

Concurrent Games with Multiple Topologies

Shaul Almagor ✉ 

Department of Computer Science, Technion, Haifa, Israel

Shai Guendelman ✉

Department of Computer Science, Technion, Haifa Israel

Abstract

Concurrent multi-player games with ω -regular objectives are a standard model for systems that consist of several interacting components, each with its own objective. The standard solution concept for such games is Nash Equilibrium, which is a “stable” strategy profile for the players.

In many settings, the system is not fully observable by the interacting components, e.g., due to internal variables. Then, the interaction is modelled by a partial information game. Unfortunately, the problem of whether a partial information game has an NE is not known to be decidable. A particular setting of partial information arises naturally when processes are assigned IDs by the system, but these IDs are not known to the processes. Then, the processes have full information about the state of the system, but are uncertain of the effect of their actions on the transitions.

We generalize the setting above and introduce Multi-Topology Games (MTGs) – concurrent games with several possible topologies, where the players do not know which topology is actually used. We show that extending the concept of NE to these games can take several forms. To this end, we propose two notions of NE: Conservative NE, in which a player deviates if she can strictly add topologies to her winning set, and Greedy NE, where she deviates if she can win in a previously-losing topology. We study the properties of these NE, and show that the problem of whether a game admits them is decidable.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory; Theory of computation → Automata over infinite objects

Keywords and phrases Concurrent games, Nash Equilibrium, Symmetry, Partial information

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2022.34

Related Version *Arxiv Version*: <https://arxiv.org/abs/2207.02596>

1 Introduction

Concurrent multi-player games of infinite duration over graphs are a standard modelling tool for representing systems that consist of several interacting components, each having its own objective. Each player in the game corresponds to a component in the interaction. In each round of the game each of the player chooses an action and the next state of the game is determined by the current state and the vector of actions chosen. A strategy for a player is then a mapping from the history of the game so far to the next action.

A strategy profile (i.e., a tuple of strategies, one for each player) induces an infinite trace of states, and the goal of each player is to direct the game into a trace that satisfies her specification. This is modeled by augmenting the game with ω -regular objectives describing the objectives of the players.

Unlike traditional zero-sum games, here the objectives of the players do not necessarily contradict each other. Accordingly, the typical questions about these games concern their stability. Specifically, the most well-known stability measure is Nash Equilibrium (NE): an NE is a strategy profile such that no single player can improve her outcome by unilaterally deviating from the profile. The problem of whether a multi-player game with ω -regular objectives has an NE was shown to be decidable in [6].

In many settings, the players only have partial information about the system, or can view only certain parts of it. This happens when e.g., the system has private and global variables, and the players model threads that can only view the global variables. To this end,



© Shaul Almagor and Shai Guendelman;
licensed under Creative Commons License CC-BY 4.0

33rd International Conference on Concurrency Theory (CONCUR 2022).

Editors: Bartek Klin, Slawomir Lasota, and Anca Muscholl; Article No. 34; pp. 34:1–34:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

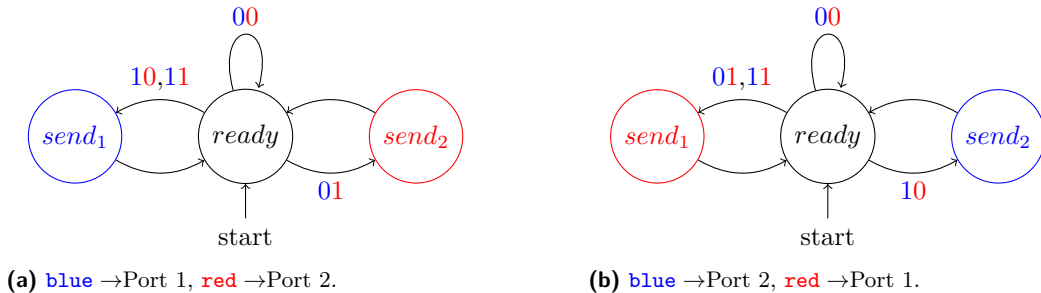
games with *partial information* have been extensively studied in various forms [3, 5, 9, 10]. However, in contrast to the full-information setting, the problem of deciding whether a partial-information multi-player game of infinite duration has a Nash equilibrium is not known to be decidable, and is known to be undecidable in the case of stochastic games [25].

In this work, we introduce and study *Multi-Topology Games (MTG)*. Intuitively, an MTG is a concurrent multi-player game with several transition functions (i.e., topologies). Then, players are fully aware of the possible topologies of the game, but do not know which topology they currently play on. Thus, MTGs capture a restricted form of partial information.

As we now demonstrate, MTGs naturally model the sort of partial information that arises in the context of *process symmetry*.

► **Example 1.** Consider a virtual router with multiple ports. When the router is initialized, several processes are plugged in. The router assigns each process to a port id, but the id is not revealed to the processes. Each process attempts to send messages, and its goal is to have its messages delivered (where some messages may be dropped due to heavy traffic). While the processes know exactly how the router works, they do not know which port they are assigned to. Therefore, their strategies must be oblivious to their port number.

As a concrete example, consider the concurrent game in Figure 1 with players {blue, red}. When both players know the port assignment, for example, blue \rightarrow Port 1 and red \rightarrow Port 2, then blue can win by always taking action 1, and red will lose in any strategy. However, if the port assignment is not known then in order for either player to win under both port assignments, the players must coordinate e.g., by taking turns trying to send a message. Thus, a-priori, the game has two possible topologies: Figure 1a and Figure 1b.



■ **Figure 1** Router game from Example 1. The players are blue and red, and the router has two ports 1, 2. In every round each player can try to *send* (action 1), or *wait* (action 0). The labels on the edges describe the actions of the players. The first is the action of the blue player, and the second is the action of the red player. From *ready*, if only the player in Port $i \in \{1, 2\}$ tries to send, the game transitions to $send_i$. If both players try to send, the router prioritizes the request from Port 1. The objective of the player Port i is to visit $send_i$ infinitely many times. Note that $send_i$ is colored according to the player that tries to reach it in each port assignment.

These type of settings are commonly referred to as *process symmetry* [12, 15, 18, 19, 1], and have been studied in several contexts (e.g., model checking with symmetry reductions). However, to our knowledge this setting has not been studied in games. In Section 3.1 we demonstrate how MTGs can model the general setting of process symmetry in games. ◻

In an MTG, a strategy for a player maps sequences of states to an action, and hence does not depend on a certain topology. Unlike standard games, a strategy profile in an MTG no longer induces a single trace, but rather a set of traces, one per topology. Thus, a player can no longer be said to be “winning” or “losing” in a strategy profile, as this may vary between topologies. In particular, it is not clear how analogues of Nash equilibrium and social optimum should be defined.

To this end, we propose two versions of Nash equilibria, corresponding to two extremities: in a Conservative NE (CNE), a player deviates if she can increase (w.r.t. containment) the set of topologies she wins in. In a Greedy NE (GNE), a player deviates if she can win in a currently-losing topology (even at the cost of losing some of the currently-winning topologies).

We study the properties of CNE and GNE and compare their strictness, showing that a GNE is also a CNE, but the converse does not hold. We also compare their properties to those of the standard notion of NE. Our main technical contribution is showing that the problem of whether a game has a CNE (resp. GNE) is decidable.

Related Work. A central work concerning NE in concurrent games is [6], where the problem of deciding whether a concurrent game admits an NE was studied for various winning conditions. Apart from establishing tight complexity bounds, this work also introduced the *suspect game* – a useful technique for reasoning about concurrent games. Interestingly, the suspect game does not seem to be adaptable to reason about MTGs, suggesting a fundamental difference between the models.

Zero-sum concurrent reachability games were studied in [13], where fundamental techniques for reasoning about them were developed. We remark that the zero-sum setting is technically very different to ours, due to the non-adversarial nature of the players.

Concurrent games can be formulated in the turn-based setting using partial information. The latter were extensively studied, e.g., in [9, 22, 10, 3, 8, 14], typically in the zero-sum setting.

Finally, the work in [3] extends strategy logic [11] with imperfect information. The authors show that, in general, the model checking problem for this logic is undecidable, but it is decidable in some special cases. Unfortunately, these cases do not readily capture MTGs.

Paper organization. In Section 2 we present the basic definitions of concurrent games. In Section 3 we formally define MTGs, introduce two notions of equilibria for them, and study their properties. In Section 4 we give our main technical result, establishing the decidability of detecting CNE in MTGs. In Section 5 we establish the decidability of detecting GNE. Finally, in Section 6 we discuss our results and some extensions, and detail future directions.

2 Preliminaries

A *concurrent parity game* is a tuple $\mathcal{G} = \langle \text{Pla}, \text{S}, s_0, \text{Act}, \delta, (\alpha_p)_{p \in \text{Pla}} \rangle$ where the components are as follows. Pla is a finite set of players, S is a finite set of states, $s_0 \in \text{S}$ is an initial state, Act is a finite set of actions. The transition function $\delta : \text{S} \times \text{Act}^{\text{Pla}} \rightarrow \text{S}$ maps a state and an *action profile* (i.e., $\mathbf{a} = (a_p)_{p \in \text{Pla}} \in \text{Act}^{\text{Pla}}$) to the next state. Every player $p \in \text{Pla}$ has a parity objective $\alpha_p \subseteq \text{S}^\omega$, as we describe below.

A *play* of \mathcal{G} is an infinite sequence of states $\rho = s_0, s_1, \dots \in \text{S}^\omega$ such that for every step $i \in \mathbb{N}$ there exists an action profile \mathbf{a} such that $s_{i+1} = \delta(s_i, \mathbf{a})$. For $k \geq 1$ we denote the length- k prefix of $\rho_{\leq k} = s_0, \dots, s_{k-1} \in \text{S}^+$. We denote by $\text{Inf}(\rho)$ the set of states that occur infinitely often in ρ . A *parity objective* is given by a function $\Omega : \text{S} \rightarrow \{0, \dots, d\}$ for some $d \in \mathbb{N}$. Then, ρ satisfies the objective if $\min\{\Omega(s) \mid s \in \text{Inf}(\rho)\}$ is even. Thus, the objective α_p is the set of all plays that satisfy the parity function of Player p . In the following, we mostly use the parity function implicitly, and so we do not include Ω in the description of \mathcal{G} .

The *description size* of \mathcal{G} , denoted $|\mathcal{G}|$ is the number of bits required to represent the components of \mathcal{G} .

► **Remark 2 (Game representation).** Note that we assume an explicit representation of the transition function as a table. In particular, we describe for every state the transition on every action profile in Act^{Pla} . Thus, the size of the transition functions is exponential in $|\text{Pla}|$.

This is in contrast with a more succinct representation, i.e., representing the transition function as a circuit. We choose this focus to eliminate the complexity effect of succinct representation.

A *history* of \mathcal{G} is a finite prefix of a play $h \in \mathcal{S}^+$. A *strategy* for Player p is a function $\sigma : \mathcal{S}^+ \rightarrow \text{Act}$ that maps a history to the next action of Player p . A *strategy profile* $\sigma = (\sigma_p)_{p \in \text{Pla}}$ is vector of strategies, one for each player. We denote the set of all strategies by $\Sigma_{\mathcal{G}}$ and the set of all strategy profiles by $\Sigma_{\mathcal{G}}^{\text{Pla}}$ (we omit the subscript \mathcal{G} when it is clear from context). A strategy profile σ can be thought as a function that maps histories to action profiles: given a history $h \in \mathcal{S}^+$ we have $\sigma(h) = (\sigma_p(h))_{p \in \text{Pla}} \in \text{Act}^{\text{Pla}}$.

For a strategy profile σ we define its *outcome* to be the infinite sequence of states (i.e. play) in \mathcal{G} that is taken when all the players follow their strategies in σ . Formally, $\text{out}_{\mathcal{G}}(\sigma) = s_0 s_1 \dots \in \mathcal{S}^{\omega}$ where s_0 is the initial state, and for every $i \geq 1$ we have $s_i = \delta(s_{i-1}, \sigma(s_0, \dots, s_{i-1}))$. Consider a play $\rho \in \mathcal{S}^{\omega}$. The set of *winners* in ρ is the set of players whose objectives are met in ρ . Formally, $\text{Win}_{\mathcal{G}}(\rho) = \{p \in \text{Pla} \mid \rho \in \alpha_p\} \subseteq \text{Pla}$. The set of winners in a strategy profile σ is then $\text{Win}_{\mathcal{G}}(\sigma) = \text{Win}_{\mathcal{G}}(\text{out}_{\mathcal{G}}(\sigma))$. Player p is said to be *losing* if she is not winning.

► **Remark 3 (Action visibility).** Note that strategies are defined to “see” only the history of visited states, and not the history of actions taken by the other players. This is a standard and natural assumption [6, 10] for concurrent models. There are, however, works (e.g., [2]) where players can view the entire action history. The latter approach is slightly easier to reason about, as players have full information on the game progress.

A strategy profile σ is a *Nash Equilibrium (NE)* if, intuitively, no single player can benefit from unilaterally changing her strategy. Since the objectives in our setting are binary, “benefiting” amounts to moving from the set of losers to the set of winners. We refer to such a change as a *beneficial deviation*. Formally, consider a strategy profile σ , a player $p \in \text{Pla}$ and a strategy $\sigma'_p \in \Sigma_{\mathcal{G}}$ for Player p . We denote by $\sigma[p \mapsto \sigma'_p] \in \Sigma_{\mathcal{G}}^{\text{Pla}}$ the strategy profile obtained from σ by replacing σ_p with σ'_p . Then, σ is an NE if for every player $p \in \text{Pla}$ and every strategy $\sigma'_p \in \Sigma_{\mathcal{G}}$ for Player p , if $p \in \text{Win}_{\mathcal{G}}(\sigma[p \mapsto \sigma'_p])$ then $p \in \text{Win}_{\mathcal{G}}(\sigma)$. Viewed contrapositively: if p loses when \mathcal{G} is played with σ , then p also loses after changing her strategy.

3 Multi-Topology Games

A *multi-topology game (MTG)* is a tuple $\mathcal{G} = \langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \text{Top}, (\delta_t)_{t \in \text{Top}}, (\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}} \rangle$ where $\text{Pla}, \mathcal{S}, s_0, \text{Act}$, are the same as in concurrent games. Top is a finite set of *topologies*, and for every $t \in \text{Top}$ we have a transition function $\delta_t : \mathcal{S} \times \text{Act}^{\text{Pla}} \rightarrow \mathcal{S}$ and objective $\alpha_{t,p} \subseteq \mathcal{S}^{\omega}$ for every player $p \in \text{Pla}$. An MTG can be thought of as a tuple of games over the same states, players and actions. That is, for $t \in \text{Top}$, we can define $\mathcal{G}_t = \langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \delta_t, (\alpha_{t,p})_{p \in \text{Pla}} \rangle$ to be the concurrent parity game obtained by fixing the transition function to δ_t and the objective for Player p to $\alpha_{t,p}$.

Crucially, the players are assumed to have no a-priori information on which topology is selected when the game is played. This is captured in the definition of strategies: a strategy for Player p is identical to the setting of concurrent parity games, i.e., $\sigma_p : \mathcal{S}^+ \rightarrow \text{Act}$. This lifts to strategy profiles and outcomes, as per Section 2. In particular, a strategy σ in \mathcal{G}

can be applied to \mathcal{G}_t for every $t \in \text{Top}$. Consider a strategy profile $\sigma \in \Sigma^{\text{Pla}}$. The *winning topologies* of Player p is the set of topologies that Player p wins in when \mathcal{G} is played with strategy profile σ . Formally, $\text{WinTop}_{\mathcal{G}}^p(\sigma) = \{t \in \text{Top} \mid p \in \text{Win}_{\mathcal{G}_t}(\sigma)\}$.

3.1 Process Symmetry in Concurrent Games

As we discuss in Section 1, a central motivation for MTGs come from settings where players plug in to the system without knowing their identity. This setting is commonly referred to as *process symmetry* [12, 15, 18, 19, 1]. Symmetry in games was studied in [24, 23, 7, 17] for *strategic form games*, which are games with a single turn. In [5, 26], symmetry in concurrent games was studied by imposing restrictions on the game structure. We consider a different setting, where processes $1, \dots, k$ log into a system described as a concurrent game, but the index of the action controlled by each process is not revealed to the processes. This setting is naturally modelled as an MTG, as follows.

Consider a concurrent game $\mathcal{G} = \langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \delta, (\alpha_p)_{p \in \text{Pla}} \rangle$ with $k \geq 2$ players, and that $\text{Pla} = \{1, \dots, k\}$. We obtain from \mathcal{G} an MTG with $k!$ topologies by letting each topology correspond to a different permutation of the players. Formally, consider a permutation $\pi \in \mathcal{S}_k$, where \mathcal{S}_k is the set of permutations over $\{1, \dots, k\}$. For an action profile $\mathbf{a} \in \text{Act}^{\text{Pla}}$ we define $\pi(\mathbf{a}) = (a_{\pi^{-1}(1)}, \dots, a_{\pi^{-1}(k)})$. That is, the action performed by Player i is taken at index $\pi(i)$. We now obtain the MTG $\mathcal{G}' = \langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \mathcal{S}_k, (\delta_\pi)_{\pi \in \mathcal{S}_k}, (\alpha_{\pi,p})_{\pi \in \mathcal{S}_k, p \in \text{Pla}} \rangle$ where \mathcal{S}_k is the set of topologies, δ_π is obtained by applying π to the action profile of the players, that is, for $s \in \mathcal{S}$ and $\mathbf{a} \in \text{Act}^{\text{Pla}}$ we have $\delta_\pi(s, \mathbf{a}) = \delta(s, \pi(\mathbf{a}))$. Finally, the objective of Player p is $\alpha_{\pi,p} = \alpha_{\pi(p)}$. Figure 1 is an example of such game.

3.2 Solution Concepts

Recall that in NE, a beneficial deviation moves a player from losing to winning. In MTGs, however, winning is no longer binary. Indeed, a strategy profile associates with each player a set of winning topologies. Thus, the meaning of “beneficial deviation” becomes context dependent. We introduce and study two notions of equilibria for MTGs that lie on two “extremities”: in the *conservative* approach, a deviation is beneficial if it strictly increases (w.r.t. containment) the set of winning topologies. In the *greedy* approach, a deviation is beneficial if a previously-losing topology becomes winning. We now turn to formally define and demonstrate these notions.

Conservative NE. A *conservative NE (CNE)* is a strategy profile σ where no player can deviate from σ and have her winning topologies be a strict superset¹ of her winning topologies when obeying σ . Formally, $\sigma \in \Sigma^{\text{Pla}}$ is a CNE if the following holds:

$$\forall p \in \text{Pla} \forall \sigma'_p \in \Sigma_{\mathcal{G}}^p \left((\forall t \in \text{Top} p \in \text{Win}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p]) \rightarrow \text{Win}_{\mathcal{G}_t}(\sigma)) \vee \right. \\ \left. (\exists t \in \text{Top} p \notin \text{Win}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p]) \wedge p \in \text{Win}_{\mathcal{G}_t}(\sigma)) \right)$$

Equivalently, this condition can be written in terms of the set of winning topologies:

$$\forall p \in \text{Pla} \forall \sigma'_p \in \Sigma_{\mathcal{G}}^p \neg(\text{WinTop}_{\mathcal{G}}^p(\sigma) \subsetneq \text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p]))$$

We refer to this notion as *conservative* since a deviating player wants to conserve her existing winning strategies.

¹ we emphasize that the relation \subsetneq means “strictly contained”.

Greedy NE. A *greedy NE (GNE)* is a strategy profile σ where no player can unilaterally deviate and win in a previously-losing topology. Formally, $\sigma \in \Sigma^{\text{Pla}}$ is a GNE if the following holds:

$$\forall p \in \text{Pla} \forall \sigma'_p \in \Sigma_G^p \forall t \in \text{Top} (p \in \text{Win}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p]) \rightarrow p \in \text{Win}_{\mathcal{G}_t}(\sigma))$$

Equivalently, this condition can also be written in terms of the set of winning topologies:

$$\forall p \in \text{Pla} \forall \sigma'_p \in \Sigma_G^p (\text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p]) \subseteq \text{WinTop}_{\mathcal{G}}^p(\sigma))$$

The latter formulation shows that in a GNE, for every player and for every deviation, the player's winning topologies when deviating are a subset of the player's winning topologies when obeying σ . It refers to this notion as *greedy* since it assumes that a player deviates if she improves her outcome in a single topology, disregarding the outcome in other topologies.

► **Example 4 (CNE and GNE).** Recall the router game from Figure 1. The strategy profile where Player **blue** repeatedly plays $(0, 0, 1, 1)^\omega$ and **red** plays $(1, 1, 0, 0)^\omega$ is a CNE, since the set of winning topologies of this profile is $\{1, 2\}$ for both players. Thus, no deviation can win in strictly more topologies.

Note that the same strategy profile is also a GNE, since every set of winning topologies is a subset of $\{1, 2\}$.

► **Remark 5 (Additional notions of NE).** CNE and GNE are based on the \subseteq preorder on the sets of topologies, 2^{Top} . In Section 6 we discuss other notions of NE in MTGs.

3.3 Properties of CNE and GNE

We start by examining some properties and relationships between the notions of CNE and GNE, as well as their relation to standard NE.

Consider an MTG $\langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \text{Top}, (\delta_t)_{t \in \text{Top}}, (\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}} \rangle$. The following observation is immediate from the definitions of GNE and CNE, since if there is only a single topology, the MTG collapses into a concurrent game.

► **Observation 6.** *If $\text{Top} = \{t\}$, i.e. there is only a single topology t , then the definitions of NE in \mathcal{G}_t coincides with that of CNE and of GNE in \mathcal{G} .*


Next, we observe that GNE is a stricter notion than CNE. Indeed, a beneficial deviation in the conservative setting (namely increasing the set of winning topologies) implies a beneficial deviation in the greedy setting (namely winning in a previously-losing topology). Contrapositively, if there is no greedy beneficial deviation, there is also no conservative beneficial deviation. We thus have the following.

► **Observation 7.** *Let \mathcal{G} be an MTG. If σ is a GNE in \mathcal{G} then σ is a CNE in \mathcal{G} .*

The following example shows that the implication of Observation 7 is strict. That is, there are MTGs with a CNE but without a GNE.

► **Example 8 (CNE without GNE).** Consider the single-player game depicted in Figure 2. The outcome of the game depends only on the first action that the player takes and the topology that the game is played in. If the player takes action 1, then the set of winning topologies is $\{t_1\}$. If the player takes action 2, then the set of winning topologies is $\{t_2\}$. Since $\{t_1\} \not\subseteq \{t_2\}$ and $\{t_2\} \not\subseteq \{t_1\}$, there is no GNE in the game, as the player can switch strategies from t_1 to t_2 and vice versa to win in a previously-losing topology.

However, since there is no strategy for the player such that the set of winning topologies is $\{t_1, t_2\}$ (the only strict superset of $\{t_1\}$ and $\{t_2\}$), then every strategy is a CNE.



■ **Figure 2** A single player MTG with two topologies, t_1 and t_2 . In both topologies, the objective of the player is to reach s_1 (it is easy to capture this using a parity objective).

► **Remark 9** (Best-response dynamics in GNE). Example 8 demonstrates that, in stark contrast to NE, an MTG might not have a GNE even when there is only a single player. This has to do, in particular, with the notion of *best-response dynamics*: in standard games, one can approach an NE by starting from some profile, and repeatedly letting players deviate to their best-response strategy, until this process converges. While this does not always converge, it does so for a large class of games (e.g., *finite-potential games* [21]).

Thus, Example 8 shows that best-response does not converge even for a single player in MTGs, whereas it does converge for a single player both for standard NE, as well as in CNE for MTGs. Indeed, the best-response of a single player in the conservative setting will increase her set of winning topologies to the maximum, and from there she will no longer have incentive to deviate.

Remark 9 reflects the intuition that a GNE must be stable in each topology separately. That is, it captures the notion “NE on all topologies”, in the following sense.

► **Observation 10.** A GNE σ is also an NE in \mathcal{G}_t for every $t \in \text{Top}$.

Indeed, if σ was not an NE in \mathcal{G}_t for some $t \in \text{Top}$, then a player that deviates from σ in \mathcal{G}_t would similarly deviate from σ in \mathcal{G} , greedily winning in the previously-losing topology t .

In contrast, we now show that CNE is a more intricate notion, and might hold even when there is no NE in the separate topologies.



■ **Figure 3** Symmetric XOR game. The players are **blue** and **red**. In topology t_1 , the objective of **blue** is to reach s_1 , and the objective of **red** is to reach s_2 . In topology t_2 the objectives of the players are swapped. The game starts from s_0 . If both players take the same action, then the game transitions to state s_1 and gets stuck there. If the players take different actions then the game transitions to s_2 and gets stuck there.

► **Example 11** (CNE without NE). Consider the Symmetric XOR game \mathcal{G} depicted in Figure 3. Note that neither \mathcal{G}_{t_1} nor \mathcal{G}_{t_2} have a NE, since if a strategy for a single player is fixed, the other player can respond to it and win.

On the other hand, any strategy profile is a CNE, since every player always wins in exactly one topology. Thus, there is no way for a player to deviate and get strict superset of winning topologies.

There are MTGs without CNE. For example, every concurrent game \mathcal{G} without an NE can be viewed as an MTG with a single topology t_1 . Since there is no NE in \mathcal{G} , then for every profile σ there exists a player p that loses with σ , which corresponds to $\text{WinTop}_{\mathcal{G}}^p(\sigma) = \emptyset$ but p can deviate and win \mathcal{G} , which corresponds to $\text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p]) = \{t_1\}$. Since $\emptyset \not\subseteq \{t_1\}$, then σ is not a CNE.

4 Existence of Conservative NE is Decidable

We now turn to our main technical contribution – showing that the existence of a CNE is a decidable property.

► **Theorem 12.** *The problem of deciding, given an MTG \mathcal{G} , whether there exists a CNE in \mathcal{G} is in 2-EXPTIME.*

The remainder of the section is devoted to proving Theorem 12. Our solution is based on a reduction to the problem of solving a restricted form of partial-information game. We then employ a result from [10], and obtain the complexity result by a careful analysis of the construction. The rest of the section is organized as follows. In Section 4.1 we present the model of partial-information games and the result of [10]. In Section 4.2 we give an overview of the reduction and in Section 4.3 we describe and analyze the reduction from our setting.

4.1 Partial-Information Games

Partial-information games (also known as *games with incomplete information*) are a ubiquitous model for settings where the players cannot fully observe the state of the game due to e.g., private/hidden variables, unknown parameters or abstractions of part of the system.

Formally, a *partial-information game* is a tuple $\mathcal{G} = \langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \delta, (\mathcal{O}_p)_{p \in \text{Pla}} \rangle$ where Pla , \mathcal{S} , s_0 , Act and δ are the same as in concurrent games. For every player $p \in \text{Pla}$, the set of *observations* $\mathcal{O}_p \subseteq 2^{\mathcal{S}}$ is a partition of \mathcal{S} . We omit the acceptance condition, and we will include it explicitly in Theorem 13 below.

Intuitively, when the play of \mathcal{G} is at state $s \in \mathcal{S}$, Player p can only observe $o \in \mathcal{O}_p$ such that $s \in o$, and needs to select an action according to o . Thus, we distinguish between *state histories*, \mathcal{S}^+ and *observation histories (of Player p)*, $(\mathcal{O}_p)^+$. For $s \in \mathcal{S}$ we define $\text{obs}_p(s) = o \in \mathcal{O}_p$ to be the unique observation of Player p such that $s \in o$. We extend obs_p to histories: let $h = s_0 s_1 \dots s_k \in \mathcal{S}^+$ be a state history, we define $\text{obs}_p(h) = \text{obs}_p(s_0) \text{obs}_p(s_1), \dots, \text{obs}_p(s_k) \in (\mathcal{O}_p)^+$ to be the corresponding observation history.

Strategies are *observation based*, that is, a *strategy* for Player p is a function $\sigma_p : \mathcal{O}_p^+ \rightarrow \text{Act}$. Since different players may have different observation sets, we denote by $\Sigma_{\mathcal{G}}^p$ the set of all strategies for Player p . We denote by $\Sigma_{\mathcal{G}}^{\text{Pla}}$ the set of all strategy profiles.

Similarly to concurrent games, a strategy profile σ can be thought of as a function that maps histories to action profiles $\sigma(h) = (\sigma_p(\text{obs}_p(h)))_{p \in \text{Pla}} \in \text{Act}^{\text{Pla}}$, and we define $\text{out}_{\mathcal{G}}(\sigma) \in \mathcal{S}^{\omega}$ similarly to concurrent games.

We say that Player $p \in \text{Pla}$ has *perfect information* if $\mathcal{O}_p = \{\{s\} \mid s \in \mathcal{S}\}$. That is, Player p can observe the exact state of the game. If all players have perfect information then the game is a *perfect information game*, and coincides with our definition of concurrent games. We say that *Player i is less informed than Player j* if \mathcal{O}_j is a refinement of \mathcal{O}_i . That is, for every $o_j \in \mathcal{O}_j$ there exists $o_i \in \mathcal{O}_i$ such that $o_j \subseteq o_i$.

Finally, consider an objective $\alpha \subseteq \mathcal{S}^{\omega}$, we say that α is *visible to Player p* if for every $\rho, \rho' \in \mathcal{S}^{\omega}$ such that $\text{obs}_p(\rho) = \text{obs}_p(\rho')$ we have that $\rho \in \alpha$ if and only if $\rho' \in \alpha$. That is, the objective can be defined according to observation sequences rather than plays.

The following theorem is a result from [10] that will serve as the target of our reduction.

► **Theorem 13.** *Let $\mathcal{G} = \langle \text{Pla}, S, s_0, \text{Act}, \delta, (\mathcal{O}_p)_{p \in \text{Pla}} \rangle$ be a partial information game, with $\text{Pla} = \{1, 2, 3\}$ where Player 1 less informed than Player 2. Let $\alpha \subseteq S^\omega$ be parity objective over S . The problem of deciding whether $\exists \sigma_1 \in \Sigma_{\mathcal{G}}^1 \forall \sigma_2 \in \Sigma_{\mathcal{G}}^2 \exists \sigma_3 \in \Sigma_{\mathcal{G}}^3 \text{out}_{\mathcal{G}}(\sigma_1, \sigma_2, \sigma_3) \in \alpha$ is 2-EXPTIME complete.*

4.2 Overview of the Reduction

We now turn to describe a reduction from the CNE existence problem to the setting of Theorem 13. We start with a high-level description. Consider an MTG \mathcal{G} . Instead of asking directly whether \mathcal{G} admits a CNE, we first fix a set of “intended” winning topologies $T_p \subseteq \text{Top}$ for each player $p \in \text{Pla}$. Then, we ask whether \mathcal{G} admits a CNE σ in which $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ for every $p \in \text{Pla}$. If we are able to answer the latter problem, we can iterate over every possible tuple $(T_p)_{p \in \text{Pla}}$ (or nondeterministically guess a set) and conclude whether \mathcal{G} admits a CNE. We remark that this approach is reminiscent of the technique in [6], where the existence of an NE in a game is decided by first guessing a “witness” path.

Once the set of intended topologies is fixed, we construct a 3-player partial-information game whose players are **Eve**, **Adam** and **Snake**, with the following roles:

- **Eve** controls the coalition of all players, and suggests a strategy profile σ by selecting the actions for all the players at each step.
- **Adam** selects a deviating player p , and the deviating strategy σ'_p for that player. In addition, **Adam** selects a set $T \subseteq \text{Top}$ in which Player p tries to win when playing σ'_p .
- **Snake** helps² **Eve** by selecting a concrete topology t from the set T picked by **Adam**.

The game starts with **Adam** and **Snake** choosing p , T and $t \in T$. It then proceeds with **Eve** and **Adam** choosing σ and σ'_p , respectively, while playing on \mathcal{G}_t . The observation sets of the players are such that both **Eve** and **Adam** can only observe the current state of the game, so **Eve** is ignorant of p , T and t , and **Adam** is ignorant of t (except knowing that $t \in T$).

The objective of **Eve** and **Snake** is then composed of three conditions:

1. **Snake** must choose a topology $t \in T$.
2. If the strategy σ'_p proposed by **Adam** does not in fact deviate from the profile σ proposed by **Eve** (dubbed “**Adam obeys Eve**”), and if $t \in T_p$, i.e., p was intended to win in t , then the outcome must be winning for Player p .
3. If **Adam** selected T to contain a topology not in T_p (i.e., Player p potentially tries to win in a superset of T_p), then the outcome must be losing for Player p .

The overall idea is that if **Eve** can find a strategy for all the players, from which any deviation choice of **Adam** can be shown to be non-beneficial by an appropriate choice by **Snake**, then there is a CNE with the intended winning topologies, and vice-versa.

There are, however, some caveats: first, in order to allow **Adam** to choose any set of topologies, the size of the game would be exponential, which is undesirable. Second, it is not immediate that the conjunction of conditions above can be captured by a small parity objective (since the parity condition does not allow conjunction without a change of state space [4]). Third, we need to separate the cases where **Adam** obeys **Eve**. In the following we give the complete construction, which overcomes these caveats.

² It is arguable whether this matches the biblical interpretation. This work makes no theological claims.

4.3 Reduction to Partial Information Game

Consider an MTG $\mathcal{G} = \langle \text{Pla}, \text{S}, s_0, \text{Act}, \text{Top}, (\delta_t)_{t \in \text{Top}}, (\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}} \rangle$. For every Player $p \in \text{Pla}$, fix $T_p \subseteq \text{Top}$ to be the intended set of winning topologies.

Game construction. We construct a 3-player partial-information game \mathcal{H} with the following components. The players are **Eve**, **Adam** and **Snake**. The states of \mathcal{H} are $Q_{\mathcal{H}} = \{q_0\} \cup Q$, where q_0 is a designated initial state and $Q \subseteq \text{S} \times \text{Pla} \times 2^{\text{Top}} \times \text{Top} \times \{\text{true}, \text{false}\}$ is described in the following. A state $(s, p, T, t, b) \in Q$ comprises $s \in \text{S}$ which tracks the state of \mathcal{G} , a player $p \in \text{Pla}$ that is controlled by **Adam**, a set $T \subseteq \text{Top}$ of topologies that **Adam** picks, $t \in \text{Top}$ is a topology picked by **Snake** and determines the topology \mathcal{G} is played in, and a bit $b \in \{\text{true}, \text{false}\}$ which tracks whether **Adam** obeys **Eve**.

In order to restrict the state space to a polynomial size in $|\mathcal{G}|$, i.e. reduce the 2^{Top} component, we define $\mathcal{T}_p = \{T_p \cup \{t\} \mid t \in \text{Top}\} \subseteq 2^{\text{Top}}$ and $\mathcal{T} = (\bigcup_{p \in \text{Pla}} \mathcal{T}_p) \cup \{\{t\} \mid t \in \text{Top}\}$. Note that $|\mathcal{T}| \leq (|\text{Pla}| + 1) \cdot |\text{Top}| \leq 2 \cdot |\text{Pla}| \cdot |\text{Top}|$. We now define $Q = \text{S} \times \text{Pla} \times \mathcal{T} \times \text{Top} \times \{\text{true}, \text{false}\}$. Intuitively, the restriction of 2^{Top} to \mathcal{T} is sound, since if a Player p is able to deviate and increase her winning topologies from T_p to some T , then she can also increase her winning topologies by just one topology, and thus we can assume $T \in \mathcal{T}_p$.

We now turn to define the transitions in \mathcal{H} . The actions are defined implicitly by the transitions.³ From q_0 , **Adam** selects a player $p \in \text{Pla}$ and a set of topologies $T \in \mathcal{T}_p$. As explained in Section 4.2, **Adam** controls Player p and attempts to show that p wins in T . Still in q_0 , **Snake** selects a topology $t \in \text{Top}$ that \mathcal{G} will be played in. Then, \mathcal{H} transitions to state $(s_0, p, T, t, \text{true}) \in Q$.

Henceforth, p, T and t remain fixed throughout the play, and **Snake** has no further effect on the play. From state $(s, p, T, t, b) \in Q$, **Eve** chooses an action profile $\mathbf{a} \in \text{Act}^{\text{Pla}}$ and **Adam** selects an action $a'_p \in \text{Act}$. Then, the game transitions to state $(s', p, T, t, b') \in Q$ such that $s' = \delta_t(s, \mathbf{a}[p \mapsto a'_p])$, and $b' = b \wedge a_p = a'_p$. That is, **Eve** chooses an action profile, **Adam** chooses a possible deviation, and the game proceeds according to \mathcal{G}_t . If **Adam** actually deviates, the bit b becomes **false** and remains so throughout the play. Adding $\{\{t\} \mid t \in \mathcal{T}\}$ to \mathcal{T} is to make sure that if Player p is supposed to win in topology t (that is, $t \in T_p$), then, the profile suggested by **Eve** must lead to player p winning in topology t . If not, **Adam** can choose $\{t\}$ and Player p at the start of the game, and obey **Eve**, falsifying one of **Eve**'s winning conditions (ψ_2).

Next, we define the observation sets of \mathcal{H} . For a state $q = (s, p, T, t, b) \in Q$ we define the *projection* of q on \mathcal{G} to be $\text{proj}(q) = s$. For every state $s \in \text{S}$ of \mathcal{G} , let $o_s = \{q \in Q \mid \text{proj}(q) = s\} \subseteq Q$. The observation sets in \mathcal{H} are $\mathcal{O}_{\text{Adam}} = \mathcal{O}_{\text{Eve}} = \mathcal{O} = \{\{q_0\}\} \cup \{o_s \mid s \in \text{S}\}$. That is, **Adam** and **Eve** can observe the initial state q_0 , and for every $q \in Q$ they can only observe $\text{proj}(q)$. **Snake** has perfect information.

This completes the construction of the game \mathcal{H} (recall that \mathcal{H} does not have an objective). We proceed to formalize the connection between \mathcal{G} and \mathcal{H} .

Correspondence between \mathcal{H} and \mathcal{G} . We lift the definition of projection to plays: for a play $\rho = q_0 q_1 q_2 \dots \in q_0 \cdot Q^\omega$ of \mathcal{H} define $\text{proj}(\rho) = \text{proj}(q_1)\text{proj}(q_2)\dots$ (note that we skip the initial state q_0). We also define the predicate $\text{obey}(\rho) = \bigwedge_{i \geq 1} b_i$, where b_i is the **true/false** bit of q_i . That is, $\text{obey}(\rho)$ is true if and only if **Adam** always takes the actions suggested by **Eve**. When $\text{obey}(\rho)$ is true, we say that **Adam** *obeys* **Eve**.

³ In the model we describe, actions are identical for all players. However, the model of [10] allows different actions as well as enabled and disabled actions in each state, so it is easy to accommodate our actions.

Since the observation of **Eve** and **Adam** correspond to states of \mathcal{G} , there is a correspondence between plays, observation-histories and strategies in \mathcal{H} to plays, histories and strategies in \mathcal{G} . We make this precise in the following. Consider the function $\gamma_{\text{obs}} : \{q_0\} \cdot \mathcal{O}^\omega \rightarrow \mathcal{S}^\omega$ defined $\gamma_{\text{obs}}(\{q_0\}, o_{s_0}, o_{s_1}, \dots) = s_0, s_1, \dots$. Since $o_s = \{q \mid \text{proj}(q) = s\}$ for every $s \in \mathcal{S}$, we have that γ_{obs} is a bijection between observation-plays of **Eve** and **Adam** in \mathcal{H} , and plays of \mathcal{G} . By looking at finite sequences, namely histories, we can refer to γ_{obs} as a bijection between observation-histories of **Adam** and **Eve** in \mathcal{H} , and histories in \mathcal{G} . Moreover, since strategies in \mathcal{H} are observation based, the following functions are also bijective:

- $\gamma_{\text{Eve}} : \Sigma_{\mathcal{H}}^{\text{Eve}} \rightarrow \Sigma_{\mathcal{G}}$ defined by $\gamma_{\text{Eve}}(\sigma_{\text{Eve}}) = \sigma_{\text{Eve}} \circ \gamma_{\text{obs}}^{-1}$.
- $\gamma_{\text{Adam}} : \Sigma_{\mathcal{H}}^{\text{Adam}} \rightarrow \bigcup_{p \in \text{Pla}} \{p\} \times \mathcal{T}_p \times \Sigma_{\mathcal{G}}^p$ defined $\gamma_{\text{Adam}}(\sigma_{\text{Adam}}) = (p, T, \sigma'_p)$ such that $\sigma_{\text{Adam}}(q_0) = (p, T)$ are the player and the set of topologies selected by **Adam** in state q_0 , and $\sigma'_p = \sigma_{\text{Adam}} \circ \gamma_{\text{obs}}^{-1}$ is the deviating strategy in \mathcal{G} induced by the deviation proposed in σ_{Adam} in \mathcal{H} .
- $\gamma_{\text{Snake}} : \Sigma_{\mathcal{H}}^{\text{Snake}} \rightarrow \text{Top}$ defined by $\gamma_{\text{Snake}}(\sigma_{\text{Snake}}) = \sigma_{\text{Snake}}(q_0)$ (recall that **Snake** only acts in q_0).

For readability, we omit the the subscript and write γ instead of $\gamma_{\text{obs}}, \gamma_{\text{Adam}}, \gamma_{\text{Eve}}, \gamma_{\text{Snake}}$. The correct subscript can be resolved from context. Intuitively, γ is the correspondence from strategies/histories/plays in \mathcal{H} to their counterpart in \mathcal{G} .

The connection between strategies and outcomes in \mathcal{H} and \mathcal{G} is formalized in the following lemma (see Appendix A.1 for the proof).

► **Lemma 14.** *Consider strategies $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$, $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ and $\sigma_{\text{Snake}} \in \Sigma_{\mathcal{H}}^{\text{Snake}}$. Let $\sigma = \gamma(\sigma_{\text{Eve}})$, $(p, T, \sigma'_p) = \gamma(\sigma_{\text{Adam}})$ and $t = \gamma(\sigma_{\text{Snake}})$. Let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}})$, $\pi' = \text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p])$, and $\pi = \text{out}_{\mathcal{G}_t}(\sigma)$. Then $\text{proj}(\rho) = \pi'$. Furthermore, if **Adam** obeys **Eve** on ρ then $\text{proj}(\rho) = \pi = \pi'$.*

Objective for \mathcal{H} . As sketched in Section 4.2, the objective α in \mathcal{H} is constructed so that **Eve** and **Snake** can win if and only if there is a CNE in \mathcal{G} with winning topologies $(T_p)_{p \in \text{Pla}}$.

We define α as a conjunction of three conditions $\alpha = \{\rho \in q_0 \cdot Q^\omega \mid \psi_1(\rho) \wedge \psi_2(\rho) \wedge \psi_3(\rho)\}$, where the conditions are defined as follows. Consider a play $\rho = q_0, (s_0, p, T, t, b_0), (s_1, p, T, t, b_1), \dots$ of \mathcal{H} .

- $\psi_1(\rho) := t \in T$. That is, ψ_1 forces **Snake** to choose a topology from the set of topologies selected by **Adam**.
- $\psi_2(\rho) := (\text{obey}(\rho) \wedge t \in T_p) \rightarrow \text{proj}(\rho) \in \alpha_{t,p}$. That is, ψ_2 is satisfied if whenever **Adam** obeys **Eve** then Player p wins in any topology $t \in T_p$ selected by **Snake**.
- $\psi_3(\rho) := T_p \subsetneq T \rightarrow \text{proj}(\rho) \notin \alpha_{t,p}$. That is, ψ_3 is satisfied if whenever **Adam** tries to win in a strict superset of T_p , then Player p loses in the topology selected by **Snake**.

As mentioned in Section 4.2, it is not clear that α can be expressed as a single parity objective over $Q_{\mathcal{H}}$. Nonetheless, we prove that this is possible. The key observation is that the “postconditions” of ψ_2 and ψ_3 contradict, hence one of them must hold vacuously. This allows us to decouple the parity conditions for each of them and obtain a single parity objective that captures both, as follows.

For each objective $\alpha_{t,p}$ in \mathcal{G} we write $\alpha_{t,p} = \text{Parity}(\Omega_{t,p})$ such that $\Omega_{t,p} : \mathcal{S} \rightarrow \{0, \dots, d\}$ is the parity ranking function, where $d \in \mathbb{N}$. We define a new ranking function $\Omega : Q_{\mathcal{H}} \rightarrow \{0, \dots, d+1\}$, and show that $\alpha = \text{Parity}(\Omega)$.

First, observe that q_0 occurs only once in each play, so its parity rank has no effect. We arbitrarily set $\Omega(q_0) = 0$. Let $\rho \in q_0 \cdot Q^\omega$ be a play of \mathcal{H} and $(s, p, T, t, b), (s', p', T', t', b') \in \text{Inf}(\rho)$. It must be that $p = p'$, $T = T'$ and $t = t'$ since those are constant throughout the play, and $b = b'$ since it is either always **true** or from some point in ρ it turns into **false** and stays that way to the rest of the play.

34:12 Concurrent Games with Multiple Topologies

Let $q = (s, p, T, t, b) \in Q$. We define $\Omega(q)$ by cases according to p, T, t, b , and show that in each case, $\rho \in \alpha$ if and only if $\rho \in \text{Parity}(\Omega)$, concluding that $\alpha = \text{Parity}(\Omega)$. For a formula of the form $\psi = \varphi_1 \rightarrow \varphi_2$, we refer to φ_1 as the *precondition* of ψ , and φ_2 as the *postcondition* of ψ .

- $t \notin T$: In this case, if $q \in \text{Inf}(\rho)$ then ρ does not satisfy ψ_1 , thus, $\rho \notin \alpha$. We set $\Omega(q) = 1$ to get $\rho \notin \text{Parity}(\Omega)$.
- $t \in T$, $b = \text{true}$, $t \in T_p$ and $T_p \subsetneq T$: In this case, if $q \in \text{Inf}(\rho)$ then ρ satisfies the preconditions of both ψ_2 and ψ_3 , but the postconditions of ψ_2 and ψ_3 contradict, thus, $\rho \notin \alpha$. We set $\Omega(q) = 1$ to get $\rho \notin \text{Parity}(\Omega)$.
- $t \in T$, $b = \text{true} \wedge t \in T_p$ and $\neg(T_p \subsetneq T)$: In this case, if $q \in \text{Inf}(\rho)$, then $\rho \in \alpha \iff \text{proj}(\rho) \in \alpha_{t,p}$. So we set $\Omega(q) = \Omega_{t,p}(s)$, to apply the objective $\alpha_{t,p}$ over $\text{proj}(\rho)$.
- $t \in T$, $\neg(b = \text{true} \wedge t \in T_p)$ and $T_p \subsetneq T$: In this case, if $q \in \text{Inf}(\rho)$, then $\rho \in \alpha \iff \text{proj}(\rho) \notin \alpha_{t,p}$. So we set $\Omega(q) = \Omega_{t,p}(s) + 1$, to apply the complement of the objective $\alpha_{t,p}$ over $\text{proj}(\rho)$.
- $t \in T$, $\neg(b = \text{true} \wedge t \in T_p)$ and $\neg(T_p \subsetneq T)$: In this case, if $q \in \text{Inf}(\rho)$ then ψ_2 and ψ_3 are vacuously satisfied, and $\rho \in \alpha$. So we set $\Omega(q) = 0$ to get that $\rho \in \text{Parity}(\Omega)$.

We are now ready to characterize the existence of a CNE in \mathcal{G} by winning strategies in \mathcal{H} .

► **Lemma 15.** *Consider an MTG $\mathcal{G} = \langle \text{Pla}, \text{S}, s_0, \text{Act}, \text{Top}, (\delta_t)_{t \in \text{Top}}, (\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}} \rangle$. Let $(T_p)_{p \in \text{Pla}}$ be sets of topologies for each player and let \mathcal{H} be the corresponding partial-information game. There exists a strategy profile σ in \mathcal{G} such that σ is a CNE and for every $p \in \text{Pla}$ we have $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ if and only if the following holds:*

$$\exists \sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}} \forall \sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}} \exists \sigma_{\text{Snake}} \in \Sigma_{\mathcal{H}}^{\text{Snake}} \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}}) \in \alpha.$$

Proof. Assume σ is a CNE in \mathcal{G} such that for every $p \in \text{Pla}$, $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$, and fix $\sigma_{\text{Eve}} = \gamma^{-1}(\sigma)$ to be the corresponding strategy for **Eve** in \mathcal{H} . Consider a strategy $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ for **Adam**, and let $(p, T, \sigma'_p) = \gamma(\sigma_{\text{Adam}})$. We show that there exists a strategy $\sigma_{\text{Snake}} \in \Sigma_{\mathcal{H}}^{\text{Snake}}$ so that the outcome satisfies α . Recall that a strategy for **Snake** amounts to choosing a topology. We divide to cases according to the choice of T by **Adam**.

- If $\neg(T_p \subsetneq T)$, then ψ_3 is satisfied vacuously. Choose $t \in T$ for **Snake**, then ψ_1 is satisfied. If **Adam** does not obey **Eve** or $t \notin T_p$ then ψ_2 is vacuously satisfied. Otherwise, if **Adam** obeys **Eve** and $t \in T_p$, let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}})$. In order to show that ψ_2 is satisfied we need to show that $\text{proj}(\rho) \in \alpha_{t,p}$. Let $\pi = \text{out}_{\mathcal{G}_t}(\sigma)$. Since $T_p = \text{WinTop}_{\mathcal{G}}^p(\sigma)$ and $t \in T_p$ we have that $\pi \in \alpha_{t,p}$. From Lemma 14 we have that $\text{proj}(\rho) = \pi$, so we get that $\text{proj}(\rho) \in \alpha_{t,p}$, as required.
- If $T_p \subsetneq T$, denote $T' = \text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p])$. Since σ is a CNE, we have that $\neg(T_p \subsetneq T')$, so $T \setminus T' \neq \emptyset$, as otherwise we would have that $T_p \subsetneq T \subseteq T'$. Choose $t \in T \setminus T'$ for **Snake**, then ψ_1 is satisfied. Let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}})$, $\pi' = \text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p])$ and $\pi = \text{out}_{\mathcal{G}_t}(\sigma)$. From Lemma 14 we have that $\text{proj}(\rho) = \pi'$ and if **Adam** obeys **Eve** then we have $\text{proj}(\rho) = \pi = \pi'$. Note that since $t \notin T' = \text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p])$ then $\pi' \notin \alpha_{t,p}$, so ψ_3 is satisfied. Finally, ψ_2 is satisfied vacuously since we cannot have $t \in T_p$ and that **Adam** obeys **Eve** simultaneously, as this would yield $T' = T_p = \text{WinTop}_{\mathcal{G}}^p(\sigma)$, but $t \notin T'$.

We conclude that in all cases $\rho \in \alpha$, as required.

Conversely, assume that $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$ is such that for every $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ there exists $\sigma_{\text{Snake}} \in \Sigma_{\mathcal{H}}^{\text{Snake}}$ such that $\text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}}) \in \alpha$. Let $\sigma = \gamma(\sigma_{\text{Eve}})$. We start by showing that for every $p \in \text{Pla}$ it holds that $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$. Indeed, let $p \in \text{Pla}$ and $t \in \text{Top}$.

If $t \in T_p$, take $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ that selects player p and $T = \{t\}$, and obeys Eve. The only strategy σ_{Snake} for Snake that satisfies ψ_1 is to select t . Let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}})$. From ψ_2 we get that $\text{proj}(\rho) \in \alpha_{t,p}$, and by Lemma 14 we have $\text{proj}(\rho) = \text{out}_{\mathcal{G}_t}(\sigma)$. Thus, $t \in \text{WinTop}_{\mathcal{G}}^p(\sigma)$.

If $t \notin T_p$, take $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ that selects Player p and $T = T_p \cup \{t\}$, and obeys Eve. Since Adam obeys Eve, in order for ψ_1 , ψ_2 and ψ_3 to be satisfied, Snake must choose t , otherwise both preconditions of ψ_2 and ψ_3 hold, which means that in order to win we must have both $\text{proj}(\rho) \in \alpha_{t,p}$ (by ψ_2) and $\text{proj}(\rho) \notin \alpha_{t,p}$ (by ψ_3), which cannot hold. Thus, Snake chooses t , and from Lemma 14 we have $\text{proj}(\rho) = \text{out}_{\mathcal{G}_t}(\sigma)$. By ψ_3 we have $\text{proj}(\rho) \notin \alpha_{t,p}$, so $\text{out}_{\mathcal{G}_t}(\sigma) \notin \alpha_{t,p}$. Thus $t \notin \text{WinTop}_{\mathcal{G}}^p(\sigma)$. Therefore, $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$.

It remains to show that σ is a CNE. Assume by way of contradiction that there exists a player $p \in \text{Pla}$ with a beneficial deviation $\sigma'_p \in \Sigma_{\mathcal{G}}^p$. That is, $T' = \text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p])$ satisfies $T_p \subsetneq T'$. We will construct a strategy of Adam such that every strategy of Snake is losing, thereby reaching a contradiction. Let $T = T_p \cup \{t'\}$ for some $t' \in T \setminus T_p$ and fix $\sigma_{\text{Adam}} = \gamma^{-1}(p, T, \sigma'_p)$. Consider a strategy σ_{Snake} , denote $t = \gamma(\sigma_{\text{Snake}})$ and let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}})$. By Lemma 14 we have $\text{proj}(\rho) = \text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p])$, and because $t \in T \subseteq \text{WinTop}_{\mathcal{G}}^p(\sigma[p \mapsto \sigma'_p])$ it holds that $\text{proj}(\rho) \in \alpha_{t,p}$. However, $T_p \subsetneq T$, so ψ_3 is violated, and $\rho \notin \alpha$, which is a contradiction. We conclude that σ is a CNE. \blacktriangleleft

Using Lemma 15 we can decide whether a given MTG \mathcal{G} has a CNE, by iterating over all possible sets of candidate winning topologies $(T_p)_{p \in \text{Pla}}$, and repeatedly applying the reduction, and using the decision procedure of Theorem 13. It remains to analyze the complexity of this procedure.

To this end, observe that the size of \mathcal{H} is polynomial in the size of \mathcal{G} . Indeed, $|Q| \leq |S| \cdot |\text{Pla}| \cdot |\mathcal{T}| \cdot |\text{Top}| \cdot 2$ where $|\mathcal{T}| \leq 2|\text{Pla}||\text{Top}|$. and the description of the actions is also polynomial in that of \mathcal{G} (note that Eve has exponentially more actions than each player in \mathcal{G} , but the overall description of the transition table in \mathcal{G} is similarly exponential, cf. Remark 2).

Finally, by Theorem 13, solving \mathcal{H} takes double-exponential time in $|\mathcal{G}|$, and we have a single-exponential number of iterations, so the overall complexity remains double-exponential time in $|\mathcal{G}|$. This completes the proof of Theorem 12.

► Remark 16 (Lower bounds and improving the upper bound). We do not have a lower bound for the 2-EXPTIME complexity of Theorem 12. Indeed, we suspect that this bound can be lowered. This is due in part to the fact that game \mathcal{H} we construct does not utilize the full scope of Theorem 13 from [10]. Unfortunately, the decision procedure in [10] goes through three nontrivial reductions, one of which involves Safra's determinization, that is notoriously difficult to analyze: The first reduction [9, 10] transforms the objective to a visible objective for Adam which involves the determinization of a parity automaton. The second reduction [10] reduces the three-player partial-information game into a two-player partial-information game. The third reduction uses the results of [22] to reduce the two-player partial-information game to a two-player perfect-information game.

Therefore, it is likely that improving the bound (if indeed possible) will involve devising an ad-hoc procedure, possibly using some key ideas from [9, 10, 22].

5 Existence of Greedy NE is Decidable

We now turn our attention to Greedy NE (GNE). Recall that a greedy beneficial deviation is one that wins in a previously-losing topology, even at the cost of losing in previously-winning topologies. That is, given an MTG $\mathcal{G} = \langle \text{Pla}, S, s_0, \text{Act}, \text{Top}, (\delta_t)_{t \in \text{Top}}, (\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}} \rangle$, a profile $\sigma \in \Sigma_{\mathcal{G}}^{\text{Pla}}$ is a GNE if for every $p \in \text{Pla}$, $\sigma'_p \in \Sigma_{\mathcal{G}}$ and $t \in \text{Top}$, if $p \in \text{Win}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p])$ then $p \in \text{Win}_{\mathcal{G}_t}(\sigma)$.

Intuitively, reasoning in the greedy approach is much less delicate than the conservative approach, since a deviating player need not concern itself with keeping the current winning topologies. As we show in the following, this allows for an exponentially faster solution.

► **Theorem 17.** *The problem of deciding, given an MTG \mathcal{G} , whether there exists a GNE in \mathcal{G} is in EXPTIME.*

Similarly to Section 4, our approach is to reduce the problem at hand to solving a partial-information game. In the greedy setting, however, it suffices to use two-player games. Specifically, we employ the following result from [9].

► **Theorem 18.** *Let $\mathcal{G} = \langle \text{Pla}, \mathcal{S}, s_0, \text{Act}, \delta, (\mathcal{O}_p)_{p \in \text{Pla}} \rangle$ with $\text{Pla} = \{1, 2\}$. Let $\alpha \subseteq \mathcal{S}^\omega$ be a parity objective. The problem of deciding whether $\exists \sigma_1 \in \Sigma_{\mathcal{G}}^1 \forall \sigma_2 \in \Sigma_{\mathcal{G}}^2 \text{out}_{\mathcal{G}}(\sigma_1, \sigma_2) \in \alpha$ is EXPTIME-complete.*

We sketch the proof of Theorem 17. The complete construction and analysis are detailed in Appendix B.

Proof sketch. As in Section 4.3, we first fix a set of “intended” winning topologies $T_p \subseteq \text{Top}$ for each player $p \in \text{Pla}$. Then, we ask whether \mathcal{G} admits a GNE σ in which $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ for every $p \in \text{Pla}$. We then construct a 2-player partial-information game whose players are Eve, Adam, where Eve again controls the coalition of all players.

The behaviour of Adam is different than in the conservative setting. Here, Adam starts by choosing a deviating player $p \in \text{Pla}$ and a *single topology* $t \in \text{Top}$ where p attempts to win. The topology t is unobservable by Eve. The observations sets of Eve and Adam are again only the current state of \mathcal{G} . Then, the game is played on topology t with Eve suggesting an action profile, and Adam possibly deviating with Player p .

The objective for Eve now comprises two conditions:

- ψ_1 requires that whenever Adam obeys Eve and $t \in T_p$, the outcome is winning for Player p in \mathcal{G}_t .
- ψ_2 requires that if $t \notin T_p$, then Player p loses in \mathcal{G}_t .

Intuitively, Adam tries to cause Player p to win in a new topology t in which Player p is not intended to win, while Eve is trying to prevent Player p from achieving this, provided that Player p is actually deviating. Note that Eve must do this without knowing which topology is chosen, nor which player deviates (if at all). ◀

6 Discussion, Extensions and Future Work

We introduced MTGs and notions of NE pertaining to them, and showed that deciding whether an MTG admits either notion is decidable (in 2-EXPTIME for CNE and in EXPTIME for GNE). We have also explored the relationships and properties of these notions of NE. We now turn to explore several extensions, and remark about future research directions.

Social optimum. A standard solution concept for concurrent games, apart from NE, is *social optimum*, namely what is the maximum welfare the player can obtain by cooperating. Since in MTGs the winning sets of topologies may be incomparable, we formulate this as follows: given sets $(T_p)_{p \in \text{Pla}}$, is there a strategy profile σ such that $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ for every $p \in \text{Pla}$?

Fortunately, the techniques we developed enable us to readily solve this problem. Indeed, we can modify the reduction used to decide the existence of GNE (Section 5) so that Adam chooses a player and a topology, but does not attempt to deviate and has no further effect on the game. Intuitively, Adam “challenges” Eve to show that the winning topologies for the players are exactly the intended ones. The complexity of this approach remains EXPTIME.

Lower bounds. As discussed in Remark 16, we do not provide lower bounds for our results. Trivial lower bounds on the existence of CNE and GNE can be obtained from those of NE existence in concurrent games, namely $P_{||}^{NP}$ -hardness [6]. This, however, is unlikely to be tight. A central open challenge is to determine the exact complexity of CNE and GNE existence in MTGs.

Additional notions of equilibria. The notions we propose, namely CNE and GNE, lie on two extremities: in the conservative setting a deviation is very strict, and in the greedy setting it is very lax. Generally, one can obtain a notion of equilibrium using any binary relation on 2^{Top} , which describes what the beneficial deviations are for each player. Moreover, different players can have different relations.

Of particular interest is a quantitative notion of NE, whereby a player deviates if she can increase the *number* of her winning topologies. This notion is fundamentally different from CNE and GNE, as it is not based on set containment, which is key to the correctness of our approach.

Succinct representation of topologies. A central motivation for MTGs, demonstrated in Example 1 and in Section 3.1 concerns process symmetry. There, from a game with k players, we construct an MTG with $k!$ topologies. However, these topologies can be succinctly represented by computing them on-the-fly. An interesting direction for future work is to determine whether we can devise a symbolic approach that is able to handle such MTGs without incurring an exponential blowup.

Logic for partial information games. Another approach to solve the CNE and GNE existence problems might be to formulate those problems with a logic for partial information games [3, 16, 20]. The most promising of those is [3], as it is the most expressive and the complexity of different fragments has been studied. From the complexity of the decision procedures for those logics, it not likely that this approach can be used to lower the EXPTIME complexity for GNE and the 2-EXPTIME complexity for CNE.

References

- 1 Shaull Almagor. Process symmetry in probabilistic transducers. In *40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, 2020.
- 2 Shaull Almagor, Guy Avni, and Orna Kupferman. Repairing multi-player games. In *26th International Conference on Concurrency Theory (CONCUR 2015)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- 3 Raphaël Berthon, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Y Vardi. Strategy logic with imperfect information. *ACM Transactions on Computational Logic (TOCL)*, 22(1):1–51, 2021.
- 4 Udi Boker. Why these automata types? In *LPAR*, volume 18, pages 143–163, 2018.
- 5 Patricia Bouyer, Nicolas Markey, and Steen Vester. Nash equilibria in symmetric graph games with partial observation. *Information and Computation*, 254:238–258, 2017.

- 6 Patricia P Bouyer, Romain Brenguier, and Nicolas N Markey. Pure nash equilibria in concurrent games. *Logical methods in computer science*, 2015.
- 7 Felix Brandt, Felix Fischer, and Markus Holzer. Equilibria of graphical games with symmetries. *Theoretical Computer Science*, 412(8-10):675–685, 2011.
- 8 Romain Brenguier, Arno Pauly, Jean-François Raskin, and Ocan Sankur. Admissibility in games with imperfect information. In *CONCUR 2017-28th International Conference on Concurrency Theory*, volume 85, pages 2:1–2:23. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2017.
- 9 Krishnendu Chatterjee and Laurent Doyen. The complexity of partial-observation parity games. In *International Conference on Logic for Programming Artificial Intelligence and Reasoning*, pages 1–14. Springer, 2010.
- 10 Krishnendu Chatterjee and Laurent Doyen. Games with a weak adversary. In *International Colloquium on Automata, Languages, and Programming*, pages 110–121. Springer, 2014.
- 11 Krishnendu Chatterjee, Thomas A Henzinger, and Nir Piterman. Strategy logic. *Information and Computation*, 208(6):677–693, 2010.
- 12 Edmund M. Clarke, Reinhard Enders, Thomas Filkorn, and Somesh Jha. Exploiting symmetry in temporal logic model checking. *Formal methods in system design*, 9(1):77–104, 1996.
- 13 Luca De Alfaro, Thomas A Henzinger, and Orna Kupferman. Concurrent reachability games. *Theoretical computer science*, 386(3):188–217, 2007.
- 14 Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Toruńczyk. Energy and mean-payoff games with imperfect information. In *International Workshop on Computer Science Logic*, pages 260–274. Springer, 2010.
- 15 E Allen Emerson and A Prasad Sistla. Symmetry and model checking. *Formal methods in system design*, 9(1):105–131, 1996.
- 16 Bernd Finkbeiner and Sven Schewe. Coordination logic. In *International Workshop on Computer Science Logic*, pages 305–319. Springer, 2010.
- 17 Nicholas Ham. Notions of anonymity, fairness and symmetry for finite strategic-form games. *arXiv preprint*, 2013. [arXiv:1311.4766](https://arxiv.org/abs/1311.4766).
- 18 C Norris Ip and David L Dill. Better verification through symmetry. In *Computer Hardware Description Languages and their Applications*, pages 97–111. Elsevier, 1993.
- 19 Anthony W Lin, Truong Khanh Nguyen, Philipp Rümmer, and Jun Sun. Regular symmetry patterns. In *International Conference on Verification, Model Checking, and Abstract Interpretation*, pages 455–475. Springer, 2016.
- 20 Bastien Maubert. *Logical foundations of games with imperfect information: uniform strategies*. PhD thesis, Université Rennes 1, 2014.
- 21 Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007. doi:10.1017/CB09780511800481.
- 22 Jean-François Raskin, Thomas A Henzinger, Laurent Doyen, and Krishnendu Chatterjee. Algorithms for omega-regular games with imperfect information. *Logical Methods in Computer Science*, 3, 2007.
- 23 Noah Daniel Stein. *Exchangeable equilibria*. PhD thesis, Massachusetts Institute of Technology, 2011.
- 24 Fernando A Tohmé and Ignacio D Viglizzo. Structural relations of symmetry among players in strategic games. *International Journal of General Systems*, 48(4):443–461, 2019.
- 25 M Ummels and DK Wojtczak. The complexity of nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 2010.
- 26 Steen Vester. *Symmetric Nash Equilibria*. PhD thesis, Master’s thesis, ENS Cachan, 2012.

A Proofs

A.1 Proof of Lemma 14

Proof. We prove by induction that for every $k \geq 1$, $\text{proj}(\rho_{\leq k+1}) = \pi'_{\leq k}$, and if **Adam** obeys **Eve** then $\text{proj}(\rho_{\leq k+1}) = \pi'_{\leq k} = \pi_{\leq k}$. For $k = 1$, $\rho_{\leq 2} = q_0, (s_0, p, t, T, b_0)$ and $\pi'_{\leq 1} = \pi_{\leq 1} = s_0$ and we have that $\text{proj}(\rho_{\leq k+1}) = \pi'_{\leq k}$. Assuming that $\text{proj}(\rho_{\leq k+1}) = \pi'_{\leq k}$ for $k \geq 1$, the next state of $\text{proj}(\rho)$ will depend on the transition function δ_t and action profile $\sigma[p \mapsto \sigma'_p](\pi'_{\leq k})$ from the way γ and the transitions of \mathcal{H} are defined, and the next state in π' will also depend on the same transition function and action profile. Thus, it holds that $\text{proj}(\rho_{\leq k+2}) = \pi'_{\leq k+1}$. Farther more, if **Adam** obeys **Eve** then in every step the action that **Adam** takes is identical to the action that **Eve** suggests for Player p , so we have that $\sigma[p \mapsto \sigma'_p](\pi'_{\leq k}) = \sigma(\pi'_{\leq k})$, and $\pi_{\leq k+1} = \pi'_{\leq k+1}$, thus, $\text{proj}(\rho_{\leq k+2}) = \pi_{\leq k+1} = \pi'_{\leq k+1}$. ◀

B Proof of Theorem 17

Consider an MTG $\mathcal{G} = \langle \text{Pla}, \text{S}, s_0, \text{Act}, \text{Top}, (\delta_t)_{t \in \text{Top}}, (\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}} \rangle$. For every Player $p \in \text{Pla}$ fix $T_p \subseteq \text{Top}$ to be the intended set of winning topologies.

Game construction. We construct a two-player partial-information game \mathcal{H} with the following components. The players are **Eve** and **Adam**. The states of \mathcal{H} are $Q_{\mathcal{H}} = \{q_0\} \cup Q$ such that q_0 is a designated initial state and $Q = \text{S} \times \text{Pla} \times \text{Top} \times \{\text{true}, \text{false}\}$ is described in the following. A state $(s, p, t, b) \in Q$ comprises of $s \in \text{S}$ which tracks the state of \mathcal{G} , a player $p \in \text{Pla}$ that is controlled by **Adam**, a topology $t \in \text{Top}$ that **Adam** picks, and a bit $b \in \{\text{true}, \text{false}\}$ which tracks whether **Adam** obeys **Eve**.

We now turn to define the transitions of \mathcal{H} . The actions are defined implicitly by the transitions. From state q_0 , **Adam** selects a player $p \in \text{Pla}$ to control and a topology $t \in \text{Top}$ that \mathcal{G} will be played in. Then, \mathcal{H} transitions to state $(s_0, p, t, \text{true}) \in Q$. Henceforth, p and t remain fixed throughout the play. From state $(s, p, t, b) \in Q$, **Eve** chooses an action profile $\alpha \in \text{Act}^{\text{Pla}}$, and **Adam** selects an action $a'_p \in \text{Act}$ and \mathcal{H} transitions to state $(s', p, t, b') \in Q$ such that $s' = \delta_t(s, \alpha[p \mapsto a'_p])$, and $b' = b \wedge (a'_p = a_p)$.

The observation sets for the players, proj and obey are defined similarly as Section 4.3. Correspondence between \mathcal{H} and \mathcal{G} , $\gamma_{\text{obs}}, \gamma_{\text{Eve}}$ is defined in the same way as in Section 4.3, and $\gamma_{\text{Adam}} : \Sigma_{\mathcal{H}}^{\text{Adam}} \rightarrow \bigcup_{p \in \text{Pla}} \{p\} \times \text{Top} \times \Sigma_{\mathcal{H}}^p$ is defined for $\gamma(\sigma_{\text{Adam}}) = (p, t, \sigma'_p)$ such that (p, t) are the player and topology selected by σ_{Adam} in state q_0 and $\sigma'_p = \sigma_{\text{Adam}} \circ \gamma_{\text{obs}}^{-1}$.

The connection between strategies and outcomes in \mathcal{H} and \mathcal{G} is formalized in the following lemma whose proof is similar to that of Lemma 14.

► **Lemma 19.** *Consider strategies $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$ and $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$. Let $\sigma = \gamma(\sigma_{\text{Eve}})$ and $(p, t, \sigma'_p) = \gamma(\sigma_{\text{Adam}})$. Let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}})$, $\pi' = \text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p])$ and $\pi = \text{out}_{\mathcal{G}_t}(\sigma)$. Then, $\text{proj}(\rho) = \pi'$. Furthermore, if **Adam** obeys **Eve** on ρ then $\text{proj}(\rho) = \pi = \pi'$.*

Objective for \mathcal{H} . Let $\rho = q_0 \cdot (s_0, p, t, b_0) \cdot (s_1, p, t, b_1) \cdot \dots$ be a play in \mathcal{H} . The objective α is such that $\rho \in \alpha \iff \psi_1(\rho) \wedge \psi_2(\rho)$, where

- $\psi_1(\rho) := (\text{obey}(\rho) \wedge t \in T_p) \rightarrow \text{proj}(\rho) \in \alpha_{t,p}$.
- $\psi_2(\rho) := t \notin T_p \rightarrow \text{proj}(\rho) \notin \alpha_{t,p}$.

α can be expressed as a parity objective as follows. For every $t \in \text{Top}$, $p \in \text{Pla}$, let $\Omega_{t,p} : \text{S} \rightarrow \{0, \dots, d_{t,p}\}$ be the priority function for the parity objective $\alpha_{t,p}$ in \mathcal{G} . We construct a priority function $\Omega : Q_{\mathcal{H}} \rightarrow \{0, \dots, d\}$ such that $d = \max\{d_{t,p} + 1 \mid t \in \text{Top}, p \in \text{Pla}\}$. We set $\Omega(q_0) = 0$ and for state $q = (s, p, t, b) \in Q$ we have

$$\Omega(q) = \begin{cases} \Omega_{t,p}(s) + 1 & t \notin T_p \\ \Omega_{t,p}(s) & b \wedge t \in T_p \\ \Omega(q) = 0 & \neg b \wedge t \in T_p \end{cases}$$

If $t \notin T_p$, then, according to α , $\rho \in \alpha$ if and only if $\text{proj}(\rho) \notin \alpha_{t,p}$. This is achieved by adding 1 to $\Omega_{t,p}$ which gives us the complement of $\alpha_{t,p}$. The case where Adam obeys Eve and $t \in T_p$ is captured in the second case, where $\rho \in \alpha$ if and only if $\text{proj}(\rho) \in \alpha_{t,p}$. This is achieved by setting Ω to be the same as $\Omega_{t,p}$. In the last case, none of the preconditions of ψ_1 and ψ_2 hold, so $\rho \in \alpha$. This is achieved by setting Ω to 0, such that every such play will satisfy the objective.

► **Lemma 20.** *There exists a GNE $\sigma \in \Sigma_{\mathcal{G}}$ in \mathcal{G} with $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ for every $p \in \text{Pla}$, if and only if $\exists \sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}} \forall \sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}} \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}) \in \alpha$.*

Proof. Let $\sigma \in \Sigma_{\mathcal{G}}$ be a GNE with $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ for every $p \in \text{Pla}$. Let $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$ be the corresponding strategy for σ , and let $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ be some strategy for Adam that corresponds to (p, t, σ'_p) . Let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}})$. If $\text{obey}(\rho) \wedge t \in T_p$, then from Lemma 19 we have that $\text{proj}(\rho) = \text{out}_{\mathcal{G}_t}(\sigma)$, and since $t \in T_p = \text{WinTop}_{\mathcal{G}}^p(\sigma)$ then $\text{out}_{\mathcal{G}_t}(\sigma) \in \alpha_{t,p}$. Thus, ψ_1 is satisfied by ρ . If $t \notin T_p$ then from Lemma 19 we have that $\text{proj}(\rho) = \text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p])$ and since Player p is losing in t when \mathcal{G} is played with σ and σ is a GNE, then $\text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p]) \notin \alpha_{t,p}$. Thus, ψ_2 is satisfied and $\rho \in \alpha$.

Conversely, let $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$ be such that for any $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ we have $\text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}) \in \alpha$. Let $\sigma \in \Sigma_{\mathcal{G}}$ correspond to σ_{Eve} . We show that σ is a GNE. First, we show that for every $p \in \text{Pla}$, $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$. Let $t \in \text{Top}$ and $p \in \text{Pla}$. Take $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ that corresponds to (p, t, σ_p) where σ_p is the strategy assigned to p in σ . Let $\rho_t = \text{out}_{\mathcal{G}_t}(\sigma)$ and $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}})$. We have that $\rho \in \alpha$. Since Adam obeys Eve on ρ , from Lemma 19 we have that $\text{proj}(\rho) = \rho_t$. If $t \in T_p$ then from ψ_1 we get that $\rho_t = \text{proj}(\rho) \in \alpha_{t,p}$, thus, $t \in \text{WinTop}_{\mathcal{G}}^p(\sigma)$. If $t \notin T_p$ then from ψ_2 we get that $\rho_t = \text{proj}(\rho) \notin \alpha_{t,p}$, thus, $t \notin \text{WinTop}_{\mathcal{G}}^p(\sigma)$. So we get that $\text{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$. Now, we show that σ is a GNE. Let $p \in \text{Pla}$, $\sigma'_p \in \Sigma_{\mathcal{G}}^p$ and $t \in \text{Top}$ such that $t \notin T_p$. Let $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ correspond to (p, t, σ'_p) , and let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}})$. We have that $\rho \in \alpha$, thus, since $t \notin T_p$ then $\text{proj}(\rho) \notin \alpha_{t,p}$. From Lemma 19 we have that $\rho'_t = \text{out}_{\mathcal{G}_t}(\sigma[p \mapsto \sigma'_p]) = \text{proj}(\rho) \notin \alpha_{t,p}$, thus, $t \notin \text{WinTop}_{\mathcal{G}_t}^p(\sigma[p \mapsto \sigma'_p]) = T_p$, so σ is a GNE. ◀

The algorithm for solving the GNE existence problem is, for each $(T_p)_{p \in \text{Pla}} \in (2^{\text{Top}})^{\text{Pla}}$ we construct \mathcal{H} from \mathcal{G} and $(T_p)_{p \in \text{Pla}}$, and check if there exists $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$ such that for every $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$, $\text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}) \in \alpha$, if there exists such σ_{Eve} , then according to Lemma 20 is corresponding strategy profile is a GNE, then we return it. If we went through all $(T_p)_{p \in \text{Pla}} \in (2^{\text{Top}})^{\text{Pla}}$, then return that there does not exist a GNE in \mathcal{G} .

The size of \mathcal{H} is polynomial in the size of \mathcal{G} . We copy each $s \in \mathcal{S}$ for every combination of $p \in \text{Pla}$, $t \in \text{Top}$, $b \in \{\text{true}, \text{false}\}$, so we get $|\mathcal{Q}_{\mathcal{H}}| = 2 \cdot |\mathcal{S}| \cdot |\text{Pla}| \cdot |\text{Top}| + 1$, which is polynomial in the size of \mathcal{G} . The number of actions in \mathcal{H} is also polynomial in the number of enabled actions in \mathcal{G} (similarly to the analysis in Section 4.3).

The algorithm performs at most $2^{|\text{Top}| \cdot |\text{Pla}|}$ iterations, which is exponential in $|\mathcal{G}|$. In each iteration we solve \mathcal{H} with size that is polynomial in $|\mathcal{G}|$, so according to Theorem 18 this takes exponential time in $|\mathcal{G}|$, so the GNE existence problem is in EXPTIME.