Concurrent Games with Multiple Topologies

Shai Guendelman

Concurrent Games with Multiple Topologies

Research Thesis

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Computer Science

Shai Guendelman

Submitted to the Senate of the Technion — Israel Institute of Technology Tishrei 5783 Haifa October 2022

This research was carried out under the supervision of Dr. Shaull Almagor, in the Faculty of Computer Science.

Some results in this thesis have been published as articles by the author and research collaborators in conferences and journals during the course of the author's master's research period, the most up-to-date versions of which being:

Shaull Almagor and Shai Guendelman. Concurrent games with multiple topologies. In 33rd International Conference on Concurrency Theory, 2022.

Acknowledgements

I would like to thank my advisor, Dr. Shaull Almagor, for his time, patience and professional help.

The Technion's funding of this research is hereby acknowledged.

Contents

Li	st of	Figures		
Al	ostra	lct	1	
1	Intr	oduction	3	
2	\mathbf{Pre}	liminaries	7	
3	Multi-Topology Games			
	3.1	Process Symmetry in Concurrent Games	9	
	3.2	Solution Concepts	10	
	3.3	Properties of CNE and GNE	11	
4	Existence of Conservative NE is Decidable			
	4.1	Partial-Information Games	15	
	4.2	Overview of the Reduction	16	
	4.3	Reduction to Partial Information Game	17	
5	Exis	stence of Greedy NE is Decidable	25	
	5.1	Reduction to Partial Information Game	25	
6	Stra	ategy Logic with Imperfect Information	29	
	6.1	Overview of SL_{ii}	29	
	6.2	$\mathrm{MTG} \ \mathrm{to} \ CGS_{ii} \ldots $	30	
	6.3	Expressing GNE Existence Problem with SL_{ii}	31	
	6.4	Expressing CNE Existence Problem with SL_{ii}	32	
7	Discussion, Extensions and Future Work		33	
Bi	bliog	graphy	35	
He	Hebrew Abstract i			

List of Figures

1.1	Router game with two players	4
3.1	A single player MTG with two topologies, t_1 and t_2 . In both topologies,	
	the objective of the player is to reach $s_1, \ldots, \ldots, \ldots, \ldots, \ldots$	12
3.2	Symmetric XOR game.	13

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, in general, undecidable. 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 above setting and introduce Multi-Topology Games (MTGs) – concurrent games with several possible transition functions, each transition function is called a *topology*. At the start of the game, a topology is chosen, without players knowing which one. 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.

Chapter 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. An example of a concurrent game can be seen in Figure 1.1. A strategy for a player is 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 [BBM15].

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, games with *partial information* have been extensively studied in various forms [BMM⁺21; BMV17; CD10; CD14]. 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 undecidable in the general case where there are 3 or more players [FGR18] or in the case of stochastic games [UW10].

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 1.0.1. 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 1.1. 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 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 1.1a and Figure 1.1b.



Figure 1.1: Router game with two players.

These type of settings are commonly referred to as *process symmetry* [CEFJ96; ES96; ID93; LNRS16; Alm20], 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.

Settings where a component might enter an interaction without knowing the exact configuration of the system are common, for example, in interactions over networks where connectivity is not known a-priory [GT07].

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, Conservative NE (CNE) and Greedy NE (GNE). In CNE we assume that players are conservative, that is, a player might deviate only if the deviation leads to a better outcome in at least one topology, without leading to a worse outcome in any topology. In this case, we only consider deviations that lead to strictly better outcomes for a player. In GNE we assume that players are greedy, that is, a player might deviate if she can improve her outcome in a single topology, regardless of how it affects other topologies. GNE is useful in cases where players might have unknown preferences over the different topologies. In this case, we want to make sure that no player has a profitable deviation, under every possible preference.

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 is decidable in 2-EXPTIME, and that the problem of whether a game has a GNE is decidable in EXPTIME.

Related Work A central work concerning NE in concurrent games is [BBM15], 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 [DHK07], 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.

In distributed computing, the notion of *anonymity* [AGM02; GR07] is similar to process symmetry, in that it considers a setting where processes are interchangeable. But, in contrast to process symmetry, all processes run the same program. In our discussion of process symmetry, players are allowed to play different strategies. *Anonymous games* [DP07] is another notion that resembles process symmetry. In this setting, players are aware of their position in the game, but their objectives only depend on

the number of actions of each type that were taken, and not on the exact action of individual players.

A concurrent game can be formulated as a turn-based partial information game, by players choosing their actions one by one, without revealing any information on the actions that were chosen until all players have selected their actions, and only then take the corresponding transition. Partial information games are more expressive than concurrent games – not every partial information game can be represented as a concurrent game. Partial information games were extensively studied, e.g., in [CD10; RHDC07; CD14; BMM⁺21; BPRS17; DDG⁺10], typically in the zero-sum setting.

Finally, the work in [BMM⁺21] extends strategy logic [CHP10] 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. Strategy logic with imperfect information can be used for reasoning about MTGs. For more details refer to Chapter 6.

Thesis organization In Chapter 2 we present the basic definitions of concurrent games. In Chapter 3 we formally define MTGs, introduce two notions of equilibria for them, and study their properties. In Chapter 4 we give our main technical result, establishing the decidability of detecting CNE in MTGs. In Chapter 5 we establish the decidability of detecting GNE. In Chapter 6 we show how strategy logic with imperfect information can be used for reasoning about MTGs. Finally, in Chapter 7 we discuss our results and some extensions, and detail future directions.

Chapter 2

Preliminaries

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

A play of \mathcal{G} is an infinite sequence of states $\rho = s_0, s_1, \ldots \in \mathsf{S}^{\omega}$ such that for every step $i \in \mathbb{N}$ there exists an action profile \boldsymbol{a} such that $s_{i+1} = \delta(s_i, \boldsymbol{a})$. For $k \ge 1$ we denote the length-k prefix of $\rho_{\le k} = s_0, \ldots, s_{k-1} \in \mathsf{S}^+$. We denote by $\mathrm{Inf}(\rho)$ the set of states that occur infinitely often in ρ .

In this work we focus on parity objectives. A parity objective α is defined by a priority function over the states of the game $\Omega : S \to \{0, \ldots, d\}$ for some $d \in \mathbb{N}$. For a state $s \in S$, $\Omega(s)$ is called the priority or rank of s. A play ρ satisfies α if the minimal priority of the states in $\text{Inf}(\rho)$ is even. The objective $\alpha = \text{Parity}(\Omega) \subseteq S^{\omega}$ is the set of plays that satisfy the above condition. We mostly use the parity function implicitly, and so we do not include Ω in the description of \mathcal{G} . We chose to focus on parity objectives since other types of ω -regular games can be translated into games with parity objectives.

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

Remark. 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 |Pla|. This is in contrast with a more succinct representation, i.e., representing the transition function as a circuit. This assumption is common in the literature [BBM15]. We take it to simplify the complexity analysis of our solution.

A history of \mathcal{G} is a finite prefix of a play $h \in S^+$. A strategy for Player p is a function $\sigma : S^+ \to Act$ that maps a history to the next action of Player p. A strategy profile $\sigma = (\sigma_p)_{p \in \mathsf{Pla}}$ 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 $\Sigma_{\mathcal{G}}^{\mathsf{Pla}}$ (we omit the subscript \mathcal{G} when it is clear from context). A strategy profile $\boldsymbol{\sigma}$ can be thought as a function that maps histories to action profiles: given a history $h \in \mathsf{S}^+$ we have $\boldsymbol{\sigma}(h) = (\sigma_p(h))_{p \in \mathsf{Pla}} \in \mathsf{Act}^{\mathsf{Pla}}$.

For a strategy profile $\boldsymbol{\sigma}$ 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 $\boldsymbol{\sigma}$. Formally, $\operatorname{out}_{\mathcal{G}}(\boldsymbol{\sigma}) = s_0 s_1 \ldots \in S^{\omega}$ where s_0 is the initial state, and for every $i \geq 1$ we have $s_i = \delta(s_{i-1}, \boldsymbol{\sigma}(s_0, \ldots, s_{i-1}))$. Consider a play $\rho \in S^{\omega}$. The set of *winners* in ρ is the set of players whose objectives are met in ρ . Formally, $\operatorname{Win}_{\mathcal{G}}(\rho) = \{p \in \mathsf{Pla} \mid \rho \in \alpha_p\} \subseteq \mathsf{Pla}$. The set of winners in a strategy profile $\boldsymbol{\sigma}$ is then $\operatorname{Win}_{\mathcal{G}}(\boldsymbol{\sigma}) = \operatorname{Win}_{\mathcal{G}}(\operatorname{out}_{\mathcal{G}}(\boldsymbol{\sigma}))$. Player pis said to be *losing* if she is not winning.

Remark. Action visibility. Note that strategies are defined to only observe the history of visited states, and not the history of actions taken by the other players. This is a standard and natural assumption [BBM15; CD14] for concurrent models. There are, however, works (e.g., [AAK15]) 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. In [AAK15] it was shown that assuming visible actions reduces the complexity of the NE existence problem for parity games from being P_{\parallel}^{NP} -complete to being NP-complete. We expect a similar effect in our setting.

A strategy profile $\boldsymbol{\sigma}$ 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 $\boldsymbol{\sigma}$, a player $p \in \mathsf{Pla}$ and a strategy $\sigma'_p \in \Sigma_{\mathcal{G}}$ for Player p. We denote by $\boldsymbol{\sigma}[p \mapsto \sigma'_p] \in \Sigma^{\mathsf{Pla}}$ the strategy profile obtained from $\boldsymbol{\sigma}$ by replacing σ_p with σ'_p . Then, $\boldsymbol{\sigma}$ is an NE if for every player $p \in \mathsf{Pla}$ and every strategy $\sigma'_p \in \Sigma_{\mathcal{G}}$ for Player p, if $p \in \operatorname{Win}_{\mathcal{G}}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ then $p \in \operatorname{Win}_{\mathcal{G}}(\boldsymbol{\sigma})$. Viewed contrapositively: if p loses when \mathcal{G} is played with $\boldsymbol{\sigma}$, then palso loses after changing her strategy.

Chapter 3

Multi-Topology Games

An MTG is a tuple $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathsf{Top}, (\delta_t)_{t \in \mathsf{Top}}, (\alpha_{t,p})_{t \in \mathsf{Top}, p \in \mathsf{Pla}} \rangle$ where $\mathsf{Pla}, \mathsf{S}, s_0$, Act, are the same as in concurrent games. Top is a finite set of *topologies*, and for every $t \in \mathsf{Top}$ we have a transition function $\delta_t : \mathsf{S} \times \mathsf{Act}^{\mathsf{Pla}} \to \mathsf{S}$ and objective $\alpha_{t,p} \subseteq \mathsf{S}^{\omega}$ for every player $p \in \mathsf{Pla}$. An MTG can be thought of as a tuple of games over the same states, players and actions. That is, for $t \in \mathsf{Top}$, we define $\mathcal{G}_t = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \delta_t, (\alpha_{t,p})_{p \in \mathsf{Pla}} \rangle$ 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, 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., $\sigma_p: S^+ \to Act$. This lifts to strategy profiles and outcomes, as per Chapter 2. In particular, a strategy σ in \mathcal{G} can be applied to \mathcal{G}_t for every $t \in \mathsf{Top}$. Although players have no information at the start of the game on which topology is played, they might reason about the set of possible topologies as the game progresses. For example, if the observed history is not possible in some topology, the player knows that this topology is not played. This is captured implicitly in the way strategies are defined. Consider a strategy profile $\sigma \in \Sigma^{\mathsf{Pla}}$. 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, WinTop^p_{\mathcal{G}}(\sigma) = { $t \in \mathsf{Top} \mid p \in Win_{\mathcal{G}_t}(\sigma)$ }.

3.1 Process Symmetry in Concurrent Games

As we discuss in Chapter 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* [CEFJ96; ES96; ID93; LNRS16; Alm20]. Symmetry in games was studied in [TV19; Ste11; BFH11; Ham13] for *strategic form games*, which are games with a single turn. In [BMV17; Ves12], symmetry in concurrent games was studied by imposing restrictions on the game structure. We consider a different setting, where processes $1, \ldots, k$ 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 $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \delta, (\alpha_p)_{p \in \mathsf{Pla}} \rangle$ with $k \geq 2$ players, and that $\mathsf{Pla} = \{1, \ldots, k\}$. 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 $\pi \in \mathcal{S}_k$, were \mathcal{S}_k is the set of permutations over $\{1, \ldots, k\}$. For an action profile $\mathbf{a} \in \mathsf{Act}^{\mathsf{Pla}}$ we define $\pi(\mathbf{a}) = (a_{\pi^{-1}(1)}, \ldots, a_{\pi^{-1}(k)})$. That is, the action performed by Player *i* is taken at index $\pi(i)$. We now obtain the MTG $\mathcal{G}_{\pi} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathcal{S}_k, (\delta_{\pi})_{\pi \in \mathcal{S}_k}, (\alpha_{\pi,p})_{\pi \in \mathcal{S}_k, p \in \mathsf{Pla}} \rangle$ 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 $s \in \mathsf{S}$ and $\mathbf{a} \in \mathsf{Act}^{\mathsf{Pla}}$ 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 1.1 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:

$$\forall p \in \mathsf{Pla} \ \forall \sigma'_p \in \Sigma^p_{\mathcal{G}} \ ((\forall t \in \mathsf{Top} \ p \in Win_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p]) \to p \in Win_{\mathcal{G}_t}(\boldsymbol{\sigma})) \lor (\exists t \in \mathsf{Top} \ p \notin Win_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p]) \land p \in Win_{\mathcal{G}_t}(\boldsymbol{\sigma})))$$

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

$$\forall p \in \mathsf{Pla} \ \forall \sigma'_p \in \Sigma^p_{\mathcal{G}} \neg (\mathsf{WinTop}^p_{\mathcal{G}}(\boldsymbol{\sigma}) \subsetneq \mathsf{WinTop}^p_{\mathcal{G}}(\boldsymbol{\sigma}[p \mapsto \sigma'_p]))$$

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

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

¹The relation \subsetneq means "strictly contained".

the following holds:

$$\forall p \in \mathsf{Pla} \ \forall \sigma'_p \in \Sigma^p_{\mathcal{G}} \ \forall t \in \mathsf{Top} \ (p \in \mathrm{Win}_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p]) \to p \in \mathrm{Win}_{\mathcal{G}_t}(\boldsymbol{\sigma}))$$

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

$$\forall p \in \mathsf{Pla} \ \forall \sigma'_p \in \Sigma^p_{\mathcal{G}} \ (\mathrm{WinTop}^p_{\mathcal{G}}(\boldsymbol{\sigma}[p \mapsto \sigma'_p]) \subseteq \mathrm{WinTop}^p_{\mathcal{G}}(\boldsymbol{\sigma}))$$

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 refer 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 3.2.1. CNE and GNE. Recall the router game from Figure 1.1. 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. Additional notions of NE. CNE and GNE are based on the \subseteq preorder on the sets of topologies, 2^{Top}. In Chapter 7 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 (Pla, S, s_0 , Act, Top, $(\delta_t)_{t \in \text{Top}}$, $(\alpha_{t,p})_{t \in \text{Top}, p \in \text{Pla}}$). 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 3.3.1. If 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 3.3.2. 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 3.3.2 is strict. That is, there are MTGs with a CNE but without a GNE.



Figure 3.1: 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 .

Example 3.3.3. CNE without GNE. Consider the single-player game depicted in Figure 3.1. 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 $\{t_1\} \notin \{t_2\}$ and $\{t_2\} \notin \{t_1\}$, 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.

Remark. Best-response dynamics in GNE. Example 3.3.3 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* [NRTV07]).

Thus, Example 3.3.3 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 4 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 3.3.4. A GNE σ is also an NE in \mathcal{G}_t for every $t \in \mathsf{Top}$.

Indeed, if σ was not an NE in \mathcal{G}_t for some $t \in \mathsf{Top}$, 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 3.2: Symmetric XOR game.

Example 3.3.5. CNE without NE. Consider the Symmetric XOR game \mathcal{G} depicted in Figure 3.2. 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. 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 $\boldsymbol{\sigma}$ there exists a player p that loses with $\boldsymbol{\sigma}$, which corresponds to WinTop^p_{\mathcal{G}}(\boldsymbol{\sigma}) = \emptyset but p can deviate and win \mathcal{G} , which corresponds to WinTop^p_{\mathcal{G}}(\boldsymbol{\sigma}[p \mapsto \boldsymbol{\sigma}'_p]) = \{t_1\}. Since $\emptyset \not\subseteq \{t_1\}$, then $\boldsymbol{\sigma}$ is not a CNE.

Chapter 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 4.1. 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 4.1. 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 [CD14], 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 [CD14]. 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 $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \delta, (\mathcal{O}_p)_{p \in \mathsf{Pla}} \rangle$ where $\mathsf{Pla}, \mathsf{S}, s_0$, Act and δ are the same as in concurrent games. For every player $p \in \mathsf{Pla}$, the set of observations $\mathcal{O}_p \subseteq 2^{\mathsf{S}}$ is a partition of S . We omit the acceptance condition, and we will include it explicitly in Theorem 4.2 below.

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

Strategies are observation based, that is, a strategy for Player p is a function σ_p : $\mathcal{O}_p^+ \to \mathsf{Act}$. 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 $\Sigma_{\mathcal{G}}^{\mathsf{Pla}}$ 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 $\sigma(h) = (\sigma_p(\text{obs}_p(h)))_{p \in \mathsf{Pla}} \in \mathsf{Act}^{\mathsf{Pla}}$, and we define $\operatorname{out}_{\mathcal{G}}(\sigma) \in \mathsf{S}^{\omega}$ similarly to concurrent games.

We say that Player $p \in \text{Pla}$ has perfect information if $\mathcal{O}_p = \{\{s\} \mid s \in S\}$. 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 $\alpha \subseteq S^{\omega}$, we say that α is visible to Player p if for every $\rho, \rho' \in S^{\omega}$ such that $obs_p(\rho) = obs_p(\rho')$ 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 [CD14] that will serve as the target of our reduction.

Theorem 4.2. Let $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \delta, (\mathcal{O}_p)_{p \in \mathsf{Pla}} \rangle$ be a partial information game, with $\mathsf{Pla} = \{1, 2, 3\}$ where Player 1 is less informed than Player 2. Let $\alpha \subseteq \mathsf{S}^{\omega}$ be parity objective over S . The problem of deciding whether $\exists \sigma_1 \in \Sigma_{\mathcal{G}}^1 \ \forall \sigma_2 \in \Sigma_{\mathcal{G}}^2 \ \exists \sigma_3 \in \Sigma_{\mathcal{G}}^3$ 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 4.2. 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 [BBM15], 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 $T \subseteq \mathsf{Top}$ in which Player p tries to win when playing σ'_p .

• Snake helps¹ 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 [Bok18]). Third, we need to separate the cases where Adam obeys Eve. In the following we give the complete construction, which overcomes these caveats.

4.3 Reduction to Partial Information Game

Consider an MTG $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathsf{Top}, (\delta_t)_{t \in \mathsf{Top}}, (\alpha_{t,p})_{t \in \mathsf{Top}, p \in \mathsf{Pla}} \rangle$. For every Player $p \in \mathsf{Pla}$, fix $T_p \subseteq \mathsf{Top}$ 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 $Q \subseteq S \times Pla \times 2^{\mathsf{Top}} \times \mathsf{Top} \times \{\mathsf{true}, \mathsf{false}\}$ is described in the following. A state $(s, p, T, t, b) \in Q$ comprises $s \in S$ which tracks the state of \mathcal{G} , a player $p \in \mathsf{Pla}$ that is controlled by Adam, a set $T \subseteq \mathsf{Top}$ of topologies that

 $^{^1\}mathrm{It}$ is arguable whether this matches the biblical interpretation. This work makes no theological claims.

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.

To make |Q| polynomial in $|\mathcal{G}|$ we restrict the component containing sets of topologies. Instead of allowing Adam to chose any $T \in 2^{\mathsf{Top}}$, that he intends to win with the selected player p, we observe that it is enough to consider only the sets that add a single topology to T_p , in case that Adam wants to show a profitable deviation, and all the sets containing a single topology, in the case that Adam wants to show that Eve fails to get the desired outcome in this topology. We define $\mathcal{T}_p = \{T_p \cup \{t\} \mid t \in \mathsf{Top}\} \subseteq 2^{\mathsf{Top}}$ and $\mathcal{T} = (\bigcup_{p \in \mathsf{Pla}} \mathcal{T}_p) \cup \{\{t\} \mid t \in \mathsf{Top}\}$. Note that $|\mathcal{T}| \leq (|\mathsf{Pla}| + 1) \cdot |\mathsf{Top}| \leq 2 \cdot |\mathsf{Pla}| \cdot |\mathsf{Top}|$. We now define $Q = \mathsf{S} \times \mathsf{Pla} \times \mathcal{T} \times \mathsf{Top} \times \{\mathsf{true}, \mathsf{false}\}$.

We now turn to define the transitions in \mathcal{H} . The actions are defined implicitly by the transitions.² From q_0 , Adam selects a player $p \in \mathsf{Pla}$ 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 $t \in \mathsf{Top}$ that \mathcal{G} will be played in. Then, \mathcal{H} transitions to state $(s_0, p, T, t, \mathsf{true}) \in Q$.

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 $a \in \operatorname{Act}^{\operatorname{Pla}}$ and Adam selects an action $a'_p \in \operatorname{Act}$. Then, the game transitions to state $(s', p, T, t, b') \in Q$ such that $s' = \delta_t(s, a[p \mapsto a'_p])$, 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.

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 $\operatorname{proj}(q) = s$. For every state $s \in S$ of \mathcal{G} , let $o_s = \{q \in Q \mid \operatorname{proj}(q) = s\} \subseteq Q$. The observation sets in \mathcal{H} are $\mathcal{O}_{Adam} = \mathcal{O}_{Eve} = \mathcal{O} = \{\{q_0\}\} \cup \{o_s \mid s \in S\}$. That is, Adam and Eve can observe the initial state q_0 , and for every $q \in Q$ they can only observe $\operatorname{proj}(q)$. 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 $\operatorname{proj}(\rho) = \operatorname{proj}(q_1)\operatorname{proj}(q_2)\dots$ (note that we skip the initial state q_0). We also define the predicate $\operatorname{obey}(\rho) = \bigwedge_{i\geq 1} b_i$, where b_i is the **true/false** bit of q_i . That is, $\operatorname{obey}(\rho)$ is true if and only if Adam always takes the actions suggested by Eve. When $\operatorname{obey}(\rho)$ is true, we say that Adam *obeys* Eve.

 $^{^{2}}$ In the model we describe, actions are identical for all players. However, the model of [CD14] 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 γ_{obs} : $\{q_0\} \cdot \mathcal{O}^{\omega} \to \mathsf{S}^{\omega}$ defined $\gamma_{obs}(\{q_0\}, o_{s_0}, o_{s_1}, \ldots) = s_0, s_1, \ldots$ Since $o_s = \{q \mid \operatorname{proj}(q) = s\}$ for every $s \in \mathsf{S}$, 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:

- $\gamma_{\text{Eve}} : \Sigma_{\mathcal{H}}^{\text{Eve}} \to \Sigma_{\mathcal{G}} \text{ defined by } \gamma_{\text{Eve}}(\sigma_{\text{Eve}}) = \sigma_{\text{Eve}} \circ \gamma_{\text{obs}}^{-1}$
- $\gamma_{\text{Adam}} : \Sigma_{\mathcal{H}}^{\text{Adam}} \to \bigcup_{p \in \text{Pla}} \{p\} \times \mathcal{T}_p \times \Sigma_{\mathcal{G}}^p$ 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} .
- $\gamma_{\text{Snake}} : \Sigma_{\mathcal{H}}^{\text{Snake}} \to \text{Top defined by } \gamma_{\text{Snake}}(\sigma_{\text{Snake}}) = \sigma_{\text{Snake}}(q_0) \text{ (recall that Snake only acts in } q_0).$

For readability, we omit the subscript and write γ instead of γ_{obs} , γ_{Adam} , γ_{Eve} , γ_{Snake} . The correct subscript can be deduced 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:

Lemma 4.3.1. 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 $\boldsymbol{\sigma} = \gamma(\sigma_{\text{Eve}})$, $(p, T, \sigma'_p) = \gamma(\sigma_{\text{Adam}})$ and $t = \gamma(\sigma_{\text{Snake}})$. Let $\rho = \text{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}})$, $\pi' = \text{out}_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$, and $\pi = \text{out}_{\mathcal{G}_t}(\boldsymbol{\sigma})$. Then $\text{proj}(\rho) = \pi'$. Furthermore, if Adam obeys Eve on ρ then $\text{proj}(\rho) = \pi = \pi'$.

Proof We prove by induction that for every $k \ge 1$, $\operatorname{proj}(\rho_{\le k+1}) = \pi'_{\le k}$, and if Adam obeys Eve then $\operatorname{proj}(\rho_{\le k+1}) = \pi'_{\le k} = \pi_{\le k}$. For k = 1, $\rho_{\le 2} = q_0$, (s_0, p, t, T, b_0) and $\pi'_{\le 1} = \pi_{\le 1} = s_0$ and we have that $\operatorname{proj}(\rho_{\le k+1}) = \pi'_{\le k}$. Assuming that $\operatorname{proj}(\rho_{\le k+1}) = \pi'_{\le k}$ for $k \ge 1$, the next state of $\operatorname{proj}(\rho)$ will depend on the transition function δ_t and action profile $\sigma[p \mapsto \sigma'_p](\pi'_{\le k})$ 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 $\operatorname{proj}(\rho_{\le k+2}) = \pi'_{\le k+1}$. 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 $\sigma[p \mapsto \sigma'_p](\pi'_{\le k}) = \sigma(\pi'_{\le k})$, and $\pi_{\le k+1} = \pi'_{\le k+1}$, thus, $\operatorname{proj}(\rho_{\le k+2}) = \pi_{\le k+1} = \pi'_{\le k+1}$. **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 $(T_p)_{p \in \mathsf{Pla}}$.

We define α as a conjunction of three conditions $\alpha = \{\rho \in q_0 \cdot Q^{\omega} \mid \psi_1(\rho) \land \psi_2(\rho) \land \psi_3(\rho)\}$, 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} .

- ψ₁(ρ) := t ∈ T. That is, ψ₁ forces Snake to choose a topology from the set of topologies selected by Adam.
- ψ₂(ρ) := (obey(ρ) ∧ t ∈ T_p) → proj(ρ) ∈ α_{t,p}. That is, ψ₂ is satisfied if whenever Adam obeys Eve then Player p wins in the topology t ∈ T_p selected by Snake.
- ψ₃(ρ) := T_p ⊊ T → proj(ρ) ∉ α_{t,p}. That is, ψ₃ 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 $\alpha_{t,p} = \operatorname{Parity}(\Omega_{t,p})$ such that $\Omega_{t,p} : \mathsf{S} \to \{0, \ldots, d\}$ is the parity ranking function, where $d \in \mathbb{N}$. We define a new ranking function $\Omega : Q_{\mathcal{H}} \to \{0, \ldots, d+1\}$, and show that $\alpha = \operatorname{Parity}(\Omega)$.

Observe that q_0 occurs only once in a play, so its priority 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.

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 ∉ T: In this case, if q ∈ Inf(ρ) then ρ does not satisfy ψ₁, thus, ρ ∉ α. We set Ω(q) = 1 to get ρ ∉ Parity(Ω).
- t ∈ T, b = true, t ∈ T_p and T_p ⊊ T: In this case, if q ∈ Inf(ρ) then ρ satisfies the preconditions of both ψ₂ and ψ₃, but the postconditions of ψ₂ and ψ₃ contradict, thus, ρ ∉ α. We set Ω(q) = 1 to get ρ ∉ Parity(Ω).
- $t \in T$, $b = \text{true} \land t \in T_p$ and $\neg(T_p \notin T)$: 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 $\text{proj}(\rho)$.

- $t \in T$, $\neg(b = \text{true} \land t \in T_p)$ and $T_p \not\subseteq T$: 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 ∈ T, ¬(b = true ∧ t ∈ T_p) and ¬(T_p ⊊ T): In this case, if q ∈ Inf(ρ) then ψ₂ and ψ₃ are vacuously satisfied, and ρ ∈ α. So we set Ω(q) = 0 to get that ρ ∈ Parity(Ω).

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

Lemma 4.3.2. Consider an MTG $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathsf{Top}, (\delta_t)_{t \in \mathsf{Top}}, (\alpha_{t,p})_{t \in \mathsf{Top}, p \in \mathsf{Pla}} \rangle$. Let $(T_p)_{p \in \mathsf{Pla}}$ be sets of topologies for each player and let \mathcal{H} be the corresponding partialinformation game. There exists a strategy profile σ in \mathcal{G} such that σ is a CNE and for every $p \in \mathsf{Pla}$ we have $\mathsf{WinTop}_{\mathcal{G}}^p(\sigma) = T_p$ if and only if the following holds:

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

Proof Assume $\boldsymbol{\sigma}$ is a CNE in \mathcal{G} such that for every $p \in \text{Pla}$, $\text{WinTop}_{\mathcal{G}}^{p}(\boldsymbol{\sigma}) = T_{p}$, and fix $\sigma_{\text{Eve}} = \gamma^{-1}(\boldsymbol{\sigma})$ 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 ¬(T_p ⊊ T), then ψ₃ is satisfied vacuously. Choose t ∈ T for Snake, then ψ₁ is satisfied. If Adam does not obey Eve or t ∉ T_p then ψ₂ is vacuously satisfied. Otherwise, if Adam obeys Eve and t ∈ T_p, let ρ = out_H(σ_{Eve}, σ_{Adam}, σ_{Snake}). In order to show that ψ₂ is satisfied we need to show that proj(ρ) ∈ α_{t,p}. Let π = out_{Gt}(σ). Since T_p = WinTop^p_G(σ) and t ∈ T_p we have that π ∈ α_{t,p}. From Lemma 4.3.1 we have that proj(ρ) = π, so we get that proj(ρ) ∈ α_{t,p}, as required.
- If $T_p \notin T$, denote $T' = \operatorname{WinTop}_{\mathcal{G}}^p(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$. Since $\boldsymbol{\sigma}$ is a CNE, we have that $\neg(T_p \notin T')$, so $T \smallsetminus T' \neq \emptyset$, as otherwise we would have that $T_p \notin T \subseteq T'$. Choose $t \in T \smallsetminus T'$ for Snake, then ψ_1 is satisfied. Let $\rho = \operatorname{out}_{\mathcal{H}}(\sigma_{\mathsf{Eve}}, \sigma_{\mathsf{Adam}}, \sigma_{\mathsf{Snake}}), \pi' = \operatorname{out}_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ and $\pi = \operatorname{out}_{\mathcal{G}_t}(\boldsymbol{\sigma})$. From Lemma 4.3.1 we have that $\operatorname{proj}(\rho) = \pi'$ and if Adam obeys Eve then we have $\operatorname{proj}(\rho) = \pi = \pi'$. Note that since $t \notin T' = \operatorname{WinTop}_{\mathcal{G}}^p(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ 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 $T' = T_p = \operatorname{WinTop}_{\mathcal{G}}^p(\boldsymbol{\sigma})$, 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 $\operatorname{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}, \sigma_{\text{Snake}}) \in \alpha$. Let $\boldsymbol{\sigma} = \gamma(\sigma_{\text{Eve}})$. We start by showing that for every $p \in \text{Pla}$ it holds that $\operatorname{WinTop}_{\mathcal{G}}^{p}(\boldsymbol{\sigma}) = T_{p}$. 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 4.3.1 we have $\text{proj}(\rho) = \text{out}_{\mathcal{G}_t}(\boldsymbol{\sigma})$. Thus, $t \in \text{WinTop}_{\mathcal{G}}^p(\boldsymbol{\sigma})$.

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 4.3.1 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 $\boldsymbol{\sigma}$ is a CNE. Assume by way of contradiction that there exists a player $p \in \mathsf{Pla}$ with a beneficial deviation $\sigma'_p \in \Sigma^p_{\mathcal{G}}$. That is, $T' = \mathrm{WinTop}^p_{\mathcal{G}}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ satisfies $T_p \notin T'$. 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 $\sigma_{\mathrm{Adam}} = \gamma^{-1}(p, T, \sigma'_p)$. Consider a strategy σ_{Snake} , denote $t = \gamma(\sigma_{\mathrm{Snake}})$ and let $\rho = \mathrm{out}_{\mathcal{H}}(\sigma_{\mathrm{Eve}}, \sigma_{\mathrm{Adam}}, \sigma_{\mathrm{Snake}})$. By Lemma 4.3.1 we have $\mathrm{proj}(\rho) = \mathrm{out}_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$, and because $t \in T \subseteq \mathrm{WinTop}^p_{\mathcal{G}}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ it holds that $\mathrm{proj}(\rho) \in \alpha_{t,p}$. However, $T_p \notin T$, so ψ_3 is violated, and $\rho \notin \alpha$, which is a contradiction. We conclude that $\boldsymbol{\sigma}$ is a CNE.

Using Lemma 4.3.2 we can decide whether a given MTG \mathcal{G} has a CNE, by iterating over all possible sets of candidate winning topologies $(T_p)_{p \in \mathsf{Pla}}$, and repeatedly applying the reduction, and using the decision procedure of Theorem 4.2. 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, $|\mathcal{Q}| \leq |\mathsf{S}| \cdot |\mathsf{Pla}| \cdot |\mathcal{T}| \cdot |\mathsf{Top}| \cdot 2$ where $|\mathcal{T}| \leq 2|\mathsf{Pla}||\mathsf{Top}|$. 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 1). From Theorem 4.2, solving \mathcal{H} takes double-exponential time in $|\mathcal{G}|$. In the worst case, we will iterate over all $2^{|\mathsf{Top}|\cdot|\mathsf{Pla}|}$ options for $(T_p)_{p\in\mathsf{Pla}}$, which is exponential in $|\mathcal{G}|$. Repeating the double-exponential procedure an exponential number of times results in a double-exponential algorithm. This completes the proof of Theorem 4.1.

Remark. Lower bounds and improving the upper bound. We do not have a lower bound for the 2-EXPTIME complexity of Theorem 4.1. 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 4.2 from [CD14]. Unfortunately, the decision procedure in [CD14] goes through three nontrivial reductions, one of which involves Safra's determinization, that is notoriously difficult to analyze: The first reduction [CD10; CD14] transforms the objective to a visible objective for Adam which involves the determinization of a parity automaton. The second reduction [CD14] reduces the three-player partial-information game into a two-player partial-information game. The third reduction uses the results of [RHDC07] to reduce the two-player partial-information game to a two-player perfect-information game.

Therefore, it is likely that improving the bound (if possible) will involve devising an ad-hoc procedure, possibly using some key ideas from [CD10; CD14; RHDC07].

Chapter 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 $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathsf{Top}, (\delta_t)_{t \in \mathsf{Top}}, (\alpha_{t,p})_{t \in \mathsf{Top}, p \in \mathsf{Pla}} \rangle$, a profile $\boldsymbol{\sigma} \in \Sigma_{\mathcal{G}}^{\mathsf{Pla}}$ is a GNE if for every $p \in \mathsf{Pla}, \sigma'_p \in \Sigma_{\mathcal{G}}$ and $t \in \mathsf{Top}$, if $p \in \mathrm{Win}_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ then $p \in \mathrm{Win}_{\mathcal{G}_t}(\boldsymbol{\sigma})$.

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

Similarly to Chapter 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 [CD10].

Theorem 5.2. Let $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \delta, (\mathcal{O}_p)_{p \in \mathsf{Pla}} \rangle$ with $\mathsf{Pla} = \{1, 2\}$. Let $\alpha \subseteq \mathsf{S}^{\omega}$ be a parity objective. The problem of deciding whether $\exists \sigma_1 \in \Sigma^1_{\mathcal{G}} \ \forall \sigma_2 \in \Sigma^2_{\mathcal{G}} \ \operatorname{out}_{\mathcal{G}}(\sigma_1, \sigma_2) \in \alpha$ is EXPTIME-complete.

5.1 Reduction to Partial Information Game

Consider an MTG $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathsf{Top}, (\delta_t)_{t \in \mathsf{Top}}, (\alpha_{t,p})_{t \in \mathsf{Top}, p \in \mathsf{Pla}} \rangle$. For every Player $p \in \mathsf{Pla}$ fix $T_p \subseteq \mathsf{Top}$ 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 $Q = S \times Pla \times Top \times \{true, false\}$ is described in the following. A state $(s, p, t, b) \in Q$ comprises of $s \in S$ which tracks the state of \mathcal{G} , a player $p \in Pla$ that is controlled by Adam, a topology $t \in Top$ that Adam picks, and a bit $b \in \{true, false\}$ 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 $p \in \mathsf{Pla}$ to control and a topology $t \in \mathsf{Top}$ that \mathcal{G} will be played in. Then, \mathcal{H} transitions to state $(s_0, p, t, \mathsf{true}) \in Q$. Henceforth, p and t remain fixed throughout the play. From state $(s, p, t, b) \in Q$, Eve chooses an action profile $a \in \mathsf{Act}^{\mathsf{Pla}}$, and Adam selects an action $a'_p \in \mathsf{Act}$ and \mathcal{H} transitions to state $(s', p, t, b') \in Q$ such that $s' = \delta_t(s, a[p \mapsto a'_p])$, and $b' = b \land (a'_p = a_p)$.

The observation sets for the players, proj and obey are defined similar to Section 4.3. Correspondence between \mathcal{H} and \mathcal{G} , γ_{obs} , γ_{Eve} is defined in the same way as in Section 4.3, and $\gamma_{\text{Adam}} : \Sigma_{\mathcal{H}}^{\text{Adam}} \to \bigcup_{p \in \text{Pla}} \{p\} \times \text{Top} \times \Sigma_{\mathcal{H}}^{p}$ 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 4.3.1.

Lemma 5.1.1. Consider strategies $\sigma_{\text{Eve}} \in \Sigma_{\mathcal{H}}^{\text{Eve}}$ and $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$. Let $\boldsymbol{\sigma} = \gamma(\sigma_{\text{Eve}})$ and $(p, t, \sigma'_p) = \gamma(\sigma_{\text{Adam}})$. Let $\rho = \operatorname{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}) \pi' = \operatorname{out}_{\mathcal{G}_t}(\boldsymbol{\sigma}[p \mapsto \sigma'_p])$ and $\pi = \operatorname{out}_{\mathcal{G}_t}(\boldsymbol{\sigma})$. Then, $\operatorname{proj}(\rho) = \pi'$. Furthermore, if Adam obeys Eve on ρ then $\operatorname{proj}(\rho) = \pi = \pi'$.

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

- $\psi_1(\rho) \coloneqq (\operatorname{obey}(\rho) \land t \in T_p) \to \operatorname{proj}(\rho) \in \alpha_{t,p}.$
- $\psi_2(\rho) \coloneqq t \notin T_p \to \operatorname{proj}(\rho) \notin \alpha_{t,p}$.

 α can be expressed as a parity objective as follows. For every $t \in \mathsf{Top}$, $p \in \mathsf{Pla}$, let $\Omega_{t,p} : \mathsf{S} \to \{0, ..., d_{t,p}\}$ be the priority function for the parity objective $\alpha_{t,p}$ in \mathcal{G} . We construct a priority function $\Omega : Q_{\mathcal{H}} \to \{0, ..., d\}$ such that $d = \max\{d_{t,p} + 1 \mid t \in \mathsf{Top}, p \in \mathsf{Pla}\}$. 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 $\operatorname{proj}(\rho) \notin \alpha_{t,p}$. 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 $\operatorname{proj}(\rho) \in \alpha_{t,p}$. 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 5.1.2. There exists a GNE $\boldsymbol{\sigma} \in \Sigma_{\mathcal{G}}$ in \mathcal{G} with $\operatorname{WinTop}_{\mathcal{G}}^{p}(\boldsymbol{\sigma}) = T_{p}$ for every $p \in \mathsf{Pla}$, if and only if $\exists \sigma_{\mathsf{Eve}} \in \Sigma_{\mathcal{H}}^{\mathsf{Eve}} \ \forall \sigma_{\mathsf{Adam}} \in \Sigma_{\mathcal{H}}^{\mathsf{Adam}} \operatorname{out}_{\mathcal{H}}(\sigma_{\mathsf{Eve}}, \sigma_{\mathsf{Adam}}) \in \alpha$.

Proof Let $\boldsymbol{\sigma} \in \Sigma_{\mathcal{G}}$ be a GNE with $\operatorname{WinTop}_{\mathcal{G}}^{p}(\boldsymbol{\sigma}) = T_{p}$ for every $p \in \operatorname{Pla}$. Let $\sigma_{\operatorname{Eve}} \in \Sigma_{\mathcal{H}}^{\operatorname{Eve}}$ be the corresponding strategy for $\boldsymbol{\sigma}$, and let $\sigma_{\operatorname{Adam}} \in \Sigma_{\mathcal{H}}^{\operatorname{Adam}}$ be some strategy for Adam that corresponds to (p, t, σ'_{p}) . Let $\rho = \operatorname{out}_{\mathcal{H}}(\sigma_{\operatorname{Eve}}, \sigma_{\operatorname{Adam}})$. If $\operatorname{obey}(\rho) \wedge t \in T_{p}$, then from Lemma 5.1.1 we have that $\operatorname{proj}(\rho) = \operatorname{out}_{\mathcal{G}_{t}}(\boldsymbol{\sigma})$, and since $t \in T_{p} = \operatorname{WinTop}_{\mathcal{G}}^{p}(\boldsymbol{\sigma})$ then $\operatorname{out}_{\mathcal{G}_{t}}(\boldsymbol{\sigma}) \in \alpha_{t,p}$. Thus, ψ_{1} is satisfied by ρ . If $t \notin T_{p}$ then from Lemma 5.1.1 we have that $\operatorname{proj}(\rho) = \operatorname{out}_{\mathcal{G}_{t}}(\boldsymbol{\sigma}[p \mapsto \sigma'_{p}])$ and since Player p is losing in t when \mathcal{G} is played with $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}$ is a GNE, then $\operatorname{out}_{\mathcal{G}_{t}}(\boldsymbol{\sigma}[p \mapsto \sigma'_{p}]) \notin \alpha_{t,p}$. 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 $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$, $\operatorname{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}}) \in \alpha$. Let $\sigma \in \Sigma_{\mathcal{G}}$ correspond to σ_{Eve} . We show that σ is a GNE. First, we show that for every $p \in \text{Pla}$, $\operatorname{WinTop}_{\mathcal{G}}^{p}(\sigma) = T_{p}$. 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 $\rho_{t} = \operatorname{out}_{\mathcal{G}_{t}}(\sigma)$ and $\rho = \operatorname{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}})$. We have that $\rho \in \alpha$. Since Adam obeys Eve on ρ , from Lemma 5.1.1 we have that $\operatorname{proj}(\rho) = \rho_{t}$. If $t \in T_{p}$ then from ψ_{1} we get that $\rho_{t} = \operatorname{proj}(\rho) \in \alpha_{t,p}$, thus, $t \in \operatorname{WinTop}_{\mathcal{G}}^{p}(\sigma)$. If $t \notin T_{p}$ then from ψ_{2} we get that $\rho_{t} = \operatorname{proj}(\rho) \notin \alpha_{t,p}$, thus, $t \notin \operatorname{WinTop}_{\mathcal{G}}^{p}(\sigma)$. So we get that $\operatorname{WinTop}_{\mathcal{G}}^{p}(\sigma) = T_{p}$. Now, we show that σ is a GNE. Let $p \in \operatorname{Pla}$, $\sigma'_{p} \in \Sigma_{\mathcal{G}}^{p}$ and $t \in \operatorname{Top}$ such that $t \notin T_{p}$. Let $\sigma_{\text{Adam}} \in \Sigma_{\mathcal{H}}^{\text{Adam}}$ correspond to (p, t, σ'_{p}) , and let $\rho = \operatorname{out}_{\mathcal{H}}(\sigma_{\text{Eve}}, \sigma_{\text{Adam}})$. We have that $\rho \in \alpha$, thus, since $t \notin T_{p}$ then $\operatorname{proj}(\rho) \notin \alpha_{t,p}$. From Lemma 5.1.1 we have that $\rho'_{t} = \operatorname{out}_{\mathcal{G}_{t}}(\sigma[p \mapsto \sigma'_{p}]) = \operatorname{proj}(\rho) \notin \alpha_{t,p}$, thus, $t \notin \operatorname{WinTop}_{\mathcal{G}_{t}}^{p}(\sigma[p \mapsto \sigma'_{p}]) = T_{p}$, so σ is a GNE.

The algorithm for solving the GNE existence problem is the following. For every $(T_p)_{p\in\mathsf{Pla}} \in (2^{\mathsf{Top}})^{\mathsf{Pla}}$ we construct \mathcal{H} from \mathcal{G} and $(T_p)_{p\in\mathsf{Pla}}$. Then, we check if there exists if there exists a winning strategy for Eve in \mathcal{H} . If there is such a strategy, then, according to Lemma 5.1.2, its corresponding strategy in \mathcal{G} is a GNE and the algorithm returns that there exists a GNE in \mathcal{G} . If we went through all the sets $(T_p)_{p\in\mathsf{Pla}} \in (2^{\mathsf{Top}})^{\mathsf{Pla}}$ without finding a GNE, then the algorithm returns that there is no GNE in \mathcal{G} .

The size of \mathcal{H} is polynomial in the size of \mathcal{G} . We copy each $s \in S$ for every combination of $p \in \mathsf{Pla}$, $t \in \mathsf{Top}$, $b \in \{\mathsf{true}, \mathsf{false}\}$, so we get $|Q_{\mathcal{H}}| = 2 \cdot |S| \cdot |\mathsf{Pla}| \cdot |\mathsf{Top}| + 1$, 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 $2^{|\mathsf{Top}|\cdot|\mathsf{Pla}|}$ 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 5.2 this takes exponential time in $|\mathcal{G}|$, so the GNE existence problem is in EXPTIME.

We now present a sketch of the proof for the correctness of the algorithm above. Then, we can conclude Theorem 5.1.

Proof sketch As in Section 4.3, we first fix a set of "intended" winning topologies $T_p \subseteq$ Top for each player $p \in \mathsf{Pla}$. Then, we ask whether \mathcal{G} admits a GNE σ in which WinTop^p_{\mathcal{G}}(\sigma) = T_p for every $p \in \mathsf{Pla}$. We then construct a 2-player partial-information game whose players are Adam and Eve, where Eve 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 $p \in \mathsf{Pla}$ and a *single topology* $t \in \mathsf{Top}$ where pattempts 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).

Chapter 6

Strategy Logic with Imperfect Information

In this section we discuss solving the GNE and CNE existence problems using strategy logic with imperfect information, SL_{ii} , introduced in [BMM⁺21]. SL_{ii} is an expressive logic that is generally undecidable, but a decidable fragment, called *hierarchical in*stances, can capture the GNE and CNE existence problems. The complexity of SL_{ii} model-checking for hierarchical instances depends on a parameter called the *simulation depth*. SL_{ii} model-checking for formulas with simulation depth up to k is (k + 1)-EXPTIME-complete, and the procedure suggested in [BMM⁺21] is (k+1)-EXPTIME. Our formulation of the GNE and CNE existence problems with SL_{ii} , has a simulation depth of 2 for both problems, resulting in 3-EXPTIME procedure for solving those problems. This is a worse complexity result than the EXPTIME and 2-EXPTIME results that we got in Chapter 5 and Chapter 4, respectively. It might be possible that there is an alternative formulation with a lower simulation depth, resulting in a lower complexity for this approach.

The section is organized as follows. In Section 6.1 we give a short overview of SL_{ii} . In Section 6.2 we discuss how to convert a multi-topology game to a model called *concurrent game structure with imperfect information* that SL_{ii} is interpreted over. Then, in Section 6.3 we formalize the GNE existence problem with SL_{ii} and compute it's simulation depth. In Section 6.4 we do the same for the CNE existence problem.

6.1 Overview of SL_{ii}

 SL_{ii} formulas are defined over a number of fixed parameters – a set of *atomic propositions* AP, a set of *players* (or *agents*) Ag, a set of *strategy variables* Var and a set of *observation symbols* Obs. SL_{ii} formulas are interpreted over a *Concurrent Game Structure with Imperfect Information*, abbreviated CGS_{ii} . A CGS_{ii} is a tuple $\mathcal{G} = \langle Ac, V, E, \mathcal{L}, v_0, \mathcal{O} \rangle$ such that Ac is a set of *actions*, V is a set of *states*, $E : V \times Ac^{Ag} \rightarrow V$ is a *transition function*, $\mathcal{L} : V \rightarrow 2^{AP}$ is a *labelling function*, $v_0 \in V$ is an *initial state* and $\mathcal{O} : Obs \rightarrow 2^{V \times V}$ is an observation interpretation, which maps each observation symbol $o \in Obs$ to an equivalence relation over the states $\mathcal{O}(o) \subseteq V \times V$. SL_{ii} has the following syntax:

$$\begin{split} \varphi &\coloneqq p \mid \neg \varphi \mid \varphi \lor \varphi \mid \langle \! \langle x \rangle \! \rangle^o \varphi \mid (a, x) \varphi \mid (a, ?) \varphi \mid \mathrm{E}\psi; \ p \in \mathsf{AP}, x \in \mathsf{Var}, a \in \mathsf{Ag} \\ \psi &\coloneqq \varphi \mid \neg \psi \mid \psi \lor \psi \mid \mathrm{X}\psi \mid \psi \mathrm{U}\psi; \end{split}$$

Formulas of type φ are called *state formulas* and formulas of type ψ are called *path formulas*. The boolean and temporal operators \neg, \lor, X, U have their usual semantics. The syntax is extended with the boolean and temporal operators \land, \rightarrow, F, G that can be expressed with the operators already in the syntax. The *existential strategy quantifier* $\langle\!\langle x \rangle\!\rangle^o \varphi$ means, "there exists a strategy x over the observations $\mathcal{O}(o)$ that satisfies φ ". The syntax is extended with a *universal strategy quantifier* defined $[\![x]\!]^o \varphi := \neg \langle\!\langle x \rangle\!\rangle^o \neg \varphi$. The binding operator (a, x) binds strategy x to player a and the *unbinding operator* (a, ?) unbinds player a from it's current strategy. The *existential outcome quantifier* $E\psi$ means "there exists an outcome of the current strategy assignment that satisfies ψ ". The syntax is extended with a *universal outcome quantifier* defined $A\psi := \neg E \neg \psi$. For a full description of the semantics of SL_{ii} we refer readers to [BMM⁺21].

An SL_{ii} instance is a pair (\mathcal{G}, Φ) where \mathcal{G} is a CGS_{ii} and Φ is an SL_{ii} state formula. In general, SL_{ii} is undecidable. But, a fragment called *hierarchical instances* is decidable. An *hierarchical instance* is such that as we go down the syntax tree of the formula, observations only get finer.

The complexity of the model-checking problem for an hierarchical SL_{ii} instance (\mathcal{G}, Φ) depends on the *simulation depth* of (\mathcal{G}, Φ) . The simulation depth is computed recursively on the formula's structure. The complexity of the model-checking procedure for an instance with simulation depth k is (k + 1)-EXPTIME. For a description of how to compute the simulation depth we refer readers to [BMM⁺21].

6.2 MTG to CGS_{ii}

In this section we show how to translate an MTG to a CGS_{ii} and a set of formulas that describe the players winning conditions.

Let $\mathcal{G} = \langle \mathsf{Pla}, \mathsf{S}, s_0, \mathsf{Act}, \mathsf{Top}, (\delta_t)_{t \in \mathsf{Top}}, (\alpha_{t,p})_{t \in \mathsf{Top}, p \in \mathsf{Pla}} \rangle$ be an MTG. We denote the players $\mathsf{Pla} = \{p_1 \dots p_n\}$. First, we fix the parameters over which the SL_{ii} formulas are defined, AP , Ag , Var and Obs . The set of atomic propositions is such that we can encode each state and each topology with a unique label (a subset of AP). This will enable us to write the LTL formula $\psi_{t,p}$ for every $t \in \mathsf{Top}$ and $p \in \mathsf{Pla}$ which means that the topology t is played and p's objective is satisfied. The set of agents is $\mathsf{Ag} = \mathsf{Pla} \cup \{T\}$ where T is the topology player that selects the topology. The set of strategy variables is $\mathsf{Var} = \{\sigma_p \mid p \in \mathsf{Pla}\} \cup \{\sigma'_p \mid p \in \mathsf{Pla}\}$. Since all players have the same observation sets (i.e., can observe the state, but not the topology), we only need a single observation symbol o. Note that every SL_{ii} instance with a single observation symbol is inherently

hierarchical.

The CGS_{ii} that we use is $\mathcal{H} = \langle \mathsf{Ac}, \mathsf{V}, \mathsf{E}, \mathcal{L}, v_0, \mathcal{O} \rangle$. The actions in \mathcal{H} are the actions in \mathcal{G} together with actions for T that enable him to select the topology in the first turn of the game. The states of \mathcal{H} are $\mathsf{V} = (\mathsf{S} \times \mathsf{Top}) \cup \{v_0\}$, where v_0 is the initial state where T selects the topology. The transition function corresponds to the transition function of \mathcal{G} , and allowing T to select the topology from the initial state v_0 . The observation symbol o is interpreted such that v_0 is distinguishable from all other states and $((s,t), (s',t')) \in \mathcal{O}(o)$ (that is, (s,t) and (s',t') are indistinguishable) if and only if s = s'.

6.3 Expressing GNE Existence Problem with SL_{ii}

The following formula expresses the GNE existence problem in SL_{ii} :

$$\langle\!\langle \boldsymbol{\sigma} \rangle\!\rangle^{o}(\mathsf{Pla}, \boldsymbol{\sigma}) \bigwedge_{p \in \mathsf{Pla}} \Biggl[\llbracket \sigma'_{p} \rrbracket^{o} \Biggl(\bigwedge_{t \in \mathsf{Top}} \Bigl(\mathsf{E} \psi_{t,p} \lor \neg (p, \sigma'_{p}) \mathsf{E} \psi_{t,p} \Bigr) \Biggr) \Biggr]$$

Where $\langle\!\langle \boldsymbol{\sigma} \rangle\!\rangle^{o} \coloneqq \langle\!\langle \sigma_{p_{1}} \rangle\!\rangle^{o} \ldots \langle\!\langle \sigma_{p_{n}} \rangle\!\rangle^{o}$ is a shorthand way of writing "there exists a strategy profile". Similarly, $(\mathsf{Pla}, \boldsymbol{\sigma}) \coloneqq (p_{1}, \sigma_{p_{1}}) \ldots (p_{n}, \sigma_{p_{n}})$ is binding the strategy profile to the players. When all players except for the topology player T are bound to a strategy, the formula $\mathrm{E}\psi_{p,t}$ means that player p wins in topology t under the given strategy assignment. After we quantify over strategy profiles, we require that for every player p in \mathcal{G} , every strategy σ'_{p} and every topology t, either player p wins topology t when players are assigned strategy profile $\boldsymbol{\sigma}$ or player p loses topology t when she changes her strategy to σ'_{p} .

Simulation depth Now, we compute the simulation depth of the instance. The computation involves two parameters – first is the current simulation depth $k \in \mathbb{N}$ and the second is a parameter that can be either nd or alt. The computation is performed according to Section 5.2 in [BMM⁺21]. Quantifying an LTL formula with E gives the simulation depth (0, nd). Thus, $sd(E\psi_{t,p}) = (0, nd)$. Binding a strategy to a player does not change the simulation depth, so we have $sd((p, \sigma'_p)E\psi_{t,p}) = (0, nd)$. Negating a formula keeps the current simulation depth the same and sets the second parameter to alt. Thus, $sd(\neg(p, \sigma'_p)E\psi_{t,p}) = (0, alt)$. Taking a disjunction between two formulas results in the maximum of each parameter of the subformulas (where nd < alt), thus, $sd(E\psi_{t,p} \lor \neg(p, \sigma'_p)E\psi_{t,p}) = (0, alt)$. The conjunction over all the topologies translates into a negation, disjunction and another negation. Since each subformula φ has $sd(\varphi) = (0, alt)$, we have that:

$$sd\left(\bigwedge_{t\in\mathsf{Top}}\left(\mathrm{E}\psi_{t,p}\vee\neg(p,\sigma'_p)\mathrm{E}\psi_{t,p}\right)\right)=(0,\mathtt{alt})$$

The universal strategy quantifier translates into a negation that does not change the simulation depth, an existential strategy quantifier that increases the first parameter by 1 and sets the second parameter to nd and another negation that sets the second parameter to alt. So we have that

$$sd\left(\llbracket\sigma_p'\rrbracket^o\left(\bigwedge_{t\in\mathsf{Top}}\left(\mathrm{E}\psi_{t,p}\vee\neg(p,\sigma_p')\mathrm{E}\psi_{t,p}\right)\right)\right)=(1,\mathtt{alt})$$

Binding the strategy profile to the players has no effect and the universal strategy quantifier increases the first parameter by 1 and sets the second to nd, thus,

$$sd\left(\langle\!\langle \boldsymbol{\sigma} \rangle\!\rangle^{o}(\mathsf{Pla},\boldsymbol{\sigma}) \bigwedge_{p \in \mathsf{Pla}} \left[\left[\!\left[\boldsymbol{\sigma}_{p}^{\prime} \right]\!\right]^{o} \left(\bigwedge_{t \in \mathsf{Top}} \left(\mathrm{E}\psi_{t,p} \lor \neg(p,\sigma_{p}^{\prime}) \mathrm{E}\psi_{t,p} \right) \right) \right] \right) = (2,\mathsf{nd})$$

Making model-checking complexity of the instance to be 3-EXPTIME.

6.4 Expressing CNE Existence Problem with SL_{ii}

The following formula expresses the CNE existence problem in SL_{ii} :

$$\langle\!\langle \boldsymbol{\sigma} \rangle\!\rangle^o(\mathsf{Pla},\boldsymbol{\sigma}) \bigwedge_{p \in \mathsf{Pla}} \xi_p$$
$$\xi_p \coloneqq \left[\!\left[\sigma_p' \right]\!\right]^o \left(\left(\bigwedge_{t \in \mathsf{Top}} \left(\mathrm{E}\psi_{t,p} \lor \neg(p,\sigma_p') \mathrm{E}\psi_{t,p} \right) \right) \lor \left(\bigvee_{t \in \mathsf{Top}} \left(\mathrm{E}\psi_{t,p} \land \neg(p,\sigma_p') \mathrm{E}\psi_{t,p} \right) \right) \right)$$

The formula for CNE is similar to the formula for GNE. We change the subformula $\wedge_{t\in \mathsf{Top}} (\mathrm{E}\psi_{t,p} \vee \neg (p, \sigma'_p) \mathrm{E}\psi_{t,p})$, which means that for every topology t, player p does not improve her outcome by switching to strategy σ'_p , by taking a disjunction with $\bigvee_{t\in\mathsf{Top}} (\mathrm{E}\psi_{t,p} \wedge \neg (p, \sigma'_p) \mathrm{E}\psi_{t,p})$, which means that there exists a topology where player p wins, and loses if she changes her strategy to σ'_p .

Simulation depth The simulation depth of the formula $\wedge_{t\in \mathsf{Top}} (\mathrm{E}\psi_{t,p} \vee \neg(p, \sigma'_p) \mathrm{E}\psi_{t,p})$ and the formula $\bigvee_{t\in \mathsf{Top}} (\mathrm{E}\psi_{t,p} \wedge \neg(p, \sigma'_p) \mathrm{E}\psi_{t,p})$ is the same and is equal to $(0, \mathsf{alt})$. Thus, the conjunction of the two results in a formula with simulation depth $(0, \mathsf{alt})$. The next steps in the computation of the simulation depth are identical to the computations for GNE, making the simulation depth be $(2, \mathsf{nd})$ and the model-checking complexity to be in 3-EXPTIME.

Chapter 7

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 explored the relationships and properties of these notions of NE. In the solution for the CNE existence problem we showed a novel technique of reducing the problem to a three player partial information game. This shows that MTGs can be seen as a restricted form of partial information games that might be more useful for modelling systems with this restricted form of partial information. 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 players can obtain by cooperating. Since in MTGs the winning sets of topologies may be incomparable, we formulate this as follows: given sets $(T_p)_{p \in \mathsf{Pla}}$, is there a strategy profile σ such that $\operatorname{WinTop}_{G}^{p}(\sigma) = T_p$ for every $p \in \mathsf{Pla}$?

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 (Chapter 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 5, 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_{\parallel}^{NP} -hardness [BBM15]. 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 1.0.1 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 In Chapter 6 we showed that logic for partial information games [BMM⁺21; FS10; Mau14] can be used for solving the GNE and CNE existence problems. It turns out, while this approach can be described with a more straightforward formula than our solution, the complexity upper bounds it gives are 3-EXPTIME for both problems. Moreover, writing the formula essentially requires an understanding of the approach we take in the paper. It remains an open question if it is possible to improve this upper bound using a more elaborate analysis. This approach might be more easily extended to other notions of NE, utilizing the expressivity of the logic.

Combinatorial topology Combinatorial topology is a useful tool for reasoning about game theoretic concepts and distributed computing [RR22]. A possible future research direction would be to investigate MTGs through the lens of combinatorial topology and to see if it offers interesting insights about the model.

Bibliography

- [AAK15] Shaull Almagor, Guy Avni, and Orna Kupferman. Repairing multi-player games. In 26th International Conference on Concurrency Theory (CON-CUR 2015). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2015.
- [AGM02] Hagit Attiya, Alla Gorbach, and Shlomo Moran. Computing in totally anonymous asynchronous shared memory systems. *Information and Com*putation, 173(2):162–183, 2002.
- [Alm20] Shaull Almagor. Process symmetry in probabilistic transducers. In 40th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, 2020.
- [BBM15] Patricia P Bouyer, Romain Brenguier, and Nicolas N Markey. Pure nash equilibria in concurrent games. *Logical methods in computer science*, 2015.
- [BFH11] Felix Brandt, Felix Fischer, and Markus Holzer. Equilibria of graphical games with symmetries. *Theoretical Computer Science*, 412(8-10):675– 685, 2011.
- [BMM⁺21] 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.
- [BMV17] Patricia Bouyer, Nicolas Markey, and Steen Vester. Nash equilibria in symmetric graph games with partial observation. Information and Computation, 254:238–258, 2017.
- [Bok18] Udi Boker. Why these automata types? In *LPAR*, volume 18, pages 143–163, 2018.
- [BPRS17] 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. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017.
- [CD10] Krishnendu Chatterjee and Laurent Doyen. The complexity of partialobservation parity games. In International Conference on Logic for Programming Artificial Intelligence and Reasoning, pages 1–14. Springer, 2010.

[CD14] Krishnendu Chatterjee and Laurent Doyen. Games with a weak adversary. In International Colloquium on Automata, Languages, and Programming, pages 110–121. Springer, 2014. [CEFJ96] 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. [CHP10] Krishnendu Chatterjee, Thomas A Henzinger, and Nir Piterman. Strategy logic. Information and Computation, 208(6):677-693, 2010. $[DDG^+10]$ 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. [DHK07] Luca De Alfaro, Thomas A Henzinger, and Orna Kupferman. Concurrent reachability games. Theoretical computer science, 386(3):188–217, 2007. [DP07] Constantinos Daskalakis and Christos Papadimitriou. Computing equilibria in anonymous games. In 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07), pages 83–93. IEEE, 2007. [ES96] E Allen Emerson and A Prasad Sistla. Symmetry and model checking. Formal methods in system design, 9(1):105-131, 1996. [FGR18] Emmanuel Filiot, Raffaella Gentilini, and Jean-François Raskin. Rational synthesis under imperfect information. In Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, pages 422–431, 2018.[FS10] Bernd Finkbeiner and Sven Schewe. Coordination logic. In International Workshop on Computer Science Logic, pages 305–319. Springer, 2010. [GR07] Rachid Guerraoui and Eric Ruppert. Anonymous and fault-tolerant sharedmemory computing. Distributed Computing, 20(3):165–177, 2007. [GT07] Fabiola Greve and Sébastien Tixeuil. Knowledge connectivity vs. synchrony requirements for fault-tolerant agreement in unknown networks. In 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07), pages 82–91. IEEE, 2007. [Ham13] Nicholas Ham. Notions of anonymity, fairness and symmetry for finite strategic-form games. arXiv preprint arXiv:1311.4766, 2013. [ID93] C Norris Ip and David L Dill. Better verification through symmetry. In Computer Hardware Description Languages and their Applications, pages 97–111. Elsevier, 1993.

- [LNRS16] 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.
- [Mau14] Bastien Maubert. Logical foundations of games with imperfect information: uniform strategies. PhD thesis, Université Rennes 1, 2014.
- [NRTV07] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors. Algorithmic Game Theory. Cambridge University Press, 2007. URL: https://doi.org/10.1017/CB09780511800481.
- [RHDC07] 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.
- [RR22] Sergio Rajsbaum and Armajac Raventós-Pujol. A distributed combinatorial topology approach to arrow's impossibility theorem. In Proceedings of the 2022 ACM Symposium on Principles of Distributed Computing, pages 471–481, 2022.
- [Ste11] Noah Daniel Stein. *Exchangeable equilibria*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [TV19] 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.
- [UW10] M Ummels and DK Wojtczak. The complexity of nash equilibria in stochastic multiplayer games. *Logical Methods in Computer Science*, 2010.
- [Ves12] Steen Vester. Symmetric Nash Equilibria. PhD thesis, Master's thesis, ENS Cachan, 2012.

בה היא מנצחת, ושיווי-משקל נאש חמדן, שבו בכדי ששחקנית תסטה מהאסטרטגיה שלה, מספיק שקיימת טופולוגיה אחת בה היא יכולה לשפר את מצבה על ידי סטיה. אנו בוחנים תכונות שונות של ההגדרות הללו לשיווי-משקל נאש, למשל ששיווי-משקל חמדן הוא מקרה פרטי של שיווי-משקל שמרן. אנו מתארים אלגוריתמים למציאת שיווי-משקל אלו במשחקים מרובי-טופולוגיות. אלגוריתמים אלו מבוססים על הרעיון העומד בבסיס האלגוריתם למציאת שיווי-משקל במשחקים מקביליים רגילים -הפיכת הבעיה למציאת אסטרטגיה מנצחת במשחק אחר. האלגוריתם למציאת שיווי-משקל חמדן הוא בעל סיבוכיות זמן מעריכית, ומסתמך על רדוקציה של הבעיה למציאת אסטרטגיה מנצחת לשחקנית במשחק ידיעה-חלקית עם שני שחקנים. האלגוריתם למציאת שיווי-משקל שמרן הוא בעל סיבוכיות זמן מעריכית-כפולה, ומסתמך על רדוקציה של הבעיה למציאת אסטרטגיה מנצחת לשחקנית זמן מעריכית-כפולה, ומסתמך על רדוקציה של הבעיה למציאת אסטרטגיה מנצחת לשחקנית זמן מעריכית-כפולה, ומסתמך על רדוקציה של הבעיה למציאת אסטרטגיה מנצחת לשחקנית מציאת אסטרטגיה מנצחת במשחק ידיעה-חלקית עם שלושה שחקנים, כאשר אחד השחקנים פועל כנגדה, והשחקן הנוסף פועל לטובתה. מציאת אסטרטגיה מנצחת במשחק ידיעה-חלקית עם שלושה שחקנים היא בעיה לא כריעה במקרה הכללי, אך תחת הנחות מסוימות שהבנייה שלנו מקיימת, קיים אלגוריתם לפתרון בעיה זו.

תקציר

משחקים מקביליים מרובי-שחקנים משמשים למידול מערכות המורכבות ממספר רכיבים המתקשרים אחד עם השני, כאשר לכל רכיב מטרה משלו, ויכולת לקבל החלטות בצורה אוטונומית. לדוגמא, ניתן למדל שירות ענן המספק שירותי חישוב באמצעות משחק מסוג זה. הרכיבים השונים הם מנהל המערכת, שהמטרה שלו היא למקסם רווח ולאפשר שירות תקין לכמה שיותר לקוחות, ולקוחות שונים המשתמשים בשירות, שהמטרה שלהם היא לבצע את החישובים שלהם בזמן הקצר ביותר. שיווי-משקל נאש הוא הפתרון הסטנדרטי למשחקים מסוג זה. שיווי-משקל נאש הוא פרופיל אסטרטגיות (אסטרטגיה לכל שחקן) "יציב" שבו לאף שחקן אין מוטיבציה לסטות מהאסטרטגיה שלו, כלומר, כל סטייה שלו מהאסטרטגיה שנתנה לו, תוביל לתוצאה טובה באותה המידה, או פחות טובה. בעבודה זו נעסוק במשחקים בעלי תנאי ניצחון אומגה-רגולרים, בהם המטרה של כל שחקן מייצגת מפרט שעל השחקן לקיים. שיווי-משקל נאש במשחקים מקביליים מרובי-שחקנים עם תנאי ניצחון אומגה-רגולרים נחקר במקרים בהם לשחקנים מידע מלא על מצב המשחק, ונמצא שהבעיה, בהינתן משחק רגולרים נחקר במקרים בהם לשחקנים מידע מלא על מצב המשחק, ונמצא שהבעיה, בהינתן משחק יעיל (פולינומיאלי) למציאת שיווי-המשקל.

במקרים רבים, רכיבים לא רואים בצורה מלאה את מצב המערכת, לדוגמה, אם מצב המערכת מכיל משתנים פנימיים שאינם חשופים כלפי חוץ. במקרה זה, האינטראקציה בין הרכיבים ממודלת על ידי משחקים עם ידיעה-חלקית. למרבה הצער, לקבוע האם קיים שיווי-משקל נאש במשחק ידיעה-חלקית כללי היא בעיה לא כריעה. עם זאת, במקרים רבים, אי-הידיעה של הרכיבים נובעת מכך שהמערכת משייכת לכל תהליך מזהה ייחודי, שאינו חשוף בפני התהליכים, במקרה זה, הרכיב יכול לראות את מצב המערכת באופן מלא, אך אינו יודע בדיוק מה תהיה ההשפעה של הפעולות שלו.

בעבודה זו, אנו מכלילים רעיון זה באמצעות מודל חדש הנקרא "משחקים מרובי טופולוגיות" - משחקים מקבילים עם מספר טופולוגיות אפשריות כך שהשחקנים משחקים את המשחק מבלי לדעת בדיוק מקבילים עם מספר טופולוגיות אפשריות כך שהשחקנים משחקים את המשחק מבלי לדעת בדיוק באיזו טופולוגיה המשחק נמצא. במשחק מרובה-טופולוגיות על השחקנים לבחור אסטרטגיה אחת שתשמש אותם לכל הטופולוגיות, וכיוצא בכך, כל פרופיל אסטרטגיות עבור השחקנים משרה תוצאה שתשמש אותם לכל הטופולוגיות, וכיוצא בכך, כל פרופיל אסטרטגיות עבור השחקנים משרה תוצאה שונה של המשחק לכל טופולוגיות, וכיוצא בכך, כל פרופיל אסטרטגיות עבור השחקנים משרה תוצאה שונה של המשחק לכל טופולוגיות, וכיוצא בכך, כל פרופיל אסטרטגיות עבור השחקנים משרה תוצאה זונה של המשחק לכל טופולוגיות, וביודה במשחק מקבילי מרובה-שחקנים רגיל. עובדה זו דורשת מאיתנו לבצע התאמות באופן שבו שיווי-משקל נאש מוגדר עבור משחקים מקביליים עבור מציאת שיווי-משקל נאש במשחקים מקביליים ענור מציאת שיווי-משקל נאש במשחקים מקביליים ענור מציאת שיווי-משקל נאש במשחקים מחקים מרובי-טופולוגיות, ומונעת מאיתנו להשתמש באלגוריתם למציאת שיווי-משקל נאש במשחקים מקביליים עבור מציאת שיווי-משקל נאש במשחקים מקביליים עבור מציאת שיווי-משקל נאש במשחקים מרובי טופולוגיות.

אנו מראים שניתן להגדיר שיווי-משקל נאש עבור משחקים אלו בכמה אופנים שונים, ובוחנים שתי אפשרויות: שיווי-משקל נאש שמרן, בו שחקנית תסטה מהאסטרטגיה שלה רק בתנאי שהסטיה תוביל לניצחון בטופולוגיות בהן היא מפסידה עם האסטרטגיה הנוכחית, מבלי להפסיד באף טופולוגיה

המחקר בוצע בהנחייתו של דר' שאול אלמגור, בפקולטה למדעי המחשב.

חלק מן התוצאות בחיבור זה פורסמו כמאמרים מאת המחבר ושותפיו למחקר בכנסים ובכתבי-עת במהלך תקופת מחקר הדוקטורט של המחבר, אשר גרסאותיהם העדכניות ביותר הינן:

Shaull Almagor and Shai Guendelman. Concurrent games with multiple topologies. In 33rd International Conference on Concurrency Theory, 2022.

תודות

אני רוצה להודות למנחה שלי, דר' שאול אלמגור, על הזמן, הסבלנות והמקצועיות שלו במהלך העבודה על התזה.

הכרת תודה מסורה לטכניון על מימון מחקר זה.

משחקים מרובי טופולוגיות

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר מגיסטר למדעים במדעי המחשב

שי גנדלמן

הוגש לסנט הטכניון – מכון טכנולוגי לישראל 2022 תשרי תשפ״ג חיפה אוקטובר

משחקים מרובי טופולוגיות

שי גנדלמן