^1)(s_1^2, \dots, s_n^2) \dots$ in H we define the objectives of the players in \mathcal{H} such that for all $j_1, \dots, j_n \in \{1, \dots, n\}$

$$\rho \in \Omega_i^l \text{ if and only if}$$

$$(H^{j_1}(s_1^1), \dots, H^{j_n}(s_n^1))(H^{j_1}(s_1^2), \dots, H^{j_n}(s_n^2)) \dots \in \Theta_i^l$$

We start by showing that \mathcal{H} as defined here is really a symmetric game structure. The arena H , the neighbour functions Nei_i and the permutations $\pi_{i,j}$ clearly satisfy the requirements. Next we need to show that \sim_i satisfies the conditions necessary for all i . First we note that every state in H can be written as $H^j(s)$ for a state s in G and an index $j \in \{1, \dots, n\}$ in a unique way. Now \sim_i is reflexive for all i since for all states s_1, \dots, s_n in G and all $j_1, \dots, j_n \in \{1, \dots, n\}$

$$\begin{aligned} (s_1, \dots, s_n) &\equiv_i (s_1, \dots, s_n) \wedge j_i = j_i \\ \Rightarrow (H^{j_1}(s_1), \dots, H^{j_n}(s_n)) &\sim_i (H^{j_1}(s_1), \dots, H^{j_n}(s_n)) \end{aligned}$$

It is symmetric for all i since for all states $s_1, \dots, s_n, s'_1, \dots, s'_n$ in G and all $j_1, \dots, j_n, m_1, \dots, m_n \in \{1, \dots, n\}$

$$\begin{aligned} (H^{j_1}(s_1), \dots, H^{j_n}(s_n)) &\sim_i (H^{m_1}(s'_1), \dots, H^{m_n}(s'_n)) \\ \Rightarrow (s_1, \dots, s_n) &\equiv_i (s'_1, \dots, s'_n) \wedge j_i = m_i \\ \Rightarrow (s'_1, \dots, s'_n) &\equiv_i (s_1, \dots, s_n) \wedge j_i = m_i \\ \Rightarrow (H^{m_1}(s'_1), \dots, H^{m_n}(s'_n)) &\sim_i (H^{j_1}(s_1), \dots, H^{j_n}(s_n)) \end{aligned}$$

It is transitive for all i since for all states $s_1, \dots, s_n, s'_1, \dots, s'_n, s''_1, \dots, s''_n$ in G and all $j_1, \dots, j_n, k_1, \dots, k_n, m_1, \dots, m_n \in \{1, \dots, n\}$

$$\begin{aligned} (H^{j_1}(s_1), \dots, H^{j_n}(s_n)) &\sim_i (H^{k_1}(s'_1), \dots, H^{k_n}(s'_n)) \wedge \\ (H^{k_1}(s'_1), \dots, H^{k_n}(s'_n)) &\sim_i (H^{m_1}(s''_1), \dots, H^{m_n}(s''_n)) \\ \Rightarrow (s_1, \dots, s_n) &\equiv_i (s'_1, \dots, s'_n) \wedge (s'_1, \dots, s'_n) \equiv_i (s''_1, \dots, s''_n) \wedge j_i = k_i \wedge k_i = m_i \\ \Rightarrow (s_1, \dots, s_n) &\equiv_i (s''_1, \dots, s''_n) \wedge j_i = m_i \\ (H^{j_1}(s_1), \dots, H^{j_n}(s_n)) &\sim_i (H^{m_1}(s''_1), \dots, H^{m_n}(s''_n)) \end{aligned}$$

This means that \sim_i is an equivalence relation for all i .

Next, suppose (t_1, \dots, t_n) is a state configuration in \mathcal{H} . Then there are

unique m_1, \dots, m_n and (s_1, \dots, s_n) such that $(t_1, \dots, t_n) = (H^{m_1}(s_1), \dots, H^{m_n}(s_n))$. For all i , all states s_1, \dots, s_n in G and all $m_1, \dots, m_n \in \{1, \dots, n\}$ we have

$$(s_1, \dots, s_n) \equiv_i (s_{\pi(1)}, \dots, s_{\pi(n)}) \wedge m_j = m_{\pi(j)} \text{ for every } \pi \text{ with } \pi(j) = j \text{ for all } j \in N(i)$$

$$\Rightarrow (H^{m_1}(s_1), \dots, H^{m_n}(s_n)) \sim_i (H^{m_{\pi(1)}}(s_{\pi(1)}), \dots, H^{m_{\pi(n)}}(s_{\pi(n)})) \text{ for every } \pi \text{ with } \pi(j) = j \text{ for all } j \in N(i)$$

$$\Rightarrow (t_1, \dots, t_n) \sim_i (t_{\pi(1)}, \dots, t_{\pi(n)}) \text{ for every } \pi \text{ with } \pi(j) = j \text{ for all } j \in N(i)$$

which means that the partial information requirement for \sim_i is satisfied too. The requirement for symmetry between the players is also satisfied since for all $k_1, \dots, k_n, m_1, \dots, m_n \in \{1, \dots, n\}$, all states $s_1, \dots, s_n, s'_1, \dots, s'_n$ in G , all players i, j and all permutations π extending $\pi_{i,j}$ we have

$$(H^{k_1}(s_1), \dots, H^{k_n}(s_n)) \sim_i (H^{m_1}(s'_1), \dots, H^{m_n}(s'_n))$$

$$\Leftrightarrow (s_1, \dots, s_n) \equiv_i (s'_1, \dots, s'_n) \wedge k_i = m_i$$

$$\Leftrightarrow (s_{\pi^{-1}(1)}, \dots, s_{\pi^{-1}(n)}) \equiv_j (s'_{\pi^{-1}(1)}, \dots, s'_{\pi^{-1}(n)}) \wedge k_{\pi^{-1}(j)} = m_{\pi^{-1}(j)}$$

$$\Leftrightarrow (H^{k_{\pi^{-1}(1)}}(s_{\pi^{-1}(1)}), \dots, H^{k_{\pi^{-1}(n)}}(s_{\pi^{-1}(n)})) \sim_j$$

$$(H^{m_{\pi^{-1}(1)}}(s'_{\pi^{-1}(1)}), \dots, H^{m_{\pi^{-1}(n)}}(s'_{\pi^{-1}(n)}))$$

We now need to show that all paths ρ of state configurations in H satisfy $\rho \in \Theta_i^l \Leftrightarrow \pi^{-1}(\rho) \in \Theta_j^l$ for all players i, j and all π which extend $\pi_{i,j}$. Let $\rho = (t_1^1, \dots, t_n^1)(t_1^2, \dots, t_n^2) \dots$ be an arbitrary legal play and let $t_k^j = H^{i_k}(s_k^j)$ for all j, k . Then we can write $\rho = (H^{i_1}(s_1^1), \dots, H^{i_n}(s_n^1))(H^{i_1}(s_1^2), \dots, H^{i_n}(s_n^2)) \dots$. Note that the i_k do not change since in a legal play in H every player will stay in the same copy of G in which they start. Now we get for all i, j, l and all π extending $\pi_{i,j}$

$$(t_1^1, \dots, t_n^1)(t_1^2, \dots, t_n^2) \dots \in \Theta_i^l$$

$$\Leftrightarrow (H^{i_1}(s_1^1), \dots, H^{i_n}(s_n^1))(H^{i_1}(s_1^2), \dots, H^{i_n}(s_n^2)) \dots \in \Theta_i^l$$

$$\Leftrightarrow (s_1^1, \dots, s_n^1)(s_1^2, \dots, s_n^2) \dots \in \Omega_i^l$$

$$\begin{aligned}
&\Leftrightarrow (s_{\pi^{-1}(1)}^1, \dots, s_{\pi^{-1}(n)}^1)(s_{\pi^{-1}(1)}^2, \dots, s_{\pi^{-1}(n)}^2) \dots \in \Omega_j^l \\
&\Leftrightarrow (H^{i_{\pi^{-1}(1)}}(s_{\pi^{-1}(1)}^1), \dots, H^{i_{\pi^{-1}(n)}}(s_{\pi^{-1}(n)}^1))(H^{i_{\pi^{-1}(1)}}(s_{\pi^{-1}(1)}^2), \dots, H^{i_{\pi^{-1}(n)}}(s_{\pi^{-1}(n)}^2)) \dots \in \\
\Theta_j^l \\
&\Leftrightarrow (t_{\pi^{-1}(1)}^1, \dots, t_{\pi^{-1}(n)}^1)(t_{\pi^{-1}(1)}^2, \dots, t_{\pi^{-1}(n)}^2) \dots \in \Theta_j^l
\end{aligned}$$

which is the final step showing that \mathcal{H} is in fact a symmetric game structure.

We wish to show that there is a Nash equilibrium in \mathcal{G} from $(s_{1,0}, \dots, s_{n,0})$ if and only if there is a symmetric Nash equilibrium in \mathcal{H} from $(H^1(s_{1,0}), \dots, H^n(s_{n,0}))$, but first we will introduce a bit of notation. The way we define the equivalence relations $(\sim_i)_{1 \leq i \leq n}$ the information sets of a player in \mathcal{H} depends on which copy of G he is in as well as which information set in \mathcal{G} the current configuration corresponds to. We denote the information sets for every player i , every copy j of G and every information set I of player i in \mathcal{G} as follows

$$H_i^j(I) = \{(H^{m_1}(s_1), \dots, H^{m_n}(s_n)) \mid (s_1, \dots, s_n) \in I \wedge m_i = j\}$$

We now start with the first direction and assume there is a Nash equilibrium σ in \mathcal{G} from $(s_{1,0}, \dots, s_{n,0})$. Then we create the strategy profile σ' in \mathcal{H} such that for all players i, j , all π extending $\pi_{i,j}$ and all information sets I_1, \dots, I_k of player i in \mathcal{G}

$$\sigma'_i(H_i^j(I_1) \dots H_i^j(I_k)) = \sigma_j(\pi^{-1}(I_1) \dots \pi^{-1}(I_k))$$

which we will prove is a symmetric Nash equilibrium. Note again that in all legal sequences of information sets, a player will stay in the same copy of G and therefore this is a full definition of a strategy for each player. To prove that it is symmetric we need that for all i, j , all π extending $\pi_{i,j}$ and all information sets I of player i in \mathcal{G} it holds that

$$\begin{aligned}
\pi^{-1}(H_i^j(I)) &= \{(H^{m_{\pi^{-1}(1)}}(s_{\pi^{-1}(1)}), \dots, H^{m_{\pi^{-1}(n)}}(s_{\pi^{-1}(n)})) \mid \\
&\quad (H^{m_1}(s_1), \dots, H^{m_n}(s_n)) \in H_i^j(I)\} \\
&= \{(H^{m_{\pi^{-1}(1)}}(s_{\pi^{-1}(1)}), \dots, H^{m_{\pi^{-1}(n)}}(s_{\pi^{-1}(n)})) \mid
\end{aligned}$$

$$\begin{aligned}
& (s_1, \dots, s_n) \in I \wedge m_i = j \} \\
= & \{ (H^{m_{\pi^{-1}(1)}}(s_{\pi^{-1}(1)}), \dots, H^{m_{\pi^{-1}(n)}}(s_{\pi^{-1}(n)})) \mid \\
& (s_{\pi^{-1}(1)}, \dots, s_{\pi^{-1}(n)}) \in \pi^{-1}(I) \wedge m_{\pi^{-1}(j)} = j \} \\
= & H_j^j(\pi^{-1}(I))
\end{aligned}$$

The requirement for σ' to be symmetric can now be reformulated as follows. For all players i, j , all π extending $\pi_{i,j}$, all π' extending $\pi_{j,j}$ and all information sets I_1, \dots, I_m of player i in \mathcal{G} we have

$$\begin{aligned}
\sigma'_i(H_i^j(I_1) \dots H_i^j(I_m)) &= \sigma'_j(\pi^{-1}(H_i^j(I_1)) \dots \pi^{-1}(H_i^j(I_m))) \\
\Leftrightarrow \sigma_j(\pi^{-1}(I_1) \dots \pi^{-1}(I_m)) &= \sigma'_j(H_j^j(\pi^{-1}(I_1)) \dots H_j^j(\pi^{-1}(I_m))) \\
\Leftrightarrow \sigma_j(\pi^{-1}(I_1) \dots \pi^{-1}(I_m)) &= \sigma_j(\pi'^{-1}(\pi^{-1}(I_1)) \dots \pi'^{-1}(\pi^{-1}(I_m)))
\end{aligned}$$

Now since $\pi_{j,j}(q) = q$ for all $q \in \text{Nei}(j)$ and $(s_1, \dots, s_n) \in I \Leftrightarrow (s_{\pi''(1)}, \dots, s_{\pi''(n)}) \in I$ for all states s_1, \dots, s_n , all information sets I of j and all π'' such that $\pi''(q) = q$ for all $q \in \text{Nei}(j)$, then it holds for π' extending $\pi_{j,j}$ in particular. From this follows that $\pi'(I) = I \Leftrightarrow \pi'^{-1}(I) = I$ for all information sets I of j . Applying this to the equation above we get equality which proves that σ' is symmetric.

To see that σ' is also a Nash equilibrium from $(H^1(s_{1,0}), \dots, H^n(s_{n,0}))$ consider the deviation of a player p from σ'_p to $\sigma'_{p,\text{dev}}$. We then look at a corresponding deviation of p from σ_p to $\sigma_{p,\text{dev}}$ in \mathcal{G} where for all sequences of information set tuples

$$\sigma_{p,\text{dev}}(I_1, \dots, I_m) = \sigma'_{p,\text{dev}}(H_p^p(I_1), \dots, H_p^p(I_m))$$

We consider the outcomes of the two profiles in the two games, denoted $\rho_{\mathcal{G}}$ and $\rho_{\mathcal{H}}$ respectively. We wish to show that $\rho_{\mathcal{G}} = \pi_1(\rho_{\mathcal{H}})$ by induction. For the base case we have

$$\rho_{\mathcal{G},=0} = (s_{1,0}, \dots, s_{n,0}) = \pi_1(H^1(s_{1,0}), \dots, H^n(s_{n,0})) = \pi_1(\rho_{\mathcal{H},=0})$$

As induction hypothesis suppose it holds for prefixes of outcomes with length

at most v . Further, let $\rho_{\mathcal{G}, \leq v+1} = (s_1^0, \dots, s_n^0) \dots (s_1^{v+1}, \dots, s_n^{v+1})$ and let I_i^j be the information set for player i in \mathcal{G} which contains s_i^j . We will need that for a move m of all the players we have for all states s_1, \dots, s_n in \mathcal{G} that $\text{Tab}((s_1, \dots, s_n), m) = \pi_1(\text{Tab}((H^1(s_1), \dots, H^n(s_n)), m))$. Then we get

$$\begin{aligned}
\rho_{\mathcal{G}, \leq v+1} &= \rho_{\mathcal{G}, \leq v} \cdot \text{Tab}(\rho_{\mathcal{G}, =v}, \sigma[\sigma_p \mapsto \sigma_{p, \text{dev}}](I(\rho_{\mathcal{G}, \leq v}))) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \text{Tab}((s_1^v, \dots, s_n^v), \sigma[\sigma_p \mapsto \sigma_{p, \text{dev}}](I(\rho_{\mathcal{G}, \leq v}))) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \pi_1(\text{Tab}((H^1(s_1^v), \dots, H^n(s_n^v)), \sigma[\sigma_p \mapsto \sigma_{p, \text{dev}}](I(\rho_{\mathcal{G}, \leq v})))) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \pi_1(\text{Tab}(\rho_{\mathcal{H}, =v}, \sigma[\sigma_p \mapsto \sigma_{p, \text{dev}}](I(\rho_{\mathcal{G}, \leq v}))))
\end{aligned}$$

Since for all players i and all information sets I_1, \dots, I_m of i in \mathcal{G} it holds that $\sigma[\sigma_p \mapsto \sigma_{p, \text{dev}}]_i(I_1, \dots, I_m) = \sigma'[\sigma'_p \mapsto \sigma'_{p, \text{dev}}]_i(H_i^i(I_1), \dots, H_i^i(I_m))$ we get

$$\begin{aligned}
\rho_{\mathcal{G}, \leq v+1} &= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \pi_1(\text{Tab}(\rho_{\mathcal{H}, =v}, \sigma[\sigma_p \mapsto \sigma_{p, \text{dev}}](I(\rho_{\mathcal{G}, \leq v})))) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \pi_1(\text{Tab}(\rho_{\mathcal{H}, =v}, \sigma'[\sigma'_p \mapsto \sigma'_{p, \text{dev}}](\\
&\quad (H_1^1(I_1^0), \dots, H_n^n(I_n^0)) \dots (H_1^1(I_1^v), \dots, H_n^n(I_n^v)))))) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \pi_1(\text{Tab}(\rho_{\mathcal{H}, =v}, \sigma'[\sigma'_p \mapsto \sigma'_{p, \text{dev}}](\rho_{\mathcal{H}, \leq v}))) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v}) \cdot \pi_1(\rho_{\mathcal{H}, =v+1}) \\
&= \pi_1(\rho_{\mathcal{H}, \leq v+1})
\end{aligned}$$

This means that when a player p can deviate from σ' in \mathcal{H} to obtain outcome $\rho_{\mathcal{H}}$ then he can deviate from σ in \mathcal{G} to obtain an outcome $\rho_{\mathcal{G}}$ with $\pi_1(\rho_{\mathcal{H}}) = \rho_{\mathcal{G}}$. The way we have defined objectives in \mathcal{H} every player will get the same payoff from a play ρ in \mathcal{H} as in $\pi_1(\rho)$ in \mathcal{G} for every play ρ . And since σ is a Nash equilibrium in \mathcal{G} from $(s_{1,0}, \dots, s_{n,0})$ where no player can deviate to improve his payoff then no player can deviate to improve his payoff from σ' in \mathcal{H} from $(H^1(s_{1,0}), \dots, H^n(s_{n,0}))$ because otherwise that player would be able to deviate from σ in \mathcal{G} to improve his payoff. Thus, σ' is a symmetric Nash equilibrium from $(H^1(s_{1,0}), \dots, H^n(s_{n,0}))$.

For the other direction we assume there is a symmetric Nash equilibrium σ' in \mathcal{H} from $(H^1(s_{1,0}), \dots, H^n(s_{n,0}))$. We now define σ in \mathcal{G} so for all infor-

mation sets I_1, \dots, I_m of player i

$$\sigma_i(I_1, \dots, I_m) = \sigma'_i(H_i^i(I_1), \dots, H_i^i(I_m))$$

and wish to show that this is a Nash equilibrium in \mathcal{G} from $(s_{1,0}, \dots, s_{n,0})$. Contrary to the previous case we consider a deviation from σ in \mathcal{G} by player p from σ_p to $\sigma_{p,\text{dev}}$ and consider a corresponding deviation in \mathcal{H} from σ' defined by

$$\sigma'_{p,\text{dev}}(H_i^i(I_1), \dots, H_i^i(I_m)) = \sigma_{p,\text{dev}}(I_1, \dots, I_m)$$

for all information sets I_1, \dots, I_m of player p in \mathcal{G} . Let the outcomes of the profiles with the deviations be $\rho_{\mathcal{G}}$ and $\rho_{\mathcal{H}}$. As in the previous case we can show that $\pi_1(\rho_{\mathcal{H}}) = \rho_{\mathcal{G}}$ which means that when a player p deviates in \mathcal{G} from σ he can do a deviation in \mathcal{H} from σ' which gives him the same payoff. Since σ' is a Nash equilibrium from $(H^1(s_{1,0}), \dots, H^n(s_{n,0}))$ no player can deviate to improve his payoff from σ' . This means that no player can deviate to improve his payoff from σ in \mathcal{G} from $(s_{1,0}, \dots, s_{n,0})$ and therefore it is a Nash equilibrium from this configuration.

In total, there is a symmetric Nash equilibrium from $(H^1(s_{1,0}), \dots, H^n(s_{1,0}))$ in \mathcal{H} if and only if there is a Nash equilibrium from $(s_{1,0}, \dots, s_{n,0})$ in \mathcal{G} . In addition the size of \mathcal{H} is polynomial in the size of \mathcal{G} \square

A.10 Proof of Theorem 21

Theorem 21. *The existence problem is undecidable for symmetric game structures*

Proof. We do a reduction of the halting problem for a deterministic two-counter machine which is undecidable. A two-counter machine M is a 3-tuple

$$M = (Q, \Delta, q_F)$$

where

- Q is a finite set of control states
- $\Delta : Q \setminus \{q_F\} \rightarrow \{\text{inc}\} \times \{c, d\} \times Q \cup \{\text{dec}\} \times \{c, d\} \times Q^2$ is an instruction function which assigns an instruction to each state.
- $q_F \in Q$ is a halting state.

A configuration of M is a 3-tuple in $Q \times \mathbb{N} \times \mathbb{N}$. A run of M is a sequence ρ of configurations

$$\rho = (q_0, c_0, d_0)(q_1, c_1, d_1) \dots$$

where (q_0, c_0, d_0) is the initial configuration and for two consecutive configurations we have

$$\Delta(q_i) = (\text{inc}, c, q_{i+1}), c_{i+1} = c_i + 1 \text{ and } d_{i+1} = d_i \text{ or}$$

$$\Delta(q_i) = (\text{inc}, d, q_{i+1}), d_{i+1} = d_i + 1 \text{ and } c_{i+1} = c_i \text{ or}$$

$$\Delta(q_i) = (\text{dec}, c, q_{i+1}, q) \text{ for some } q, c_{i+1} = c_i - 1 \geq 0 \text{ and } d_{i+1} = d_i \text{ or}$$

$$\Delta(q_i) = (\text{dec}, d, q_{i+1}, q) \text{ for some } q, d_{i+1} = d_i - 1 \geq 0 \text{ and } c_{i+1} = c_i \text{ or}$$

$$\Delta(q_i) = (\text{dec}, c, q, q_{i+1}) \text{ for some } q, c_{i+1} = c_i = 0 \text{ and } d_{i+1} = d_i \text{ or}$$

$$\Delta(q_i) = (\text{dec}, d, q, q_{i+1}) \text{ for some } q, d_{i+1} = d_i = 0 \text{ and } c_{i+1} = c_i.$$

The run is infinite if there is no i so $q_i = q_F$ and otherwise it is finite with q_F being the halting state in the final configuration of ρ . The problem of deciding if the run of a two-counter machine has a halting run from a configuration (q_0, c_0, c_d) is undecidable and we wish to reduce an instance of this problem to the existence problem in symmetric game structures.

Let M be a deterministic two-counter machine and let (q_0, c_0, d_0) be an initial configuration. From this we create a symmetric game structure with 2 players $\mathcal{G} = (G, (\text{Nei}_i), (\pi_{i,j}), (\equiv_i), (\Omega_i^i))$ where $\text{Nei}_1 = (2), \text{Nei}_2 = (1), \pi_{1,2} = (2, 1), \pi_{2,1} = (2, 1)$ and $(s_1, s_2) \equiv_i (s'_1, s'_2)$ iff $s_1 = s'_1$ and $s_2 = s'_2$ for $i = 1, 2$.

The arena G consists of two disconnected parts. It is shown in Figure 12, but without G' .

The idea is that player 1 will start in $s_{1,0}$ and player 2 will start in $s_{2,0}$. They start by playing a matching penny game and afterwards player 2 will play in G' which simulates the counter machine M . Then we will design the objectives so player 2 wins if he acts according to the rules of the counter machine and reaches a halting state. If he does not reach a halting state he wins if the two players choose different coins initially and otherwise player 1 wins. In this way, if there is a legal, halting run of the counter machine then there is a Nash equilibrium where player 2 wins and player 1 loses. If there is no legal halting run then the game is essentially reduced to a matching

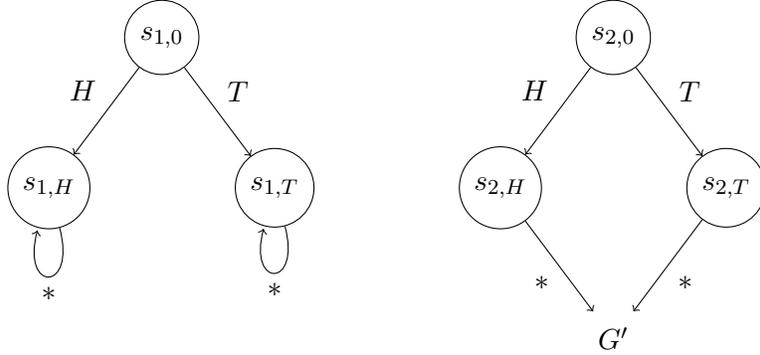


Figure 12: Illustration of G

penny game which has no Nash equilibrium.

Formally, the way we do this is let G' consist of the control states of M with the state connected to $s_{2,H}$ and $s_{2,T}$ being q_0 . Then for all i, j there will be an action $C+$ taking the play from q_i through an intermediate state $C+_{ij}$ to q_j if $\Delta(q_i) = (inc, C, q_j)$ for some counter $C \in \{c, d\}$ as illustrated in Figure 13.

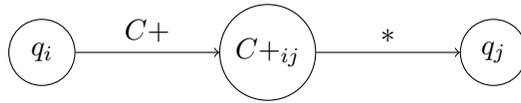


Figure 13: Construction of incrementation module

In addition for all i, j, k there will be an action $C-$ and an action $C0$ respectively taking the play from q_i through intermediate states $C-_{ij}$ to q_j and $C0_{ik}$ to q_k if $\Delta(q_i) = (dec, C, q_j, q_k)$ for some counter $C \in \{c, d\}$ as illustrated in Figure 14.

Additionally, we add a self-loop to the halting state q_F . For a finite path ρ we now define

$$C_\rho = |\{k \mid \rho_k = C +_{ij} \text{ for some } i, j\}| - |\{k \mid \rho_k = C -_{ij} \text{ for some } i, j\}|$$

When given an initial value c_0 and d_0 of the counters we then define the objectives such that player 2 loses in all plays ρ that contains a prefix $\rho_{\leq k}$ such that state $\rho_k = C -_{ij}$ and $C_0 - C_{\rho_{\leq k}} < 0$ for some C, i and j to make

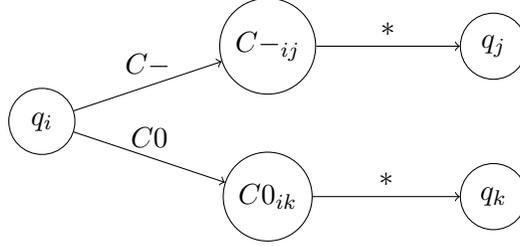


Figure 14: Construction of decrementation module

sure player 2 plays according to the rules of the counter machine and does not subtract from a counter with value zero. In addition, he loses in all plays ρ that contains a prefix $\rho_{\leq k}$ such that $\rho_k = C0_{ij}$ and $C_0 - C_{\rho_{\leq k}} \neq 0$ to make sure player 2 does not follow the true branch of a zero test when the value of the counter being tested is not 0. Finally, player 2 wins if he does not violate any of these restrictions and reaches q_F . He also wins if he wins the matching penny game (no matter if he violates the restrictions) and player 1 wins whenever player 2 doesn't win.

In total this means that there is a Nash equilibrium where player 2 wins if M halts with initial counter values c_0 and d_0 . If M doesn't halt with initial values c_0 and d_0 the game is reduced to a matching penny game which has no Nash equilibrium. Thus, there is a Nash equilibrium in \mathcal{G} if and only if M halts implying that the existence problem is undecidable. \square

A.11 Proof of Theorem 25

Theorem 25. *Suppose \mathcal{L} is a logic interpreted over infinite words. If model checking a formula $\varphi \in \mathcal{L}$ in a rooted, interpreted transition system (S, T, L, s_0) is decidable in time $f(|S|, |T|, |L|, |\varphi|)$ then the (symmetric) existence problem for m -bounded suffix strategies in a symmetric game structure $\mathcal{G} = (G, (\text{Ne}_i), (\pi_{i,j}), (\equiv_i), (\Omega_i^i))$ is decidable in time*

$$O(n \cdot |\text{States}|^{n^2 \cdot m \cdot |\text{Act}|} \cdot f((|\text{States}|^n + 1)^{n \cdot m}, |\text{Act}| \cdot (|\text{States}|^n + 1)^{n \cdot m}, |\text{Prop}|, |\varphi|))$$

when every player i has an objective given by a formula $\varphi_i \in \mathcal{L}$ where the proposition symbols occurring are Prop and $|\varphi| = \max |\varphi_i|$.

Proof. Let \mathcal{L} be a logic interpreted over infinite words where model checking is decidable in time $f(|S|, |T|, |L|, |\varphi|)$ for an interpreted transition system

(S, T, L, s_0) and \mathcal{G} be a symmetric game structure where every player i has an objective given by a formula $\varphi_i \in \mathcal{L}$. Since there are at most $|\text{States}|^n$ information sets for each player and at most $|\text{Act}|$ actions for each state there are at most

$$\sum_{i=1}^m |\text{States}|^{n \cdot i \cdot |\text{Act}|} = O(|\text{States}|^{n \cdot m \cdot |\text{Act}|})$$

different m -bounded suffix strategies. Thus, there are at most $O(|\text{States}|^{n^2 \cdot m \cdot |\text{Act}|})$ different m -bounded suffix strategy profiles. Our algorithm will check all these strategy profiles to see if any of them is a Nash equilibrium.

We define an m -bounded info matrix as an $n \times m$ matrix where each row contains a sequence of m info sets for each of the n players. We have a special null element $\mathbf{0}$ for sequences of information sets shorter than m . For a finite path ρ with $I(\rho) = (I_1^0, \dots, I_n^0) \dots (I_1^{|\rho|}, \dots, I_n^{|\rho|})$ we define $M_m(\rho)$ as the m -bounded info matrix

$$M_m(\rho) = \begin{bmatrix} I_1^{|\rho|-m+1} & \dots & I_1^{|\rho|} \\ \dots & & \dots \\ I_n^{|\rho|-m+1} & \dots & I_n^{|\rho|} \end{bmatrix}$$

where we define $I_i^j = \mathbf{0}$ for all i when $j < 0$. There are at most $(|\text{States}|^n + 1)^{n \cdot m}$ different m -bounded info matrices. The outcome of an m -bounded strategy profile will have the form $\pi \cdot \tau^\omega$ with $|\pi| + |\tau| \leq (|\text{States}|^n + 1)^{n \cdot m}$ since an m -bounded info matrix of a path uniquely identifies the last state configuration of the path and since the players will repeat the same play when reaching a repeated info matrix.

For a given m -bounded suffix strategy profile σ we can now check in time $O(n \cdot f((|\text{States}|^n + 1)^{n \cdot m}, (|\text{States}|^n + 1)^{n \cdot m}, |\text{Prop}|, |\varphi|))$ which players win and which players lose. For the players losing in σ we check if they can improve from σ in which case σ is not be a Nash equilibrium. For a given player i who loses in σ we do the check as follows. We create an interpreted transition system $\mathcal{T} = (S, T, L, s_0)$ where the set of states S is the set of different info matrices and the transition relation T is defined so $(M, M') \in T$ if and only if there is an action for player i taking the play from info matrix M to info matrix M' given that all other players play according to σ . The labeling of each info matrix is the labeling of the corresponding information set for player i , in other words $L(M) = L(M_{i,m})$. The initial state of the transition system is s_0 which is the info matrix corresponding to the initial state configuration of \mathcal{G} . This transition system represents exactly

the possibilities that player i has to deviate. Note that this is because σ is an m -bounded suffix strategy profile and therefore the players will always use the same action at a given info matrix. Since player i can improve if and only if there is a path in \mathcal{T} satisfying φ_i we can do this check in time $f((|\text{States}|^n + 1)^{n \cdot m}, |\text{Act}| \cdot (|\text{States}|^n + 1)^{n \cdot m}, |\text{Prop}|, |\varphi_i|)$. This check is done for every losing player to check if σ is a Nash equilibrium.

In total we have $O(|\text{States}|^{n^2 \cdot m \cdot |\text{Act}|})$ different strategy profiles to check, which can all be checked in time $n \cdot f((|\text{States}|^n + 1)^{n \cdot m}, |\text{Act}| \cdot (|\text{States}|^n + 1)^{n \cdot m}, |\text{Prop}|, |\varphi|)$ where $|\varphi| = \max |\varphi_i|$, which means that the problem is decidable in time

$$O(n \cdot |\text{States}|^{n^2 \cdot m \cdot |\text{Act}|} \cdot f((|\text{States}|^n + 1)^{n \cdot m}, |\text{Act}| \cdot (|\text{States}|^n + 1)^{n \cdot m}, |\text{Prop}|, |\varphi|))$$

Note that the algorithm solves the symmetric problem as well, since we can just try all strategy profiles of one player and extend them to symmetric strategy profiles, which is done in a unique way. \square

B Discussion of Symmetric Normal-Form Games

Here, we question the definition of symmetric normal-form games introduced in [8] which is as follows

Definition 31. *A symmetric normal-form game is a tuple $(\{1, \dots, n\}, S, (u_i)_{1 \leq i \leq n})$ where $\{1, \dots, n\}$ is the set of players, S is a finite set of strategies and $u_i : S^n \rightarrow \mathbb{R}$ are utility functions such that for all strategy vectors $(a_1, \dots, a_n) \in S^n$, all permutations π of $(1, \dots, n)$ and all i it holds that*

$$u_i(a_1, \dots, a_n) = u_{\pi(i)}(a_{\pi(1)}, \dots, a_{\pi(n)})$$

In [8] and other sources where symmetric games are defined [6, 7, 10, 13, 14, 17] the basic intuition of what a symmetric normal-form game should be is a game where all players have the same set of actions and the same utility functions in the sense that the utility of player should depend exactly on which action he chooses himself and on the number of other players choosing any particular action. However, the definition above does not seem to capture the scenario at hand and might even be erroneous. For two players we agree with the definition, and indeed in [13, 17] it is only defined for two players. In [6, 7, 10, 14] however it is defined such that $u_i(a_i, a_{-i}) = u_j(a_j, a_{-j})$ whenever $a_i = a_j$ and $a_{-i} = a_{-j}$. Here, a_i means the action taken by player i and a_{-i} means the set of actions taken by players other than player i . This definition seems to agree with the intuition about the games which we wish to represent. Another argument against Definition 31 is that it has the following consequence

Theorem 32. *Using Definition 31 for games with more than 2 players we have for all pairs of players i, j and all strategy profiles (a_1, \dots, a_n)*

$$u_i(a_1, \dots, a_n) = u_j(a_1, \dots, a_n)$$

In particular, for zero-sum games, we have $u_i(a_1, \dots, a_n) = 0$ for all players i and all strategy profiles (a_1, \dots, a_n) .

Proof. We start by looking at games with $n \geq 3$. First, using the permutation $(2, 3, 1, \dots)$ we get

$$u_1(a_1, a_2, a_3, \dots) = u_2(a_2, a_3, a_1, \dots)$$

Next, using $(1, 3, 2, \dots)$ we get

$$u_3(a_2, a_1, a_3, \dots) = u_2(a_2, a_3, a_1, \dots)$$

Using (2, 1, 3, ...) gives us

$$u_2(a_2, a_3, a_1, \dots) = u_1(a_3, a_2, a_1, \dots)$$

Putting all these together, we have $u_1(a_1, a_2, a_3, \dots) = u_1(a_3, a_2, a_1, \dots)$. This means that Player 1 can switch action with Player 3 without changing his utility. This can be done for arbitrary opponents by symmetry and since we can switch actions of all other players without changing the utility using permutations where $\pi(1) = 1$ we have that $u_1(a_1, \dots, a_n) = u_1(a_{\pi(1)}, \dots, a_{\pi(n)})$ for every permutation π . And by symmetry, that for every i and every permutation π

$$u_i(a_1, \dots, a_n) = u_i(a_{\pi(1)}, \dots, a_{\pi(n)})$$

Now consider any two players i and j as well as a permutation π such that $\pi(i) = j$. Then we have

$$u_i(a_1, \dots, a_n) = u_{\pi(i)}(a_{\pi(1)}, \dots, a_{\pi(n)}) = u_j(a_{\pi(1)}, \dots, a_{\pi(n)}) = u_j(a_1, \dots, a_n)$$

using the result above. This means that for any choice of actions of the players, every player will get the same utility. \square

One could define a class of games with this property, however for the definition at hand the property is not true for games with 2 players, which is a quite strange property of a class of games. We next present our definition, which is equivalent to the definitions from [6, 7, 10, 14] but defined in a form resembling Definition 31

Definition 33. *A symmetric normal-form game is a tuple $(\{1, \dots, n\}, S, (u_i)_{1 \leq i \leq n})$ where $\{1, \dots, n\}$ is the set of players, S is a finite set of strategies and $u_i : S^n \rightarrow \mathbb{R}$ are utility functions such that for all strategy vectors $(a_1, \dots, a_n) \in S^n$, all permutations π of $(1, \dots, n)$ and all i it holds that*

$$u_i(a_1, \dots, a_n) = u_{\pi^{-1}(i)}(a_{\pi(1)}, \dots, a_{\pi(n)})$$

The intuition behind this definition is as follows. Suppose we have a strategy profile $\sigma = (a_1, \dots, a_n)$ and a strategy profile where the actions of the players have been rearranged by permutation π , $\sigma_\pi = (a_{\pi(1)}, \dots, a_{\pi(n)})$. We would prefer that the player j using the same action in σ_π as player i

does in σ gets the same utility. Since j uses $a_{\pi(j)}$ this means that $\pi(j) = i \Rightarrow j = \pi^{-1}(i)$. Now, from this intuition we have that $u_i(a_1, \dots, a_n) = u_j(a_{\pi(1)}, \dots, a_{\pi(n)}) = u_{\pi^{-1}(i)}(a_{\pi(1)}, \dots, a_{\pi(n)})$. Apart from this intuition, the new definition can be shown to be equivalent to the one from [6, 7, 10, 14]. The reason that we agree with Definition 31 for 2 players is that $\pi = \pi^{-1}$ for all permutations π of 2 elements.