

Games with recurring certainty*

Dietmar Berwanger

Laboratoire Spécification et Vérification
CNRS & ENS Cachan, France
dwb@lsv.ens-cachan.fr

Anup Basil Mathew

Institute of Mathematical Sciences
Chennai, India
anupbasil@imsc.res.in

Infinite games where several players seek to coordinate under imperfect information are known to be intractable, unless the information flow is severely restricted. Examples of undecidable cases typically feature a situation where players become uncertain about the current state of the game, and this uncertainty lasts forever.

Here we consider games where the players attain certainty about the current state over and over again along any play. For finite-state games, we note that this kind of *recurring* certainty implies a stronger condition of *periodic* certainty, that is, the events of state certainty ultimately occur at uniform, regular intervals. We show that it is decidable whether a given game presents recurring certainty, and that, if so, the problem of synthesising coordination strategies under ω -regular winning conditions is solvable.

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 [10, 9, 8]. Absent such restrictions, however, the problem of coordinating two or more components of a distributed system with non-terminating executions is generally undecidable [11, 2].

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 the current position and the moves of the other players. The outcome of a play is an 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: that the play corresponds to a correct execution with respect to the system specification, no matter how Nature moves. 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, which was shown to be undecidable already in [12], for the basic setting of two players with a reachability condition, and in [7], for more complex winning conditions.

The cited undecidability arguments share a basic scenario: two players – he and she – become uncertain about the current state of the game, due to moves of Nature. As her (partial) knowledge of the state differs from his, and their actions need to respect the uncertainty of both, she needs to keep track not only of what she or he knows about the game state, but also, e.g., of what he knows about what she

*This work was partly supported by European project Cassting (FP7-ICT-601148).

knows that he knows, and so on. The scenario, set up so that the uncertainty never vanishes, leads to undecidability as the knowledge hierarchies grow unboundedly while the play proceeds [3].

The *information fork* criterion of [5] identifies distributed system architectures that may allow the knowledge of players to develop differently, for an unbounded number of rounds. Nevertheless, information forks may not cause undecidability in every context, for instance, if the “forked knowledge” is irrelevant for enforcing the winning condition, or if the effect of forking can be undone within a few rounds every time it occurs.

In this paper, we consider n -player games with imperfect information where the uncertainty of players about the game state cannot last forever. Our intuition of *recurring certainty* is that, whenever players are uncertain about the state of the game during a play, it takes only finitely many rounds until they can deduce the current state with certainty, and it becomes common knowledge among them. A faithful formalisation of this common knowledge property would most likely be undecidable. Thus, we resort to a weakening which intuitively states that the current state is evident to all players.

We show that the following two questions are decidable:

- Given an n -player game structure with imperfect information, does it satisfy the condition of recurring certainty?
- Given a game with recurring certainty and an ω -regular winning condition, does the grand coalition have a winning strategy?

Towards this, we first prove that, under recurring certainty, the intervals where the current state of the game is not common knowledge are bounded uniformly. We call this periodic certainty. Then, we show that the perfect-information *tracking* [4] of a game with periodic certainty is finite. This allows to solve the synthesis problem.

Acknowledgement. The authors thank Marie Van den Bogaard for useful discussions on related topics and for proof-reading this paper.

2 Coordination games with imperfect information

Our game model is close to that of concurrent games [1]. There are n players $1, \dots, n$ and a distinguished agent called nature. The *grand coalition* is the set $N = \{1, \dots, n\}$ of all players. We refer to a list of elements $x = (x^i)_{i \in N}$, 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 structure* $G = (V, \Delta, (\beta^i)_{i \in N})$ consists of a finite set V of *states*, a relation $\Delta \in V \times A \times V$ of simultaneous *moves* labelled by action profiles, and a profile of *observation functions* $\beta^i : V \rightarrow B^i$. We assume that each state has at least one outgoing move for every action profile, i.e., $\Delta(v, a) \neq \emptyset$, for all $v \in V$ and all $a \in A$.

Plays start at an initial state $v_0 \in V$ known to all players, and proceed in rounds. In a round, all players i choose an action $a^i \in A^i$ simultaneously, then nature chooses a successor state $v' \in \Delta(v, a)$ and each player i receives the observation $\beta^i(v')$. Notice that the players are not directly informed about the action chosen by other players nor the state chosen by nature. However, we assume that the player’s own action is part of his observation at the target state. Formally, a *play* is an infinite sequence $\pi = v_0, a_0, v_1, a_1, \dots$ alternating between positions and action profiles with $(v_\ell, a, v_{\ell+1}) \in \Delta$, for all $\ell \geq 0$. A *history* is a prefix $v_0, a_0, \dots, a_{\ell-1}, v_\ell$ of a play. The observation function extends from states to histories and plays $\pi = v_0, a_0, v_1, a_1, \dots$ as $\beta^i(\pi) = \beta^i(v_0), \beta^i(v_1), \dots$. 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 *strategy* for Player i is a mapping $s^i : (VA)^*V \rightarrow A^i$ 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 with S^i and the set of all strategy profiles by S . A history or play $\pi = v_0, a_0, v_1, a_1, \dots$ follows the strategy $s^i \in S^i$, if $a_\ell^i = s^i(v_0, a_0, v_1, \dots, a_{\ell-1}, v_\ell)$ for every $\ell > 0$. For the grand coalition, the play π follows a strategy profile 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 *winning condition* over a game structure G is a set $W \subseteq (VA)^\omega$ of plays. A *game* $\mathcal{G} = (G, W)$ consists of a game structure 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. To describe winning conditions, we use a colouring function $\gamma : V \rightarrow C$ with a finite range C of colours, and refer to the set $W \subseteq C^\omega$ of all plays $v_0, a_0, v_1, a_1, \dots$ with $\gamma(v_0), \gamma(v_1), \dots \in W$. In this paper, we assume that the colouring is *observable* to each player i , that is, $\beta^i(v) \neq \beta^i(v')$ whenever $\gamma(v) \neq \gamma(v')$.

We consider coordination games over finite game structures where the winning condition is given by finite-state automata. (See [6], for a comprehensive background.) Given such a game \mathcal{G} , we are interested in the following questions: (1) Does the grand coalition have a winning strategy profile in \mathcal{G} ? and (2) How to synthesise (distributed) winning strategies, if they exists?

3 Recurring certainty

We consider a class of games where the uncertainty of players about the current state is temporary and vanishes after a finite number of rounds.

To explain our notion of certainty, we introduce a fictitious player, let us call him Player 0, who is less informed than any actual player. He does not contribute to joint actions (i.e., his action set A^0 is a singleton), and his observation function is a coarsening of all observations of other players: for any pair v, v' of game states, $\beta^0(v) = \beta^0(v')$ whenever $\beta^i(v) = \beta^i(v')$ for some player i . Thus, for histories π, π' , we have $\pi \sim^0 \pi'$, whenever $\pi \sim^i \pi'$ for some player i (the converse does not hold, in general).

For a given game structure G , we say that the grand coalition *attains certainty* at history $\pi = v_0, a_0, \dots, a_{\ell-1}, v_\ell$, if any indistinguishable history $\pi' \sim^0 \pi$ ends at the same state v_ℓ . An infinite play π has *recurring certainty*, if the grand coalition attains certainty at infinitely many of its histories. Finally, we say that the game structure G has recurring certainty, if this is the case for every play in G .

As a simple example of a game with recurring certainty, consider the infinite repetition of a finite extensive game with imperfect information where the root is a perfect-information node, i.e., it is distinguishable from any other node, for every player. Likewise, games on graphs with the property that every cycle passes through a perfect-information state have recurring certainty.

We will also encounter the following stronger property. A game structure G has *periodic certainty* if there exists a uniform bound $t \in \mathbb{N}$ such that for every play π in G , every history ρ of π has a continuation ρ' by at most t rounds in π , such that the grand coalition attains certainty at ρ' .

3.1 Recognising games with recurring certainty

Our first result states that recurring certainty is a regular property of plays in finite game structures.

Lemma 3.1. *For any finite game structure, the set of plays where the grand coalition has recurring certainty is recognisable by a finite-state automaton.*

Proof. Let us fix a finite game structure G . First, we construct a word automaton \mathcal{A} over the alphabet AV that recognises histories ρ at which the grand coalition does not attain certainty. To witness this, the

automaton guesses a second history ρ' (of the same length) that is \sim^0 -indistinguishable from ρ and ends at a different state.

The state space of \mathcal{A} consists of pairs of game states in V , plus a sink. The first component of the automaton state keeps track of the input history and the second one of the uncertainty witness that is guessed nondeterministically. The transition function ensures that both components evolve according to the moves available in the game structure and yield the same observation to all players; otherwise, they lead to the sink. Accepting states are those where the first and the second component differ.

By complementing the automaton \mathcal{A} , we obtain an automaton $\overline{\mathcal{A}}$ that accepts the set of histories at which the grand coalition attains certainty (plus sequences that do not correspond to histories, which can be excluded easily by intersection with the unravelling of G). Next, we determinise $\overline{\mathcal{A}}$ and view the outcome as a deterministic Büchi automaton \mathcal{B} which accepts the input word, if it hits the set of final states infinitely often. Thus, \mathcal{B} accepts all plays where the grand coalition has recurring certainty. \square

The synchronous product of the deterministic Büchi automaton \mathcal{B} constructed above with the game structure G is universal, i.e. accepts every play of G , if and only if, G has recurring certainty.

Theorem 1. *The question whether a given game structure has recurring certainty is decidable.*

A further consequence of the automaton construction is that we obtain a uniform bound on the distance between two rounds at which the grand coalition attains certainty.

Theorem 2. *Every game with recurring certainty also has periodic certainty.*

Proof. Let G be a game structure with recurring certainty, \mathcal{B} the deterministic Büchi automaton constructed for G as above, and let t be the number of states in \mathcal{B} plus one. Towards a contradiction, suppose there exists a play π in G with a collection of $t > |\mathcal{B}|$ many consecutive histories $\rho_0, \rho_1, \dots, \rho_t$ at which the grand coalition does not attain certainty. Accordingly, the uniquely determined run of \mathcal{B} on input π hits no accepting state of the automaton \mathcal{B} while reading the continuation of ρ_0 up to ρ_t . On the other hand, as $t > |\mathcal{B}|$, there exists a state in \mathcal{B} that is reached by two different histories, say ρ_k and ρ_ℓ , with $0 \leq k < \ell \leq t$. Now we consider the play π' on G that begins with ρ_ℓ , and then repeats the continuation of ρ_k up to ρ_ℓ forever. Thus, the run of \mathcal{B} on π' will finally not hit any accepting state and be rejected, in contradiction to our assumption that G has recurring certainty. \square

3.2 Winner determination and strategy synthesis

Theorem 3. *Let \mathcal{G} be a coordination game with an ω -regular winning condition. If \mathcal{G} has recurring certainty, then the question whether the grand coalition has a winning strategy profile is decidable and the strategy synthesis problem is effectively solvable.*

Our argument relies on the tracking construction proposed in [4] that eliminates imperfect information in n -player games by an unravelling process that generates epistemic models of the player's information along the stages of a play. An *epistemic model* for a game structure \mathcal{G} is a Kripke structure $\mathcal{K} = (K, (Q_v)_{v \in V}, (\sim^i)_{i \in N})$ over a set K of histories in \mathcal{G} , equipped with predicates Q_v designating the histories that end in state $v \in V$ and the players' indistinguishability relations \sim^i . The construction keeps track of how the knowledge of players is updated by generating, for each epistemic model \mathcal{K} , a set of successor models along tuples $(a_k)_{k \in K}$ of action profiles $a_k \in A$ compatible with the player's current knowledge, i.e. for every $i \in N$ and for all $k, k' \in K$ with $k \sim^i k'$, we have $a_k^i = a_{k'}^i$. This leads to a possibly disconnected epistemic model with universe $K' = \{ka_kv \mid k \in K, k \in Q_w \text{ and } (w, a_k, v) \in \Delta\}$ with $Q_v = \{ka_kv \mid ka_kv \in K'\}$ and $ka_kv \sim_i k'a_kv' \iff k \sim_i^{\mathcal{K}} k' \text{ and } v \sim_i^{\mathcal{G}} v'$. By taking the connected

components of this model under the coarsening $\sim^U := \bigcup_{i=0}^{n-1} \sim_i$, we obtain the set of epistemic successor models. When starting from the trivial model that consists only of the initial node of the game, and successively applying the update, one unravels a tree labelled with epistemic models, which corresponds to a two-player game of perfect information where the strategies of one player translate to coordination strategies of the grand coalition in the original game, and vice versa. This tree structure, which in general may contain infinitely many distinct labels for its nodes (the undecidable game in [3], for example), is called the *tracking* of the game structure.

The main result of [4] shows that, whenever two nodes of the unravelling tree carry homomorphically equivalent labels, they can be identified without changing the (winning or losing) status of the game. This holds for all imperfect-information games with ω -regular winning conditions that are *observable*. Consequently, the strategy synthesis problem is decidable for a subclass of such games, whenever the unravelling process is guaranteed to generate only finitely many epistemic models, up to homomorphic equivalence.

Let us now consider the tracking of a game \mathcal{G} with an observable ω -regular winning condition. We claim that every history where the grand coalition attains certainty leads to an epistemic model that is homomorphically equivalent to the trivial structure consisting of a singleton labelled with the (certain) state at which the history ends. This is because every \sim^U -connected component is also \sim^0 -connected, and all histories in such a component end at the same state. On the other hand, when updating an epistemic model, the successor models can be at most exponentially larger (for fixed action space). The property of periodic certainty implied by recurring certainty, allows us to conclude that the number of updating rounds in which the models can grow is bounded by the certainty period of G . Therefore, games with recurring certainty have finite tracking. By [4], this implies that the winner determination problem is decidable for such games, and finite-state winning strategies can be effectively synthesised whenever the grand coalition has a winning strategy.

References

- [1] Rajeev Alur, Thomas A. Henzinger & Orna Kupferman (2002): *Alternating-time temporal logic*. *J. ACM* 49(5), pp. 672–713, doi:10.1145/585265.585270.
- [2] André Arnold & Igor Walukiewicz (2007): *Nondeterministic controllers of nondeterministic processes*. In: *Logic and Automata*, 2, Amsterdam University Press.
- [3] Dietmar Berwanger & Łukasz Kaiser (2010): *Information Tracking in Games on Graphs*. *Journal of Logic, Language and Information* 19(4), pp. 395–412, doi:10.1007/s10849-009-9115-8.
- [4] Dietmar Berwanger, Łukasz Kaiser & Bernd Puchala (2011): *Perfect-Information Construction for Coordination in Games*. In: *Proceedings of the 31st Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'11)*, *Leibniz International Proceedings in Informatics* 13, Leibniz-Zentrum für Informatik, Mumbai, India, pp. 387–398, doi:10.4230/LIPIcs.FSTTCS.2011.387.
- [5] B. Finkbeiner & S. Schewe (2005): *Uniform Distributed Synthesis*. In: *Proc. of LICS '05*, IEEE, pp. 321–330, doi:10.1109/LICS.2005.53.
- [6] Erich Grädel, Wolfgang Thomas & Thomas Wilke, editors (2002): *Automata, Logics, and Infinite Games*. LNCS 2500, Springer-Verlag, doi:10.1007/3-540-36387-4.
- [7] David Janin (2007): *On the (High) Undecidability of Distributed Synthesis Problems*. In: *Proc. of Theory and Practice of Computer Science (SOFSEM 2007)*, *Lecture Notes in Computer Science* 4362, Springer, pp. 320–329, doi:10.1007/978-3-540-69507-3_26.
- [8] Łukasz Kaiser (2006): *Game Quantification on Automatic Structures and Hierarchical Model Checking Games*. In: *Proc. of CSL '06*, LNCS 4207, Springer, pp. 411–425, doi:10.1007/11874683_27.

- [9] Orna Kupferman & Moshe Y. Vardi (2001): *Synthesizing Distributed Systems*. In: *Proc. of LICS '01*, IEEE Computer Society Press, pp. 389–398, doi:10.1109/LICS.2001.932514.
- [10] Amir Pnueli & Roni Rosner (1989): *On the synthesis of a reactive module*. In: *Proceedings of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, POPL '89*, ACM Press, pp. 179–190, doi:10.1145/75277.75293.
- [11] Amir Pnueli & Roni Rosner (1990): *Distributed Reactive Systems Are Hard to Synthesize*. In: *Proceedings of the 31st Annual Symposium on Foundations of Computer Science, FoCS '90*, IEEE Computer Society Press, pp. 746–757, doi:10.1109/FSCS.1990.89597.
- [12] John H. Reif (1984): *The Complexity of Two-Player Games of Incomplete Information*. *Journal of Computer and Systems Sciences* 29(2), pp. 274–301, doi:10.1016/0022-0000(84)90034-5.