

Dynamic Epistemic Logic Games with Epistemic Temporal Goals

Bastien Maubert,¹ Aniello Murano,¹ Sophie Pinchinat,² Francois Schwarzenruber,³ Silvia Stranieri¹

Abstract. Dynamic Epistemic Logic (DEL) is a logical framework in which one can describe in great detail how actions are perceived by the agents, and how they affect the world. DEL games were recently introduced as a way to define classes of games with imperfect information where the actions available to the players are described very precisely. This framework makes it possible to define easily, for instance, classes of games where players can only use public actions or public announcements. These games have been studied for reachability objectives, where the aim is to reach a situation satisfying some epistemic property expressed in epistemic logic; several (un)decidability results have been established.

In this work we show that the decidability results obtained for reachability objectives extend to a much more general class of winning conditions, namely those expressible in the epistemic temporal logic LTLK. To do so we establish that the infinite game structures generated by DEL public actions are regular, and we describe how to obtain finite representations on which we rely to solve them.

1 Introduction

Strategic reasoning in multi-agent systems refers to a number of important issues for settings where a team of agents have to take decisions in order to achieve some goals, while evolving in an environment that may pursue different objectives. Application domains are numerous (economics, robotics, distributed computing systems, web services, etc). For instance, drones patrolling an area may have to decide which trajectory to take so that the status (safe or unsafe) of each zone in this area is always known to at least one of them, while antagonistic agents try to keep the status of some areas secret. It is a real challenge to automatically compute adequate individual strategies for the agents. In this work we consider the *distributed strategy synthesis problem*, in which a team of agents collaborates towards a common goal, while the environment is purely antagonistic.

Because agents typically have a local view of the system, such situations are usually modelled as imperfect information game arenas, i.e., graphs whose nodes represent positions of the game, edges are the possible actions, and equivalence relations capture indistinguishability of positions. To reflect imperfect information, strategies must prescribe the same action in indistinguishable situations; such strategies are classically called *uniform* or *observation-based* strategies [35, 19, 7, 28]. Also the goal, or *winning condition*, is often expressed in some logical language such as LTL [32, 31] or LTLK, its extension with knowledge operators [39, 40]. For the patrolling example, one could consider the LTLK formula

$G \bigwedge_{\text{zone}} \bigvee_{\text{agent}} (K_{\text{agent}} \text{safe}_{\text{zone}} \vee K_{\text{agent}} \neg \text{safe}_{\text{zone}})$ which says that always, for every zone, some agent knows whether it is safe or not.

Distributed strategy synthesis is known to be undecidable [27, 31], but the numerous literature on the topic has identified two main decidable cases: the case where actions in the games are *public* (known to all agents) [39, 40, 34, 6, 5, 12], and the case of *hierarchical information* (the set of agents can be totally ordered so that what is known propagates along this order) [21, 31, 18, 26, 17, 29, 36, 8].

However, the state explosion problem makes game structures often very large, making distributed synthesis intractable. In order to circumvent this difficulty, we promote a *planning approach* using Dynamic Epistemic Logic (DEL) [41]: instead of representing explicitly the game structure, we consider implicit descriptions by means of *DEL presentations*. These consist in a finite initial *epistemic model* that reflects the initial knowledge of the agents, and a finite set of *epistemic actions* available to them and the other agents in the environment. Such implicit descriptions make it easier for the modeller to add, modify or remove actions. Also, since DEL action models were introduced to represent in detail how events are observed by agents, this setting is very convenient to define various types of actions such as public actions or semi-private announcements, and study how restricting to such actions can make distributed synthesis easier.

While DEL presentations have been widely used in epistemic planning (finding a finite succession of events that achieves some epistemic property) [11], only recently have adversarial aspects been considered in this setting, along with strategic problems such as distributed strategy synthesis. In [25], agents are split into two antagonistic teams Agt_{\exists} and Agt_{\forall} , and agents in Agt_{\exists} pursue some goal while agents in Agt_{\forall} try to prevent them from winning (these are zero-sum games). For reachability objectives where the team Agt_{\exists} wants to reach a situation that satisfies some epistemic property, it is shown in [25] that, as in the setting of explicit game arenas, distributed strategy synthesis is undecidable, but decidability can be recovered for the case of public actions and hierarchical information.

In this work we lift these decidability results from reachability goals to the much larger family of winning conditions expressible in the temporal epistemic logic LTLK, that blends temporal operators and epistemic modalities. This logic can express the reachability objectives from [25] (they correspond to formulas of the form $F\varphi$, where F is the “Finally” temporal operator and φ is an epistemic formula), but also safety, liveness properties, and many more (see for instance [40] for a detailed security example).

In all our decidability results, a crucial step is to show that the game arena induced by a DEL game presentation, which is in general infinite, can be represented finitely. This was already known for the case of actions whose preconditions do not involve knowledge but are purely propositional formulas (so-called *propositional ac-*

¹ Università degli Studi di Napoli “Federico II”, Italy

² Université de Rennes, France

³ ENS Rennes, France

Type of DEL game presentation	Logic	Finite game structure construction	Complexity
propositional actions and hierarchical information	LTLK	finitely many valuations	decidable (Th. 3)
public actions	LTLK	by quotienting	2EXPTIME-complete (Th. 4)
public announcements	LTL(U)K ₀	a poly-depth game tree	PSPACE-complete (Th. 5)

Table 1. Summary of our contribution.

tions) [23, 15], and we use it to transfer an existing result for distributed synthesis in explicit game arenas with hierarchical information. The main technical contribution of this work is to prove that the infinite game generated from a DEL game presentation is regular also in the case of public actions. This is done by observing that, modulo isomorphism, such actions can only generate finitely many different epistemic models from the initial one, thus allowing us to get an equivalent finite game as the quotient of the infinite one. This, combined with a recent result on game arenas with public actions [6], yields a procedure that runs in doubly exponential time, just as the case of LTL games [32]. Additionally, for public announcements (a special case of public actions) and the syntactic fragment of LTLK without next operator and local knowledge properties only, we show an even stronger characteristic of game arenas that allows us to reduce to a polynomial-length horizon game and to derive an optimal PSPACE procedure. Table 1 sums up our results.

Related work. The only work on DEL games that we are aware of is [25], and it only considers reachability objectives. However our results relate to the many aforementioned results for distributed synthesis in explicit game structures with either public actions or hierarchical information. Those dealing with epistemic temporal logic are the closest to ours and can be found in [33, 24] for hierarchical information, and [39, 40, 6, 5] for public actions.

Plan. Section 2 recalls games with imperfect information and the logic LTLK. Section 3 recalls DEL game presentations as defined in [25]. The central sections 4, 5, 6 describe our contributions for the cases of propositional actions, public actions and public announcements, respectively. We discuss our results in Section 7.

2 Preliminaries

In this section we recall basics about games with imperfect information and the epistemic temporal logic LTLK.

2.1 Games with imperfect information

We consider multiplayer game arenas with imperfect information in the spirit of, e.g., [37, 20, 8]. Since the DEL games we define in the next section are turn-based, i.e., the agents play in turns and not concurrently, we define turn-based arenas instead of the more general concurrent ones usually considered in the aforementioned works.

For the rest of the paper let us fix a countable set of *atomic propositions* AP and a finite set of *agents* Agt that is partitioned into two antagonistic teams, Agt_{\exists} and Agt_{\forall} .

Definition 1. A *game arena* $G = (V, V_I, Act, \delta, t, (\approx_a)_{a \in Agt}, \lambda)$ is a tuple where:

- V is a non-empty set of *positions*,
- $V_I \subseteq V$ is the set of *initial positions*,
- Act is a non-empty set of *actions*,
- $\delta : V \times Act \rightarrow V$ is a partial *transition function*,
- $t : V \rightarrow Agt$ is a *turn function*,
- $\approx_a \subseteq V \times V$ is an equivalence relation called *indistinguishability relation* for agent a , and
- $\lambda : V \rightarrow 2^{AP}$ is a *labelling* or *valuation function*.

In a position v , agent $t(v)$ chooses an action α such that $v' = \delta(v, \alpha)$ is defined, and the game proceeds similarly from position v' . We let $Act(v) = \{\alpha \mid (v, \alpha) \in \text{dom}(\delta)\}$, where dom denotes the domain, and we assume that $Act(v) \neq \emptyset$ for every position v .

A *play* $\pi = v_0 v_1 v_2 \dots$ in G is an infinite sequence of positions such that for all $i \in \mathbb{N}$, there exists $\alpha \in Act$ such that $v_{i+1} = \delta(v_i, \alpha)$. We let $\pi_i = v_i$ and $\pi_{\leq i} = v_0 v_1 \dots v_i$. We also let Plays^G denote the set of plays in G . A *history* $h = v_0 v_1 \dots v_n$ is a finite nonempty prefix of a play, $\text{last}(h) = v_n$ is the last position in h and Hist^G is the set of histories in G . We may write $t(h)$, $Act(h)$ and $\lambda(h)$ for, respectively, $t(\text{last}(h))$, $Act(\text{last}(h))$ and $\lambda(\text{last}(h))$.

The indistinguishability relation \approx_a is an equivalence relation between positions of the game arena that represents how agent a observes them: $v \approx_a v'$ means that agent a cannot distinguish between positions v and v' . As a result, we assume that if $v \approx_a v'$ for some agent a , then $t(v) = t(v')$, which means that agents know whose turn it is to play. In addition, if $t(v) = a$ and $v \approx_a v'$, we assume that $Act(v) = Act(v')$, meaning that the agent who has to make a move knows which actions are available.

We consider agents that have *synchronous perfect recall*, i.e. they remember the whole sequence of observations they made, and know how many moves have been made. To model this, each indistinguishability relation \approx_a is lifted to histories as follows: $h \approx_a h'$ if $|h| = |h'|$ and for every $i < |h|$ it holds that $h_i \approx_a h'_i$.

Remark 1. Some works require in addition that actions in equivalent histories be identical [38] or indistinguishable [33, 12]. While this can make a difference for some strategic problems such as existence of perfect-subgame equilibria, it does not for the distributed synthesis problem that we consider in this work [12].

A *strategy* for agent a is a partial function $\sigma : \text{Hist}^G \rightarrow Act$ such that for every h with $t(h) = a$, it holds that $\sigma(h) \in Act(h)$. Because agents can only base their decisions on what they observe, their strategies must assign the same action to indistinguishable situations: a strategy σ for agent a is *uniform* if, for all histories h, h' such that $t(h) = t(h') = a$ and $h \approx_a h'$, it holds that $\sigma(h) = \sigma(h')$. We say that a play π follows a strategy σ for agent a if for every $i \in \mathbb{N}$ such that $t(\pi_{\leq i}) = a$, it holds that $\pi_{i+1} = \delta(\pi_i, \sigma(\pi_{\leq i}))$. A *distributed strategy* for a group of agents $A \subseteq Agt$ is a tuple $\sigma_A = (\sigma_a)_{a \in A}$, and we write $\text{Out}(\sigma_A)$ the set of *outcomes* of σ_A , i.e. the set of plays that start in V_I and follow each σ_a for $a \in A$.

A game $\mathcal{G} = (G, \text{Win})$ is a game arena G with a winning condition $\text{Win} \subseteq \text{Plays}^G$. Team Agt_{\exists} wins a game \mathcal{G} if there is a distributed strategy $\sigma_{\text{Agt}_{\exists}}$ such that every play in $\text{Out}(\sigma_{\text{Agt}_{\exists}})$ is in Win .

Definition 2. Let $G = (V, V_I, \text{Act}, \delta, t, (\approx_a)_{a \in \text{Agt}}, \lambda)$ be a game arena. We define the *unfolding* of G as the game arena $G^{\text{unf}} = (V', V'_I, \text{Act}, \delta', t', (\approx'_a)_{a \in \text{Agt}}, \lambda')$ where $V' = \text{Hist}^G$, $V'_I = V_I$, for every $h \in \text{Hist}^G$, $t'(h) = t(\text{last}(h))$ and $\lambda'(h) = \lambda(\text{last}(h))$, \approx'_a is the synchronous perfect-recall lifting of \approx_a to histories, and

$$\delta'(h, \alpha) = \begin{cases} h \cdot v & \text{if } \delta(\text{last}(h), \alpha) = v \\ \text{undefined} & \text{if } \delta(\text{last}(h), \alpha) \text{ is undefined} \end{cases}$$

The natural bijection between plays of G and plays of G^{unf} induces a winning condition Win^{unf} over arena G^{unf} . Additionally, because of the natural bijection between strategies in G and strategies in G^{unf} , Agt_{\exists} wins (G, Win) if, and only if, Agt_{\exists} wins $(G^{\text{unf}}, \text{Win}^{\text{unf}})$. We say that two game structures \mathcal{G} and \mathcal{G}' are *equivalent* whenever their unfoldings are isomorphic⁴.

In this work we are interested in winning conditions expressed in the logic of knowledge and time called LTLK (standing for linear temporal logic with knowledge), which extends LTL with knowledge operators for each agent.

2.2 Linear temporal logic with knowledge

The syntax of LTLK is given by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi U \varphi \mid K_a\varphi$$

where $p \in AP$ and $a \in \text{Agt}$. The formula $X\varphi$ reads as “at the next step, φ holds”, $\varphi U \varphi'$ reads as “ φ holds until φ' holds”, and $K_a\varphi$ is read “agent a knows that φ is true”.

The size $|\varphi|$ of a formula φ is the number of symbols in it.

We exhibit two important syntactic fragments of LTLK: *Epistemic Logic* \mathcal{L}_{EL} obtained by removing temporal operators X and U , i.e., generated by grammar

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid K_a\varphi,$$

and *Propositional logic* $\mathcal{L}_{\text{Prop}}$ obtained by also removing the knowledge modality.

The logic LTLK is interpreted over a moment $i \in \mathbb{N}$ along a play $\pi \in \text{Plays}^G$ in a game arena $G = (V, \text{Act}, \delta, t, (\approx_a)_{a \in \text{Agt}}, \lambda)$. We write $G, \pi, i \models \varphi$, and read “formula φ holds at moment i along play π of game arena G ”, defined inductively over φ as follows.

$$\begin{aligned} G, \pi, i \models p & \quad \text{if } p \in \lambda(\pi_i) \\ G, \pi, i \models \neg\varphi & \quad \text{if } G, \pi, i \not\models \varphi \\ G, \pi, i \models \varphi_1 \vee \varphi_2 & \quad \text{if } G, \pi, i \models \varphi_1 \text{ or } G, \pi, i \models \varphi_2 \\ G, \pi, i \models X\varphi & \quad \text{if } G, \pi, i+1 \models \varphi \\ G, \pi, i \models \varphi_1 U \varphi_2 & \quad \text{if } \exists i' \geq i \text{ s.t. } G, \pi, i' \models \varphi_2 \text{ and} \\ & \quad \forall i'' \text{ s.t. } i \leq i'' < i', G, \pi, i'' \models \varphi_1 \\ G, \pi, i \models K_a\varphi & \quad \text{if } \forall \pi' \in \text{Plays}^G \text{ s.t. } \pi'_i \approx_a \pi_i, \\ & \quad G, \pi', i \models \varphi \end{aligned}$$

An LTLK formula φ naturally denotes a winning condition:

$$\text{Win}_{\varphi} = \{\pi \in \text{Plays}^G \mid G, \pi, 0 \models \varphi\}.$$

⁴ Some looser notion of bisimulation between games could also be considered but isomorphism fits here.

3 DEL games

In this section we recall the definition of DEL games as recently introduced in [25]. We start with definitions for epistemic models and DEL event models.

3.1 The classic DEL setting

In the usual possible-worlds semantics of epistemic logic, models are Kripke structures with interpretations for atomic propositions [16].

Definition 3. An *epistemic model* $\mathcal{M} = (W, (\approx_a)_{a \in \text{Agt}}, \lambda)$ is a tuple where

- W is a non-empty finite set of possible *worlds* (or situations),
- $\approx_a \subseteq W \times W$ is an *indistinguishability relation* for agent a , and
- $\lambda : W \rightarrow 2^{AP}$ is a *valuation function*.

We may write $w \in \mathcal{M}$ for $w \in W$. As for games with imperfect information introduced in the previous section, we assume that indistinguishability relations \approx_a are equivalence relations. The valuation function λ defines which atomic propositions hold in a world. A pair (\mathcal{M}, w) where $w \in \mathcal{M}$ is called a *pointed epistemic model*, and we define $|\mathcal{M}| = |W| + \sum_{a \in \text{Agt}} |\approx_a| + \sum_{w \in W} |\lambda(w)|$, the *size* of \mathcal{M} . We will only consider finite models, i.e. we assume that W is finite and $\lambda(w)$ is finite for all worlds w .

Definition 4. We define $\mathcal{M}, w \models \varphi$, read as “formula φ holds in the pointed epistemic model (\mathcal{M}, w) ”, by induction on φ , as follows:

- $\mathcal{M}, w \models p$ if $p \in \lambda(w)$;
- $\mathcal{M}, w \models \neg\varphi$ if $\mathcal{M}, w \not\models \varphi$;
- $\mathcal{M}, w \models \varphi_1 \vee \varphi_2$ if $\mathcal{M}, w \models \varphi_1$ or $\mathcal{M}, w \models \varphi_2$;
- $\mathcal{M}, w \models K_a\varphi$ if for all u such that $w \approx_a u$, $\mathcal{M}, u \models \varphi$.

Dynamic Epistemic Logic (DEL) relies on *action models* (also called “event models”). These models specify how agents perceive the occurrence of an action as well as its effects on the world.

Definition 5. An *action model* $\mathcal{A} = (A, (\approx_a^{\mathcal{A}})_{a \in \text{Agt}}, \text{pre}, \text{post})$ is a tuple where:

- A is a non-empty finite set of possible *actions*,
- $\approx_a^{\mathcal{A}} \subseteq A \times A$ is the *indistinguishability relation* for agent a ,
- $\text{pre} : A \rightarrow \mathcal{L}_{\text{EL}}$ the *precondition function*, and
- $\text{post} : A \times AP \rightarrow \mathcal{L}_{\text{Prop}}$ is the *postcondition function*.

A *pointed action model* is a pair (\mathcal{A}, α) where α represents the actual action. We let $|\mathcal{A}|$ be the *size* of \mathcal{A} , which we define as $|\mathcal{A}| := |A| + \sum_{a \in \text{Agt}} |\approx_a^{\mathcal{A}}| + \sum_{\alpha \in A} |\text{pre}(\alpha)| + \sum_{\alpha \in A, p \in AP} |\text{post}(\alpha, p)|$.

An action α is *executable* in a world w of an epistemic model \mathcal{M} if $\mathcal{M}, w \models \text{pre}(\alpha)$, and in that case we define $\lambda^{\text{update}}(w, \alpha) := \{p \in AP \mid \mathcal{M}, w \models \text{post}(\alpha, p)\}$, the set of atomic propositions that hold after occurrence of action α in world w .

Types of actions An action model \mathcal{A} is *propositional* if all pre- and postconditions of actions in \mathcal{A} belong to $\mathcal{L}_{\text{Prop}}$. A *public action* is a pointed action model \mathcal{A}, α such that for each agent a , $\approx_a^{\mathcal{A}}$ is the identity relation.

The effect of the execution of an action in an epistemic model is captured by the update product [4]:

Definition 6. Let $\mathcal{M} = (W, (\approx_a)_{a \in \text{Agt}}, \lambda)$ be an epistemic model, and $\mathcal{A} = (A, (\approx_a^{\mathcal{A}})_{a \in \text{Agt}}, \text{pre}, \text{post})$ be an action model. The *product* of \mathcal{M} and \mathcal{A} is defined as $\mathcal{M} \otimes \mathcal{A} = (W', (\approx_a')_{a \in \text{Agt}}, \lambda')$ where:

- $W' = \{(w, \alpha) \in W \times A \mid \mathcal{M}, w \models \text{pre}(\alpha)\}$,
- $(w, \alpha) \approx'_a (w', \alpha')$ if $w \approx_a w'$ and $\alpha \approx_a^A \alpha'$, and
- $\lambda'((w, \alpha)) = \lambda^{\text{update}}(w, \alpha)$.

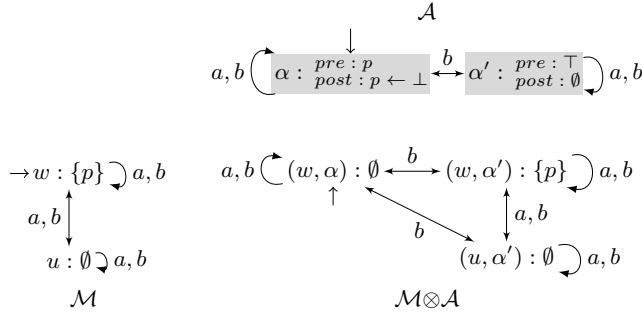


Figure 1. Example of DEL product. Symbol \emptyset indicates the trivial postcondition that leaves valuations unchanged.

Example 1. Figure 1 shows the pointed model \mathcal{M}, w that represents a situation in which p is true and both agents a and b do not know it. The pointed action model \mathcal{A}, α describes the action where agent a learns that p was true but that it is now set to false, while agent b does not learn anything: indeed she considers it possible that instead of α it was action α' that was performed, and this action has trivial pre- and postcondition. In the product epistemic model $\mathcal{M} \otimes \mathcal{A}, (w, \alpha)$, agent a now knows that p is false, while b still does not know the truth value of p , or whether agent a knows it.

3.2 Defining DEL games

We recall the definition of DEL games as introduced in [25], except that we consider more general winning conditions.

The initial situation is described by an epistemic model \mathcal{M} , and the set of possible actions by an action model \mathcal{A} . Because the update product in DEL can only model execution of a single action at a time, the games that we define are turn-based. We use the variable turn , ranging over the set of agents Agt , to represent whose turn it is to play. We require that postconditions for variable turn do not depend on the current epistemic situation, but instead the next value of turn is only determined by the action that is played. When the precondition $\text{pre}(\alpha)$ of some action α is satisfied, we may say that this action is *available*. Without loss of generality, we assume that there always is at least one action available.

The game thus starts in some world $w \in \mathcal{M}$, and the agent a such that $\mathcal{M}, w \models \text{turn} = a$ chooses some available action α which is executed. The new epistemic situation $\mathcal{M} \otimes \mathcal{A}, (w, \alpha)$ is given by the update product, and the game goes on. After n rounds, the epistemic situation is described by a pointed epistemic model of the form $\mathcal{MA}^n, (w, \alpha_1, \dots, \alpha_n)$, where \mathcal{MA}^n is defined by letting $\mathcal{MA}^0 = \mathcal{M}$ and $\mathcal{MA}^{n+1} = \mathcal{MA}^n \otimes \mathcal{A}$. In the following we may write $w\alpha_1 \dots \alpha_n$ instead of $(w, \alpha_1, \dots, \alpha_n)$, and call it a history. Given that the model \mathcal{MA}^n to which a history $w\alpha_1 \dots \alpha_n$ belongs is determined by the length of the history, we may omit it and write, e.g., $w\alpha_1 \dots \alpha_n \models \varphi$ instead of $\mathcal{MA}^n, w\alpha_1 \dots \alpha_n \models \varphi$.

In order to obtain proper games of imperfect information, we will require the following hypotheses to hold in the epistemic and event models defining DEL games:

Hypotheses on \mathcal{M} and \mathcal{A}

- (H1) **The starting player is known:** there is a player a such that for all $w \in W$, it holds that $\mathcal{M}, w \models \text{turn} = a$;
- (H2) **The turn stays known:** for all actions α, α' and agent a , if $\alpha \approx_a^A \alpha'$, then α and α' assign the same value to turn .
- (H3) **Players know their available actions:** if $w\alpha_1 \dots \alpha_n \models \text{turn} = a$ and $w\alpha_1 \dots \alpha_n \approx_a w'\alpha'_1 \dots \alpha'_n$, then the same actions are available in $w\alpha_1 \dots \alpha_n$ and in $w'\alpha'_1 \dots \alpha'_n$.

We can now define DEL games.

Definition 7. A DEL game presentation $(\mathcal{M}, \mathcal{A}, I)$ consists of an initial epistemic model \mathcal{M} and an action model \mathcal{A} that satisfy hypotheses H1, H2 and H3, together with a set of initial worlds I .

We now describe how a DEL game presentation $(\mathcal{M}, \mathcal{A}, I)$ induces a game arena $G_{\mathcal{M}, \mathcal{A}, I}$ as per Definition 1.

Definition 8. Given a DEL game presentation $(\mathcal{M}, \mathcal{A}, I)$, we define the game arena $G_{\mathcal{M}, \mathcal{A}, I} = (V, V_I, \text{Act}, \delta, t, (\approx_a)_{a \in \text{Agt}}, \lambda)$ where, letting $\mathcal{MA}^n = (W^n, (\approx_a^n)_{a \in \text{Agt}}, \lambda^n)$ for every n :

- $V = \bigcup_{n \in \mathbb{N}} W^n$,
- $V_I = I$,
- Act is the set of actions in \mathcal{A} ,
- $\delta(w\alpha_1 \dots \alpha_n, \alpha_{n+1}) = \begin{cases} w\alpha_1 \dots \alpha_n \alpha_{n+1} & \text{if } w\alpha_1 \dots \alpha_n \models \text{pre}(\alpha_{n+1}) \\ \text{undefined} & \text{otherwise} \end{cases}$
- $t(w\alpha_1 \dots \alpha_n) = a$ if $w\alpha_1 \dots \alpha_n \models \text{turn} = a$
- $\approx_a = \bigcup_{n \in \mathbb{N}} \approx_a^n$ for each agent a , and
- $\lambda(w\alpha_1 \dots \alpha_n) = \lambda^n(w\alpha_1 \dots \alpha_n)$.

Observe that this game arena is infinite: the set of positions V is the set of histories. In the following sections we will see that in some cases they admit finite representations that we can use to decide the existence of winning strategies.

A DEL game $\mathcal{G}^{\text{DEL}} = (\mathcal{M}, \mathcal{A}, I, \text{Win})$ consists of a DEL game presentation $(\mathcal{M}, \mathcal{A}, I)$ together with a winning condition $\text{Win} \subseteq \text{Plays}^{G_{\mathcal{M}, \mathcal{A}, I}}$ on the induced game arena $G_{\mathcal{M}, \mathcal{A}, I}$. We consider the following decision problem:

Definition 9 (Distributed strategy synthesis for LTLK objectives).

- Input: A DEL game $\mathcal{G}^{\text{DEL}} = (\mathcal{M}, \mathcal{A}, I, \varphi)$ with $\varphi \in \text{LTLK}$;
- Question: Does team Agt_{\exists} win the game $(G_{\mathcal{M}, \mathcal{A}, I}, \varphi)$?

Note that the games studied in [25] correspond to the class of DEL games where the winning condition is given by LTLK formulas of the form $F\varphi$, where φ is purely epistemic.

Remark 2. When we evaluate whether a tuple of strategies $(\sigma_a)_{a \in \text{Agt}_{\exists}}$ is winning for an LTLK formula φ , the semantics of the knowledge operators in φ does not depend on the strategies σ_a . In particular, it is not restricted to indistinguishable histories that follow these strategies, but instead it considers all indistinguishable histories in the game. This semantics models situations in which agents do not know the strategies of agents Agt_{\exists} , and it is the one also used in [25] and in DEL epistemic planning [9, 1, 10, 11], where the agents do not know which plan is being executed. This semantics is called *uninformed semantics* in [24], contrary to the *informed* one. Note that the choice of semantics affects the complexity of distributed strategy synthesis: for instance it is known that when the game is given explicitly as a graph, the existence of distributed strategies for LTLK on games with hierarchical information is decidable for the uninformed semantics [33] (and we use it to establish Theorem 3), but it

is undecidable for the informed semantics [40]. It is thus likely that Theorem 3 would not hold for the informed semantics. We refer the reader to [33, 23] for more discussions on the matter.

3.3 Discussion on initial positions

One subtlety that arises when formalising existence of winning strategies under imperfect information is in defining what having a winning strategy means. For instance, are we satisfied with the agents in Agt_{\exists} having a distributed strategy that is winning from the initial position of the game, even if they do not *know* that it is winning, in the sense that there is a world that some agent in Agt_{\exists} considers like a possible initial position and from which the distributed strategy is not winning? Or instead do we want everybody in the team to know that the team's strategy is winning? These two notions have sometimes been called *objective winning* and *subjective winning*, respectively (see [13] and also [20] for similar considerations). We could also ask whether there is *distributed knowledge* or *common knowledge* [16] that the distributed strategy is winning.

Note that we can model all of these notions by tuning the set of initial worlds in the definition of a DEL game, as this defines the set of outcomes that we consider, i.e., the set of plays from which the strategies should be winning. Assume we have an initial epistemic model \mathcal{M} with an initial world w_I . If we are interested in distributed strategies that are objectively winning from w_I , we simply set $I = \{w_I\}$ in the DEL game. If instead we want subjectively winning strategies, i.e., strategies that not only are winning, but such that everybody in the team Agt_{\exists} knows that they are winning, then we let the set of initial worlds in the DEL game be

$$I^{\exists} = \{w \in \mathcal{M} \mid w \approx_a w_I \text{ for some } a \in \text{Agt}_{\exists}\}.$$

Objective distributed strategy synthesis was studied in [25] for reachability epistemic objectives, i.e., when an epistemic objective is given as an epistemic formula $\varphi \in \mathcal{L}_{\text{EL}}$, and Win is defined as the sets of plays $\pi \in \text{Plays}_{\mathcal{MA}^*}$ for which there exists $i \in \mathbb{N}$ such that $\mathcal{MA}^*, \pi_{\leq i} \models \varphi$. Note that such winning conditions can be specified by LTLK formulas $F\varphi$. It is shown that objective distributed strategy synthesis is undecidable for propositional actions, already for a team of two players and reachability epistemic objectives. Since the problem we study is more general, we inherit this undecidability result.

Theorem 1. *Distributed strategy synthesis for LTLK objectives is undecidable already for propositional actions and formulas of the form $F\varphi$, where φ is purely epistemic.*

In the rest of the paper we describe various cases in which decidability can be recovered.

4 DEL games with propositional actions

Extending a result from [23, 2] in the case of DEL planning, it was proved in [25] that in the case of propositional actions, games generated by DEL game presentations are regular, and one can compute an equivalent finite game arena:

Proposition 1 ([25]). *Given a DEL game presentation $(\mathcal{M}, \mathcal{A})$ where \mathcal{A} is propositional, one can construct a finite game arena G equivalent to $G_{\mathcal{M}, \mathcal{A}, I}$ such that $|G| \leq |\mathcal{M}| + |\mathcal{A}| \times 2^m$, where m is the number of atomic propositions in \mathcal{M} and \mathcal{A} .*

From the latter result, if the winning condition Win is given as a formula $\varphi \in \text{LTLK}$, then the same winning condition on G yields

a multiplayer epistemic game that is equivalent to the original DEL game G^{DEL} in terms of existence of distributed winning strategies. And because such games can be decided in the case of hierarchical information, we obtain our result.

More precisely, we say that a DEL game $(\mathcal{M}, W_I, \mathcal{A}, \text{Win})$ presents *hierarchical information* if the set of agents Agt_{\exists} can be totally ordered $(a_1 < \dots < a_n)$ so that $\approx_{a_i} \subseteq \approx_{a_{i+1}}$ and $\approx_{a_i}^{\mathcal{A}} \subseteq \approx_{a_{i+1}}^{\mathcal{A}}$, for each $1 \leq i < n$.

Theorem 2 ([33, 24]). *In multiplayer epistemic games with hierarchical information and epistemic temporal objectives, distributed strategy synthesis is decidable for the uninformed semantics.*

Proposition 1 together with Theorem 2 imply that:

Theorem 3. *Distributed strategy synthesis for LTLK objectives with propositional actions and hierarchical information is decidable.*

Remark 3. The results in [33, 24] are established for games with a unique initial position, i.e. when V_I is a singleton $\{v_I\}$. However it is easy to see that distributed synthesis with multiple initial positions V_I can be reduced to the case of a unique initial position: one only needs to add a fresh position v_I that is used as initial position, from which one can attain all positions in V_I , and only these. It does not matter who v_I belongs to or how the agents observe it.

5 DEL games with public actions

In this section, we show that when all actions are public, the distributed strategy synthesis problem is decidable for LTLK winning conditions. Towards this end, we first prove a result similar to Proposition 1: we show that given a DEL game presentation $(\mathcal{M}, \mathcal{A}, I)$ where all actions $\alpha \in \mathcal{A}$ are public, the infinite game arena $G_{\mathcal{M}, \mathcal{A}, I}$ is regular and can be folded back into a finite game arena. This allows us to reduce the distributed strategy synthesis problem to a distributed synthesis problem on explicit game arenas, for which a solution is known in the case of public actions and LTLK objectives [40].

Note that the decidability result for reachability DEL games with public actions in [25] does not rely on this kind of construction, but instead is proved by providing a direct alternating algorithm. There is a problem in the way this algorithm forces strategies to be uniform. In the case of public actions and unique initial world considered there it can be easily corrected, as there is in fact no need to check for uniformity of strategies (see Remark 4). But for our more general setting with multiple initial worlds, this approach was not sound.

Proposition 2. *Given a DEL game presentation $(\mathcal{M}, \mathcal{A}, I)$ where all actions in \mathcal{A} are public, one can compute a finite game arena G equivalent to $G_{\mathcal{M}, \mathcal{A}, I}$ such that $|G| \leq m(2^p + 1)^m$, with p the number of atomic propositions in $(\mathcal{M}, \mathcal{A})$ and m the number of worlds in \mathcal{M} .*

Proof. For every position $w\alpha_1 \dots \alpha_n$ in $G_{\mathcal{M}, \mathcal{A}, I}$ we define its *attached epistemic model* $\mathcal{M}_{w\alpha_1 \dots \alpha_n}$ as the connected component of \mathcal{MA}^n that contains $w\alpha_1 \dots \alpha_n$. Since all actions in \mathcal{A} are public, for all positions $w\alpha_1 \dots \alpha_n$ and $w\alpha_1 \dots \alpha_n\alpha_{n+1}$ in $G_{\mathcal{M}, \mathcal{A}, I}$ we have that $\mathcal{M}_{w\alpha_1 \dots \alpha_{n+1}}$ is no bigger than $\mathcal{M}_{w\alpha_1 \dots \alpha_n}$: indeed, the application of a public action can only remove worlds from $\mathcal{M}_{w\alpha_1 \dots \alpha_n}$ (those that do not satisfy the precondition) and change the valuations of the remaining worlds. As a result there is only a finite number of different positions $w\alpha_1 \dots \alpha_n$ in $G_{\mathcal{M}, \mathcal{A}, I}$, up to isomorphism of their attached models. We write \equiv the equivalence relation on positions of $G_{\mathcal{M}, \mathcal{A}, I}$ defined by letting two positions be equivalent if

their attached models are isomorphic, and we let $[w\alpha_1 \dots \alpha_n]_{\equiv}$ be the equivalence class of position $w\alpha_1 \dots \alpha_n$ for this relation.

Let us write $G_{\mathcal{M}, \mathcal{A}, I} = (V, V_I, Act, \delta, t, \{\approx_a\}_{a \in Agt}, \lambda)$. The finite game arena G is the quotient of $G_{\mathcal{M}, \mathcal{A}, I}$ with \equiv . More precisely, $G_{\mathcal{M}, \mathcal{A}, I/\equiv} = (V', V'_I, Act', \delta', t', \{\approx'_a\}_{a \in Agt}, \lambda')$, where:

- $V' = \{[w\alpha_1 \dots \alpha_n]_{\equiv} \mid w\alpha_1 \dots \alpha_n \in V\}$,
- $V'_I = \{[w]_{\equiv} \mid w \in V_I\}$,
- $Act' = Act$,
- $\delta'([v]_{\equiv}, \alpha) = \begin{cases} [\delta(v, \alpha)]_{\equiv} & \text{if } \delta(v, \alpha) \text{ is defined} \\ \text{undefined} & \text{otherwise,} \end{cases}$
- $[v]_{\equiv} \approx'_a [v']_{\equiv}$ if $v \approx_a v'$, and
- $\lambda'([v]_{\equiv}) = \lambda(v)$.

To see that δ' is well defined, observe that if $v \equiv v'$ then $\delta(v, \alpha)$ is defined if, and only if, so is $\delta(v', \alpha)$, and in this case $\delta(v, \alpha) \equiv \delta(v', \alpha)$. The fact that \approx'_a and λ' are well defined follows directly from isomorphism of attached models.

To construct $G_{\mathcal{M}, \mathcal{A}, I/\equiv}$, one can enumerate all possible attached models $\{\mathcal{N}_1, \dots, \mathcal{N}_k\}$ (modulo isomorphism) as follows: for each world w in the original model \mathcal{M} , decide first whether there is some world of the form $w\alpha_1 \dots \alpha_n$ in \mathcal{N} ; if there is, there is only one, because all actions are public, and thus any position of the form $w\alpha'_1 \dots \alpha'_n$ (with for some i , $\alpha'_i \neq \alpha_i$) is not related to $w\alpha_1 \dots \alpha_n$ and thus does not appear in \mathcal{N} . Then, one chooses the valuation over the atomic propositions involved in the problem for each world in \mathcal{N} . Indistinguishability relations are inherited from \mathcal{M} : $w\alpha_1 \dots \alpha_n \approx_a w'\alpha'_1 \dots \alpha'_n$ if, and only if, $w \approx_a w'$. The number of such different attached models is bounded by $\sum_{i=1}^m \binom{m}{i} 2^{p_i} = (2^p + 1)^m$, and each one has at most m worlds. Positions $[w]_{\equiv}$ of $G_{\mathcal{M}, \mathcal{A}, I/\equiv}$ can be represented as pairs (\mathcal{N}, w) where $\mathcal{N} \in \{\mathcal{N}_1, \dots, \mathcal{N}_k\}$ is one of the representants of classes of isomorphic attached models enumerated above, and w is a world in \mathcal{N} . We thus have at most $m(2^p + 1)^m$ positions in V' . It remains to build the function δ' and relations \approx'_a as described in the definition of $G_{\mathcal{M}, \mathcal{A}, I/\equiv}$.

In order to build δ' one has to compute for each position (\mathcal{N}, w) and for each action α , the attached model $\mathcal{N} \otimes \mathcal{A}, w\alpha$, and then identify the pointed attached model (\mathcal{N}_i, u) in $G_{\mathcal{M}, \mathcal{A}, I/\equiv}$ to which $\mathcal{N} \otimes \mathcal{A}, w\alpha$ is isomorphic. Since testing graph isomorphism can be done in time $2^{\sqrt{m \log m}}$ for graphs with at most m vertices [3], and all attached models have at most m worlds, the whole construction can be done in exponential time. \square

Proposition 2 ensures that from a DEL game presentation $(\mathcal{M}, \mathcal{A})$ with public actions we can construct an equivalent finite game arena of exponential size. Moreover, in this game arena, all actions are public in the sense of [6]. In this latter work, model checking ATL^* with epistemic operators (ATL_K^*) on game arenas with public actions is proved in 2EXPTIME. More precisely, the proposed procedure takes time doubly exponential in the size of the formula, but only exponential time in the size of the game structure. Combined with our exponential construction from Theorem 2, we obtain a procedure to solve our distributed strategy synthesis problem for public actions in doubly exponential time, both in the size of the DEL game presentation and in the size the LTLK winning condition φ .

To make our argument more precise, we briefly recall the syntax and semantics of ATL_K^* . The syntax of ATL_K^* is given by the following grammar:

$$\begin{aligned} \varphi &::= p \mid \neg \varphi \mid \varphi \vee \varphi \mid K_a \varphi \mid \langle A \rangle \psi \\ \psi &::= \varphi \mid \neg \psi \mid \psi \vee \psi \mid X\psi \mid \psi U \psi \end{aligned}$$

where $p \in AP$ and $a \in Agt$. Formulas of type φ are called *history formulas*, while those of type ψ are called *path formulas*. Note that, in addition, the authors of [6] consider epistemic operators for common and distributed knowledge. We omit them from the syntax as we do not consider such operators in this work.

The semantics of ATL_K^* is defined in [6] on concurrent-game structures. We instead define it on our turn-based game structures, which can be seen as a particular case. Let us first recall the notion of public actions for game arenas considered in [6]:

Definition 10. A game arena $G = (V, V_I, Act, \delta, t, (\approx_a)_{a \in Agt}, \lambda)$ has *only public actions* if, for all $v, v' \in V$ and $\alpha, \alpha' \in Act$ such that $\alpha \neq \alpha'$, we have $\delta(v, \alpha) \not\approx_a \delta(v', \alpha')$.

The semantics of ATL_K^* formulas is defined with respect to a game arena G together with a history h in case of a history formula, or a play π and a point in time $i \in \mathbb{N}$ in case of a path formula (we omit boolean and temporal cases).

$$\begin{aligned} G, h &\models p && \text{if } p \in \lambda(h) \\ G, h &\models K_a \varphi && \text{if } \forall h' \in \text{Hist}^G \text{ s.t. } h' \approx_a h, G, h' \models \varphi \\ G, h &\models \langle A \rangle \psi && \text{if there exists } \sigma_A \text{ s.t. } \forall \pi \in \text{Out}(h, \sigma_A), \\ &&& G, \pi, |h| - 1 \models \psi \\ G, \pi, i &\models \varphi && \text{if } G, \pi_{\leq i} \models \varphi \end{aligned}$$

In the above definition, $\text{Out}(h, \sigma_A)$ is the set of plays that extend h by following σ_A , which corresponds to the objective semantics discussed in Section 3.3. Thus, it is easy to see that for an LTLK formula ψ and a game arena G with a singleton set of initial positions $V_I = \{v_I\}$, it holds that Agt_{\exists} wins (G, φ) if, and only if, $G, v_I \models \langle Agt_{\exists} \rangle \psi$ (technically, ψ should be modified by replacing each occurrence of knowledge modality K_a by $K_a(\emptyset)$ in order to have a history formula).

As explained in [6], their procedure can be adapted to the subjective semantics, i.e. when the distributed strategy σ_A is required to be winning from all positions that are equivalent to v_I for some agent in A . It is not hard to see that this adaptation would also work for any set V_I of initial positions, and that it does not change the complexity of the procedure.

Theorem 4. For public actions, distributed strategy synthesis for LTLK objectives is 2EXPTIME-complete.

Proof. Let $(\mathcal{M}, \mathcal{A}, I, \varphi)$ be a DEL game with $\varphi \in \text{LTLK}$ and such that all actions are public. By Proposition 2 we can compute in exponential time a finite game arena $G = G_{\mathcal{M}, \mathcal{A}, I/\equiv}$ equivalent to $G_{\mathcal{M}, \mathcal{A}, I}$ with $|G| \leq m(2^p + 1)^m$, where m is the number of worlds in \mathcal{M} and p is the number of atomic propositions in $(\mathcal{M}, \mathcal{A})$.

We show that this finite game arena G has only public actions in the sense of Definition 10. Write $G = (V, Act, \delta, t, \{\approx_a\}_{a \in Agt}, \lambda)$, and recall that $V = \{[w\alpha_1 \dots \alpha_n]_{\equiv} \mid w\alpha_1 \dots \alpha_n \in G_{\mathcal{M}, \mathcal{A}, I}\}$. Take two positions $v = [w\alpha_1 \dots \alpha_i]_{\equiv}$ and $v' = [w'\alpha'_1 \dots \alpha'_j]_{\equiv}$ in V , and $\alpha, \alpha' \in Act$ such that $\alpha \neq \alpha'$. By definition, $\delta(v, \alpha) = [w\alpha_1 \dots \alpha_i \alpha]_{\equiv}$ and $\delta(v', \alpha') = [w'\alpha'_1 \dots \alpha'_j \alpha']_{\equiv}$. Since α and α' are public actions and $\alpha \neq \alpha'$, necessarily $w\alpha_1 \dots \alpha_i \alpha \not\approx_a w'\alpha'_1 \dots \alpha'_j \alpha'$, entailing $\delta(v, \alpha) \not\approx_a \delta(v', \alpha')$.

We can thus use the model-checking procedure from [6] to evaluate whether G satisfies $\langle Agt_{\exists} \rangle \varphi$. This procedure takes time doubly exponential in the size of φ and exponential in the size of G , which is itself exponential in the size of the DEL game presentation $(\mathcal{M}, \mathcal{A})$, hence the upper bound. The lower bound is obtained by reduction from LTL synthesis, which is 2EXPTIME-complete [32]. \square

Remark 4. We point out that in the context of public actions, the case of a unique initial position is in fact equivalent to the perfect-information simplification of the problem, in which the uniformity requirement for strategies is dropped. Indeed in this case the uniformity constraint is trivial to satisfy: assume there is a winning distributed strategy $(\sigma_a)_{a \in \text{Agt}}$ where the strategies are not necessarily uniform. Take two histories h and h' that are equivalent to some agent $a \in \text{Agt}$. If one of them, say h' , does not start in w_I , then it does not matter how strategies $(\sigma_a)_{a \in \text{Agt}}$ are defined on h' , because they are not required to be winning from worlds other than w_I ; one can thus change the definition of the strategies on h' to make them uniform. Otherwise, if both start in w_I , because actions are public, $h = h'$ so that the strategies are already uniform on these histories.

6 DEL games with public announcements

We now investigate DEL games with *public announcements*, which are public actions with no effect besides epistemic ones. We assume that the winning conditions are restricted to $\text{LTL}(U)\text{K}_0$, the syntactic fragment of LTLK objectives with no next modality (X) and with no temporal operator (X or U) under the scope of a knowledge modality (K_a). We also assume that the games are round-robin, i.e. the turn goes from an agent to the next in a circular order, and we assume a unique initial world. We show that, in this context, deciding the existence of a winning strategy for the team Agt in a DEL game $\mathcal{G}^{\text{DEL}} = (\mathcal{M}, \mathcal{A}, I, \varphi)$ and $\varphi \in \text{LTL}(U)\text{K}_0$ is PSPACE-complete. Because reachability goals are definable in $\text{LTL}(U)\text{K}_0$, this result generalises the PSPACE-completeness result established in [25].

Formally, a *public announcement* is a public action (\mathcal{A}, α) such that $\text{post}(\alpha, p) = p$, for each variable p but variable turn. This is the natural generalisation of public announcements as defined in [30, 41]. As a consequence of the product update, either an announcement is *non-informative* and the updated epistemic model remains the same (modulo variable turn), or it is *informative* and yields an epistemic model with strictly less worlds.

Theorem 5. *In round-robin DEL games with unique initial world and public announcements, distributed strategy synthesis for $\text{LTL}(U)\text{K}_0$ winning conditions is PSPACE-complete.*

The rest of this section is dedicated to the proof of Theorem 5. The problem is already PSPACE-hard for reachability goals [25], therefore it is still PSPACE-hard for $\text{LTL}(U)\text{K}_0$ objectives. Regarding the membership in PSPACE, the two main ideas are:

1. From an initial epistemic model $\mathcal{M} = (W, (\approx_a)_{a \in \text{Agt}}, \lambda)$, there are at most $|W|$ informative announcements;
2. To limit the length of plays, we can shorten, as depicted in Figure 2, sequences of non-informative announcements: from a strategy σ we show how to extract an *eager* strategy σ^{eager} that performs all informative announcements eventually recommended by σ as early as possible. Thus, any sequence of non-informative announcements followed by an informative one is of length at most the number $|\text{Agt}|$ of agents: if an agent wants to perform an informative event in the future, she can do so as soon as it is her turn to play. This, in a round-robin game, happens in at most $|\text{Agt}|$ steps.

As a result of these two points, we can search for eager strategies via a depth-first search in $G_{\mathcal{M}, \mathcal{A}, I}$ up to depth $|\text{Agt}| \times |\mathcal{M}|$.

We now describe how to extract eager strategies. In the following we call *states* the attached epistemic models⁵ in $G_{\mathcal{M}, \mathcal{A}, I}$, writing

them s, s_1, \dots . These are mere submodels of the initial model \mathcal{M} , if we ignore variable turn. We then write s^k for the sequence with k consecutive s 's (only variable turn is changing).

Given a distributed strategy $\sigma = (\sigma_a)_{a \in \text{Agt}}$, we let the eager distributed strategy $\sigma^{\text{eager}} = (\sigma_a^{\text{eager}})_{a \in \text{Agt}}$ be defined by $\sigma_a^{\text{eager}}(h) := \sigma_a(\text{look}_a(h))$ where $\text{look}_a(h)$ is a history called *look ahead*. This $\text{look}_a(h)$ is, when it exists, a history that follows σ , in which it is a 's turn to play and $\sigma_a(\text{look}_a(h))$ is informative. Also, h is a stuttering-equivalent subsequence of $\text{look}_a(h)$ where agents bypass non-informative announcements and perform informative ones prescribed by σ as soon as possible. There might be no such $\text{look}_a(h)$ if agents are not eager in h . We define $\text{look}_a(h)$ by induction:

- $\text{look}_a(\epsilon) := \epsilon$ (base case);
- if $h = h's$ where h' is a history and s is a state with either $h' = \epsilon$ or $|s| < |\text{last}(h')|$ (an informative announcement has been made), $\text{look}_a(h's^k) = \text{look}_a(h')s^\ell$, such that $\text{look}_a(h')s^\ell$ follows σ and ℓ is
 - i. if it exists, the smallest integer such that $\text{turn}(\text{look}_a(h')s^\ell) = a$, and $\sigma(\text{look}_a(h')s^\ell)$ is informative at $\text{look}_a(h')s^\ell$.
 - ii. otherwise take $\ell = k$.

Lemma 1. *Any outcome of σ^{eager} is of the form $s_1^{k_1} \dots s_n^{k_n} s^\omega$ where $|s_1| < |s_2| < \dots < |s_n|$ and $k_i < |\text{Agt}|$.*

As informative announcements prescribed by $(\sigma_a^{\text{eager}})_{a \in \text{Agt}}$ coincide with the ones prescribed by $(\sigma_a)_{a \in \text{Agt}}$, the outcomes of σ^{eager} are *stuttering equivalent* to some outcome of σ . Recall that two paths are stuttering equivalent if omitting repetitions of states in both of them yields the same sequence of states. For instance, $s_1 s_2 s_2 s_3$ and $s_1 s_1 s_2 s_3 s_3$ are stuttering equivalent (see [22]).

Lemma 2. *For any outcome of σ^{eager} , there exists a stuttering equivalent outcome of σ .*

We now design a polynomial space algorithm that decides whether there exists such an eager strategy σ_a^{eager} by performing a depth-first-search (minmax-like approach) in the unfolding of $G_{\mathcal{M}, \mathcal{A}, I}$ at polynomial depth $|\text{Agt}| \times |\mathcal{M}|$.




Figure 2. Strategy σ^{eager} bypasses the non-informative announcements in σ and makes the same informative announcement (here p) but eagerly.

Every time a leaf s is reached, it is considered as the attached epistemic model in which the game stays forever with no more informative announcement (i.e. s^ω). We then evaluate the winning condition $\text{LTL}(U)\text{K}_0$ -formula by model checking the path carried by the branch in this tree. Now, model checking a path against $\text{LTL}(U)\text{K}_0$ is a problem in P: because we require that no temporal operator occur under the scope of knowledge modalities, epistemic subformulas occurring in the challenged $\text{LTL}(U)\text{K}_0$ -formula can be evaluated locally on the path so that these subformulas become mere propositions. It remains to model check an LTL-formula on this marked path which can be done in polynomial time (see for example [14, Section 6.4.3]).

Notice that while running this depth-first-search, one needs to remember the current branch (needed for backtracking in the minmax

⁵ see proof of Proposition 2, page 5.

algorithm) as well as the information used by the $LTL(U)K_0$ path model-checking procedure, which yields a poly-size information, so that the algorithm runs in polynomial space.

We prove that this algorithm is correct. If the algorithm accepts the input, then we have a winning strategy σ^{eager} , and we are done.

Conversely, assume there exists a winning strategy σ . Because any $LTL(U)$ -formula (no X operator) is stuttering-invariant (see for example [14, Th. 6.6.5 p. 184]), and because in our logic $LTL(U)K_0$, epistemic subformulas are evaluated locally in states, just as propositions, the outcomes of σ^{eager} do also satisfy the winning conditions by Lemma 2. Now because of Lemma 1 strategy σ^{eager} will be found by the algorithm, which concludes.

7 Conclusion

We generalised the setting defined in [25] for distributed synthesis in DEL games, moving from reachability winning conditions to ones expressed in LTLK, and allowing for multiple initial positions, which allows us to capture various semantics of strategic ability but also makes the problem harder in the case of public actions.

We showed that the main results established in [25] can be lifted to this more general setting: of course the problem remains undecidable, but decidability is retrieved in the case of public actions, as well as propositional actions together with hierarchical information.

In the latter case the problem is, as usual, nonelementary, as each agent in the team with a different observation of the game adds an exponential to the cost of solving it [31, 26]. But for public actions we proved that the problem is in 2EXPTIME, which is optimal as this is already the complexity of solving LTL synthesis [32]. A central technical result was to establish the regularity of infinite DEL game arenas generated from public actions. We conjecture that our techniques could extend to even more expressive winning conditions, such as ones expressible in epistemic mu-calculus.

Regarding public announcements, we showed that the distributed synthesis problem is PSPACE-complete for winning conditions in the fragment $LTL(U)K_0$, when games are round-robin and have a unique initial world. The complexity of generalisations such as several initial positions or winning conditions beyond $LTL(U)K_0$ is still open.

REFERENCES

- [1] G. Aucher and T. Bolander, ‘Undecidability in epistemic planning’, in *IJCAI’13*, (2013).
- [2] G. Aucher, B. Maubert, and S. Pinchinat, ‘Automata techniques for epistemic protocol synthesis’, in *SR’14*, pp. 97–103, (2014).
- [3] László Babai and Eugene M Luks, ‘Canonical labeling of graphs’, in *STOC’83*, pp. 171–183, (1983).
- [4] A. Baltag, L. S Moss, and S. Solecki, ‘The logic of public announcements, common knowledge, and private suspicions’, in *TARK’98*, pp. 43–56, (1998).
- [5] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin, ‘Verification of broadcasting multi-agent systems against an epistemic strategy logic’, in *IJCAI’17*, pp. 91–97, (2017).
- [6] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin, ‘Verification of multi-agent systems with imperfect information and public actions’, in *AAMAS’17*, pp. 1268–1276, (2017).
- [7] D. Berwanger, K. Chatterjee, M. De Wulf, L. Doyen, and T. A. Henzinger, ‘Strategy construction for parity games with imperfect information’, *Inf. Comput.*, **208**(10), 1206–1220, (2010).
- [8] Dietmar Berwanger, Anup Basil Mathew, and Marie van den Bogaard, ‘Hierarchical information and the synthesis of distributed strategies’, *Acta Inf.*, **55**(8), 669–701, (2018).
- [9] T. Bolander and M.B. Andersen, ‘Epistemic planning for single and multi-agent systems’, *J. Appl. Non-Class. Log.*, **21**(1), 9–34, (2011).
- [10] T. Bolander, M.H. Jensen, and F. Schwarzenrüber, ‘Complexity results in epistemic planning’, in *IJCAI*, pp. 2791–2797, (2015).
- [11] Thomas Bolander, ‘A gentle introduction to epistemic planning: The DEL approach’, in *M4M*, pp. 1–22, (2017).
- [12] P. Bouyer, ‘Games on graphs with a public signal monitoring’, in *FOS-SACS’18*. Springer, (2018).
- [13] Nils Bulling and Wojciech Jamroga, ‘Comparing variants of strategic ability: how uncertainty and memory influence general properties of games’, *Auton. Agents Multi Agent Syst.*, **28**(3), 474–518, (2014).
- [14] Stéphane Demri, Valentin Goranko, and Martin Lange, *Temporal logics in computer science: finite-state systems*, volume 58, Cambridge University Press, 2016.
- [15] G. Douéneau-Tabot, S. Pinchinat, and F. Schwarzenrüber, ‘Chain-monadic second order logic over regular automatic trees and epistemic planning synthesis’, in *AiML’18*, (2018).
- [16] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi, *Reasoning about Knowledge*, MIT Press, 1995.
- [17] B. Finkbeiner and S. Schewe, ‘Uniform distributed synthesis’, in *LICS*, pp. 321–330, (2005).
- [18] Paul Gastin, Nathalie Sznajder, and Marc Zeitoun, ‘Distributed synthesis for well-connected architectures’, *FMSD*, **34**(3), 215–237, (2009).
- [19] A. Herzig and N. Troquard, ‘Knowing how to play: uniform choices in logics of agency’, in *AAMAS 2016*, pp. 209–216. ACM, (2006).
- [20] Wojciech Jamroga and Thomas Ågotnes, ‘Constructive knowledge: what agents can achieve under imperfect information’, *Journal of Applied Non-Classical Logics*, **17**(4), 423–475, (2007).
- [21] O. Kupferman and M.Y. Vardi, ‘Synthesizing distributed systems’, in *LICS’01*, pp. 389–398, (2001).
- [22] Leslie Lamport, ‘What good is temporal logic?’, in *IFIP congress*, volume 83, pp. 657–668, (1983).
- [23] B. Maubert, *Logical foundations of games with imperfect information : uniform strategies.*, Ph.D. dissertation, University of Rennes 1, 2014.
- [24] B. Maubert and A. Murano, ‘Reasoning about knowledge and strategies under hierarchical information’, in *KR’18*, pp. 530–540, (2018).
- [25] Bastien Maubert, Sophie Pinchinat, and François Schwarzenrüber, ‘Reachability games in dynamic epistemic logic’, in *IJCAI’19*, pp. 499–505, (2019).
- [26] G. Peterson, J. Reif, and S. Azhar, ‘Decision algorithms for multiplayer noncooperative games of incomplete information’, *Comput Math Appl*, **43**(1), 179–206, (2002).
- [27] G.L. Peterson and J.H. Reif, ‘Multiple-person alternation’, in *FOCS’79*, (1979).
- [28] J. Pilecki, M. A. Bednarczyk, and W. Jamroga, ‘SMC: synthesis of uniform strategies and verification of strategic ability for multi-agent systems’, *J. Log. Comput.*, **27**(7), 1871–1895, (2017).
- [29] Sophie Pinchinat and Stéphane Riedweg, ‘A decidable class of problems for control under partial observation’, *JPL*, **95**(4), 454–460, (2005).
- [30] Jan Plaza, ‘Logics of public communications’, *Synthese*, **158**(2), 165–179, (2007).
- [31] A. Pnueli and R. Rosner, ‘Distributed reactive systems are hard to synthesize’, in *FOCS’90*, pp. 746–757, (1990).
- [32] Amir Pnueli and Roni Rosner, ‘On the synthesis of a reactive module’, in *POPL*, pp. 179–190, (1989).
- [33] Bernd Puchala, ‘Asynchronous omega-regular games with partial information’, in *MFCS’10*, pp. 592–603, (2010).
- [34] R. Ramanujam and S. Simon, ‘A communication based model for games of imperfect information’, in *CONCUR’10*, pp. 509–523, (2010).
- [35] John H. Reif, ‘The complexity of two-player games of incomplete information’, *J. Comput. Syst. Sci.*, **29**(2), 274–301, (1984).
- [36] Sven Schewe and Bernd Finkbeiner, ‘Distributed synthesis for alternating-time logics’, in *ATVA’07*, pp. 268–283, (2007).
- [37] Pierre-Yves Schobbens, ‘Alternating-time logic with imperfect recall’, *ENTCS*, **85**(2), 82–93, (2004).
- [38] Nikolay V Shilov, Natalya Olegovna Garanina, and K-M Choe, ‘Update and abstraction in model checking of knowledge and branching time’, *Fundamenta Informaticae*, **72**(1-3), 347–361, (2006).
- [39] R. van der Meyden and M.Y. Vardi, ‘Synthesis from knowledge-based specifications’, in *CONCUR’98*, 34–49, Springer, (1998).
- [40] R. van der Meyden and T. Wilke, ‘Synthesis of distributed systems from knowledge-based specifications’, in *CONCUR’05*, pp. 562–576, (2005).
- [41] H. van Ditmarsch, W. van der Hoek, and B. Kooi, *Dynamic Epistemic Logic*, Springer, Dordrecht, 2008.