Infinite games with finite knowledge gaps

Dietmar Berwanger[a], Anup Basil Mathew[a,b]

[a]*LSV, CNRS & ENS Cachan, Université Paris-Saclay, France*
[b]*Institute of Mathematical Sciences, Chennai, India*

Abstract

Infinite games where several players seek to coordinate under imperfect information are deemed to be undecidable, unless the information is hierarchically ordered among the players.

We identify a class of games for which joint winning strategies can be constructed effectively without restricting the direction of information flow. Instead, our condition requires that the players attain common knowledge about the actual state of the game over and over again along every play.

We show that it is decidable whether a given game satisfies the condition, and prove tight complexity bounds for the strategy synthesis problem under ω-regular winning conditions given by deterministic parity automata.

Key words: games on graphs, infinite games, imperfect information, distributed synthesis, coordination
2010 MSC: 05C57, 68M14, 91A06, 91A28, 93B50

1. Introduction

Automated synthesis of systems that are correct by construction is a persistent ambition of computational engineering. One major challenge consists in controlling components that have only partial information about the global system state. Building on automata and game-theoretic foundations, significant progress has been made towards synthesising finite-state components that interact with an uncontrollable environment either individually, or in coordination with other controllable components — provided the information they have about the global system is distributed hierarchically [1, 2]. For the general case, however, it was shown that the problem of coordinating two or more components of a distributed system with non-terminating executions is undecidable [3, 4].

The distributed synthesis problem can be formulated alternatively in terms of games between n players (the components) that move along the edges of a finite graph (the state-transitions of the global system) with imperfect information about the current position and the moves of the other players. The

Email addresses: dwb@lsv.fr (Dietmar Berwanger), anupbasil@imsc.res.in (Anup Basil Mathew)

outcome of a play is a possibly infinite path (system execution) determined by the joint actions of the players and moves of Nature (the uncontrollable environment). The players have a common winning condition: to form a path that corresponds to a correct execution with respect to the system specification, no matter how Nature moves. Winning conditions may be specified by finite-state automata, temporal logics, or in the canonical form of parity conditions. Thus, distributed synthesis under partial information corresponds to the problem of constructing a winning profile of finite-state strategies in a coordination game with imperfect information. This problem was shown to be undecidable by Peterson and Reif [5], already for the basic setting of two players with a reachability condition; infinitary winning conditions, which lead to higher degrees of undecidability, have been studied by Janin [6].

As in the case of distributed systems, decidable classes of coordination games rely on restrictions of the information flow according to an order among the players [7, 8, 9]. In their survey article on the complexity of multiplayer games, Azhar, Peterson, and Reif conclude that "[i]n general, multiplayer games of incomplete information can be undecidable, unless the information is hierarchically arranged"[7, p. 991].

The undecidability arguments cited above share a basic scenario: two players become uncertain about the current state of the game, due to moves of Nature. The structure of the game requires them to take into account not only their first-order uncertainty about the actual state, but also the higher-order uncertainty of one player about the knowledge of the other. Finally, the players can win only by attaining common knowledge about a property of the actual history, which requires them to maintain knowledge hierarchies of increasing height, as the play proceeds. The scenario is set up so that the uncertainty never vanishes and the knowledge hierarchies grow unboundedly, which leads to undecidability [10].

One systematic approach to characterising undecidable classes of problems that involve multiple players with imperfect information is the *information fork* criterion formulated by Finkbeiner and Schewe [11]. The criterion applies to the distributed synthesis problem in the basic setting of Pnueli and Rosner [3], for fixed architectures with synchronous communication channels. Intuitively, an architecture has an information fork if it allows for two players to reach a situation in which neither one can infer the observation received by the other player from his own observation. Under this condition, distributed architectures may allow the knowledge of players to diverge over an unbounded number of rounds for particular specifications. Conversely, all architectures that do not contain an information fork admit an information ordering among the players and are therefore decidable.

When applied to games, however, the information fork criterion yields only a coarse classification for decidability, as the parametrisation over architectures has no natural correspondent in terms of game graphs. Indeed, the set of game instances obtained by modelling a given family of distributed systems may be solvable, in spite of possible information forks in the underlying architecture. This can occur, for instance, if the information flow affected by the fork is inessential to the players, or if the divergence between the knowledge of players vanishes after few rounds.

In this article, we identify a new condition for the decidability of coordination games with imperfect information that does not rely on the hierarchical arrangement of information. Similar to the information fork approach, our focus is on situations in which the knowledge of players diverges. We use the term *knowledge gap* to describe an interval of rounds at which the players do not attain common knowledge about the actual state. Essentially, our condition requests that all knowledge gaps of a game are finite, or in other words, that the players attain common knowledge of the actual game state infinitely often, along every play. In this case we say that the game allows for *recurring common knowledge of the state*.

Questions about common knowledge in infinite runs are typically hard. In their study of the model checking complexity of epistemic temporal logics [12], van der Meyden and Shilov point out that already the problem of determining whether the players attain common knowledge about an atomic property is undecidable, even for synchronous models as we consider here. Indeed, it turns out that it is undecidable whether the players attain common knowledge of the state at any history within a given part of the game graph (Proposition 4.3).

Surprisingly, the situation improves when we look at the recurring formulation relevant for our characterisation: We are able to show that the question of whether the common-knowledge property holds infinitely often, on every play in a game is decidable with low complexity. This has several favourable consequences for solving infinite coordination games with imperfect information.

Our results are summarised as follows:

(1) The question of whether a game for n players with imperfect information satisfies the condition of recurring common knowledge of the state is decidable in NLogSpace.

(2) If a coordination game for n players with imperfect information satisfies the condition of recurring common knowledge of the state, then the problem of whether a joint winning strategy exists is decidable, and it is NExpTime-complete.

(3) If there exists a joint winning strategy in a game with recurring common knowledge of the state, then there also exists a profile of finite-state strategies of exponential size, which can be synthesised in 2ExpTime.

The conclusions rely on three key arguments. Firstly, we show that under recurring common knowledge of the state, the intervals where the current state of the game is not common knowledge are bounded uniformly. This implies that the *perfect-information tracking* of such a game is finite, which yields decidability of the strategy synthesis problem as a consequence of a metatheorem from [13]. Secondly, we characterise recurring *common* knowledge in terms of recurring *mutual* knowledge. This allows us to establish tight complexity bounds. Finally, we prove that the problem of solving imperfect-information games with recurring common knowledge of the state can be reduced to solving parity games with perfect information, at a relatively low cost in terms of complexity.

2. Basic notions

2.1. Coordination games with imperfect information

Our game model is close to that of concurrent games [14]. There are n players 1, ..., n and a distinguished agent called Nature. The *grand coalition* is the set $\{1, \ldots, n\}$ of all players. We refer to a list of elements $x = (x^i)_{1 \leq i \leq n}$, one for each player, as a *profile*.

For each player i, we fix a set A^i of *actions* and a set B^i of *observations*, finite unless stated otherwise. The *action space* A consists of all action profiles. A *game graph* $G = (V, E, (\beta^i)_{1 \leq i \leq n})$ consists of a finite set V of nodes called *states*, an edge relation $E \subseteq V \times A \times V$ representing simultaneous *moves* labelled by action profiles, and a profile of *observation* functions $\beta^i : V \to B^i$ that label every state with an observation, for each player. We assume that for every state v and every action profile a there is an outgoing move $(v, a, v') \in E$. For convenience, we will include special sink states from which any outgoing move is a self loop.

Plays start at an initial state $v_0 \in V$ known to all players and proceed in rounds where each player i chooses an action $a^i \in A^i$, then Nature chooses a successor state v' reachable via a move $(v, a, v') \in E$, and each player i receives the observation $\beta^i(v')$. Notice that the players are not informed about the action chosen by other players, nor about the state chosen by Nature.

Formally, a *play* is an infinite sequence $\pi = v_0, a_1, v_1, a_2, v_2, \ldots$ alternating between positions and action profiles with $(v_\ell, a_{\ell+1}, v_{\ell+1}) \in E$, for all $\ell \geq 0$. A *history* is a prefix $\pi = v_0, a_1, v_1, \ldots, a_\ell, v_\ell$ of a play; we refer to ℓ as the *length* of the history. For convenience, we omit commas in the sequence notation and write plays and histories as words $\pi = v_0\, a_1 v_1\, a_2 v_2 \ldots$. Whenever we refer to a finite prefix ρ of a play or history π, we mean a history $\rho \in V(AV)^*$ such that $\pi = \rho\tau$ for some τ in $(AV)^*$ or $(AV)^\omega$; further, we call π a *prolongation* and τ a *continuation* of ρ.

The observation function is extended from states to histories and plays by setting $\beta^i(\pi) := \beta^i(v_0)\, a_1^i \beta^i(v_1) \ldots$ We say that two histories π, π' are *indistinguishable* to Player i, and write $\pi \sim^i \pi'$, if $\beta^i(\pi) = \beta^i(\pi')$. This is an equivalence relation, and its classes are called the *information sets* of Player i. A game (graph) with *perfect information* is one where all information sets are singletons. In general, we do not assume that this is the case, so we speak about games with *imperfect information*.

When viewed as a distributed system in the taxonomy of Halpern and Vardi [15], our game model belongs to the class of *synchronous* systems with *perfect recall*. This is implicit in our definition of observation functions: the players are able to distinguish between histories of different length (synchronicity), and if two histories are indistinguishable for a player i at round ℓ, then so are they at any previous round $r < \ell$ (perfect recall).

A *strategy* for Player i is a mapping $s^i : V(AV)^* \to A^i$ from histories to actions such that $s^i(\pi) = s^i(\pi')$, for any pair $\pi \sim^i \pi'$ of indistinguishable histories. We denote the set of all strategies of Player i by S^i and the set of all strategy profiles by S. A history or play $\pi = v_0\, a_1 v_1 \ldots$ *follows* the strategy

$s^i \in S^i$ if $a^i_{\ell+1} = s^i(v_0\, a_1 v_1 \dots a_\ell v_\ell)$, for every $\ell \geq 0$. For the grand coalition, the play π follows a strategy profile $s \in S$ if it follows all strategies s^i. The set of possible *outcomes* of a strategy profile s is the set of plays that follow s.

A general *winning condition* over a game graph G is a set $W \subseteq (VA)^\omega$ of plays. A *coordination game* $\mathcal{G} = (G, W)$ is described by a game graph and a winning condition. We say that a play π on G is winning in $\mathcal{G}$ if $\pi \in W$. A strategy profile s is winning in $\mathcal{G}$, if all its possible outcomes are so. In this case, we refer to s as a *joint winning strategy.*

With a view to effective algorithms for synthesising strategies, we focus on finitely presented games where the winning condition is described by a colouring function $\gamma : V \to C$ with a finite range of colours, and an ω-regular set $W \subseteq C^\omega$ given, e.g., by finite-state automaton. Then, a play $v_0\, a_1 v_1 \dots$ is winning if $\gamma(v_0)\gamma(v_1) \dots \in W$. We generally assume that the colouring is *observable* to each player i, that is, $\beta^i(v) \neq \beta^i(v')$ whenever $\gamma(v) \neq \gamma(v')$. Given such a game, the *distributed synthesis problem* consists of two tasks: (1) to decide whether there exists a joint winning strategy, and (2) to construct a winning profile of finite-state strategies, if any exist. These are strategies implemented by automata that input observations and output actions. For more background on finite-state strategy synthesis we refer to the expository article of Thomas [16].

For lower bounds, we refer to simple safety conditions which require the players to avoid an observable sink $\ominus$. The technical results on upper bounds are formulated in terms of *parity* winning conditions represented by a coloring function $\gamma : V \to \mathbb{N}$ that maps every state to a number called *priority*: A play is winning if the least priority seen infinitely often along a play is even. Parity conditions provide a canonical form for observable ω-regular winning conditions, in the sense that each game with a regular condition can be reduced to one with a parity condition such that the solution of the synthesis problem is preserved. The reduction for the perfect-information setting described by Thomas in the handbook chapter [17] generalises to imperfect-information games with observable winning conditions, as pointed out in [13].

2.2. Domino tiling problems

As a tool for proving lower complexity bounds, we use domino tiling problems, which allow a more transparent representation of combinatorial problems than encoding machine models. Our exposition follows the notation of Börger, Grädel, and Gurevich [18].

A *domino system* $\mathcal{D} = (D, E_H, E_V)$ is described by a finite set D of *dominoes* together with horizontal and vertical compatibility relations $E_H, E_V \subseteq D \times D$. The generic domino tiling problem is to determine, for a given system $\mathcal{D}$, whether copies of the dominoes can be arranged to cover a region $Z \subseteq \mathbb{Z} \times \mathbb{Z}$, such that any two vertically or horizontally adjacent dominoes are compatible. Here we consider finite rectangular regions $Z(\ell, m) := \{0, \dots, \ell+1\} \times \{0, \dots, m\}$ where the first and the last column, and the bottom row are distinguished as border areas to be tiled with special dominoes $\#$ and $\Box$. The concrete question is whether there exists a *tiling* $\tau : Z(\ell, m) \to D$ that assigns to every point in the region a domino, subject to the border constraints:

- $\tau(x, 0) = \Box$, for all $x = 1, \ldots, \ell$, and
- $\tau(0, y) = \tau(\ell + 1, y) = \#$, for all $y = 0, \ldots, m$,

and the compatibility constraints, for all $x \leq \ell$ and $y < m$:

- if $\tau(x, y) = d$ and $\tau(x + 1, y) = d'$ then $(d, d') \in E_H$, and
- if $\tau(x, y) = d$ and $\tau(x, y + 1) = d'$ then $(d, d') \in E_V$.

In addition, we may specify constraints on the *frontier* of the tiling, that is, the sequence $\tau(1, m), \tau(2, m), \ldots, \tau(\ell, m)$ of dominoes in the top row. To ensure that border dominoes do not appear at the interior of a correct tiling and to avoid the trivial tiling, we generally assume that $E_V \subseteq D \times (D \setminus \{\#, \Box\}) \cup \{(\#, \#)\}$ and $(\Box, \Box) \notin E_H$.

We use three variants of the domino problem. Firstly, the Corridor Tiling problem takes as input a domino system $\mathcal{D}$ together with a frontier constraint $w \in D^\ell$ and asks whether there exists a height m such that the region $Z(\ell, m)$ allows a tiling τ that additionally satisfies:

- $\tau(i, m) = w_i$, for all $i = 1, \ldots, \ell$.

Secondly, the Corridor Universality problem takes as input a domino system $\mathcal{D}$ together with a subset of dominoes $\Sigma \subseteq D$ and asks whether for *all* frontier constraints $w \in \Sigma^\ell$ of arbitrary length $\ell > 0$, there exists a height m such that the region $Z(\ell, m)$ allows a corridor tiling.

The basic variant of corridor tiling is a well-known PSpace-complete problem [19]. One way to explain the complexity of both variants is via the correspondence between context-sensitive languages and domino systems, established by Latteux and Simplot [20, 21]. The *frontier language* of a domino system $\mathcal{D}$ is the set $L(\mathcal{D})$ of words $w \in D^*$ such that $(\mathcal{D}, w)$ yields a positive instance of the Corridor Tiling problem. We refer to standard notions on context sensitive languages as found, for instance, in the handbook [22, Chapter 3].

Theorem 2.1 ([20, 21])**.** *For every context-sensitive language $L \subseteq \Sigma^*$ given as a linear bounded automaton, one can construct in polynomial time a domino system $\mathcal{D}$ over a set of dominoes $D \supseteq \Sigma$, such that $L(\mathcal{D}) \cap \Sigma^* = L$.*

Via this correspondence, the membership problem for context-sensitive language, which is PSpace-complete, reduces to Corridor Tiling and the universality problem for context-sensitive languages, which is undecidable, reduces to Corridor Universality. Converse reductions from domino tiling to context-sensitive language problems can also be done in polynomial time, by translating domino systems into linear-bounded automata.

Theorem 2.2 ([19], [20, 21]).

(i) Corridor Tiling *is* PSpace-*complete.*

(ii) Corridor Universality *is undecidable.*

Finally, the Exp-Square Tiling problem takes as input a domino system together with a number $\ell \in \mathbb{N}$ in binary encoding, and asks whether the region $Z(\ell, \ell)$ allows a correct tiling. The problem was first studied by Fürer [23].

Theorem 2.3 ([23]). Exp-Square Tiling *is* NExpTime-*complete.*

2.3. Common knowledge

We use the notion of knowledge in the sense of having information. That Player i knows proposition φ at history π should mean that, from the structure of the game graph and the sequence $\beta^i(\pi)$ of observations she received, it can be inferred that φ holds. Specifically, we are interested in propositions about play histories. To formalise knowledge and uncertainty, we rely on the standard semantic model introduced by Aumann [24] and follow the treatment of Osborne and Rubinstein [25, Chapter 5]. For an extensive account of distributed knowledge in computational systems, we refer the reader to the book of Fagin, Halpern, Moses, and Vardi [26, Chapters 10, 11] and, as a standard reference on common knowledge, to the handbook chapter of Geanakoplos [27]. The enlightening article of [28] addresses foundational issues about the formalisation of common-knowledge.

Let us fix a game graph G and denote by Ω the set of histories. The *possibility* correspondence $P^i : \Omega \to \mathcal{P}(\Omega)$ associates to each history π its information set:

$$P^i(\pi) := \{\pi' \in \Omega \mid \pi' \sim^i \pi\}, \quad \text{for every player } i.$$

Thus, at history π, Player i knows that the actual history is in $P^i(\pi)$, but he may be uncertain which one it is. The sets $P^i(\pi)$ form a partition of Ω. Observe that each information set $P^i(\pi)$ consists of histories of the same length as π, hence it is finite.

An *event* is a subset $F \subseteq \Omega$. We say that F *occurs* at history π if $\pi \in F$. The *knowledge* operator $K^i : \mathcal{P}(\Omega) \to \mathcal{P}(\Omega)$ associates to every event F the set of histories at which Player i knows that F occurs:

$$K^i(F) := \{\pi \in \Omega \mid P^i(\pi) \subseteq F\}, \quad \text{for every player } i.$$

Note that $K^i(F)$ is itself an event. If $\pi \in K^i(F)$, then (the occurrence of) F is private knowledge of Player i at π: For any $\pi' \sim^i \pi$, it holds that $\pi' \in F$. If, moreover, $\pi \in K^i(F)$, for all players i, we say that F is *mutual knowledge* among the players at π. In this case, every player knows that F occurs, although one player may be uncertain about whether another player knows it.

An event $F \subseteq \Omega$ is *common knowledge* at π if for every sequence of histories $\pi_1, \ldots, \pi_k$ and players $i_1, \ldots, i_k$ such that $\pi \sim^{i_1} \pi_1 \sim^{i_2} \cdots \sim^{i_k} \pi_k$, it is the case that $\pi_k \in F$. In other words, π belongs to the image of F under every iteration $K^{i_1}(K^{i_2}(\dots K^{i_k}(F)\dots))$ of knowledge operators.

We will use an alternative characterisation in terms of shared information. An event F is *self-evident* if it is mutual knowledge among the players at every history in F, that is, if $F \subseteq K^i(F)$ for all players i. As the converse inclusion $K^i(F) \subseteq F$ always holds, this amounts to saying that F is a simultaneous fixed point of the player's knowledge operators. Self-evident sets allow an interpretation of common knowledge that coincides with the iterated-knowledge interpretation, if the situation model is sufficiently simple (see Barwise [28], for a thorough analysis), particularly if the sample space Ω is finite. Although in our setting Ω is infinite, for histories π of length ℓ, only the finite space Ω_ℓ of histories of the same length matters: An event $F \subseteq \Omega$ is mutual or common knowledge at π if, and only if, this holds for the event $F \cap \Omega_\ell$. Therefore, the argument for the finite setting given, for instance, in the handbook chapter of Geanakoplos [27, Section 6], justifies the following characterisation.

Theorem 2.4 ([24])**.** *An event $C \subseteq \Omega$ is common knowledge at history π, if and only if, there exists a self-evident event $F \subseteq C$ with $\pi \in F$.*

3. Uncertainty and coordination

Under perfect information, coordination games are trivial to solve, by considering the two-player zero-sum game where the observations and the actions of the grand coalition are attributed to the first player and the role of Nature is played by the second player. Then, any winning strategy of the first player can be viewed as a joint winning strategy and vice versa. Intuitively, players of the coalition can act as one because each player knows the actual history when he chooses an action. Unlike the case considered in [29], where players choose their strategy independently, in the setting of distributed strategies there is no risk of discoordination due to strategic uncertainty: The joint strategy is centrally planned, it is common knowledge among the players.

Under imperfect information, the problem is more complex, because players may not know where they are in the game. Strategies need to adjust to the uncertainty around the current history, which is induced by moves of Nature. The prescribed actions should be suitable in any contingency of the unobservable state of nature. In the interaction between a single player and Nature (or, equivalently, between two players with strictly conflicting interests), the knowledge relevant to this task is of first order: it regards only the set of possible contingencies, that is, the information set.

Yet, in games among several players, whether an action of one player is suitable or not at a particular history may depend on whether another player chooses a matching action at the same history. He, the one player, should thus be certain that she, the other player, would play the matching action, according to her commonly known strategy, which, however, responds to the observations she received on her own side. In other words, to avoid discoordination, he needs

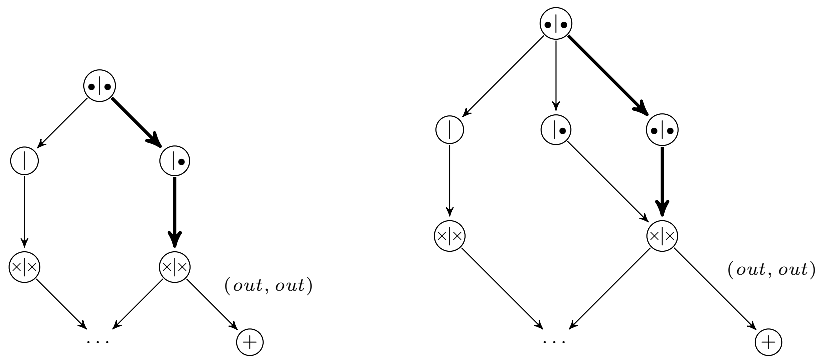

(a) first-order

(b) second order

Figure 1: Lacking knowledge to coordinate

to know about what she knows about the current history. In contrast to the one-player or the two-player conflict case, here it is relevant to consider higher-order knowledge, i.e., knowledge about the knowledge of other players.

The role of knowledge is particularly obvious in coordination games where the actions of players must agree at every history. Formally, we call a *consensus game* a coordination game with a set of actions that is common to all players, and where every move $(v, a, w) \in E$ in which two players $i \neq j$ disagree on their actions $a^i \neq a^j$ leads to a special sink state $\ominus$ from which no play is winning. A necessary condition for a strategy profile s to be winning in a consensus game is that, for every history π that follows s, all components prescribe the same action, that is, $s^i(\pi) = s^j(\pi)$, for all players i, j. When speaking about a winning strategies in such a game, we may therefore identify any strategy s^i of an individual player with the profile s of strategies where all components are equal to s^i, without loss of generality.

Figures 1 and 2 show examples of consensus games for two players, Player 1 (he) and 2 (she), with actions *in* and *out*; unlabelled arcs represent moves where both players choose *in*. The observations $\bullet$, $\circ$, and $\times$ are indicated in split states: he sees the left side, she the right side. Apart from the unsafe sink $\ominus$ that also collects the moves along any unrepresented action profiles, there is a safe sink $\oplus$, from which all plays are winning; the sinks are observable to both players. The dots on the bottom stand for an arbitrary continuation. For each of these games, we consider the situation where the actual history corresponds to the rightmost path, marked by thicker arrows leading to the *out* state where playing (out, out) would lead to an immediate win. Along the examples, we will discuss the question of whether the information that players have at the marked history allows them to infer with certainty that (out, out) is a safe action.

In Figure 1(a), at the marked history, Player 2 knows about being in the *out* state which requires the action (out, out) to win. However, Player 1 cannot distinguish the current history from the one along the left path $(\bullet|\bullet)(\circ|\circ)(\times|\times)$, where playing *out* would be losing. So it occurs that at the current state (out, out) is the right joint move, but Player 1 lacks first-order knowledge about it, whereas Player 2 has the relevant first-order knowledge, but not the second-order knowledge to ascertain that Player 1 will play *out*. At best, the players could coordinate on (in, in), based on their common knowledge that this leads to continuing the game, and not straight to the $\ominus$ sink.

In Figure 1(b), both players know that they are in the *out* state. Nevertheless, Player 2 is uncertain about whether Player 1 knows it, because according to her observation, the current history may be $(\bullet|\bullet)(\circ|\bullet)(\times|\times)$, in which case, just as in Figure 1(a), Player 1 would consider the history $(\bullet|\bullet)(\circ|\circ)(\times|\times)$ possible, and thus not play *out*. So both players have first-order knowledge about being at the *out* state, Player 1 even has the second-order knowledge that Player 2 knows it, and still they cannot coordinate with certainty, because Player 2 does not know that Player 1 knows it. Moreover, in Figure 2(a), both players know that the play is at the *out* state and each of them knows that they both know it. But Player 1 regards it as possible that Player 2 observed $\bullet\bullet\circ\times$, again raising the uncertainty of Figure 1(b). Here, the reason why the players cannot coordinate is that Player 1 does not know that Player 2 knows that Player 1 knows about being in a *out* state.

The argument can be lifted to arbitrary levels of the knowledge hierarchy. This is illustrated in Figure 2(b), where the loop around the observation $(\bullet|\bullet)$ may be unravelled n times to obtain an instance where coordination on the winning action fails in spite of the players having mutual knowledge of order n about being in a state where this action is safe.

Indeed, the examples embody the *coordinated attack problem*, a parable that illustrates the intricacy of coordination via unreliable communication. The story features two generals camped with their armies on two hills surrounding a fortification that they plan to attack. As either one alone would lose the battle, they need to agree on attacking simultaneously. However, they can only communicate by sending messengers, which may be captured on the way. The challenge is to attain common knowledge about being in a state where the attack can occur after a finite history of message exchanges, given that whenever a general receives a message, he is uncertain of whether the sender knows that he received it. Proofs that this cannot be achieved have been given for different settings, e.g., by Gray [30], and Halpern and Moses [31], in the distributed-systems literature, and by Rubinstein [32] in game theory. In our setting, it is Nature that induces nondeterministically a possible loss of one message between the two generals who would attack if, and only if, they had common knowledge of being in the *out*-state. The analyses put forward the paradigm that that if common knowledge is not attainable, then coordination is impossible in spite of an arbitrarily high level of mutual knowledge.

As our examples suggest, already in the simple case of consensus games with a safety condition (avoid the $\ominus$-sink), coordination games with imperfect information are sensitive to common knowledge, and thus vulnerable to problems caused by its inapproximability through finite levels of mutual knowledge. Still, one may argue that the problem of synthesising a joint winning strategy does

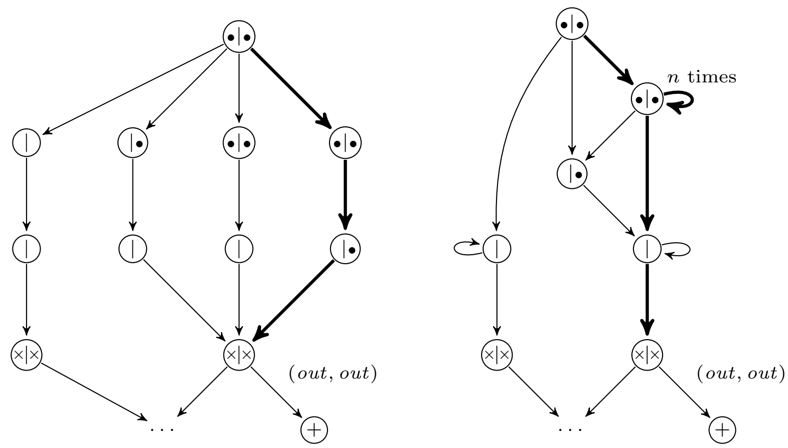

(a) third order

(b) arbitrary order n

Figure 2: Coordinated attack

not invoke the reasoning process of individual players. In the end, strategies only rely on first-order knowledge. Nevertheless, we will show in the remainder of this article that the problem of attaining common knowledge about certain events — namely, the game state at the actual history — is relevant for solving coordination games, in the following sense.

(i) There exists a family of games that admit a solution if, and only if, the players attain common knowledge about the state at a particular history.

(ii) Every game of infinite duration where common knowledge of the actual state is attained infinitely often along every play can be solved effectively.

4. Common knowledge of the state

Let G be a game graph, and let Ω be the set of histories in G. For a history π, we denote by $Ter(\pi) \subseteq \Omega$ the set of all histories that end at the same state as π. We say that the players attain *common knowledge of the state* (CKS) at history π, if $Ter(\pi)$ is common knowledge at π.

To develop familiarity with the notion, we first show that attaining CKS at a particular history in a game graph is equivalent to having a joint winning strategy in an associated coordination game, more precisely, a consensus game. For simplicity, we detail the case where Nature controls all moves in the original game graph, i.e., the players have only one, trivial action; the general case requires only notational changes.

Given a game G, let π be a history that starts at the initial state v_0 and ends at some state $z \in V$. We construct a consensus game $\mathcal{G}_\pi$ on the the disjoint union of G and (the unravelling of) π as follows. The players are the same as in G, and they have a common set $\{in, out\}$ of actions. The state set of $\mathcal{G}_\pi$ consists of copies of the states in G and a fresh state $\hat{\tau}$ for every history τ in π — note that the copy of the initial state $v_0 \in V_0$ is distinct from the initial history $\hat{v}_0$. For each state copy, the observations get inherited from the corresponding state in G, and for each history from its last state. Likewise, the moves from G and along π get inherited with the action label *in* for all players (in consensus). The initial-history state $\hat{v}_0$ is designated as the initial state of $\mathcal{G}_\pi$, and we add *in* moves from $\hat{v}_0$ to every successor of v_0 in G. There is an unsafe sink $\ominus$ and a safe sink $\oplus$; the winning condition requires to avoid $\ominus$. Finally, we add *out* moves (in consensus) to the safe sink $\oplus$ from the copy of state z in G and from the state $\hat{\pi}$ corresponding to the history π. Moves with any other action profile lead to the unsafe sink $\ominus$, namely any action profile that is not in consensus, the action *out* from any state other than z or $\hat{\pi}$, and the action *in* from $\hat{\pi}$.

Proposition 4.1. *For a game graph G and a history π, the players have a joint winning strategy in the game $\mathcal{G}_\pi$ if, and only if, they attain common knowledge of the state at π in G.*

Proof. To see that winning in $\mathcal{G}_\pi$ implies attaining CKS at π, suppose that there exists a joint winning strategy s in $\mathcal{G}_\pi$. Since this is a consensus game, we may assume that all components of s are equal and can identify the profile s and its component strategies. Now, let C be the set of histories ρ in G that follow s and are assigned $s(\rho) = out$. We argue that C satisfies the conditions of Theorem 2.4 to witness that the players attain CKS at π:

- C is a self-evident event, for each player i: for every history $\rho \in C$, any indistinguishable history $\rho' \sim^i \rho$ follows s and is assigned the same action $s(\rho') = s(\rho) = out$, which means $\rho' \in C$.

- $\pi \in C$: the action *in* is losing at $\hat{\pi}$, and since no winning strategy can avoid $\hat{\pi}$, we must have $s(\hat{\pi}) = out$. As π and its copy ending at $\hat{\pi}$ are indistinguishable to all players, it follows that $s(\pi) = out$.

- $C \subseteq Ter(\pi)$: the action *out* is losing at all states except for z and $\hat{\pi}$. As we assumed that s is a winning strategy, all histories in C must end at z.

For the converse, assume that, at the history π in G, the players attain common knowledge of $Ter(\pi)$. We define a function s that associates to any history ρ in G the action $s(\rho) := out$ if ρ has the same length as π and $Ter(\pi)$ is common knowledge at ρ, and otherwise $s(\rho) := in$.

First, let us verify that s is a valid strategy on the game graph G, for each player i. Notice that, if an event is common knowledge at a history ρ, then it is also common knowledge at every indistinguishable history $\rho' \sim^i \rho$ (each history that is accessible from ρ' via a sequence of pairwise indistinguishable histories is also accessible from ρ via the same sequence preceded by $\rho \sim^i \rho'$). In particular, whenever $s(\rho) = out$, for a history ρ, that is, when the event $Ter(\pi)$ is common knowledge at ρ, it is also common knowledge at every history $\rho' \sim^i \rho$, hence $s(\rho') = out$. Consequently $s(\rho) = s(\rho')$ for every pair $\rho \sim^i \rho'$.

Now, we extend the strategy s on G to the graph of $\mathcal{G}_\pi$ in the only consistent way, by assigning to every history that ends at a state $\hat{\rho}$ corresponding to a history ρ in π, the action $s(\rho)$ prescribed in G. We argue that s is a winning strategy in the consensus game $\mathcal{G}_\pi$: The unsafe sink $\ominus$ can only be reached by taking a wrong move in one of the following two situations: either choosing *in* at state $\hat{\pi}$, or choosing *out* at a state different from z and $\hat{\pi}$. The former situation is excluded, as $s(\hat{\pi}) = s(\pi) = out$ holds by definition of s. The latter situation cannot occur either: On the one hand, at all histories ρ in G with $s(\rho) = out$, the players attain common knowledge of $Ter(\pi)$, so in particular, $\rho \in Ter(\pi)$ (by definition of the knowledge operator, players can only know an event if it actually occurs). On the other hand, among the histories along the copy of π, only the one ending at $\hat{\pi}$ has the same length as π, which is necessary to be asigned *out*. In conclusion, all plays that follow s are winning in $\mathcal{G}_\pi$. □

The argument illustrates that the need for (common) knowledge about the actual game state is a source of computational complexity in coordination games with imperfect information. In general, there may be further sources. One class of games, where the issue is exactly whether the players attain common knowledge about reaching a certain state set, are *consensus acceptors* investigated in [33]. Essentially, these are consensus games with a simple safety condition (avoid the sink $\ominus$) where the players have only one nontrivial decision in every play. Which decision to take thus depends on the common knowledge of the players about the actual state. The games $\mathcal{G}_\pi$ from Proposition 4.1 as well as those represented in Figures 1 of the previous section are examples of consensus acceptor games. The complexity analysis for consensus acceptors in [33] sheds light on the problem of attaining common knowledge of the state in an arbitrary game graph. For completeness, we reproduce the part of the analysis relevant for our setting.

Proposition 4.2. *Given a game graph and a history π, the problem of deciding whether the players attain common knowledge of the state at π is* PSpace-*complete.*

Proof. For membership, consider the procedure that takes a game graph and a history π as input, and iterates the following loop: guess nondeterministically a player i and a history $\rho \sim^i \pi$; accept if π and ρ end at different states, otherwise repeat with ρ as the new value of π. [1] As any two indistinguishable histories have the same length, the procedure requires only linear space. It accepts if, and only if, there exists a sequence $i_1, \dots, i_k$ of players and histories $\rho_1 \sim^{i_1} \rho_2 \sim^{i_2} \cdots \sim^{i_k} \rho_k$ with $\pi = \rho_1$ and $\rho_k \notin Ter(\pi)$, that is, if $Ter(\pi)$ is not

[1] We adopt the convention that machines reject by looping; Hopcroft and Ullman [34] showed that lower space bounds above NLogSpace do not change when dropping the halting assumption.

common knowledge at π. Hence, we have a nondeterministic PSPACE procedure for deciding the complement problem which asks whether the players do *not* have CKS common knowledge of the state at the given history. Since nondeterministic and deterministic PSPACE are equal, it follows that the original problem can be solved in PSPACE. Even more, the argument shows that for any game there exists a linear-bounded automaton that recognises the set of histories at which the players attain CKS.

To prove hardness, we describe a reduction from (the complement of) the CORRIDOR TILING problem. Given a domino system $\mathcal{D} = (D, E_V, E_H)$ with a frontier constraint $w \in D^\ell$ we construct, in polynomial time, a game graph G for two players and a history π, such that the players attain CKS at π if, and only if, the domino-problem instance $(\mathcal{D}, w)$ is negative, i.e., there does not exist a height $m \in \mathbb{N}$ such that the rectangle $Z(\ell, m)$ can be tiled with w in the top row.

The two players in G have one trivial action and their observations correspond to the dominoes in D. The set of states consists of *singleton* states $d \in D \setminus \{\#\}$, *pair* states $(d, b) \in E_V$, an initial state v_0, and two sinks $\oplus$ and $\ominus$. At each singleton state d, both players receive the same observation d, whereas at each pair state (d, b), the first player observes d and the second player b; at the initial state and the two sinks, both players receive the observation $\#$ corresponding to the vertical border domino.

The singleton states are connected by moves $d \to d'$ for every (d, d') in E_H, and the pair states by moves $(d, b) \to (d', b')$ whenever (d, d') and (b, b') are in E_H. From the initial state v_0, there are moves to all singleton states d with $(\#, d)$ in E_H, and all pair states (d, b) with $(\#, d)$ and $(\#, b)$ in E_H. Conversely, the sink $\oplus$ is reachable from all singleton states d with $(d, \#)$ in E_H, and from all pair states (d, b) with $(d, \#)$ and $(b, \#)$ in E_H; the sink $\ominus$ is reachable only from the singleton bottom-domino state $\Box$. Clearly, the game graph G can be constructed from $\mathcal{D}$ in linear time.

Note that any sequence $x = d_1, d_2, \ldots, d_\ell \in D^\ell$ that forms a horizontally consistent row (omitting the borders) in a corridor tiling corresponds to a history $\pi_x = v_0 d_1 d_2 \ldots d_\ell \oplus$ in the game graph. For the special case of the bottom row *bot* with $d_1 = d_2 = \cdots = d_\ell = \Box$, apart of π_{bot}, we also have the history $\hat{\pi}_{bot} = v_0 d_1 d_2 \ldots d_\ell \ominus$. On the other hand, every history in G that ends at a sink corresponds either to one consistent row, in case Nature chooses a singleton state in the first move, or to two rows, in case Nature chooses a pair state. Moreover, a row $x = d_1, d_2, \ldots, d_\ell$ can appear below a row $y = b_1, b_2, \ldots, b_\ell$ in a correct tiling if, and only if, there exists a history ρ in G such that $\pi_x \sim^1 \rho \sim^2 \pi_y$, namely $\rho = v_0\,(d_1, b_1)\,(d_2, b_2)\ldots(d_\ell, b_\ell)\,\oplus$.

Now, we claim that the players attain CKS at the history π_w corresponding to the frontier constraint $w \in D^\ell$ if, and only if, there exists no corridor tiling for the instance $(\mathcal{D}, w)$. According to our observation, if there exists a correct tiling of the corridor, then there exists a sequence of rows corresponding to histories $\pi_1, \ldots, \pi_m$, and a sequence of witnessing histories $\rho_1, \ldots, \rho_{m-1}$ such that

$$\pi_w = \pi_1 \sim^1 \rho_1 \sim^2 \pi_2 \cdots \sim^1 \rho_{m-1} \sim^2 \pi_m = \pi_{bot}.$$

However, the history π_{bot} is indistinguishable from $\hat{\pi}_{bot}$, for both players. As these two histories end at different states, it follows that $Ter(\pi_w)$ is not common knowledge at π_w. Conversely, if $Ter(\pi_w)$ is not common knowledge at π_w, then there exists a sequence of pairwise indistinguishable histories that leads from π_w to $\hat{\pi}_{bot}$, from which we can extract a correct corridor tiling.

Corridor Tiling is PSpace-hard, according to Theorem 2.2. Thus, the reduction shows that the problem of deciding whether the players attain CKS at a given history is Co-PSpace hard, and since PSpace is closed under complement, it is hard for PSpace. $\square$

The above argument can be extended to prove that it is undecidable whether the players can ever attain CKS along a given set of plays in a game. We say that a history π is *within* a subset $S \subseteq V$ of states if all states that occur in π belong to S.

Proposition 4.3. *It is undecidable whether, for a given game graph with a designated subset $S \subseteq V$ of states, there exists a nontrivial history within S at which the players attain common knowledge of the state.*

Proof. We proceed by reduction from Corridor Universality: Given a domino system $\mathcal{D}$ with a designated subset $\Sigma \subseteq D$, we construct a game graph G with a subset S of states such that the players attain CKS at some nontrivial history π within S if, and only if, there exists a frontier constraint $w \in \Sigma^\ell$ that does not allow a corridor tiling with $\mathcal{D}$; actually, the sequence of observations along π will yield such a constraint.

The construction of G is as in the proof of Proposition 4.2 except that we take two disjoint copies of the game graph associated to $\mathcal{D}$ and identify the two copies of the initial state v_0, and those of the sinks $\ominus$, $\oplus$, respectively. In this way, no player knows the actual state of any history that does not reach a sink. The set S consists of the initial state, the two sinks, and the states corresponding to singleton dominoes in Σ (from both copies).

Now, every nontrivial history π that ends at a sink corresponds either to a single row or to a pair of horizontally consistent rows, depending on whether π proceeds through singletons or pair states. This holds by the same argument as in the proof of Proposition 4.2. Likewise, it follows that the two players do not attain CKS at a history π if, and only if, the row corresponding to π, or either one of the two rows, appears in the frontier of some correct corridor tiling with $\mathcal{D}$. For histories within S, when π is of the form $\#w\#$ for some (nonempty) word $w \in \Sigma^*$, this means that the players do not attain CKS if, and only if, there exists a correct corridor tiling with w in the frontier.

In conclusion, there exists a history within S at which the players attain CKS in G if, and only if, the Corridor Universality instance $(\mathcal{D}, \Sigma)$ at the outset is negative — an undecidable problem, by Theorem 2.2. $\square$

The results of Latteux and Simplot stated in Theorem 2.1 describe a correspondence between context-sensitive languages and domino systems. Via our construction for proving the lower bounds in Propositions 4.2 and 4.3, this correspondence is extended to game graphs, as detailed in [33] for the more general case of consensus acceptor games. For our setting, it implies that any context-sensitive language $L \subseteq \Sigma^*$ can be translated into a game graph G such that a word w belongs to L if, and only if, the players do not attain CKS at a particular history π_w in G which yields w as an observation to both players. Conversely, the nondeterministic linear-space procedure witnessing the upper bound in Propositions 4.2 shows that the set of histories at which the players do not attain CKS is recognisable by a nondeterministic linear-bounded automaton, and it is hence context-sensitive. Thus, we can formulate the following corollary which implies, in particular, that common knowledge of the state is not a finite-state property in arbitrary games, i.e., the set of histories at which the players attain CKS is not regular.

Corollary 4.4 ([33])**.** (i) *In any game, the set of histories at which the players attain* CKS *forms a context-sensitive language.*

(ii) *For every context-sensitive language $L \subseteq \Sigma^*$, we can construct a game in which an observation history belongs to L if, and only if, the players attain* CKS *at the history.*

5. Recurring common knowledge

Let us now turn to the use of common knowledge in infinite plays. We say that a play π allows for *recurring* common knowledge of the state (ω-CKS) if there are infinitely many histories in π at which the players attain CKS. Likewise, we say that a game graph G, or a game over G, allows for ω-CKS if this is true for every play in G.

A *knowledge gap* in a play π is an interval $[\![\ell, t]\!]$ with $t \geq \ell > 0$, such that the players do not attain CKS in π at any round in $[\![\ell, t]\!]$. The *length* of the gap is $t - \ell + 1$. Hence, a play allows for ω-CKS if the length of every knowledge gap in it is finite. The *gap size* (for CKS) of a play π is the least upper bound on the length of knowledge gaps in π. Likewise, the gap size of a game (graph) is the least upper bound on the gap size of its plays.

Intuitively, the uncertainty of players about the game state progresses along knowledge gaps, and it vanishes at every history at which CKS is attained. If we reindex the histories of a game G by forgetting any prefix history at which the players attain CKS, the knowledge of players about the game state is preserved. Concretely, for the case of private knowledge, let π, π' be two histories at which the players attain CKS, and let v be the state at which they end. Then, for any two continuations τ, τ' and every player i, we have $\pi\tau \sim^i \pi'\tau'$ if, and only if, $v\tau \sim^i v\tau'$ in the game G with v as initial state. The preservation for the case of common knowledge follows immediately; we formulate it for further reference.

Lemma 5.1. *For a game G, let π be a history at which the players attain* CKS*, and let v be its last state. Then, the players attain* CKS *at a prolongation history*

$\pi\tau$ in G if, and only if, they attain CKS *at history $v\tau$ in the game G with initial state v.*

Notice that for a play in an arbitrary game, the length of knowledge gaps may be unbounded, even if the play allows for ω-CKS; its gap size is then infinite. Nevertheless, we show that, if a *game* allows for ω-CKS, then there exists a uniform, finite bound on the length of the knowledge gaps in its plays.

Proposition 5.2. *If a game graph allows for recurring common knowledge of the state, then its gap size is finite.*

Proof. Let G be a game graph that allows for ω-CKS. Without loss of generality, we assume that all states are reachable from the initial state v_0.

For each state $v \in V$, we construct a tree T_v that may be understood as the unravelling of G from v, up to common knowledge. The nodes of T_v correspond to the histories in G, v that have no strict, nontrivial prefix at which the players attain CKS. The edges are labelled with action profiles and correspond to moves in G: for any history ρ in the domain of T_v at which the players do not attain CKS, or for $\rho = v$, we have an edge $(\rho, a, \rho a w)$ whenever $(u, a, w) \in E$, for the last state u of ρ. The leaves of T_v thus correspond to the histories in G, v at which the players attain CKS for the first time (not counting the initial history). Finally, we associate to every history the observations of its last state.

Notice that each of the constructed trees has finite branching and all its paths are finite, according to our assumption that all plays allow for ω-CKS. Hence by König's lemma, every tree in the collection $(T_v)_{v\in V}$ is finite. We claim that the maximal height of a tree in this collection is an upper bound for the length of knowledge gaps in the plays of G, v_0.

To show this, we construct a game graph G^{CK} over the disjoint union of all unravelling trees T_v, where we identify every leaf history with the root of the tree associated to its last state. Formally, in each tree T_v, we replace every edge (ρ, a, π), where π is a leaf history ending at w, with an edge (ρ, a, w) leading to the root of the tree T_w. This induces a natural bijection h between histories of G, v_0 and G^{CK}, v_0, which is also a bisimulation — clearly, the two game graphs have the same infinite unravelling. The bijection h preserves CKS: By the reindexing argument of Lemma 5.1, the players attain CKS at a history π in G, v_0, if, and only if, they attain CKS at the image $h(\pi)$ in G^{CK}, v_0. As a consequence it follows that, on the one hand, every history in G^{CK}, v_0 at which the players attain CKS ends at the root of some tree T_v, and on the other hand, for every knowledge gap, i.e., every sequence of consecutive histories $\pi^1, \pi^2, \dots, \pi^t$ in G, v_0 at which the players do not attain CKS, the image $h(\pi^1), h(\pi^2), \dots, h(\pi^t)$ describes a sequence of consecutive histories in G^{CK} that never visit the root of any tree T_v. Hence, the length t of such a sequence is bounded by the maximal length of a path in any of the trees $(T_v)_{v\in V}$. This concludes the proof. □

The insight that games with recurring common knowledge of the state have finitely bounded knowledge gaps allows us to conclude that these games are decidable via a generic argument which, however, does not yield meaningful complexity bounds.

Theorem 5.3. *For games that allow for recurring common knowledge of the state, with observable ω-regular winning conditions,*

(i) *it is decidable whether there exists a joint winning strategy, and*

(ii) *if it is the case, there also exists a finite-state winning strategy, which can be constructed effectively.*

Proof. The argument relies on the tracking construction from [13] that reduces the problem of solving coordination games with imperfect information for n players against Nature to that of solving a zero-sum game for two players with perfect information. The construction proceeds via an unravelling process that generates epistemic models of the player's information along the rounds of a play, and thus encapsulates their uncertainty.

This process described as "epistemic unfolding" in the paper [13, Section 3] is outlined as follows. An *epistemic model* for a game graph G with the usual notation, is a Kripke structure $\mathcal{K} = (K, (Q_v)_{v \in V}, (\sim^i)_{1 \le i \le n})$ over a set K of histories of the same length in in G, equipped with predicates Q_v designating the histories that end in state $v \in V$ and with the players' indistinguishability relations $\sim^i$. The construction keeps track of how the knowledge of players about the actual history is updated during a round, by generating for each epistemic model $\mathcal{K}$ a set of new models, one for each assignment of an action profile a_k to each history $k \in K$ such that the action assigned to any player i is compatible with his knowledge, i.e. for all $k, k' \in K$ with $k \sim^i k'$, we have $a^i_k = a^i_{k'}$. The update of a model $\mathcal{K}$ with such an action assignment $(a_k)_{k \in K}$ leads to a new, possibly disconnected epistemic model $\mathcal{K}'$ over the universe

$$K' = \{k a_k w \mid k \in K \cap Q_v \text{ and } (v, a_k, w) \in E\},$$

with predicates Q_w designating the histories $k a_k w \in K'$, and with $k a_k w \sim^i k' a_k w'$ whenever $k \sim^i k'$ in $\mathcal{K}$ and $w \sim^i w'$ in $\mathcal{G}$. By taking the connected components of this updated model under the coarsening $\sim := \bigcup_{i=1}^n \sim^i$, we obtain the set of epistemic successor models of $\mathcal{K}$ in the unfolding. The tracking construction starts from the trivial model that consists only of the initial state of the game $\mathcal{G}$. By successively applying the update, it unfolds a tree labelled with epistemic models, which corresponds to a two-player game $\mathcal{G}'$ of perfect information where the strategies of one player translate into joint strategies of the grand coalition in $\mathcal{G}$ and vice versa, such that a strategy in $\mathcal{G}'$ is winning if and only if the corresponding joint strategy in $\mathcal{G}$ is so [13, Theorem 5].

The construction can be exploited algorithmically if the perfect-information tracking of a game can be folded back into a finite game. A homomorphism from an epistemic model $\mathcal{K}$ to $\mathcal{K}'$ is a function $f : K \to K'$ that preserves the state predicates and the indistinguishability relations, that is, $Q_v(k) \Rightarrow Q_v(f(k))$ and $k \sim^i k' \Rightarrow f(k) \sim^i f(k')$. The main result of [13] shows that, whenever two nodes of the unfolded tree carry homomorphically equivalent labels, they can be identified without changing the (winning or losing) status of the game [13, Theorem 9]. This holds for all imperfect-information games with ω-regular winning conditions that are observable. Consequently, the strategy synthesis problem is decidable for a class of such games, whenever the unravelling process of any game in the class is guaranteed to generate only finitely many epistemic models, up to homomorphic equivalence.

Let us now consider the tracking of a coordination game $\mathcal{G}$ with observable ω-regular winning condition that allows for ω-CKS. We claim that every history π where the players attain CKS corresponds to an epistemic model that is homomorphically equivalent to one with a single element labelled with the (commonly known) state at which the history π ends. This is because, by our hypothesis of CKS, in the $\sim$-connected component of any epistemic model containing π, all histories end at the same state. On the other hand, when updating an epistemic model $\mathcal{K}$, there are only finitely many successor models and each of them can be at most exponentially larger than $\mathcal{K}$, for any fixed action space. Accordingly, the number of updating rounds in which the models can grow is bounded by the gap size of $\mathcal{G}$, which is finite, according to Proposition 5.2.

Therefore, every game with ω-CKS has a finite tracking quotient under homomorphic equivalence. By [13, Theorems 9 and 11], this implies that the winner determination problem is decidable for such games, and finite-state winning strategies can be effectively synthesised whenever the players have a joint winning strategy. □

6. Characterisation via mutual knowledge

Proposition 4.3 leaves little hope for deciding whether a game allows for ω-CKS by checking that the property holds within parts of the game graphs or on individual plays. In general, histories at which the players do not attain CKS may be connected to arbitrarily long chains of indistinguishable histories that end at the same state, before reaching one with a different end state to witness the lack of CKS. Fortunately, there is a way around this obstacle. It turns out that in any game that allows for ω-CKS we can find, for each play π, an associated play π' that aligns witnesses for all the histories in π that lack CKS; in particular, whenever the players lack common knowledge of the state at some round in π, there is one player that lacks first-order knowledge of the state in π'. This will allow us to characterise games with recurring common knowledge of the state as those where mutual knowledge of the state is attained over and over again, along every play.

We say that the players attain *mutual knowledge of the state* (MKS) at a history π in a game if $Ter(\pi)$ is mutual knowledge at π, that is, if all indistinguishable histories $\rho \sim^i \pi$ end at the same state as π, for all players i. A play π allows for recurring mutual knowledge of the state (ω-MKS) if the players attain MKS at infinitely many histories along π, and a game (graph) G allows for ω-MKS if all plays in G do.

The link between common and mutual knowledge is made by the notion of connected ambiguous histories. We say that two histories, or plays, π and π'

are *connected* if there exists a sequence of histories or plays $\pi_1, \dots, \pi_k$ and a sequence of players $i_1, \dots, i_{k+1}$ such that

$$\pi \sim^{i_1} \pi_1 \sim^{i_2} \cdots \sim^{i_k} \pi_k \sim^{i_{k+1}} \pi'.$$

In the special case when two histories π and π' end at the same state v and are connected via a sequence of histories that also end at v, we say that π and π' are *twins*. Clearly, the relations of connectedness and twins are equivalences between histories. Moreover, if two histories π and ρ that end at the same state v are connected or twins, then for every move $(v, a, v') \in E$ the prolongation histories $\pi a v'$ and $\rho a v'$ are in the same relation.

A history π is *ambiguous* if the players do not attain MKS at π, that is, if there exists an indistinguishable history $\rho \sim^i \pi$, for some player i, which ends at a different state. In this case, we refer to ρ is an *ambiguity witness* for π and we say that ρ, π is an *ambiguous pair*. Notice that the players do not attain CKS at a history π if, and only if, there exists an ambiguous twin of π.

Our goal is to show that every play in which the players do not attain recurring common knowledge of the state is witnessed by one where they do not attain recurring mutual knowledge of the state. Towards this, we first prove that if the players never attain CKS in a play, there exists a witnessing play in which all histories are ambiguous.

Lemma 6.1. *For any game, if there exists a play π along which the players never attain* common *knowledge of the state (except for the initial state), then there also exists a play π' along which they never attain* mutual *knowledge of the state. Moreover, the plays π and π' are connected.*

Proof. For an arbitrary game G and a play $\pi = v_0\, a_1 v_1 \dots$, we consider the set T_π of all histories τ in G such that every nontrivial prefix history of τ is ambiguous and connected to the history of the same length in π. As the set T_π is closed under prefix histories, we can view it as a finitely branching tree. We wish to show that if the players do not attain CKS along π, then every history in π is connected to some history in T_π, and therefore T_π contains an infinite play in G along which the players never attain MKS.

We prove a stronger property, for every history π_ℓ of lenght $\ell \geq 1$ in π: If the players do not attain CKS along π_ℓ (except for the trivial history), then for every ambiguous pair $\tau \sim^i \rho$ connected to π_ℓ there exists a pair $\tau' \sim^i \rho'$, such that

(i) $\tau' \in T_\pi$ is a twin of τ, and

(ii) ρ' ends at the same state as ρ.

For the base case with $\ell = 1$, if the players do not attain CKS at the history $\pi_1 := v_0\, a_1 v_1$, then there exist ambiguous histories connected to π_1, and they all belong to T_π, because the only preceding history is trivial. Hence, for any ambiguous pair $\tau \sim^i \rho$, already $\tau' = \tau$ and $\rho' = \rho$ witness the statement.

For the induction step, suppose the statement holds for $\ell \geq 1$ and assume that the players do not attain CKS up to (and including) the history $\pi_{\ell+1}$ of length $\ell + 1$. In particular, this means that there exist ambiguous histories connected to $\pi_{\ell+1}$; among these, let us pick an ambiguous pair $\tau av \sim^i \rho cw$, with $v \neq w$. Due to perfect recall, we have $\tau \sim^i \rho$.

We distinguish two cases. (1) If τ and ρ end at different states $v' \neq w'$, by induction hypothesis, there exists a twin $\tau' \in T_\pi$ of τ and a history $\rho' \sim^i \tau'$ that ends at w'. On the one hand $\tau' av$ is a twin of τav. On the other hand, by definition of the observation function, $\tau' av \sim^i \rho' cw$. Since $\tau' \in T_\pi$ and $v \neq w$, this also implies $\tau' av \in T_\pi$. (2) Otherwise, suppose τ and ρ end at the same state. As the histories are connected to π_ℓ, the players do not attain CKS at τ. Hence, there exists an ambiguous twin τ' of τ, and by induction hypothesis, we can choose $\tau' \in T_\pi$. On the one hand, $\tau' av$ is a twin of τav. On the other hand, as τ' and ρ end at the same state, so $\tau' cw$ is a valid history in G, and we have $\tau' av \sim^i \tau' cw$. Again, since $\tau' \in T_\pi$, and $v \neq w$ it follows $\tau' av \in T_\pi$. This completes the induction argument.

In conclusion, for a play π in which the players do not attain CKS at any round, there exist histories in T_π that are connected to arbitrarily long histories of π. As the tree T_π is finitely branching, it follows from König's Lemma that it has an infinite path π'. By construction, each nontrivial prefix of π' is an ambiguous history, and it is connected to the history of π of the same length. Hence, π' describes a play connected to π along which the players never attain mutual knowledge of the state. □

We are now ready to formulate our characterisation result that will be instrumental for the algorithmics of games with ω-CKS.

Theorem 6.2. *A game allows for recurring common knowledge of the state if, and only if, it allows for recurring mutual knowledge of the state.*

Proof. The *only if* direction is trivial: common knowledge of an event implies mutual knowledge.

For the converse, let us consider a game G that does not allow for ω-CKS. Then, there exists a play π in which the players attain CKS at some round ℓ, but not at any later history. Accordingly, in the game G, v starting from the (commonly known) state v that is reached in round ℓ of π, there exists a play along which the player never attain CKS, except for the initial state. Then, by Lemma 6.1, there exists a play π' in G, v along which the players never attain mutual knowledge of the state. Furthermore, in the play that follows π for the first ℓ rounds and, upon reaching v, proceeds like π', the players do not attain MKS at the infinitely many histories from round ℓ onwards. Hence, the game G does not allow for ω-MKS, which concludes the proof. □

Before turning to algorithmic questions, let us state the following corollary of arguments from the proofs of Lemma 6.1 and Theorem 6.2, which will be useful for bounding the gap size of games in Section 7.

Corollary 6.3. *For any game G, if the players do not attain common knowledge of the state in a play π along a sequence of rounds $\ell+1, \dots, \ell+t$, then there exists a play π' in G that is connected to π and on which the players do not attain mutual knowledge of the state along the rounds $\ell+1, \dots, \ell+t$.*

Proof. Let G be a game graph and let π be a play with the stated property, for some $\ell, t > 0$. We assume, without loss of generality, that the players attain CKS at round ℓ in π. For the game G, v starting at the state v reached in this round, we consider the suffix τ of π from round ℓ onwards, and construct the tree T_τ of hereditarily ambiguous histories connected to τ, as in the proof of Lemma 6.1. The induction argument from the proof then shows that the history of length t in τ is connected to some ambiguous history $\tau' \in T_\tau$. The histories of τ' from round 1 to t are ambiguous and each of them is connected to the history of the same length in τ. Hence, the play π' that follows π for the first ℓ rounds, then proceeds like τ' for t rounds, and then again follows π satisfies the required properties: π' is connected to π and the players do not attain MKS along the rounds $\ell+1, \dots, \ell+t$. □

7. Recognising recurring mutual and common knowledge

An automaton for recognising the plays that allow for ω-MKS could easily be designed using the powerset construction described by Reif [35] for solving one-player games with imperfect information. This would yield a PSPACE-procedure for deciding whether a game allows for ω-MKS and thus for ω-CKS. To obtain a sharper complexity bound, we will show that ambiguity witnesses along a play can be represented efficiently, by a tree of very low width, which allows to reduce the complexity to NLOGSPACE.

Let us fix an arbitrary game graph G. A *fork tree* for a play π is a prefix-closed set T of histories that contains, for every level $\ell \geq 0$,

(i) the history π_ℓ of π in round ℓ, and

(ii) at most one history $\rho_\ell \neq \pi_\ell$ with $\rho_\ell \sim^i \pi_\ell$, for some player i.

A fork tree T is *complete*, if it additionally satisfies, for every level ℓ:

(iii) if π_ℓ is ambiguous, then T contains an ambiguity witness ρ_ℓ of π_ℓ.

We can view fork trees as induced subtrees in the unravelling of G that contain π as a central branch and have width at most two, that is, at most two elements on each level. For convenience, we let ρ_ℓ refer to π_ℓ whenever T contains only π_ℓ at level ℓ. In case π_ℓ and ρ_ℓ end at different states, we say that the level ℓ is a *doubleton*, else it is a *singleton*.

If we consider an arbitrary family of ambiguity witnesses to the histories of a play, the subtree induced in the game unravelling can have unbounded width. Nevertheless, the following lemma states that every play π admits a complete family of witnesses ρ_ℓ, one for each of its ambiguous histories π_ℓ, that forms a fork tree.

Lemma 7.1. *For every play in an arbitrary game there exists a complete fork tree.*

Proof. It is convenient to extend the notion of ambiguity witness to knowledge gaps in histories. For a history π and an interval $[\![\ell, t]\!]$, we say that a history π' is an ambiguity witness *along the gap* $[\![\ell, t]\!]$ if π and π' have length at least t, and π'_r is an ambiguity witness for π_r, for every round $\ell \leq r \leq t$. Likewise, for a play π, we say that a play π' is an ambiguity witness *from round ℓ onwards* if π'_r is an ambiguity witness for π_r for every $r \geq \ell$.

Now, consider an arbitrary game G and a play π. By induction on the number of rounds ℓ, we construct a finite or infinite sequence of trees T_ℓ that satisfy the fork-tree conditions (i) and (ii) for the first ℓ levels, and, in addition, the following strengthening of the completeness condition (iii) for the last level ℓ:

(iii)* If, for some $t \geq \ell$, there exists a history in G that is an ambiguity witness for π along the gap $[\![\ell, t]\!]$, then there also exists a prolongation history of ρ_ℓ that is such a witness.

In particular, this implies that whenever the history π_ℓ is ambiguous, the level ℓ in T_ℓ is a doubleton.

Each tree $T_{\ell+1}$ is finite and extends its predecessor T_ℓ by one level, except if the sequence ends at some stage $\ell+1$, in which case $T_{\ell+1}$ extends T_ℓ with the (infinite) prolongation of π_ℓ to π and with a prolongation play ρ of either π_ℓ or ρ_ℓ that is an ambiguity witness for π from round ℓ onwards.

For the base case, we take the tree T_0 consisting only of the initial history v_0. For the induction step, suppose that a tree T_ℓ with ℓ levels satisfying the conditions (i), (ii), and (iii)* has been constructed. To extend it to $T_{\ell+1}$, we look at the set R of histories τ that prolong either π_ℓ or ρ_ℓ, and are ambiguity witnesses for π along the gap $[\![\ell+1, t]\!]$ up to the length t of τ. Now we distinguish three cases. (1) If R is empty, we set $\rho_{\ell+1} := \pi_{\ell+1}$, that is, $\ell+1$ is a singleton level. (2) If R is nonempty, but finite, we pick a history $\tau \in R$ of maximal length, and add $\rho_{\ell+1} := \tau_{\ell+1}$ together with $\pi_{\ell+1}$ as a new level to T_ℓ. (3) Finally, if R is infinite, there exists an infinite play τ in G such that all its histories from round ℓ onwards are in R. This follows from König's Lemma, since the histories in R form an infinite tree that is finitely branching (indeed, a subtree of the unravelling of G). In this case, we add the histories π_r and $\rho_r := \tau_r$, for all levels $r > \ell$ and terminate the sequence with this infinite tree $T_{\ell+1}$.

In any case, $\rho_{\ell+1}$ is a history in G and is indistinguishable from $\pi_{\ell+1}$ which is also contained on level $\ell+1$. Condition (iii)* holds trivially in case (3), we shall verify that it is also maintained in case (1) and (2).

For case (1) assume, towards a contradiction, that R is empty and there exists a history π' of length $\ell+1$ that is an ambiguity witness for $\pi_{\ell+1}$. If π'_ℓ ends at the same state as π_ℓ, then the last action-state pair $(a'_{\ell+1} v'_{\ell+1})$ of π' yields a prolongation $\tau = \pi_\ell a'_{\ell+1} v'_{\ell+1}$ that should be included in R, a contradiction. Else, if π'_ℓ ends at a different state than π_ℓ, by perfect recall, π'_ℓ is an ambiguity witness for π_ℓ along the gap $[\![\ell, \ell+1]\!]$, which, by induction hypothesis, implies that there also exists such a witness that prolongs ρ_ℓ and is thus contained in R, again in contradiction to our assumption that $R = \emptyset$.

For case (2), consider a history π' of length $t > \ell$ that is an ambiguity witness for π along the gap $[\![\ell + 1, t]\!]$. We claim that there also exists a prolongation of $\rho_{\ell+1}$ with this property. There are two situations to distinguish: If π'_ℓ reaches the same state as π_ℓ, then the history π'' that follows π until round ℓ and then continues like π' belongs to R, and is at most as long as the witness τ chosen to construct $\rho_{\ell+1}$. Hence, τ prolongs $\rho_{\ell+1}$ and is an ambiguity witness for π along the gap $[\![\ell + 1, t]\!]$. Otherwise, if π'_ℓ reaches a different state than π_ℓ, then, by perfect recall, we have $\pi'_\ell \sim^i \pi_\ell$, for some player i, and hence π' is already an ambiguity witness for π along the gap $[\![\ell, t]\!]$. By induction hypothesis, there exists an ambiguity witness π'' for π along the gap $[\![\ell, t]\!]$ that prolongs ρ_ℓ. Hence, $\pi'' \in R$ and, as the history $\tau \in R$ chosen to construct $\rho_{\ell+1}$ is of maximal length, τ prolongs $\rho_{\ell+1}$ and is also an ambiguity witness for π along the gap $[\ell + 1, t]$.

Clearly, each tree T_ℓ constructed along the induction satisfies the conditions of a complete fork tree and agrees with its successor $T_{\ell+1}$, up to level ℓ. In conclusion, the sequence converges and the infinite tree T obtained at the limit is a complete fork tree for π. □

Fork trees for a fixed play π can be represented by ω-words. We say that a word $\tau \in V(AV)^\omega$ is a *fork sequence* for a play π if it starts with $\tau_0 = v_0$ and there exists a fork tree T for π such that τ_ℓ is the last action-state pair of ρ_ℓ in T, for every $\ell > 0$. In the following we construct, for any arbitrary game, an ω-word automaton that takes as input an (infinite) play π in G and guesses in every run a fork sequence for π. The automaton is equipped with a co-Büchi acceptance condition: a run is accepting if it visits only finitely many non-accepting states. For background on the automaton model, we refer to the survey [36, Chapter 1].

Proposition 7.2. *For any game with m states, the set of plays that do not allow for recurring mutual knowledge of the state is recognisable by a nondetermistic co-Büchi automaton with m^2 states.*

Proof. Let us fix an arbitrary game graph G. We construct an ω-word automaton A with co-Büchi acceptance condition that recognises the set of histories π in G, for which there exists a fork tree with only finitely many singleton levels. To witness this, the automaton guesses non-deterministically a fork sequence τ for π and accepts if the states at τ_ℓ and π_ℓ are different, for all but finitely many rounds ℓ.

The states of the automaton are pairs of game states from V: the first component keeps track of the input play, the second one is used for guessing the fork sequence τ. The transition function ensures that the two components evolve according to the moves available in the game graph and that the current input symbol yields the same observation as the second component to some player i.

Concretely, the co-Büchi automaton A is defined over the input alphabet $A \times V$ on the state set $V \times V$ with initial state (v_0, v_0) and transitions from state (u, u') on input (a, v) to state (v, v') whenever $(u, a, v) \in E$ and $\beta^i(v') = \beta^i(v)$, for some player i, and either $(u', a', v') \in E$ or $(u, a', v') \in E$, for some action $a' \in A$ with $a'^i = a^i$ for this player. The set of final states is $Q \setminus \{(v, v) \mid v \in V\}$; the automaton accepts an infinite input word if all states that occur infinitely often in a run are final.

We claim that an input word $\pi \in V(AV)^\omega$ is accepted by A if, and only if, π corresponds to a play in G, and the players never attain mutual knowledge of the state along π, from some round onwards.

For the *if* direction, consider a play π along which the players never attain mutual knowledge of the state from some round onwards. By Lemma 7.1, there exists a complete fork tree T for π, in which all but finitely many levels are doubletons. Let τ be the fork sequence associated to T. Then, the sequence $((\pi_\ell, \tau_\ell))_{\ell<\omega}$ describes a run of A on input π in which non-final states (v, v) occur only at the finitely many positions ℓ corresponding to singleton levels in T, thus witnessing that π is accepted.

For the converse, inputs that do not correspond to histories in G are rejected, by construction of A. Furthermore, if an input word π corresponds to a play with infinitely many histories π_ℓ at which the players attain mutual knowledge of the state, then every run of the automaton visits a non-final state whenever such an input prefix π_ℓ is read. As this occurs infinitely often, the input π is rejected. $\square$

Theorem 7.3. *The problem of whether a game graph allows for recurring common knowledge, or equivalently, recurring mutual knowledge of the state is* NLogSpace*-complete.*

Proof. According to the Characterisation Theorem 6.2, a game graph G allows for recurring common knowledge of the state if, and only if, it allows for recurring mutual knowledge of the state. Our problem thus reduces to checking whether the language recognised by the co-Büchi automaton A constructed for G in Proposition 7.2 is non-empty. It is well known that the non-emptiness test for co-Büchi automata is in NLogSpace (see, for instance, Vardi and Wolper [37]).

Concretely, a nondeterministic procedure can guess a run of A that leads to a cycle included in the set of final states. This requires only pointers to three states of the automaton: two for the current transition and one for storing a state to verify that a cycle is formed. As each state of the automaton is formed by two states of the game, the overall space requirement is logarithmic in the size of the game graph G. Accordingly, the problem of determining whether a game graph allows for common knowledge of state is in NLogSpace.

Hardness for NLogSpace follows via a straightforward reduction from directed graph acyclicity, shown to be NLogSpace-hard by Jones in [38]: Given a directed graph G, we construct a game graph G' for one player by taking two disjoint copies of G and assigning all non-terminal nodes with the same observation; each terminal node is assigned with a distinct observation and equipped with a self-loop. Finally, we add a fresh initial state to G', with moves to all other states. Clearly, the game graph G' can be constructed using logarithmic space, and the player has recurring (mutual, common) knowledge of the state in G' if, and only if, the directed graph G is acyclic.

This shows that the problem of determining whether a game graph allows for common knowledge of the state, or equivalently, for mutual knowledge of the state, is NLogSpace-complete. □

Theorem 7.4. *The gap size of any game with m states that allows for recurring common knowledge of the state is bounded by m^2.*

Proof. Consider a game G with m states that allows for ω-CKS. Towards a contradiction, suppose that in G there exists a play with gap size greater than m^2, that is, the players do not attain CKS along a sequence of consecutive rounds $r, \ldots, r + m^2$, for some r. Due to Corollary 6.3, there also exists a play π in G such that the players do not attain MKS in π along these rounds. Let T be a complete fork tree for π, according to Lemma 7.1, and let τ be the associated fork sequence.

As G allows for ω-MKS, the automaton A constructed in Proposition 7.2 recognises the empty language, in particular it rejects the run on π described by (π, τ). But A has at most m^2 states, so there must be a cycle in the transition graph that is visited by this run, say from position $\ell \geq r$ to $t \leq r+m^2$. Along the interval $[\![\ell, t]\!]$, the players do not attain MKS in π, therefore the corresponding levels in the fork tree T are doubletons, and the states on the cycle visited in the run (π, τ) from position ℓ to t are final.

Consider now the sequences π', and τ' that follow π and τ, respectively, until position t and then loop from ℓ to t forever. Then, the pair (π', τ') describes a run in A that eventually cycles through final states, hence, the input π' is accepted. But this means that π' is a play in G that does not allow for ω-MKS, in contradiction to our assumption that all plays in G allow for ω-CKS. □

We observe that the quadratic bound on the gap size is tight. Consider, for instance, the game graph G_m for one player with two observations, black and white, depicted in Figure 3, for an arbitrary number $m > 1$. There is only one bit of uncertainty induced by the choice of Nature at the initial state, where it can either move up, into the cycle with $m - 1$ white states followed by a black one, or down, to the path consisting of m white states with selfloops, each followed by a black state, except for the last one which leads to the black state on the cycle. Consider the play π where Nature moves into the cycle (and stays there forever). Along π, every nontrivial history up to round m^2 is indistinguishable from the one where Nature moves initially down to the path and loops on each white state precisely $m - 1$ times. For the first m^2 rounds in π, the player does therefore not know the current state, which means that the gap size of the game is at least m^2. On the other hand, notice that all histories that are distinguishable from π are non-ambiguous, and that from round $m^2 + 1$ onwards, any history that is indistinguishable from π leads to the same state

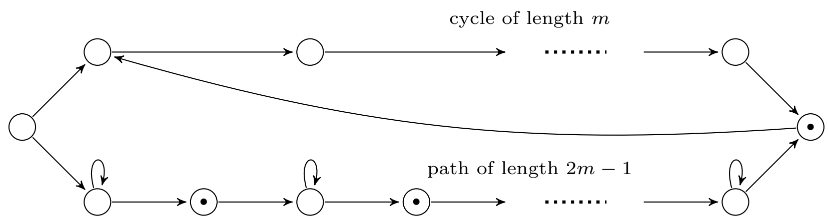

Figure 3: A game graph with 3m states and gap size m^2

as π itself. Accordingly the game graph G_m with $3m$ states allows for ω-CKS and its gap size is m^2.

One consequence of Theorem 7.4 is that the knowledge hierarchy for any game of size m that allows for ω-CKS collapses to the level $|B|^{m^2}$: By the Reindexing Lemma 5.1, whenever two histories are connected, they are connected via a sequence of histories that differ only on the last knowledge gap. As the gap size is bounded by m^2, there are at most $|B|^{m^2}$ different observation histories along a gap. Hence, the players attain common knowledge about an event $F \subseteq \Omega$ at a history π if, and only if, they attain MKS of order $|B|^{m^2}$ about F. In particular, this holds for the event $Ter(\pi)$.

For any game, an automaton that recognises the set of histories at which the players attain MKS of a fixed order can be constructed as in the proof of Proposition 7.2. Accordingly, for any game that allows for ω-CKS, the set of histories at which the players attain CKS is regular. This is in contrast with the general situation of games that may not allow for ω-CKS where, by Corollary 4.4, the set of histories at which the players attain CKS form arbitrary context-sensitive sets.

8. Strategy synthesis

We are now ready to establish complexity bounds for the basic algorithmic questions on games with recurring common knowledge of the state. Our analysis focuses on the canonical case of parity condition. At the end of the section, we explain how the results apply to observable ω-regular conditions.

Theorem 8.1. *For games that allow for recurring common knowledge of the state, with parity winning conditions,*

(i) *the problem of deciding whether there exists a joint winning strategy is* NExpTime*-complete;*

(ii) *if joint winning strategies exist, there also exists a winning profile of finite-state strategies of at most exponential size, which can be synthesised in* 2-ExpTime*.*

The lower bound for the decision problem (i) follows from NExpTime-hardness of the corresponding problem for two-player reachability or safety games of finite horizon. These are games where the underlying graph is acyclic except for having self-loops at observable sinks — hence, the simplest examples of games that allow for cks. The original proof, due to Azhar, Peterson, and Reif [7, Section 5], is by reduction from the time-bounded halting problem via a variant of QBF with dependency quantifiers.

For the sake of completeness, we outline a direct reduction from the Exp-Square Tiling problem to synthesis problem for safety games with finite horizon, similar to the one described by Bernstein, Zilberstein, and Immerman in [39] for decentralised planning in partially observable Markov decision processes of finite horizon.

Given a domino system $\mathcal{D}$ and the logarithm ℓ of the square size, we construct a two-player game with the following scenario: First, Nature sends to each player i privately a pair $(x^i, y^i) \in Z(2^\ell, 2^\ell)$ of coordinates in binary encoding over ℓ bits, such that either

(i) $(x^2, y^2) = (x^1, y^1)$, or

(ii) $(x^2, y^2) = (x^1 + 1, y^1)$, or

(iii) $(x^2, y^2) = (x^1, y^1 + 1)$.

Then, each player i produces a domino d^i. The play is winning if the produced dominoes are consistent with the relative position of the recieved coordinates, that is, if $d^1 = d^2$ in case (i) , $(d^1, d^2) \in E_H$ in case (ii), and $(d^1, d^2) \in E_V$ in case (iii). This can be formulated either as an observable reachability or safety condition: reach $\oplus$ or avoid $\ominus$. If a tiling of the exponential square exists, then the strategy to produce the domino placed at the received coordinates guarantees a joint win. Conversely, any winning strategy can be turned into a correct tiling of the exponential square.

The game is of finite horizon: After $\ell + 1$ rounds (ℓ for observing the coordinates and one for producing the domino), each play reaches either the safe sink $\oplus$ or the unsafe sink $\ominus$, which are observable to both players. Hence, the game trivially allows for ω-cks. The construction can be done in time $O(\ell + |\mathcal{D}|)$.

In summary, we have a linear-time reduction from Exp-Square Tiling, to the problem of deciding whether there exists a joint winning strategy in a game that allows for ω-cks. According to Theorem 2.2, this shows that our problem is NExpTime-hard.

For the upper bound and the strategy-construction procedure, it would be inconvenient to rely on the tracking construction used in the decidability proof of Theorem 5.3, as the number of epistemic structures (over histories of quadratic length that are relevant here) is already doubly exponential in the size of the game graph. Instead, we introduce an auxiliary representation of the game which retains the histories at which players attain cks and is only simply exponential in the size of the input game graph.

8.1. The abridged game

For the proof in the reminder of the section, let us fix a coordination game $\mathcal{G} = (G, \gamma)$ for n players over a game graph G that allows for ω-CKS, with a parity condition over a set of priorities $C = \{1, \dots, |C|\}$ described by the colouring function $\gamma : V \to C$. Recall that a play is winning under the parity condition if the least priority seen infinitely often along a play is even. The assumption that the players can observe the the prioriy coloring is inessential for parity conditions; it is only needed for ω-regular conditions that will be discussed at the end of the section.

The *abridged* game $\widehat{\mathcal{G}}$ of $\mathcal{G}$ is a game with perfect information for one player against Nature. Intuitively, $\widehat{\mathcal{G}}$ is obtained by contracting knowledge gaps and recording only the most significant priority seen between two consecutive histories where the players attain CKS.

Concretely, the states of the abridged game graph $\widehat{G}$ are pairs (v, c) of states $v \in V$ and priorities $c \in C$; for convenience, we also include a sink $\ominus$. We shall refer to the states of $\widehat{G}$ as *positions*, to avoid confusion with the ones of $\mathcal{G}$. The initial position $(v_0, |C|)$ corresponds to the initial state of $\mathcal{G}$ labelled with the least significant priority. The set of actions consists of all nonempty subsets $U \subseteq V \times C$ of positions. The player has perfect information, so the observation function is the identity on $V \times C$.

To define the moves, we look at the unravelling G^{CK} up to common knowledge of the game graph G, as constructed in the proof of Proposition 5.2. Recall that G^{CK} is built from a disjoint collection of trees $(T_v)_{v \in V}$, which are then connected by identifying all leaves with the corresponding roots. For every state $v \in V$ and any joint strategy t over the tree component T_v of G^{CK}, we define the set $\text{outcome}_v(t)$ of pairs $(u, d) \in V \times C$, for which there exists a history τ in T_v that follows t, such that τ ends at u, and the most significant priority that occurs along τ is d. Now, the set of available moves is defined as follows. For an action $U \subseteq V \times C$ there are moves from a position (v, c) to every position $(u, d) \in U$ if there exists a joint strategy t in T_v with $\text{outcome}_v(t) = U$. Otherwise, the action leads to the $\ominus$-sink. Notice that the moves depend only on the first component of the position, that is, on the state and not on the priority.

At last, we define a parity condition on $\mathcal{G}$, by assigning to every position $(v, c) \in V \times C$ the priority c.

The plays of G and $\widehat{G}$ are related via their summaries. Intuitively, this is the sequence of states reached when the players attain CKS in a play, together with the most significant priority seen along the last knowledge gap. More precisely, for a play $\pi = v_0\, a_1 v_1\, \dots$ in G, we look at the subsequence of rounds $t_0, t_1, t_2, \dots$ such that, for all $\ell \geq 0$, the players attain CKS at round t_ℓ in π, but not at any round t in between $t_\ell < t < t_{\ell+1}$. Next, we associate to each index $\ell > 0$, the most significant colour that occurred in the gap between t_ℓ and $t_{\ell+1}$, setting $c_{\ell+1} := \min\{\gamma(v_t) \;:\; t_\ell < t \leq t_{\ell+1}\}$. Now, the *summary* of π is the sequence $[\pi] := v_0, (v_{t_1}, c_1), (v_{t_2}, c_2) \dots$ Notice that for every play π in G, the summary $[\pi]$ corresponds to a sequence of states in $\widehat{G}$, which is infinite, since we assume that π allows for ω-CKS.

The notion of summary is defined analogously for histories, and it also applies to plays $\hat{\pi}$ in $\widehat{G}$. Indeed, $[\hat{\pi}]$ is obtained simply by dropping the actions in $\hat{\pi}$. We say that a play π in G *matches* a play $\hat{\pi}$ in $\widehat{G}$ if they have the same summary: $[\hat{\pi}] = [\pi]$.

The winning or losing status is preserved among matching plays.

Lemma 8.2. *If a play π of $\mathcal{G}$ matches a play $\hat{\pi}$ of $\hat{\mathcal{G}}$, then π is winning if, and only if, $\hat{\pi}$ is winning.*

Proof. Let c be the least priority that appears infinitely often in π. As each knowledge gap in π is finite, c appears in infinitely many knowledge gaps in π, hence it is recorded infinitely often in the summary $[\pi]$. Conversely, all priorities that appear infinitely often in the summary $[\pi]$, also appear infinitely often in π, so c is minimal among them. In conclusion, the least priority appearing infinitely often in the summaries $[\pi] = [\hat{\pi}]$ is the same as in the plays π and $\hat{\pi}$. □

8.2. Reduction to parity games with perfect information

To use results from the standard literature on parity games, it is convenient to view the abridged game $\widehat{\mathcal{G}}$ formally as a turn-based game between two players, *Coordinator* and *Nature*. In contrast to before, we shall hence regard Nature as an actual player with proper positions, moves, and strategies.

Towards this, we view the game graph $\widehat{G}$ as a bipartite graph, with one partition $V \times C$ controlled by Coordinator, and a second one formed by the nonempty subsets of $V \times C$, controlled by Nature. The initial position $(v_0, |C|)$ is unchanged. Coordinator can move from every position $(v, c) \in V \times C$ to a position $U \subseteq V \times C$, if $U = \text{outcome}_v(t)$ for some joint strategy t on T_v, whereas Nature can move from every position $U \subseteq V \times C$ to any element $(u, d) \in U$. The new positions from $U \subseteq V \times C$ receive the least significant priority $|C|$, whereas position $(v, c) \in V \times C$ have priority c, as before.

A fundamental result about parity games is that they enjoy *positional determinacy*. A strategy is positional if the choice prescribed at a history π depends only on the last position in π. The following theorem was first proved by Emerson and Jutla [40], a comprehensive exposition can be found in the survey of Zielonka in [41].

Theorem 8.3 ([40])**.** *For every parity game with perfect information, one of the two players has a positional winning strategy.*

For our setting, positional determinacy means that in the abridged game $\widehat{\mathcal{G}}$, either Coordinator or Nature has a winning strategy defined on the set of positions. This yields witnesses of manageable size for determining which player wins the abridged game.

In the following, we argue that positional strategies for the abridged game $\widehat{\mathcal{G}}$ can be translated effectively into strategies on $\mathcal{G}$, such that the resulting plays match in the sense of Lemma 8.2.

Proposition 8.4. *Let $\mathcal{G}$ be a game that allows for ω-CKS, and let $\widehat{G}$ be the abridged game.*

(i) *For every positional Coordinator strategy $\hat{s}$ in $\widehat{\mathcal{G}}$, we can effectively construct a strategy profile s for the grand coalition in $\mathcal{G}$ such that, for every play π that follows s, there exists a matching play $\hat{\pi}$ that follows $\hat{s}$.*

(ii) *For every positional Nature strategy $\hat{r}$ in $\widehat{\mathcal{G}}$, and every strategy profile s for the coalition in $\mathcal{G}$, there exists a play π in $\mathcal{G}$ that follows s with a matching play $\hat{\pi}$ that follows $\hat{r}$.*

Proof. (i) Let $\hat{s} : V \times C \to 2^{V \times C}$ be a positional strategy for Coordinator in the abridged game $\widehat{\mathcal{G}}$. We construct a strategy s for the unravelling $\mathcal{G}^{\text{CK}}$ of $\mathcal{G}$ up to common knowledge. As the two game graphs have the same unravelling, s is also a strategy for $\mathcal{G}$.

We can assume that the strategy $\hat{s}$ prescribes for every state $v \in V$ the same choice at all positions (v, c) independently of the priority. This is without loss of generality: Recall that all positions (v, c) in $\widehat{\mathcal{G}}$ have the same set of successors U. If we add a fresh position z_v, of least significant priority, from which Coordinator can move to every position in U, and replace the outgoing moves from each position (v, c) with a move to z_v, the game remains essentially unchanged. Whenever Coordinator has a winning strategy for the original game, he has one for the modified game. Then, due to positional determinacy, he also has a positional winning strategy and its choice at the new position z_v can be transferred as a uniform choice to all positions (v, c) in the original game, still yielding a winning strategy.

To transfer the given strategy from $\widehat{\mathcal{G}}$ to $\mathcal{G}^{\text{CK}}$, we consider for each state $v \in V$ the tree component T_v of $\mathcal{G}^{\text{CK}}$ separately. For an arbitrarily chosen colour c, we look at the set $U := \hat{s}(v, c)$ and pick a joint strategy t_v on T_v with $\text{outcome}_v(t_v) = U$. Now, for every history π that ends in T_v, we take the suffix π_v contained in T_v, that is, we forget the prefix history up to entering the tree, and set $s(\pi) = t_v(\pi_v)$. This is a valid strategy profile, due to the reindexing argument for private knowledge undelying Lemma 5.1.

With s constructed this way, every play π in G that follows s has the same summary $[\pi] = v_0\,(v_1, c_1)(v_2, c_2)\ldots$ as the play $\hat{\pi} = v_0\,a_1(v_1, c_1)\,a_2(v_2, c_2)\ldots$ in $\widehat{G}$ with actions $a_\ell = \hat{s}(v_\ell, c_\ell)$. Hence, $\hat{\pi}$ follows $\hat{s}$ and matches π, as required.

(ii) For the converse, let $\hat{r} : 2^{V \times C} \to V \times C$ be a positional strategy for Nature in $\widehat{G}$ and let s be an arbitrary strategy for Coordinator in G. We construct a pair of plays π in G, and $\hat{\pi}$ in $\widehat{\mathcal{G}}$ with the desired properties.

The construction is by induction on the number of knowledge gaps in π: For every ℓ, we construct a history π_ℓ in G with ℓ knowledge gaps that follows s and ends at some state v, where the players attain CKS. At the same time, we construct a matching history $\hat{\pi}_\ell$ that follows $\hat{r}$ and ends at a position (v, c) in $\widehat{G}$, associated with the same state v.

For the base case, both histories π_0 and $\hat{\pi}_0$ are set to v_0. For the induction step, suppose that the two histories π_ℓ and $\hat{\pi}_\ell$ satisfy the hypothesis, and that they end at state v and position (v, c), respectively. We construct a prolongation $\pi_{\ell+1}$ that follows s over the $\ell + 1$st knowledge gap and matches a one-round prolongation $\hat{\pi}_{\ell+1}$ of $\hat{\pi}_\ell$. Towards this, we consider the strategy t_v induced by s in the set of histories $\pi_\ell T_v$, that is, the prolongations of π_ℓ into the tree component T_v of G^{CK}. For $U := \text{outcome}_v(t_v)$ and $(u, d) := \hat{r}(U)$, there exists a history τ in T_v that ends at u and has d as most significant priority after the initial state v. Now, we update $\pi_{\ell+1} := \pi_\ell \tau$ and $\hat{\pi}_{\ell+1} := \hat{\pi}_\ell U(u, d)$. This way, $\pi_{\ell+1}$ follows s and the players attain CKS, and $\hat{\pi}_{\ell+1}$ follows $\hat{r}$. Moreover, the two plays have the same summary, and $\hat{\pi}_{\ell+1}$ ends at a position corresponding to the last state of $\pi_{\ell+1}$.

For the infinite plays π and $\hat{\pi}$ obtained at the limit, we have: π follows s and matches $\hat{\pi}$ which follows $\hat{r}$, as required.

□

The correspondence between strategies in the abridged game and in the original game allows us to draw the following conclusion.

Proposition 8.5. *Let $\mathcal{G}$ be a coordination game that allows for recurring common knowledge of the state, with m states and a parity winning condition over d priorities.*

(i) *The grand coalition has a joint winning strategy for $\mathcal{G}$ if, and only if, Coordinator has a positional winning strategy for the abridged game $\widehat{\mathcal{G}}$, that is a perfect-information parity game with $md + 2^{md}$ positions and d priorities.*

(ii) *If the grand coalition has a joint winning strategy in $\mathcal{G}$, then there exists a winning profile of finite-state strategies with $2^{\mathcal{O}(m^2 \log m)}$ states.*

Proof. (i) If Coordinator has a positional winning strategy $\hat{s}$ in $\widehat{\mathcal{G}}$, then the corresponding profile s according to Proposition 8.4(i) is winning in $\mathcal{G}$, because every play π that follows s has a matching play in $\mathcal{G}$ that follows $\hat{s}$ and is hence winning, which implies that π is also winning, by Lemma 8.2.

Conversely, assume that there exists a joint winning strategy s in G. By Proposition 8.4(ii), for any arbitrary positional strategy $\hat{r}$ of Nature, there exists a play $\hat{\pi}$ that follows $\hat{r}$ and matches some play π in G which follows s and thus wins. Hence, $\hat{\pi}$ is also winning for Coordinator, by Lemma 8.2 which means that $\hat{r}$ in not winning for Nature. By positional determinacy, it follows that Coordinator has a positional winning strategy in $\widehat{G}$.

The state space of the abridged game is $V \times C \cup 2^{V \times C}$, it has $md + 2^{md}$ positions; the number d of priorities is as in $\mathcal{G}$.

(ii) Let $\hat{s} : V \times C \to 2^{V \times C}$ be a winning strategy for Coordinator in the abridged game $\widehat{\mathcal{G}}$. As in the proof of Proposition 8.4(i), we assume, without loss of generality, that the strategy prescribes the same move $U = \hat{s}(v, c)$, at all positions corresponding to v, independently of the priority c; for each of the m states $v \in V$, the move $\hat{s}(v, c)$ is translated into an imperfect-information strategy t_v on the tree component T_v of G^{CK}. We use these local strategies to construct a joint winning strategy s for the grand coalition in G as follows.

For each player i, the component strategy s^i is implemented by a reactive procedure that maintains, along the infinite sequence of input observations, a record (v, ρ^i) of the last state v about which the players attained common knowledge and the observation history ρ^i along the subsequent knowledge gap. Initially v is set to v_0 and ρ^i to $\beta^i(v_0)$. In each step, the procedure returns the action $a^i := t^i_v(\rho)$, inputs the next observation b^i, and repeats with $\rho^i a^i b^i$ as a new value for ρ^i, unless this corresponds to a history in G_v at which the players attain common knowledge of the current state v'. In that case, the root v is replaced with v' and the new value of ρ^i becomes $b^i = \beta^i(v')$. Each local strategy t^i_v can be represented by a (tree shaped) automaton that outputs actions in response to observation sequences along knowledge gaps — of length at most m^2, by Theorem 7.4. Since there are no more observations than game states, m^{m^2} automaton states are sufficent to store these responses, as well as the set of histories at which the players attain CKS. Globally, the strategy s^i of each player i combines m local strategies t^i_v. Hence, we need at most $m \cdot m^{m^2} = 2^{\mathcal{O}(m^2 \log m)}$ many states to represent each component of the profile s by a strategy automaton. □

8.3. Complexity

A nondeterministic procedure for deciding whether there exists a joint winning strategy in a game $\mathcal{G}$ with ω-CKS, according to Proposition 8.5, can guess the abridged game $\widehat{G}$ and determine whether Coordinator has a winning strategy in the obtained parity game with perfect information. The complexity is dominated by the verification of the transition relations between Coordinator positions (v, c) and Nature positions $U \subseteq V \times C$, which involves guessing a witnessing strategy profile t_v over the tree T_v such that $\text{outcome}_v(t_v) = U$. As we pointed out in the proof of Proposition 8.5, for a game $\mathcal{G}$ of size m, such a strategy t_v can be represented by a collection of n trees of size $2^{\mathcal{O}(m^2 \log m)}$, one for every player. Once the local strategy trees t^i are guessed, the verification that $\text{outcome}_v(t) = U$ is done in time linear in their size. Given the abridged game, a winning strategy for Coordinator can be guessed and verified in nondeterministic linear time with respect to the size $md + 2^{md}$ of $\widehat{\mathcal{G}}$ where d is the number of priorites. Overall, the procedure runs in $\text{NTIME}(2^{\mathcal{O}(m^2 \log m)})$, that is, nondeterministic exponential time.

With a deterministic procedure, the abridged game can be constructed by exhaustive search over witnessing strategies over the component trees in G^{CK} in time $2^{2^{\mathcal{O}(m^2 \log m)}}$. Once this is done, winning strategies for the obtained parity game $\widehat{\mathcal{G}}$ can be constructed in time $\mathcal{O}(2^{md^2})$ using the basic iterative algorithm presented by Zielonka in [41]. This concludes the proof of Theorem 8.1.

8.4. Observable ω-regular conditions

In view of applying our results to the practice of automated verification and design, we briefly outline a procedure for synthesing distributed winning strategies in games with winning conditions expressed by standard specification formalisms rather than by parity conditions.

It is well known that every ω-regular language of infinite words can be recognised by a deterministic automaton with parity acceptance condition [17]. Given a game $\mathcal{G} = (G, W)$ with an ω-regular winning condition W represented by a deterministic automaton $\mathcal{A}$ over attributes (colours) of the game states, we construct a parity game $\mathcal{G}'$ on the game graph obtained as the synchronised product of G with $\mathcal{A}$ and define a priority colouring that associates to every state of the product graph G' the priority of the automaton state in its second component. Informally, this corresponds to running the automaton along the plays in G to monitor the ω-regular winning condition over the (colouring of) game states by a parity condition over the automaton states. Then, the synthesis problem for the original game $\mathcal{G}$ reduces to the synthesis problem for the game $\mathcal{G}'$ with a parity condition. This transformation works as in the case of perfect-information games detailed in [36, Chapter 2].

Now, let us assume that the game G at the outset allows for ω-CKS and that the winning condition W is expressed by an observable colouring. The latter assumption implies that, at every history in the product game G', the current state of the automaton monitoring the winning condition is common knowledge among the players. Since along every play in G the players attain CKS infinitely often, and there are only finitely many automaton states, it follows that the players also attain common knowledge of the product state in G' infinitely often. Hence, the product game graph G' allows for ω-CKS. In conclusion, the synthesis problem for games that allow for ω-CKS and with observable ω-regular conditions represented by deterministic automata can be solved with the same generic complexity as games with parity conditions.

Corollary 8.6. *For games that allow for recurring common knowledge of the state, with observable ω-regular winning conditions represented by a deterministic parity automata,*

(i) *the problem of deciding whether there exists a joint winning strategy is* NEXPTIME-*complete;*

(ii) *if joint winning strategies exist, there also exists a winning profile of finite-state strategies of at most exponential size, which can be synthesised in* 2-EXPTIME.

To obtain precise upper bounds, we need to take into account that the product construction increases the size of the game graph by a factor corresponding to the size of the deterministic automaton. More generally, for winning conditions specified in common verification formalisms, e.g., ω-regular expressions, PDL, LTL, or nondeterministic automata, we can apply the standard techniques for transforming the specification into deterministic parity automaton to establish the complexity of the synthesis problem for games with ω-CKS.

9. Conclusion

We identified a new class of games with imperfect information for which the distributed synthesis problem can be solved effectively: It is decidable whether distributed winning strategies exist and, if so, a profile of finite-state winning strategies can be computed. Our procedure for solving the distributed synthesis problem for infinite games with ω-CKS under parity winning condition matches the lower complexity bounds for solving the particular case of multi-player safety games of finite horizon. Whether a game belongs to the class is decidable efficiently.

Known decidable classes from the distributed-systems literature rely on decomposing the global synthesis problem into separate instances, each involving only one player and the environment (Nature), that can be solved by automata-theoretic techniques for zero-sum games. The approach proved successful in several cases where the dependencies between the behaviour of players are restricted, typically by a hierarchical information order among them. Prominent examples are weakly-ordered architectures [2] and doubly-flanked pipelines [42], both subsumed by Coordination Logic [43], or, in the asynchronous setting, well-connected architectures [44]. As the synthesis procedures rely on solving nested instances for all players, these classes generally display nonelementary complexity. A class of more moderate NExpTime complexity was recently proposed by Chatterjee et al. [45]. Here, the winning conditions are restricted to ensure an even higher degree of independence: Essentially, each player can achieve her part of the global objective independently of the others.

Our approach is orthogonal to the idea of decomposing games into two-player zero-sum instances. Instead of restricting the order of information or the game objective, our decidability condition requires that players attain common knowledge of the game state infinitely often along every play. Intuitively, this allows to decompose the game tree into a sequence of time slices (the gaps) that can be solved independently, rather than reducing it to parallel zero-sum instances for each individual players. As a most simple case, our class subsumes repeated safety games of finite horizon with imperfect information, where the initial state (re-entered at each repetition) is observable. Since safety games of finite horizon are already NExpTime hard to solve, this justifies the lower bound for our solution procedure. Nevertheless, it is rewarding to see that the synthesis problem for arbitrary games with ω-CKS has a matching NExpTime upper bound, in spite of covering much more general examples of games.

For instance, the class captures the interaction scenarios that proceed in phases where imperfect information can arise and evolve in any form, provided that at the outcome of each phase the participants synchronise in some state that will become common knowledge among them. This may be guaranteed explicitely, by restrictin to phase games on acyclic graphs with observable exit states, or implicitely, by ensuring that the players attain CKS due to the structure of game graph. We believe that designs of distributed systems tend to follow such patterns naturally, as developers introduce breakpoints or synchronisation barriers for monitoring their system. One challenge is to develop concrete communication schemes for distributed systems that yield games with ω-CKS. As an application to decentralised control, it will be interesting to identify conditions under which common knowledge of an event can be approximated by (a finite degree of) mutual knowledge.

Apart from direct applications, we hope that our contribution may help to demystify the subject of imperfect information in multiplayer games haunted by the discouraging complexity results for the general case. As pointed out by Muscholl and Walukiewicz in their concise survey on distributed synthesis [46], currently we have no evidence that the constructions causing undecidability in general game models would arise in real-world systems, but we have no systematic justification for ruling them out either. We may paraphrase the insights of the present article by concluding that, if imperfect information is admitted only as a temporary perturbation of perfect information, then we can rule out undecidable situations. It remains to investigate whether this intuition applies to further relevant forms of perturbation in the information structure of games as, for instance, communication delays or incomplete but perfect information.

Acknowledgements. This work was supported by the European Union Seventh Framework Programme under Grant Agreement 601148 (CASSTING) and by the Indo-French Formal Methods Lab (LIA Informel). A preliminary report [47] was presented at the 2014 Workshop on Strategic Reasoning.

References

[1] A. Pnueli, R. Rosner, On the synthesis of a reactive module, in: Proceedings of the 16th Annual ACM Symposium on Principles of Programming Languages (POPL 1989), ACM Press, 1989, pp. 179–190.

[2] O. Kupferman, M. Y. Vardi, Synthesizing distributed systems, in: Proceedings of the 18th IEEE Symposium on Logic in Computer Science (LICS 2001), IEEE Computer Society Press, 2001, pp. 389–398.

[3] A. Pnueli, R. Rosner, Distributed reactive systems are hard to synthesize, in: Proceedings of the 31st Annual Symposium on Foundations of Computer Science (FOCS 1990), IEEE Computer Society Press, 1990, pp. 746–757.

[4] A. Arnold, I. Walukiewicz, Nondeterministic controllers of nondeterministic processes, in: Logic and Automata, Vol. 2, Amsterdam University Press, 2008, pp. 29–52.

[5] G. L. Peterson, J. H. Reif, Multiple-Person Alternation, in: Proceedings of the 20th Annual Symposium on Foundations of Computer Science, (FOCS 1979), IEEE, 1979, pp. 348–363.

[6] D. Janin, On the (high) undecidability of distributed synthesis problems, in: Proceedings of the 33rd Conference on Theory and Practice of Computer Science (SOFSEM 2007), Vol. 4362 of LNCS, Springer, 2007, pp. 320–329.

[7] S. Azhar, G. Peterson, J. Reif, Lower bounds for multiplayer noncooperative games of incomplete information, Journal of Computers and Mathematics with Applications 41 (2001) 957–992.

[8] Ł. Kaiser, Game quantification on automatic structures and hierarchical model checking games, in: Proceedings of the 15th Annual Conference of the Europan Association for Computer Science Logic (CSL 2006), Vol. 4207 of LNCS, Springer, 2006, pp. 411–425.

[9] D. Berwanger, A. B. Mathew, M. Van den Bogaard, Hierarchical information patterns and distributed strategy synthesis, in: Proceedings of the 13th International Symposium on Automated Technology for Verification and Analysis (ATVA 2015), LNCS, Springer, 2015, to appear.

[10] D. Berwanger, Ł. Kaiser, Information tracking in games on graphs, Journal of Logic, Language and Information 19 (4) (2010) 395–412.

[11] B. Finkbeiner, S. Schewe, Uniform distributed synthesis, in: Proceedings of the 20th IEEE Symposium on Logic in Computer Science (LICS 2005), IEEE, 2005, pp. 321–330.

[12] R. van der Meyden, N. V. Shilov, Model checking knowledge and time in systems with perfect recall, in: Proceedings of Foundations of Software Technology and Theoretical Computer Science (FSTTCS 1999), Vol. 1738 of LNCS, 1999, pp. 432–445. doi:10.1007/3-540-46691-6_35.

[13] D. Berwanger, Ł. Kaiser, B. Puchala, A perfect-information construction for coordination in games, in: Proceedings of Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2011), Vol. 13 of LIPICS, Leibniz-Zentrum für Informatik, 2011, pp. 387–398.

[14] R. Alur, T. A. Henzinger, O. Kupferman, Alternating-time temporal logic, Journal of the ACM 49 (5) (2002) 672–713.

[15] J. Y. Halpern, M. Y. Vardi, The complexity of reasoning about knowledge and time. I. Lower bounds, Journal of Computer and System Sciences 38 (1) (1989) 195 – 237. doi:10.1016/0022-0000(89)90039-1.

[16] W. Thomas, On the synthesis of strategies in infinite games, in: Proceedings of the 12th Annual Symposium on Theoretical Aspects of Computer Science, STACS '95, Vol. 900 of LNCS, Springer, 1995, pp. 1–13.

[17] W. Thomas, Automata on infinite objects, in: Handbook of Theoretical Computer Science, Vol. B: Formal Models and Semantics, Elsevier, 1990, pp. 133–192.

[18] E. Börger, E. Grädel, Y. Gurevich, The Classical Decision Problem, Perspectives in Mathematical Logic, Springer, 1997.

[19] P. Van Emde Boas, The convenience of tilings, in: Complexity, Logic, and Recursion Theory, Vol. 18 of Lecture Notes in Pure and Applied Mathematics, Marcel Dekker Inc, 1997, pp. 331–363.

[20] M. Latteux, D. Simplot, Context-sensitive string languages and recognizable picture languages, Information and Computation 138 (2) (1997) 160–169. doi:10.1006/inco.1997.2659.

[21] M. Latteux, D. Simplot, Recognizable picture languages and domino tiling, Theoretical Computer Science 178 (12) (1997) 275–283. doi:10.1016/S0304-3975(96)00283-6.

[22] A. Salomaa, Formal Languages, Academic Press, 1973.

[23] M. Fürer, The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems), in: Logic and Machines: Decision Problems and Complexity, Vol. 171 of LNCS, Springer, 1984, pp. 312–319. doi:10.1007/3-540-13331-3_48.

[24] R. J. Aumann, Agreeing to disagree, The Annals of Statistics 4 (6) (1976) 1236–1239. doi:10.1214/aos/1176343654.

[25] M. J. Osborne, A. Rubinstein, A Course in Game Theory, MIT Press, 1994.

[26] R. Fagin, J. Y. Halpern, Y. Moses, M. Y. Vardi, Reasoning about Knowledge, MIT Press, 1995.

[27] J. Geanakoplos, Common knowledge, in: R. Aumann, S. Hart (Eds.), Handbook of Game Theory with Economic Applications, 1st Edition, Vol. 2, Elsevier, 1994, Ch. 40, pp. 1437–1496.

[28] J. Barwise, Three views of common knowledge, in: Proceedings of the 2nd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK 1988), Morgan Kaufmann, 1988, pp. 365–379.

[29] D. Berwanger, Infinite coordination games, in: Logic and the Foundations of Game and Decision Theory, Vol. 6006 of LNCS, Springer, 2010, pp. 1–19. doi:10.1007/978-3-642-15164-4_1.

[30] J. Gray, Notes on data base operating systems, in: Operating Systems, An Advanced Course, Springer, 1978, pp. 393–481.

[31] J. Y. Halpern, Y. Moses, Knowledge and common knowledge in a distributed environment, Journal of the ACM (JACM) 37 (3) (1990) 549–587.

[32] A. Rubinstein, The Electronic Mail Game: Strategic Behavior under "Almost Common Knowledge", American Economic Review 79 (3) (1989) 385–391.

[33] D. Berwanger, M. Van den Bogaard, Consensus game acceptors, in: Proceedings of the 17th International Conference on Developments in Language Theory (DLT 2015), Vol. 9168 of LNCS, Springer, 2015, pp. 108–119.

[34] J. E. Hopcroft, J. D. Ullman, Some results on tape-bounded Turing machines, J. ACM 16 (1) (1969) 168–177. doi:10.1145/321495.321508.

[35] J. H. Reif, The complexity of two-player games of incomplete information, Journal of Computer and Systems Sciences 29 (2) (1984) 274–301.

[36] E. Grädel, W. Thomas, T. Wilke (Eds.), Automata, Logics, and Infinite Games, no. 2500 in LNCS, Springer-Verlag, 2002.

[37] M. Y. Vardi, P. Wolper, Reasoning about infinite computations, Information and Computation 115 (1994) 1–37.

[38] N. D. Jones, Space-Bounded Reducibility among Combinatorial Problems, Journal of Computer and System Sciences 11 (1) (1975) 68–85.

[39] D. S. Bernstein, R. Givan, N. Immerman, S. Zilberstein, The complexity of decentralized control of Markov decision processes, Mathematics of Operations Research 27 (4) (2002) 819–840. doi:10.1287/moor.27.4.819.297.

[40] E. A. Emerson, C. S. Jutla, Tree automata, mu-calculus and determinacy (extended abstract), in: Proceedings of the 32nd Annual Symposium on Foundations of Computer Science (FOCS 1991), IEEE Computer Society Press, 1991, pp. 368–377.

[41] W. Zielonka, Infinite games on finitely coloured graphs with applications to automata on infinite trees, Theoretical Computer Science 200 (1–2) (1998) 135–183.

[42] P. Madhusudan, P. Thiagarajan, Distributed controller synthesis for local specifications, in: Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001), Vol. 2076 of LNCS, Springer, 2001, pp. 396–407.

[43] B. Finkbeiner, S. Schewe, Coordination logic, in: Proceedings of the 19th Annual Conference on Computer Science Logic (CSL 2010), Vol. 6247 of LNCS, Springer, 2010, pp. 305–319. doi:10.1007/978-3-642-15205-4_25.

[44] P. Gastin, N. Sznajder, M. Zeitoun, Distributed synthesis for well-connected architectures, Formal Methods in System Design 34 (3) (2009) 215–237. doi:10.1007/s10703-008-0064-7.

[45] K. Chatterjee, T. Henzinger, J. Otop, A. Pavlogiannis, Distributed synthesis for LTL fragments, in: Formal Methods in Computer-Aided Design (FMCAD 2013), 2013, pp. 18–25. doi:10.1109/FMCAD.2013.6679386.

[46] A. Muscholl, I. Walukiewicz, Distributed synthesis for acyclic architectures, in: Proceedings of Foundation of Software Technology and Theoretical Computer Science (FSTTCS 2014), Vol. 29 of LIPIcs, Schloss Dagstuhl, 2014, pp. 639–651. doi:10.4230/LIPIcs.FSTTCS.2014.639.

[47] D. Berwanger, A. B. Mathew, Games with recurring certainty, in: Proceedings 2nd International Workshop on Strategic Reasoning, (SR 2014), Vol. 146 of EPTCS, 2014, pp. 91–96. doi:10.4204/EPTCS.146.12.