

Strategy Logic with Imperfect Information

Raphaël Berthon*, Bastien Maubert†, Aniello Murano†, Sasha Rubin† and Moshe Y. Vardi‡

*École Normale Supérieure de Rennes, Rennes, France

†Università degli Studi di Napoli Federico II, Naples, Italy

‡Rice University, Houston, Texas, USA

Abstract—We introduce an extension of Strategy logic for the imperfect-information setting, called SL_{ii} , and study its model-checking problem. As this logic naturally captures multi-player games with imperfect information, the problem turns out to be undecidable. We introduce a syntactical class of “hierarchical instances” for which, intuitively, as one goes down the syntactic tree of the formula, strategy quantifications are concerned with finer observations of the model. We prove that model-checking SL_{ii} restricted to hierarchical instances is decidable. This result, because it allows for complex patterns of existential and universal quantification on strategies, greatly generalises previous ones, such as decidability of multi-player games with imperfect information and hierarchical observations, and decidability of distributed synthesis for hierarchical systems.

To establish the decidability result, we introduce and study $QCTL_{ii}^*$, an extension of QCTL (itself an extension of CTL with second-order quantification over atomic propositions) by parameterising its quantifiers with observations. The simple syntax of $QCTL_{ii}^*$ allows us to provide a conceptually neat reduction of SL_{ii} to $QCTL_{ii}^*$ that separates concerns, allowing one to forget about strategies and players and focus solely on second-order quantification. While the model-checking problem of $QCTL_{ii}^*$ is, in general, undecidable, we identify a syntactic fragment of hierarchical formulas and prove, using an automata-theoretic approach, that it is decidable. The decidability result for SL_{ii} follows since the reduction maps hierarchical instances of SL_{ii} to hierarchical formulas of $QCTL_{ii}^*$.

I. INTRODUCTION

Logics for strategic reasoning are a powerful tool for expressing correctness properties of multi-player graph-games, which in turn are natural models for reactive systems and discrete event systems. In particular, ATL^* and its related logics were introduced to capture the realisability/synthesis problem for open systems with multiple components. These logics were designed as extensions of branching-time logics such as CTL^* that allow one to write alternating properties directly in the syntax, *i.e.*, statements of the form “there exist strategies, one for each player in A , such that for all strategies of the remaining players, the resulting play satisfies φ ”. Strategy logic (SL) [1] generalises these by treating strategies as first-order objects x that can be quantified $\langle\langle x \rangle\rangle$ (read “there exists a strategy x ”) and bound to players (a, x) (read “player a uses strategy x ”). This syntax has flexibility very similar to

first-order logic, and thus allows one to directly express many solution concepts from game-theory, *e.g.*, the SL formula $\langle\langle x_1 \rangle\rangle \langle\langle x_2 \rangle\rangle (a_1, x_1)(a_2, x_2) \wedge_{i=1,2} [\langle\langle y_i \rangle\rangle (a_i, y_i) goal_i] \rightarrow goal_i$ expresses the existence of a Nash equilibrium in a two-player game (with individual Boolean objectives).

An essential property of realistic multi-player games is that players only have a limited view of the state of the system. This is captured by introducing partial-observability into the models, *i.e.*, equivalence-relations o (called *observations*) over the state space that specify indistinguishable states. In the formal-methods literature it is typical to associate observations to players. In this paper, instead, we associate observations to strategies. Concretely, we introduce an extension SL_{ii} of SL that annotates the strategy quantifier $\langle\langle x \rangle\rangle$ by an observation o , written $\langle\langle x \rangle\rangle^o$. Thus, both the model and the formulas mention observations o . This novelty allows one to express, in the logic, that a player’s observation changes over time.

Our logic SL_{ii} is extremely powerful: it is an extension of SL (which is in the perfect-information setting), and of the imperfect-information strategic logics $ATL_{i,R}^*$ [2] and $ATL_{sc,i}^*$ [3]. A canonical specification in multi-player game of partial observation is that the players, say a_1, \dots, a_n , can form a coalition and beat the environment player, say b . This can be expressed in SL_{ii} as $\Phi_{SYNTH} := \langle\langle x_1 \rangle\rangle^{o_1} \dots \langle\langle x_n \rangle\rangle^{o_n} \llbracket y \rrbracket^o (a_1, x_1) \dots (a_n, x_n)(b, y)\varphi$, where φ is quantifier- and binding-free. Also, SL_{ii} can express more complicated specifications by alternating quantifiers, binding the same strategy to different agents and rebinding (these are inherited from SL), as well as changing observations.

The complexity of SL_{ii} is also visible from an algorithmic point of view. Its satisfiability problem is undecidable (this is already true of SL), and its model-checking problem is undecidable (this is already true of $ATL_{i,R}^*$, in fact, even for the single formula $\langle\langle a, b \rangle\rangle \mathbf{F}p$ in systems with three players [4]). In fact, similar undecidability occurs in foundational work in multi-player games of partial observation, and in distributed synthesis [5], [6]. Since then, the formal-methods community has spent much effort finding restrictions and variations that ensure decidability [7]–[15]. The common thread in these approaches is that the players’ observations (or what they can deduce from their observations) are hierarchical.

Motivated by the problem of finding decidable extensions of strategy logic in the imperfect-information setting, we introduce a syntactical class of “hierarchical instances” of SL_{ii} , *i.e.*, formula/model pairs, and prove that the model-checking problem on this class of instances is decidable. Intuitively,

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 709188. It was also supported in part by NSF grants CCF-1319459 and IIS-1527668, and by NSF Expeditions in Computing project “ExCAPE: Expeditions in Computer Augmented Program Engineering”. Finally, this work was done while Sasha Rubin was a Marie Curie fellow of the Istituto Nazionale di Alta Matematica.

an instance of SL_{ii} is hierarchical if, as one goes down the syntactic tree of the formula, the observations annotating strategy quantifications can only become finer. Although the class of hierarchical instances refers not only to the syntax of the logic but also to the model, the class is syntactical in the sense that it depends only on the structure of the formula and the observations in the model. Moreover, it is easy to check (*i.e.*, in linear time) if an instance is hierarchical or not.

The class of hierarchical instances generalises some existing approaches and supplies new classes of systems and properties that can be model-checked. For instance, suppose that there is a total order \preceq among the players such that $a \preceq b$ implies player b 's observation is finer than player a 's observation — such games are said to yield “hierarchical observation” in [15]. In such games it is known that synthesis against CTL^* specifications is decidable (in fact, this holds for ω -regular specifications [8], [15]). This corresponds to hierarchical instances of SL_{ii} in which the observations form a total order in the model and the formula is of the form Φ_{SYNTH} (mentioned above). On the other hand, in hierarchical instances of SL_{ii} , the ordering on observations can be a pre partial-order (*i.e.*, not just total), and one can arbitrarily alternate quantifiers in the formulas. For instance, hierarchical instances allow one to decide if a game that yields hierarchical information has a Nash equilibrium. For example, assuming observations p_1, p_2, o_1, o_2 with p_i finer than o_2 (for $i = 1, 2$), and o_2 finer than o_1 , the formula $\langle\langle x_1 \rangle\rangle^{o_1} \langle\langle x_2 \rangle\rangle^{o_2} (a_1, x_1)(a_2, x_2) \wedge_{i=1,2} [\langle\langle y_i \rangle\rangle^{p_i} (a_i, y_i) goal_i] \rightarrow goal_i$ expresses that there exists a strategy profile (x_1, x_2) of uniform strategies, such that neither player has an incentive to deviate using a strategy that observes at least as much as the observations that both players started with. Observe that this formula is in fact *equivalent* to the existence of a Nash equilibrium, *i.e.*, to the same formula in which $p_i = o_i$. This shows we can decide the existence of Nash equilibria in games that yield hierarchical observation.

As a tool to study SL_{ii} we introduce $QCTL_{ii}^*$, an extension to the imperfect-information setting of $QCTL^*$ [16], itself an extension of CTL^* by second-order quantifiers over atoms. This is a low-level logic that does not mention strategies and into which one can effectively compile instances of SL_{ii} . States of the models of the logic $QCTL_{ii}^*$ have internal structure, much like the multi-player game structures from [17] and distributed systems [18]. Model-checking $QCTL_{ii}^*$ is also undecidable (indeed, we show how to reduce from the MSO-theory of the binary tree extended with the equal-length predicate, known to be undecidable [19]). We introduce the syntactical class $QCTL_{i,\subseteq}^*$ of hierarchical formulas as those in which innermost quantifiers observe more than outermost quantifiers, and prove that model-checking is decidable using an extension of the automata-theoretic approach for branching-time logics (our decision to base models of $QCTL_{ii}^*$ on local states greatly eases the use of automata). Moreover, the compilation of SL_{ii} into $QCTL_{ii}^*$ preserves being hierarchical, thus establishing our main contribution, *i.e.*, that model checking the hierarchical instances of SL_{ii} is decidable.

Related work. Formal methods for reasoning about reactive

systems with multiple components have been studied mainly in two theoretical frameworks: a) multi-player graph-games of partial-observation [6], [8], [15] and b) synthesis in distributed architectures [5], [7], [9], [11], [20] (the relationship between these two frameworks is discussed in [15]). All of these works consider the problem of synthesis, which (for objectives in temporal logics) can be expressed in SL_{ii} using the formula Φ_{SYNTH} mentioned above. Limited alternation was studied in [12] that, in the language of SL_{ii} , considers the model-checking problem of formulas of the form $\langle\langle x_1 \rangle\rangle^{o_1} \langle\langle x_2 \rangle\rangle^{o_2} \langle\langle x_3 \rangle\rangle^{o_3} (a_1, x_1)(a_2, x_2)(a_3, x_3)\varphi$, where φ is an ω -regular objective. They prove that this is decidable in case player 3 has perfect observation and player 2 observes at least as much as player 1.

In contrast to all these works, formulas of SL_{ii} can express much more complex specifications by alternating quantifiers, sharing strategies, rebinding, and changing observations.

Recently, [15] generalised the classic result of [8]: it weakens the assumption of hierarchical observation to hierarchical information (which are both static notions), and then, further to dynamic hierarchical information which allows for the hierarchy amongst players' information to change along a play. However, it only considers the synthesis problem.

We are aware of two papers that (like we do) give simultaneous structural constraints on both the formula and the model that result in decidability: in the context of synthesis in distributed architecture with process delays, [20] considers CTL^* specifications that constrain external variables by the input variables that may effect them in the architecture; and in the context of asynchronous perfect-recall, [10] considers a syntactical restriction on instances for Quantified μ -Calculus with partial observation (in contrast, we consider the case of synchronous perfect recall).

The work closest to ours is [21] which introduces a decidable logic CL in which one can encode many distributed synthesis problems. However, CL is close in spirit to our $QCTL_{i,\subseteq}^*$, and is more appropriate as a tool than as a high-level specification logic like SL_{ii} . Furthermore, by means of a natural translation we derive that CL is strictly included in the hierarchical instances of SL_{ii} (Section II-E). In particular, we find that hierarchical instances of SL_{ii} can express non-observable goals, while CL does not. Non-observable goals arise naturally in problems in distributed synthesis [5].

Finally, our logic SL_{ii} is the first generalisation of SL to include strategies with partial observation and, unlike CL , to generalise previous logics with partial-observation strategies, *i.e.*, $ATL_{i,R}^*$ [2] and $ATL_{sc,i}^*$ [3]. A comparison of SL_{ii} to SL , $ATL_{i,R}^*$, $ATL_{sc,i}^*$ and CL is given in Section II-E.

Plan. The definition of SL_{ii} and of hierarchical instances, and the discussion about Nash equilibria, are in Section II. The definition of $QCTL_{ii}^*$, and its hierarchical fragment $QCTL_{i,\subseteq}^*$, are in Section III. The proof that model checking $QCTL_{i,\subseteq}^*$ is decidable, including the required automata preliminaries, are in Section IV. The translation of SL_{ii} into $QCTL_{ii}^*$, and the fact that this preserves hierarchy, are in Section V.

II. SL WITH IMPERFECT INFORMATION

In this section we introduce SL_{ii} , an extension of SL [1] to the imperfect-information setting with synchronous perfect-recall. Our logic presents two original features: first, observations are not bound to players (as is done in extensions of ATL by imperfect information [22] or logics for reasoning about knowledge [23]), and second, we have syntactic observations in the language, which need to be interpreted.

We follow the presentation of SL in [1], except that we make some changes that simplify their presentation but do not change their semantics, and some that allow us to capture imperfect information. We introduce the syntax and semantics of SL_{ii} , carefully detailing these changes. We first fix some notation used throughout the paper.

A. Notation

Let Σ be an alphabet. A *finite* (resp. *infinite*) *word* over Σ is an element of Σ^* (resp. Σ^ω). Words are written $w = w_0w_1w_2\dots$, *i.e.*, indexing begins with 0. The *length* of a finite word $w = w_0w_1\dots w_n$ is $|w| := n + 1$, and $\text{last}(w) := w_n$ is its last letter. Given a finite (resp. infinite) word w and $0 \leq i \leq |w|$ (resp. $i \in \mathbb{N}$), we let w_i be the letter at position i in w , $w_{\leq i}$ is the prefix of w that ends at position i and $w_{\geq i}$ is the suffix of w that starts at position i . We write $w \preceq w'$ if w is a prefix of w' , and w^{\prec} is the set of finite prefixes of word w . Finally, the domain of a mapping f is written $\text{dom}(f)$, and for $n \in \mathbb{N}$ we let $[n] := \{i \in \mathbb{N} : 1 \leq i \leq n\}$. The literature sometimes refers to “imperfect information” and sometimes to “partial observation”; we will use the terms interchangeably.

B. Syntax

The syntax of SL_{ii} is similar to that of strategy logic SL as defined in [1]: the only difference is that we annotate strategy quantifiers $\langle\langle x \rangle\rangle$ by *observation symbols* o . For the rest of the paper, for convenience we fix a number of parameters for our logics and models: AP is a finite set of *atomic propositions*, Ag is a finite set of *agents* or *players*, Var is a finite set of *variables* and Obs is a finite set of *observation symbols*. When we consider model-checking problems, these data are implicitly part of the input.

Definition 1 (SL_{ii} Syntax). *The syntax of SL_{ii} is defined by the following grammar:*

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \langle\langle x \rangle\rangle^o \varphi \mid (a, x)\varphi$$

where $p \in \text{AP}$, $x \in \text{Var}$, $o \in \text{Obs}$ and $a \in \text{Ag}$.

We use abbreviations $\top := p \vee \neg p$, $\perp := \neg\top$, $\varphi \rightarrow \varphi' := \neg\varphi \vee \varphi'$, $\varphi \leftrightarrow \varphi' := \varphi \rightarrow \varphi' \wedge \varphi' \rightarrow \varphi$ for boolean connectives, $\mathbf{F}\varphi := \top \mathbf{U}\varphi$, $\mathbf{G}\varphi := \neg\mathbf{F}\neg\varphi$ for temporal operators, and finally $\llbracket x \rrbracket^o \varphi := \neg\langle\langle x \rangle\rangle^o \neg\varphi$.

The notion of free variables and sentences are defined as for SL . We recall these for completeness: A variable x appears *free* in a formula φ if it appears out of the scope of a strategy quantifier, and a player a appears free in φ if a temporal operator (either \mathbf{X} or \mathbf{U}) appears in φ out of the scope of any binding for player a . We let $\text{free}(\varphi)$ be the set of variables

and players that appear free in φ . If $\text{free}(\varphi)$ is empty, φ is a *sentence*.

C. Semantics

The models of SL_{ii} are like those of SL , *i.e.*, concurrent game structures, but extended by a finite set of observations Obs and, for each $o \in \text{Obs}$, by an equivalence-relation $\mathcal{O}(o)$ over positions that represents what a player using a strategy with that observation can see. That is, $\mathcal{O}(o)$ -equivalent positions are indistinguishable to a player using a strategy associated with observation o .

Definition 2 (CGS_{ii}). *A concurrent game structure with imperfect information (or CGS_{ii} for short) is a structure $\mathcal{G} = (\text{Ac}, V, E, \ell, v_i, \mathcal{O})$ where*

- Ac is a finite non-empty set of actions,
- V is a finite non-empty set of positions,
- $E : V \times \text{Ac}^{\text{Ag}} \rightarrow V$ is a transition function,
- $\ell : V \rightarrow 2^{\text{AP}}$ is a labelling function,
- $v_i \in V$ is an initial position,
- $\mathcal{O} : \text{Obs} \rightarrow 2^{V \times V}$ is an observation interpretation mapping observations to equivalence relations on positions.

We may write \sim_o for $\mathcal{O}(o)$, and $v \in \mathcal{G}$ for $v \in V$.

The notions of joint actions, plays, strategies and assignments are similar to those for SL . We will recall these for completeness. Since the main difference between the semantics of SL and SL_{ii} is in the strategy-quantification case where we require strategies to be consistent with observations, we define synchronous perfect recall and o -strategies before defining the semantics of SL_{ii} . Finally, we mention that we make some (inconsequential) simplifications to the semantics of SL as proposed in [1]: i) we dispense with partial assignments and only consider complete assignments, ii) our semantics are w.r.t. a finite play instead of a position (these simplifications do not change the expressive power of SL).

Joint actions. In a position $v \in V$, each player a chooses an action $c_a \in \text{Ac}$, and the game proceeds to position $E(v, \mathbf{c})$, where $\mathbf{c} \in \text{Ac}^{\text{Ag}}$ stands for the *joint action* $(c_a)_{a \in \text{Ag}}$. Given a joint action $\mathbf{c} = (c_a)_{a \in \text{Ag}}$ and $a \in \text{Ag}$, we let \mathbf{c}_a denote c_a . For each position $v \in V$, $\ell(v)$ is the set of atomic propositions that hold in v .

Plays and strategies. A *finite* (resp. *infinite*) *play* is a finite (resp. infinite) word $\rho = v_0 \dots v_n$ (resp. $\pi = v_0 v_1 \dots$) such that $v_0 = v_i$ and for all i with $0 \leq i < |\rho| - 1$ (resp. $i \geq 0$), there exists a joint action \mathbf{c} such that $E(v_i, \mathbf{c}) = v_{i+1}$. We let Plays be the set of finite plays. A *strategy* is a function $\sigma : \text{Plays} \rightarrow \text{Ac}$, and the set of all strategies is denoted Str .

Assignments. An *assignment* is a function $\chi : \text{Ag} \cup \text{Var} \rightarrow \text{Str}$, assigning a strategy to each player and variable. For an assignment χ , player a and strategy σ , $\chi[a \mapsto \sigma]$ is the assignment that maps a to σ and is equal to χ on the rest of its domain, and $\chi[x \mapsto \sigma]$ is defined similarly, where x is a variable.

Outcomes. For an assignment χ and a finite play ρ , we let $\text{out}(\chi, \rho)$ be the only infinite play that starts with ρ and is then extended by letting players follow the strategies assigned by χ .

Formally, $\text{out}(\chi, \rho) := \rho \cdot v_1 v_2 \dots$ where, for all $i \geq 0$, $v_{i+1} = E(v_i, \mathbf{c})$, where $v_0 = \text{last}(\rho)$ and $\mathbf{c} = (\chi(a)(\rho \cdot v_1 \dots v_i))_{a \in \text{Ag}}$.

Synchronous perfect recall. In this work we consider players with *synchronous perfect recall*, meaning that each player remembers the whole history of a play, a classic assumption in games with imperfect information and logics of knowledge and time. Each observation relation is thus extended to finite plays as follows: $\rho \sim_o \rho'$ if $|\rho| = |\rho'|$ and $\rho_i \sim_o \rho'_i$ for every $i \in \{0, \dots, |\rho| - 1\}$. For $o \in \text{Obs}$, an *o-strategy* is a strategy $\sigma : V^+ \rightarrow \text{Ac}$ such that $\sigma(\rho) = \sigma(\rho')$ whenever $\rho \sim_o \rho'$. The latter constraint captures the essence of imperfect information, which is that players can base their strategic choices only on the information available to them. For $o \in \text{Obs}$ we let Str_o be the set of all *o*-strategies.

Definition 3 (SL_{ii} Semantics). The semantics $\mathcal{G}, \chi, \rho \models \varphi$ is defined inductively, where φ is an SL_{ii}-formula, $\mathcal{G} = (\text{Ac}, V, E, \ell, v_i, \mathcal{O})$ is a CGS_{ii}, ρ is a finite play, and χ is an assignment:

$$\begin{aligned} \mathcal{G}, \chi, \rho \models p & \quad \text{if } p \in \ell(\text{last}(\rho)) \\ \mathcal{G}, \chi, \rho \models \neg \varphi & \quad \text{if } \mathcal{G}, \chi, \rho \not\models \varphi \\ \mathcal{G}, \chi, \rho \models \varphi \vee \varphi' & \quad \text{if } \mathcal{G}, \chi, \rho \models \varphi \text{ or } \mathcal{G}, \chi, \rho \models \varphi' \\ \mathcal{G}, \chi, \rho \models \langle\langle x \rangle\rangle^o \varphi & \quad \text{if } \exists \sigma \in \text{Str}_o \text{ s.t. } \mathcal{G}, \chi[x \mapsto \sigma], \rho \models \varphi \\ \mathcal{G}, \chi, \rho \models (a, x)\varphi & \quad \text{if } \mathcal{G}, \chi[a \mapsto \chi(x)], \rho \models \varphi \end{aligned}$$

and, writing $\pi = \text{out}(\chi, \rho)$:

$$\begin{aligned} \mathcal{G}, \chi, \rho \models \mathbf{X}\varphi & \quad \text{if } \mathcal{G}, \chi, \pi_{\leq |\rho|} \models \varphi \\ \mathcal{G}, \chi, \rho \models \varphi \mathbf{U}\varphi' & \quad \text{if } \exists i \geq 0 \text{ s.t. } \mathcal{G}, \chi, \pi_{\leq |\rho|+i-1} \models \varphi' \\ & \quad \text{and } \forall j \text{ s.t. } 0 \leq j < i, \\ & \quad \mathcal{G}, \chi, \pi_{\leq |\rho|+j-1} \models \varphi \end{aligned}$$

To explain the temporal operators, we remind the reader that positions begin at 0 and thus π_n is the $(n+1)$ -st position of π . Clearly, the satisfaction of a sentence is independent of the assignment. For an SL_{ii} sentence φ we thus let $\mathcal{G}, \rho \models \varphi$ if $\mathcal{G}, \chi, \rho \models \varphi$ for some assignment χ , and we write $\mathcal{G} \models \varphi$ if $\mathcal{G}, v_i \models \varphi$.

D. Model checking and hierarchical instances

Model Checking. We now introduce the main decision problem of this paper, *i.e.*, the model-checking problem for SL_{ii}. An SL_{ii}-instance is a formula/model pair (Φ, \mathcal{G}) where $\Phi \in \text{SL}_{ii}$ and \mathcal{G} is a CGS_{ii}. The *model-checking problem* for SL_{ii} is the decision problem that, given an SL_{ii}-instance (Φ, \mathcal{G}) , returns ‘yes’ if $\mathcal{G} \models \Phi$, and ‘no’ otherwise.

It is well known that deciding the existence of winning strategies in multi-player games with imperfect information is undecidable for reachability objectives [17]. Since this problem is easily reduced to the model-checking problem for SL_{ii}, we get the following result:

Theorem 1. *The model-checking problem for SL_{ii} is undecidable.*

Hierarchical instances. We isolate a sub-problem obtained by restricting attention to *hierarchical instances*. Intuitively, an SL_{ii}-instance (Φ, \mathcal{G}) is hierarchical if, as one goes down

a path in the syntactic tree of Φ , the observations tied to quantifications become finer.

Definition 4 (Hierarchical instances). An SL_{ii}-instance (Φ, \mathcal{G}) is hierarchical if for all subformulas φ_1, φ_2 of Φ of the form $\varphi_2 = \langle\langle x \rangle\rangle^{o_2} \varphi'_2$ and $\varphi_1 = \langle\langle y \rangle\rangle^{o_1} \varphi'_1$ where φ_1 is a subformula of φ'_2 , it holds that $\mathcal{O}(o_1) \subseteq \mathcal{O}(o_2)$.

If $\mathcal{O}(o_1) \subseteq \mathcal{O}(o_2)$ we say that o_1 is *finer* than o_2 in \mathcal{G} , and that o_2 is *coarser* than o_1 in \mathcal{G} . Intuitively, this means that a player with observation o_1 observes game \mathcal{G} no worse than, *i.e.*, is not less informed, a player with observation o_2 .

Example 1 (Security levels). We illustrate hierarchical instances in a “security levels” scenario, where higher levels have access to more data (*i.e.*, can observe more). Assume that the CGS_{ii} \mathcal{G} has $\mathcal{O}(o_3) \subseteq \mathcal{O}(o_2) \subseteq \mathcal{O}(o_1)$ (level 3 is the highest security clearance, level 1 is the lowest). Let $\varphi = (a_1, x_1)(a_2, x_2)(a_3, x_3)\mathbf{G}p$. The SL_{ii} formula $\Phi := \langle\langle x_1 \rangle\rangle^{o_1} \llbracket x_2 \rrbracket^{o_2} \langle\langle x_3 \rangle\rangle^{o_3} \varphi$ and \mathcal{G} together form a hierarchical instance. It expresses that player a_1 (with lowest clearance) can collude with player a_3 (with highest clearance) to ensure a safety property p , even in the presence of an adversary a_2 (with intermediate clearance), as long as the strategy used by a_3 can depend on the strategy used by a_2 .

On the other hand, formula $\langle\langle x_1 \rangle\rangle^{o_1} \langle\langle x_3 \rangle\rangle^{o_3} \llbracket x_2 \rrbracket^{o_2} \varphi$, which is similar to Φ except that the strategy used by a_3 cannot depend on the adversarial strategy used by a_2 , does not form a hierarchical instance with \mathcal{G} .

Here is the main contribution of this paper:

Theorem 2. *The model-checking problem for SL_{ii} restricted to the class of hierarchical instances is decidable.*

This is proved in Section V by reducing it to the model-checking problem of the hierarchical fragment of a logic called QCTL_{ii}^{*}, which we introduce, and prove decidable, in Section III. We now give an important corollary of Theorem 2.

A Nash equilibrium in a game is a tuple of strategies such that no player has the incentive to deviate. Assuming that $\text{Ag} = \{a_i : i \in [n]\}$ and goals are written in SL_{ii}, say goal_i for $i \in [n]$, the following formula of SL_{ii} expresses the existence of a Nash equilibrium:

$$\Phi_{\text{NE}} := \langle\langle x_1 \rangle\rangle^{o_1} \dots \langle\langle x_n \rangle\rangle^{o_n} (a_1, x_1) \dots (a_n, x_n) \Psi_{\text{NE}}$$

where $\Psi_{\text{NE}} := \bigwedge_{i \in [n]} [\langle\langle y_i \rangle\rangle^{o_i} (a_i, y_i) \text{goal}_i \rightarrow \text{goal}_i]$.

A CGS_{ii} \mathcal{G} is said to *yield hierarchical observation* [15] if the “finer-than” relation is a total ordering, *i.e.*, if for all $o, o' \in \text{Obs}$, either $\mathcal{O}(o) \subseteq \mathcal{O}(o')$ or $\mathcal{O}(o') \subseteq \mathcal{O}(o)$.

Note that even if \mathcal{G} yields hierarchical information, the instance $(\Phi_{\text{NE}}, \mathcal{G})$ is *not* hierarchical (unless $\mathcal{O}(o_i) = \mathcal{O}(o_j)$ for all $i, j \in [n]$). *Nonetheless*, we can decide if a game that yields hierarchical observation has a Nash equilibrium:

Corollary 1. *The following problem is decidable: given a CGS_{ii} that yields hierarchical observation, whether $\mathcal{G} \models \Phi_{\text{NE}}$.*

Proof. The main idea is to use the fact that in a one-player game of partial-observation (such a game occurs when all but

one player have fixed their strategies, as in the definition of Nash equilibrium), the player has a strategy enforcing some goal iff the player has a uniform strategy enforcing that goal. Here are the details. Let $\mathcal{G} = (\text{Ac}, V, E, \ell, v_\ell, \mathcal{O})$ be a CGS_{ii} that yields hierarchical observation. Suppose the observation set is Obs. To decide if $\mathcal{G} \models \Phi_{\text{NE}}$ first build a new CGS_{ii} $\mathcal{G}' = (\text{Ac}, V, E, \ell, v_\ell, \mathcal{O}')$ over observations $\text{Obs}' := \text{Obs} \cup \{o_p\}$ such that $\mathcal{O}'(o) = \mathcal{O}(o)$ and $\mathcal{O}'(o_p) = \{(v, v) : v \in V\}$, and consider the sentence

$$\Phi' := \langle\langle x_1 \rangle\rangle^{o_1} \dots \langle\langle x_n \rangle\rangle^{o_n} (a_1, x_1) \dots (a_n, x_n) \Psi'$$

where $\Psi' := \bigwedge_{i \in [n]} [(\langle\langle y_i \rangle\rangle^{o_p} (a_i, y_i) \text{goal}_i) \rightarrow \text{goal}_i]$.

Then (Φ', \mathcal{G}') is a hierarchical instance, and by Theorem 2 we can decide $\mathcal{G}' \models \Phi'$. We claim that $\mathcal{G}' \models \Phi'$ iff $\mathcal{G} \models \Phi_{\text{NE}}$. To see this, it is enough to establish that:

$$\mathcal{G}', \chi, v_\ell \models \langle\langle y_i \rangle\rangle^{o_p} (a_i, y_i) \text{goal}_i \leftrightarrow \langle\langle y_i \rangle\rangle^{o_i} (a_i, y_i) \text{goal}_i,$$

for every $i \in [n]$ and every assignment χ such that $\chi(x_i) = \chi(a_i)$ is an o_i -uniform strategy.

To this end, fix i and χ . The right-to-left implication is immediate (since o_p is finer than o_i). For the converse, let σ be a p -uniform strategy (i.e., perfect-information) such that $\mathcal{G}', \chi[y_i \mapsto \sigma, a_i \mapsto \sigma], v_\ell \models \text{goal}_i$. Let $\pi := \text{out}(\chi[y_i \mapsto \sigma, a_i \mapsto \sigma], v_\ell)$. Construct an o_i -uniform strategy σ' that agrees with σ on prefixes of π . This can be done as follows: if $\rho \sim_{o_i} \pi_{\leq j}$ for some j then define $\sigma'(\rho) = \sigma(\pi_{\leq j})$ (note that this is well-defined since if there is some such j then it is unique), and otherwise define $\sigma'(\rho) = a$ for some fixed action $a \in \text{Ac}$. \square

E. Comparison with other logics

The main difference between SL and ATL-like strategic logics is that in the latter a strategy is always bound to some player, while in the former bindings and quantifications are separated. This separation adds expressive power, e.g., one can bind the same strategy to different players. Extending ATL with imperfect-information is done by giving each player an indistinguishability relation that its strategies must respect [2]. Our extension of SL by imperfect information, instead, assigns each strategy x an indistinguishability relation o when it is quantified $\langle\langle x \rangle\rangle^o$. Thus $\langle\langle x \rangle\rangle^o \varphi$ means “there exists a strategy with observation o such that φ holds”. Associating observations in this way, i.e., to strategies rather than players has two consequences. First, it is a clean generalisation of SL in the perfect information setting [1]. Define the *perfect-information fragment* of SL_{ii} to be the logic SL_{ii} assuming that $\text{Obs} = \{o\}$ and $\mathcal{O}(o) = \{(v, v) : v \in \mathcal{G}\}$ for every CGS_{ii} \mathcal{G} ; also let us assimilate such structures with classic perfect-information concurrent game structures (CGS), which are the models of SL. Finally, let $\text{tr}_1 : \text{SL} \rightarrow \text{SL}_{\text{ii}}$ be the trivial translation that annotates each strategy quantifier $\langle\langle x \rangle\rangle$ with observation o . The next proposition says that the perfect-information fragment of SL_{ii} is a notational variant of SL.

Proposition 1. *For every SL sentence φ and every CGS \mathcal{G} , it holds that $\mathcal{G} \models \varphi$ iff $\mathcal{G} \models \text{tr}_1(\varphi)$.*

Second, SL_{ii} subsumes imperfect-information extensions of ATL* that associate observations to players.

We recall that an ATL_{i,R}* formula¹ $\langle A \rangle \psi$ reads as “there are strategies for players in A such that ψ holds whatever players in $\text{Ag} \setminus A$ do”. Consider the translation $\text{tr}_2 : \text{ATL}_{i,R}^* \rightarrow \text{SL}_{\text{ii}}$ that replaces each subformula of the form $\langle A \rangle \psi$, where $A = \{a_1, \dots, a_k\} \subset \text{Ag}$ is a coalition of players and $\text{Ag} \setminus A = \{a_{k+1}, \dots, a_n\}$, with formula $\langle\langle x_1 \rangle\rangle^{o_1} \dots \langle\langle x_k \rangle\rangle^{o_k} \llbracket x_{k+1} \rrbracket^{o_p} \dots \llbracket x_n \rrbracket^{o_p} (a_1, x_1) \dots (a_n, x_n) \psi'$, where $\psi' = \text{tr}_2(\psi)$. Also, for every CGS_{ii} as considered in the semantics of ATL_i, i.e., where each agent is assigned an equivalence relation on positions (let us refer to such structures as ATL-CGS_{ii}), define the CGS_{ii} \mathcal{G}' by interpreting each o_i as the equivalence relation for agent a_i in \mathcal{G} , and interpreting o_p as the identity relation.

Proposition 2. *For every ATL_{i,R}* formula φ and ATL-CGS_{ii} \mathcal{G} , it holds that $\mathcal{G} \models \varphi$ iff $\mathcal{G}' \models \text{tr}_2(\varphi)$.*

Third, SL_{ii} also subsumes the imperfect-information extension of ATL* with strategy context (see [3] for the definition of ATL_{sc}* with partial observation, which we refer to as ATL_{sc,i}*). The only difference between ATL_{sc,i}* and ATL_{i,R}* is the following: in ATL_{i,R}*, when a subformula of the form $\langle A \rangle \psi$ is met, we quantify existentially on strategies for players in A , and then we consider all possible outcomes obtained by letting other players behave however they want. Therefore, if any player in $\text{Ag} \setminus A$ had previously been assigned a strategy, it is forgotten. In ATL_{sc,i}* on the other hand, these strategies are stored in a *strategy context*, which is a *partial* assignment χ , defined for the subset of players currently bound to a strategy. A strategy context allows one to quantify universally only on strategies of players who are not in A and who are not already bound to a strategy. It is then easy to define a translation $\text{tr}_3 : \text{ATL}_{\text{sc,i}}^* \rightarrow \text{SL}_{\text{ii}}$ by adapting translation tr_2 from Proposition 2, with the strategy context as parameter. For an ATL-CGS_{ii} \mathcal{G} , the CGS_{ii} \mathcal{G}' is defined as for Proposition 2.

Proposition 3. *For every ATL_{sc,i}* formula φ and ATL-CGS_{ii} \mathcal{G} , it holds that $\mathcal{G} \models \varphi$ iff $\mathcal{G}' \models \text{tr}_3(\varphi)$.*

Fourth, there is a natural and simple translation of instances of the model-checking problem of CL [21] into the hierarchical instances of SL_{ii}. Moreover, the image of this translation consists of instances of SL_{ii} with a very restricted form, i.e., atoms mentioned in the SL_{ii}-formula are *observable* for all observations of the CGS_{ii}, i.e., players know the truth value of all atoms in all positions, for any observation they are assigned.

Proposition 4. *There is an effective translation that, given a CL-instance (\mathcal{S}, φ) produces a hierarchical SL_{ii}-instance (\mathcal{G}, Φ) such that*

- 1) $\mathcal{S} \models \varphi$ iff $\mathcal{G} \models \Phi$,
- 2) For all atoms p in Φ , observations $o \in \text{Obs}$ and positions $v, v' \in \mathcal{G}$, if $v \sim_o v'$ then $p \in \ell(v) \leftrightarrow p \in \ell(v')$.

¹See [2] for the definition of ATL_{i,R}*, where subscript i refers to “imperfect information” and subscript R to “perfect recall”. Also, we consider the so-called *objective semantics* for ATL_{i,R}*.

To do this, one first translates CL into (hierarchical) $\text{QCTL}_{\text{ii}}^*$, the latter is defined in the next section. This step is a simple reflection of the semantics of CL in that of $\text{QCTL}_{\text{ii}}^*$. Then one translates $\text{QCTL}_{\text{ii}}^*$ into SL_{ii} by a simple adaptation of the translation of QCTL^* into ATL_{sc}^* [24].

III. QCTL^* WITH IMPERFECT INFORMATION

In this section we introduce an imperfect-information extension of QCTL^* [16], [25]–[29]. In order to introduce imperfect information, instead of considering equivalence relations between states as in concurrent game structures, we will enrich Kripke structures by giving internal structure to their states, *i.e.*, we see states as n -tuples of local states. This way of modelling imperfect information is inspired from Reif’s multi-player game structures [17] and distributed systems [18], and we find it very suitable to application of automata techniques, as discussed in Section III-C.

The syntax of $\text{QCTL}_{\text{ii}}^*$ is similar to that of QCTL^* , except that we annotate second-order quantifiers by subsets $\mathbf{o} \subseteq [n]$. The idea is that quantifiers annotated by \mathbf{o} can only “observe” the local states indexed by $i \in \mathbf{o}$. We define the tree-semantics of $\text{QCTL}_{\text{ii}}^*$: this means that we interpret formulas on trees that are the unfoldings of Kripke structures (this will capture the fact that players in SL_{ii} have synchronous perfect recall).

We then define the syntactic class of *hierarchical formulas* and prove, using an automata-theoretic approach, that model checking this class of formulas is decidable.

A. $\text{QCTL}_{\text{ii}}^*$ Syntax

The syntax of $\text{QCTL}_{\text{ii}}^*$ is very similar to that of QCTL^* : the only difference is that we annotate quantifiers by a set of indices that defines the “observation” of that quantifier.

Definition 5 ($\text{QCTL}_{\text{ii}}^*$ Syntax). *Fix $n \in \mathbb{N}$. The syntax of $\text{QCTL}_{\text{ii}}^*$ is defined by the following grammar:*

$$\begin{aligned} \varphi &:= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{E}\psi \mid \exists^{\mathbf{o}}p. \varphi \\ \psi &:= \varphi \mid \neg\psi \mid \psi \vee \psi \mid \mathbf{X}\psi \mid \psi \mathbf{U}\psi \end{aligned}$$

where $p \in \text{AP}$ and $\mathbf{o} \subseteq [n]$.

Formulas of type φ are called *state formulas*, those of type ψ are called *path formulas*, and $\text{QCTL}_{\text{ii}}^*$ consists of all the state formulas defined by the grammar. We use standard abbreviation $\mathbf{A}\psi := \neg\mathbf{E}\neg\psi$. We also use $\exists p. \varphi$ as a shorthand for $\exists^{[n]}p. \varphi$, and we let $\forall p. \varphi := \neg\exists p. \neg\varphi$.

Given a $\text{QCTL}_{\text{ii}}^*$ formula φ , we define the set of *quantified propositions* $\text{AP}_{\exists}(\varphi) \subseteq \text{AP}$ as the set of atomic propositions p such that φ has a subformula of the form $\exists^{\mathbf{o}}p. \varphi$. We also define the set of *free propositions* $\text{AP}_f(\varphi) \subseteq \text{AP}$ as the set of atomic propositions that appear out of the scope of any quantifier of the form $\exists^{\mathbf{o}}p$. Observe that $\text{AP}_{\exists}(\varphi) \cap \text{AP}_f(\varphi)$ may not be empty in general, *i.e.*, a proposition may appear both free and quantified in (different places of) a formula.

B. $\text{QCTL}_{\text{ii}}^*$ tree-semantics

We define the semantics on structures whose states are tuples of local states.

Local states. Let $\{L_i\}_{i \in [n]}$ denote $n \in \mathbb{N}$ disjoint finite sets of *local states*. For $I \subseteq [n]$, we let $L_I := \prod_{i \in I} L_i$ if $I \neq \emptyset$, and $L_{\emptyset} := \{\mathbf{0}\}$ where $\mathbf{0}$ is a special symbol.

Concrete observations. A set $\mathbf{o} \subseteq [n]$ is called a *concrete observation* (to distinguish it from observations o in the definitions of SL_{ii}).

Compound Kripke structures. These are like Kripke structures except that the states are elements of $L_{[n]}$. A *compound Kripke structure*, or CKS, over AP, is a tuple $\mathcal{S} = (S, R, s_{\ell}, \ell)$ where $S \subseteq L_{[n]}$ is a set of *states*, $R \subseteq S \times S$ is a left-total² *transition relation*, $s_{\ell} \in S$ is an *initial state*, and $\ell : S \rightarrow 2^{\text{AP}}$ is a *labelling function*.

A *path* in \mathcal{S} is an infinite sequence of states $\lambda = s_0 s_1 \dots$ such that for all $i \in \mathbb{N}$, $(s_i, s_{i+1}) \in R$. For $s \in S$, we let $\text{Paths}(s)$ be the set of all paths that start in s . A *finite path* is a finite non-empty prefix of a path. We may write $s \in \mathcal{S}$ for $s \in S$. Since we will interpret $\text{QCTL}_{\text{ii}}^*$ on unfoldings of CKS, we now define infinite trees.

Trees. In many works, trees are defined as prefix-closed sets of words with the empty word ϵ as root. Here trees represent unfoldings of Kripke structures, and we find it more convenient to see a node u as a sequence of states and the root as the initial state. Let X be a finite set (typically a set of states). An *X-tree* τ is a nonempty set of words $\tau \subseteq X^+$ such that:

- there exists $r \in X$, called the *root* of τ , such that each $u \in \tau$ starts with r ($r \preceq u$);
- if $u \cdot x \in \tau$ and $u \neq \epsilon$, then $u \in \tau$, and
- if $u \in \tau$ then there exists $x \in X$ such that $u \cdot x \in \tau$.

The elements of a tree τ are called *nodes*. If $u \cdot x \in \tau$, we say that $u \cdot x$ is a *child* of u . The *depth* of a node u is $|u|$. An *X-tree* τ is *complete* if $u \in \tau, x \in X$ implies $u \cdot x \in \tau$. A *path* in τ is an infinite sequence of nodes $\lambda = u_0 u_1 \dots$ such that for all $i \in \mathbb{N}$, u_{i+1} is a child of u_i , and $\text{Paths}(u)$ is the set of paths that start in node u . An *AP-labelled X-tree*, or (AP, X)-tree for short, is a pair $t = (\tau, \ell)$, where τ is an *X-tree* called the *domain* of t and $\ell : \tau \rightarrow 2^{\text{AP}}$ is a *labelling*. For a labelled tree $t = (\tau, \ell)$ and an atomic proposition $p \in \text{AP}$, we define the *p-projection* of t as the labelled tree $t \downarrow_p := (\tau, \ell \downarrow_p)$, where for each $u \in \tau$, $\ell \downarrow_p(u) := \ell(u) \setminus \{p\}$. For a set of trees \mathcal{L} , we let $\mathcal{L} \downarrow_p := \{t \downarrow_p \mid t \in \mathcal{L}\}$. Finally, two labelled trees $t = (\tau, \ell)$ and $t' = (\tau', \ell')$ are *equivalent modulo p*, written $t \equiv_p t'$, if $t \downarrow_p = t' \downarrow_p$ (in particular, $\tau = \tau'$).

Narrowing. Let X and Y be two finite sets, and let $(x, y) \in X \times Y$. The *X-narrowing* of (x, y) is $(x, y) \downarrow_X := x$. This definition extends naturally to words over $X \times Y$ (pointwise), and thus to $X \times Y$ -trees.

For $J \subseteq I \subseteq [n]$ and $z = (l_i)_{i \in I} \in L_I$, we also define $z \downarrow_J := z \downarrow_{L_J}$, where z is seen as a pair $z = (z_1, z_2) \in L_J \times L_{I \setminus J}$. This is well defined because having taken sets L_i to be disjoint, the ordering of local states in z is indifferent. We also extend this definition to words and trees.

²*i.e.*, for all $s \in S$, there exists s' such that $(s, s') \in R$.

Observe that when narrowing a tree, nodes with same narrowing are merged. In particular, for every L_I -tree τ , $\tau \downarrow_{\emptyset}$ is the only L_{\emptyset} -tree, $\mathbf{0}^\omega$.

Quantification and uniformity. In $\text{QCTL}_{\mathbf{ii}}^*$ the intuitive meaning of $\exists^{\mathbf{o}}p. \varphi$ in a tree t is that there is some equivalent tree t' modulo p such that t' is \mathbf{o} -uniform in p and satisfies φ . Intuitively, a tree is \mathbf{o} -uniform in p if it is uniformly labelled by p , i.e., if every two nodes that are indistinguishable when projected onto the local states indexed by $\mathbf{o} \subseteq [n]$ agree on their labelling of p .

Definition 6 (\mathbf{o} -indistinguishability and \mathbf{o} -uniformity in p). Fix $\mathbf{o} \subseteq [n]$ and $I \subseteq [n]$.

- Two tuples $x, x' \in L_I$ are \mathbf{o} -indistinguishable, written $x \approx_{\mathbf{o}} x'$, if $x \downarrow_{I \cap \mathbf{o}} = x' \downarrow_{I \cap \mathbf{o}}$.
- Two words $u = u_0 \dots u_i$ and $u' = u'_0 \dots u'_j$ over alphabet L_I are \mathbf{o} -indistinguishable, written $u \approx_{\mathbf{o}} u'$, if $i = j$ and for all $k \in \{0, \dots, i\}$ we have $u_k \approx_{\mathbf{o}} u'_k$.
- A tree t is \mathbf{o} -uniform in p if for every pair of nodes $u, u' \in \tau$ such that $u \approx_{\mathbf{o}} u'$, we have $p \in \ell(u)$ iff $p \in \ell(u')$.

Finally, we inductively define the satisfaction relation \models for the semantics on trees, where $t = (\tau, \ell)$ is a 2^{AP} -labelled L_I -tree, u is a node and λ is a path in τ :

$t, u \models p$	if	$p \in \ell(u)$
$t, u \models \neg \varphi$	if	$t, u \not\models \varphi$
$t, u \models \varphi \vee \varphi'$	if	$t, u \models \varphi$ or $t, u \models \varphi'$
$t, u \models \mathbf{E}\psi$	if	$\exists \lambda \in \text{Paths}(u)$ s.t. $t, \lambda \models \psi$
$t, u \models \exists^{\mathbf{o}}p. \varphi$	if	$\exists t' \equiv_p t$ s.t. t' is \mathbf{o} -uniform in p and $t', u \models \varphi$.
$t, \lambda \models \varphi$	if	$t, \lambda_0 \models \varphi$
$t, \lambda \models \neg \psi$	if	$t, \lambda \not\models \psi$
$t, \lambda \models \psi \vee \psi'$	if	$t, \lambda \models \psi$ or $t, \lambda \models \psi'$
$t, \lambda \models \mathbf{X}\psi$	if	$t, \lambda_{\geq 1} \models \psi$
$t, \lambda \models \psi \mathbf{U} \psi'$	if	$\exists i \geq 0$ s.t. $t, \lambda_{\geq i} \models \psi'$ and $\forall j$ s.t. $0 \leq j < i$, $t, \lambda_{\geq j} \models \psi$

We write $t \models \varphi$ for $t, r \models \varphi$, where r is the root of t .

Example 2. Consider the following CTL formula:

$$\text{border}(p) := \mathbf{AF}p \wedge \mathbf{AG}(p \rightarrow \mathbf{AXAG}\neg p).$$

This formula holds in a labelled tree if and only if each path contains exactly one node labelled with p . Now, consider the following $\text{QCTL}_{\mathbf{ii}}^*$ formula:

$$\text{level}(p) := \exists^{\mathbf{o}}p. \text{border}(p).$$

For a blind quantifier, two nodes of a tree are indistinguishable if and only if they have same depth. Therefore, this formula holds on a tree iff the p 's label all and only the nodes at some fixed depth. This formula can thus be used to capture the equal level predicate on trees. Actually, just as QCTL^* captures MSO, one can prove that $\text{QCTL}_{\mathbf{ii}}^*$ with tree semantics subsumes MSO with equal level [19], [30], [31]. In Theorem 3

we make use of a similar observation to prove that model-checking $\text{QCTL}_{\mathbf{ii}}^*$ is undecidable.

Model-checking problem for $\text{QCTL}_{\mathbf{ii}}^*$ under tree semantics.

For the model-checking problem, we interpret $\text{QCTL}_{\mathbf{ii}}^*$ on unfoldings of CKSs.

Tree unfoldings $t_{\mathcal{S}}(s)$. Let $\mathcal{S} = (S, R, s_i, \ell)$ be a compound Kripke structure over AP, and let $s \in S$. The *tree-unfolding of \mathcal{S} from s* is the (AP, S) -tree $t_{\mathcal{S}}(s) := (\tau, \ell')$, where τ is the set of all finite paths that start in s , and for every $u \in \tau$, $\ell'(u) := \ell(\text{last}(u))$. Given a CKS \mathcal{S} , a state $s \in S$ and a $\text{QCTL}_{\mathbf{ii}}^*$ formula φ , we write $\mathcal{S}, s \models \varphi$ if $t_{\mathcal{S}}(s) \models \varphi$. Write $\mathcal{S} \models \varphi$ if $t_{\mathcal{S}}(s_i) \models \varphi$.

The *model-checking problem for $\text{QCTL}_{\mathbf{ii}}^*$ under tree-semantics* is the following decision problem: given an instance (φ, \mathcal{S}) where \mathcal{S} is a CKS, and φ is a $\text{QCTL}_{\mathbf{ii}}^*$ formula, return ‘Yes’ if $\mathcal{S} \models \varphi$ and ‘No’ otherwise.

C. Discussion of the definition of $\text{QCTL}_{\mathbf{ii}}^*$

We now motivate in detail some aspects of $\text{QCTL}_{\mathbf{ii}}^*$.

Modelling of imperfect information. We model imperfect information by means of local states (rather than equivalence relations) since this greatly facilitates the use of automata techniques. More precisely, in our decision procedure of Section IV, we make extensive use of an operation on tree automata called *narrowing*, which was introduced in [32] to deal with imperfect-information in the context of distributed synthesis for temporal specifications. Given an automaton \mathcal{A} that works on $X \times Y$ -trees, where X and Y are two finite sets, and assuming that we want to model an operation performed on trees while observing only the X component of each node, this narrowing operation allows one to build from \mathcal{A} an automaton \mathcal{A}' that works on X -trees, such that \mathcal{A}' accepts an X -tree if and only if \mathcal{A} accepts its widening to $X \times Y$ (see Section IV for details). One can then make this automaton \mathcal{A}' perform the desired operation, which will by necessity be performed uniformly with regards to the partial observation, since the Y component is absent from the input trees.

With our definition of compound Kripke structures, their unfoldings are trees over the Cartesian product $L_{[n]}$. To model a quantification $\exists^{\mathbf{o}}p$ with observation $\mathbf{o} \subseteq [n]$, we can thus use the narrowing operation to forget about components L_i , for $i \in [n] \setminus \mathbf{o}$. We then use the classic projection of nondeterministic tree automata to perform existential quantification on atomic proposition p . Since the choice of the p -labelling is made directly on $L_{\mathbf{o}}$ -trees, it is necessarily \mathbf{o} -uniform.

Choice of the tree semantics. QCTL^* is obtained by adding to CTL^* second-order quantification on atomic propositions. Several semantics have been considered. The two most studied ones are the *structure semantics*, in which formulas are evaluated directly on Kripke structures, and the *tree semantics*, in which Kripke structures are first unfolded into infinite trees. Tree semantics thus allows quantifiers to choose the value of a quantified atomic proposition in each *finite path* of the model, while in structure semantics the choice is only made in each state. When QCTL^* is used to express existence

of strategies, existential quantification on atomic propositions labels the structure with strategic choices; in this kind of application, structure semantics reflects so-called *positional* or *memoryless* strategies, while tree semantics captures *perfect-recall* or *memoryfull* strategies. Since in this work we are interested in perfect-recall strategies, we only consider the tree semantics.

D. Model checking $\text{QCTL}_{\text{ii}}^*$

We now prove that the model-checking problem for $\text{QCTL}_{\text{ii}}^*$ under tree semantics is undecidable. This comes as no surprise since, as we will show, $\text{QCTL}_{\text{ii}}^*$ can express the existence of winning strategies in imperfect-information games.

Theorem 3. *The model-checking problem for $\text{QCTL}_{\text{ii}}^*$ under tree-semantics is undecidable.*

Proof. Let MSO_{eq} denote the extension of the logic MSO by a binary predicate symbol eq . Formulas of MSO_{eq} are interpreted on trees, and the semantics of $\text{eq}(x, y)$ is that x and y have the same depth in the tree. There is a translation of MSO -formulas to QCTL^* -formulas that preserves satisfaction [16]. This translation can be extended to map MSO_{eq} -formulas to $\text{QCTL}_{\text{ii}}^*$ -formulas using the formula $\text{level}(\cdot)$ from Example 2 to help capture the equal-length predicate. Our result follows since the MSO_{eq} -theory of the binary tree is undecidable [19]. \square

IV. A DECIDABLE FRAGMENT OF $\text{QCTL}_{\text{ii}}^*$: HIERARCHY ON OBSERVATIONS

The main result of this section is the identification of an important decidable fragment of $\text{QCTL}_{\text{ii}}^*$.

Definition 7 (Hierarchical formulas). *A $\text{QCTL}_{\text{ii}}^*$ formula φ is hierarchical if for all subformulas φ_1, φ_2 of the form $\varphi_1 = \exists^{o_1} p_1. \varphi'_1$ and $\varphi_2 = \exists^{o_2} p_2. \varphi'_2$ where φ_2 is a subformula of φ'_1 , we have $o_1 \subseteq o_2$.*

In other words, a formula is hierarchical if innermost propositional quantifiers observe at least as much as outermost ones. We let $\text{QCTL}_{\text{ii}, \subseteq}^*$ be the set of hierarchical $\text{QCTL}_{\text{ii}}^*$ formulas.

Theorem 4. *Model checking $\text{QCTL}_{\text{ii}, \subseteq}^*$ under tree semantics is non-elementary decidable.*

Since our decision procedure for the hierarchical fragment of $\text{QCTL}_{\text{ii}}^*$ is based on an automata-theoretic approach, we recall some definitions and results for alternating tree automata.

A. Alternating parity tree automata

We briefly recall the notion of alternating (parity) tree automata. Because it is sufficient for our needs and simplifies definitions, we assume that all input trees are complete trees. For a set Z , $\mathbb{B}^+(Z)$ is the set of formulas built from the elements of Z as atomic propositions using the connectives \vee and \wedge , and with $\top, \perp \in \mathbb{B}^+(Z)$. An *alternating tree automaton* (ATA) on (AP, X) -trees is a structure $\mathcal{A} = (Q, \delta, q_i, C)$ where Q is a finite set of states, $q_i \in Q$ is an initial state,

$\delta : Q \times 2^{\text{AP}} \rightarrow \mathbb{B}^+(X \times Q)$ is a transition function, and $C : Q \rightarrow \mathbb{N}$ is a colouring function. To ease reading we shall write atoms in $\mathbb{B}^+(X \times Q)$ between brackets, such as $[x, q]$. A *nondeterministic tree automaton* (NTA) on (AP, X) -trees is an ATA $\mathcal{A} = (Q, \delta, q_i, C)$ such that for every $q \in Q$ and $a \in 2^{\text{AP}}$, $\delta(q, a)$ is written in disjunctive normal form and for every direction $x \in X$ each disjunct contains exactly one element of $\{x\} \times Q$. Acceptance is defined as usual (see, e.g., [33]), and we let $\mathcal{L}(\mathcal{A})$ be the set of trees accepted by \mathcal{A} .

We recall three classic results on tree automata. The first one is that nondeterministic tree automata are closed under projection, and was established by Rabin to deal with second-order monadic quantification:

Theorem 5 (Projection [34]). *Given an NTA \mathcal{N} and an atomic proposition $p \in \text{AP}$, one can build an NTA $\mathcal{N} \downarrow_p$ such that $\mathcal{L}(\mathcal{N} \downarrow_p) = \mathcal{L}(\mathcal{N}) \downarrow_p$.*

Because it will be important to understand the automata construction for our decision procedure in Section IV, we briefly recall that the projected automaton $\mathcal{N} \downarrow_p$ is simply automaton \mathcal{N} with the only difference that when it reads the label of a node, it can choose whether p is there or not: if δ is the transition function of \mathcal{N} , that of $\mathcal{N} \downarrow_p$ is $\delta'(q, a) = \delta(q, a \cup \{p\}) \vee \delta(q, a \setminus \{p\})$, for any state q and label $a \in 2^{\text{AP}}$. Another way of seeing it is that $\mathcal{N} \downarrow_p$ first guesses a p -labelling for the input tree, and then simulates \mathcal{N} on this modified input. To prevent $\mathcal{N} \downarrow_p$ from guessing different labels for a same node in different executions, it is crucial that \mathcal{N} be nondeterministic, which is the reason why we need the next classic result: the crucial simulation theorem, due to Muller and Schupp.

Theorem 6 (Simulation [35]). *Given an ATA \mathcal{A} , one can build an NTA \mathcal{N} such that $\mathcal{L}(\mathcal{N}) = \mathcal{L}(\mathcal{A})$.*

The third result was established by Kupferman and Vardi to deal with imperfect information aspects in distributed synthesis. To state it we need to define a widening operation on trees which simply expands the directions in a tree.

Widening [32]. Let X and Y be two finite sets, let t be an X -tree with root x , and let $y \in Y$. The Y -widening of t with root (x, y) is the $X \times Y$ -tree

$$\tau \uparrow_y^Y := \{u \in (x, y) \cdot (X \times Y)^* \mid u \downarrow_X \in \tau\}.$$

For an (AP, X) -tree $t = (\tau, \ell)$, we let $t \uparrow_y^Y := (\tau \uparrow_y^Y, \ell')$ where $\ell'(u) := \ell(u \downarrow_X)$.

Similarly to the narrowing operation, we extend this definition to tuples of local states by letting, for $J \subseteq I \subseteq [n]$, τ an L_J -tree and $z' \in L_{I \setminus J}$,

$$\tau \uparrow_{z'}^I := \tau \uparrow_{z'}^{L_{I \setminus J}},$$

and similarly for a labelled L_J -tree t ,

$$t \uparrow_{z'}^I := t \uparrow_{z'}^{L_{I \setminus J}}.$$

Recall that because the sets of local states L_i are disjoint, the order of local states in a tuple does not matter and we can identify L_I with $L_J \times L_{I \setminus J}$.

The rough idea of the narrowing operation on ATA is that, if one just observes X , uniform p -labellings on $X \times Y$ -trees can be obtained by choosing the labellings directly on X -trees, and then lifting them to $X \times Y$.

Theorem 7 (Narrowing [32]). *Given an ATA \mathcal{A} on $X \times Y$ -trees one can build an ATA $\mathcal{A} \downarrow_X$ on X -trees such that for all $y \in Y$, $t \in \mathcal{L}(\mathcal{A} \downarrow_X)$ iff $t \uparrow_y^{X \times Y} \in \mathcal{L}(\mathcal{A})$.*

B. Translating $\text{QCTL}_{i,\subseteq}^*$ to ATA

In order to prove Theorem 4 we need some more notations and a technical lemma that contains the automata construction.

For every $\varphi \in \text{QCTL}_{ii}^*$, we let

$$I_\varphi := \bigcap_{o \in \text{Obs}(\varphi)} o \subseteq [n],$$

where $\text{Obs}(\varphi)$ is the set of concrete observations that occur in φ , with the intersection over the empty set defined as $[n]$. We also let $L_\varphi := L_{I_\varphi}$ (recall that for $I \subseteq [n]$, $L_I = \prod_{i \in I} L_i$).

Our construction, that transforms a $\text{QCTL}_{i,\subseteq}^*$ formula φ and a CKS \mathcal{S} into an ATA, builds upon the classic construction from [36], which builds hesitant ATA for CTL^* formulas. Since our aim here is to establish decidability and that the hesitant condition is only used to improve complexity, we do not consider it. However we need to develop an original technique to implement quantifiers with imperfect information thanks to automata narrowing and projection.

The classical approach to model checking via tree automata is to build an automaton that accepts all tree models of the input formula, and check whether it accepts the unfolding of the model [36]. We now explain how we adapt this approach.

Narrowing of non-uniform trees. Quantification on atomic propositions is classically performed by means of automata projection (see Theorem 5). But in order to obtain a labelling that is uniform with regards to the observation of the quantifier, we need to make use of the narrow operation (see Theorem 7). Intuitively, to check that a formula $\exists^o p. \varphi$ holds in a tree t , we would like to work on its narrowing $t' := t \downarrow_o$, guess a labelling for p on this tree thanks to automata projection, thus obtaining a tree t'_p , take its widening $t''_p := t'_p \uparrow^{[n]}$, obtaining a tree with an o -uniform labelling for p , and then check that φ holds on t''_p . The problem here is that, while the narrowing $\tau \downarrow_o$ of an unlabelled tree τ is well defined (see Section III-B), that of a labelled tree $t = (\tau, \ell)$ is undefined: indeed, unless t is o -uniform in every atomic proposition in AP , there is no way to define the labelling of $\tau \downarrow_o$ without losing information. This implies that we cannot feed (a narrowing of) the unfolding of the model to our automata. Still, we need an input tree to be successively labelled and widened to guess uniform labellings.

Splitting quantified from free propositions. To address this problem, we remark that since we are interested in model checking a QCTL_{ii}^* formula φ on a CKS \mathcal{S} , the automaton that we build for φ can depend on \mathcal{S} . It can thus guess paths in \mathcal{S} , and evaluate free occurrences of atomic propositions in \mathcal{S} without reading the input tree. The input tree is thus no longer used to represent the model. However we use it to carry

labellings for quantified propositions $\text{AP}_\exists(\varphi)$: we provide the automaton with an input tree whose labelling is initially empty, and the automaton, through successive narrowing and projection operations, decorates it with uniform labellings for quantified atomic propositions.

We remark that this technique allows one to go beyond **CL** [21]: by separating between quantified atomic propositions (that need to be uniform) and free atomic propositions (that state facts about the model), we manage to remove the restriction present in **CL**, that requires that all facts about the model are known to every strategy/agent (see Proposition 4).

To do this we assume without loss of generality that propositions that are quantified in φ do not appear free in φ , i.e., $\text{AP}_\exists(\varphi) \cap \text{AP}_f(\varphi) = \emptyset$. If necessary, for every $p \in \text{AP}_\exists(\varphi) \cap \text{AP}_f(\varphi)$, we take a fresh atomic proposition p' and replace all quantified occurrences of p in φ with p' . We obtain an equivalent formula φ' on $\text{AP}' := \text{AP} \cup \{p' \mid p' \in \text{AP}_\exists(\varphi) \cap \text{AP}_f(\varphi)\}$ such that $\text{AP}_\exists(\varphi') \cap \text{AP}_f(\varphi') = \emptyset$. Observe also that given a formula φ such that $\text{AP}_\exists(\varphi) \cap \text{AP}_f(\varphi) = \emptyset$, a CKS \mathcal{S} and a state $s \in \mathcal{S}$, the truth value of φ in \mathcal{S}, s does not depend on the labelling of \mathcal{S} for atomic propositions in $\text{AP}_\exists(\varphi)$, which can thus be forgotten.

As a consequence, henceforth we assume that an instance (φ, \mathcal{S}) of the model-checking problem for QCTL_{ii}^* is such that $\text{AP}_\exists(\varphi) \cap \text{AP}_f(\varphi) = \emptyset$, and \mathcal{S} is a CKS over $\text{AP}_f(\varphi)$.

Merging the decorated input tree and the model. To state the correctness of our construction, we will need to merge the labels for quantified propositions, carried by the input tree, with those for free propositions, carried by CKS \mathcal{S} . Because, through successive widenings, the input tree (represented by t in the definition below) will necessarily be a complete tree, its domain will always contain the domain of the unfolding of \mathcal{S} (represented by t' below), hence the following definition.

Definition 8 (Merge). *Let $t = (\tau, \ell)$ be a complete (AP, X) -tree and $t' = (\tau', \ell')$ an (AP', X) -tree, where $\text{AP} \cap \text{AP}' = \emptyset$. We define the merge of t and t' as the $(\text{AP} \cup \text{AP}', X)$ -tree $t \bowtie t' := (\tau \cap \tau' = \tau', \ell'')$, where $\ell''(u) = \ell(u) \cup \ell'(u)$.*

We now describe our automata construction and establish the following lemma, which is our main technical contribution.

Lemma 1 (Translation). *Let (Φ, \mathcal{S}) be an instance of the model-checking problem for $\text{QCTL}_{i,\subseteq}^*$. For every subformula φ of Φ and state s of \mathcal{S} , one can build an ATA \mathcal{A}_s^φ on $(\text{AP}_\exists(\Phi), L_\varphi)$ -trees such that for every $(\text{AP}_\exists(\Phi), L_\varphi)$ -tree t rooted in $s \downarrow_{I_\varphi}$,*

$$t \in \mathcal{L}(\mathcal{A}_s^\varphi) \quad \text{iff} \quad t \uparrow_y^{[n]} \bowtie t_{\mathcal{S}}(s) \models \varphi, \quad \text{where } y = s \downarrow_{[n] \setminus I_\varphi}.$$

For an L_I -tree t , from now on $t \uparrow^{[n]} \bowtie t_{\mathcal{S}}(s)$ stands for $t \uparrow_y^{[n]} \bowtie t_{\mathcal{S}}(s)$, where $y = s \downarrow_{[n] \setminus I}$: the missing local states in the root of t are filled with those from s .

Proof sketch of Lemma 1. Let (Φ, \mathcal{S}) be an instance of the model-checking problem for $\text{QCTL}_{i,\subseteq}^*$. Let $\text{AP}_\exists = \text{AP}_\exists(\Phi)$ and $\text{AP}_f = \text{AP}_f(\Phi)$, and recall that \mathcal{S} is labelled over AP_f . For each state $s \in \mathcal{S}$ and each subformula φ of Φ (note that all

subformulas of Φ are also hierarchical), we define by induction on φ the ATA \mathcal{A}_s^φ on $(\text{AP}_\exists, L_\varphi)$ -trees.

$\varphi = p$:

We let \mathcal{A}_s^p be the ATA over $L_{[n]}$ -trees with one unique state q_ι , with transition function defined as follows:

$$\delta(q_\iota, a) = \begin{cases} \top & \text{if } \begin{array}{l} p \in \text{AP}_f \text{ and } p \in \ell_S(s) \\ \text{or} \\ p \in \text{AP}_\exists \text{ and } p \in a \\ p \in \text{AP}_f \text{ and } p \notin \ell_S(s) \end{array} \\ \perp & \text{if } \begin{array}{l} \text{or} \\ p \in \text{AP}_\exists \text{ and } p \notin a \end{array} \end{cases}$$

$\varphi = \neg\varphi'$:

We obtain \mathcal{A}_s^φ by dualising $\mathcal{A}_s^{\varphi'}$, a classic operation.

$\varphi = \varphi_1 \vee \varphi_2$:

Because $I_\varphi = I_{\varphi_1} \cap I_{\varphi_2}$, and each $\mathcal{A}_s^{\varphi_i}$ works on L_{φ_i} -trees, we first narrow them so that they work on L_φ -trees: for $i \in \{1, 2\}$, we let $\mathcal{A}_i := \mathcal{A}_s^{\varphi_i} \downarrow_{I_\varphi}$. We then build \mathcal{A}_s^φ by taking the disjoint union of \mathcal{A}_1 and \mathcal{A}_2 and adding a new initial state that nondeterministically chooses which of \mathcal{A}_1 or \mathcal{A}_2 to execute on the input tree, so that $\mathcal{L}(\mathcal{A}_s^\varphi) = \mathcal{L}(\mathcal{A}_1) \cup \mathcal{L}(\mathcal{A}_2)$.

$\varphi = \mathbf{E}\psi$:

The aim is to build an automaton \mathcal{A}_s^φ that works on L_φ -trees and that on input t , checks for the existence of a path in $t \uparrow^{[n]} \bowtie t_S(s)$ that satisfies ψ . Observe that a path in $t \uparrow^{[n]} \bowtie t_S(s)$ is a path in $t_S(s)$, augmented with the labelling for atomic propositions in AP_\exists carried by t .

To do so, \mathcal{A}_s^φ guesses a path λ in (S, s) . It remembers the current state in S , which provides the labelling for atomic propositions in AP_f , and while it guesses λ it follows its L_φ -narrowing in its input tree t (which always exists since input to tree automata are complete trees), reading the labels to evaluate propositions in AP_\exists .

Let $\text{max}(\psi) = \{\varphi_1, \dots, \varphi_n\}$ be the set of maximal state sub-formulas of ψ . In a first step we see these maximal state sub-formulas as atomic propositions. The formula ψ can thus be seen as an LTL formula, and we can build a nondeterministic parity word automaton $\mathcal{W}^\psi = (Q^\psi, \Delta^\psi, q_\iota^\psi, C^\psi)$ over alphabet $2^{\text{max}(\psi)}$ that accepts exactly the models of ψ [37].³ We define the ATA \mathcal{A} that, given as input a $(\text{max}(\psi), L_\varphi)$ -tree t , nondeterministically guesses a path λ in $t \uparrow^{[n]} \bowtie t_S(s)$ and simulates \mathcal{W}^ψ on it, assuming that the labels it reads while following $\lambda \downarrow_{I_\varphi}$ in its input correctly represent the truth value of formulas in $\text{max}(\psi)$ along λ . Recall that $S = (S, R, s_\iota, \ell_S)$; we define $\mathcal{A} := (Q, \delta, q_\iota, C)$, where

- $Q = Q^\psi \times S$,
- $q_\iota = (q_\iota^\psi, s)$,
- $C(q^\psi, s') = C^\psi(q^\psi)$, and

³Note that, as usual for nondeterministic word automata, we take the transition function of type $\Delta^\psi : Q^\psi \times 2^{\text{max}(\psi)} \rightarrow 2^{Q^\psi}$.

- for each $(q^\psi, s') \in Q$ and $a \in 2^{\text{max}(\psi)}$,

$$\delta((q^\psi, s'), a) = \bigvee_{q' \in \Delta^\psi(q^\psi, a)} \bigvee_{s'' \in R(s')} [s'' \downarrow_{I_\varphi}, (q', s'')].$$

The intuition is that \mathcal{A} reads the current label, chooses nondeterministically which transition to take in \mathcal{W}^ψ , chooses a next state in S and proceeds in the corresponding direction in X_φ . Thus, \mathcal{A} accepts a $\text{max}(\psi)$ -labelled L_φ -tree t iff there is a path in t that is the L_φ -narrowing of some path in $t_S(s)$, and that satisfies ψ , where maximal state formulas are considered as atomic propositions.

Now from \mathcal{A} we build the automaton \mathcal{A}_s^φ over L_φ -trees labelled with ‘‘real’’ atomic propositions in AP_\exists . In each node it visits, \mathcal{A}_s^φ guesses what should be its labelling over $\text{max}(\psi)$, it simulates \mathcal{A} accordingly, and checks that the guess it made is correct. If the path being guessed in $t_S(s)$ is currently in node u ending with state s' , and \mathcal{A}_s^φ guesses that φ_i holds in u , it checks this guess by starting a simulation of automaton $\mathcal{A}_s^{\varphi_i}$ from node $v = u \downarrow_{I_\varphi}$ in its input t .

For each $s' \in S$ and each $\varphi_i \in \text{max}(\psi)$ we first build $\mathcal{A}_s^{\varphi_i}$, and we let $\mathcal{A}_s^i := \mathcal{A}_s^{\varphi_i} = (Q_{s'}^i, \delta_{s'}^i, q_{s'}^i, C_{s'}^i)$. We also let $\overline{\mathcal{A}}_s^i = (\overline{Q}_{s'}^i, \overline{\delta}_{s'}^i, \overline{q}_{s'}^i, \overline{C}_{s'}^i)$ be its dualisation, and we assume w.l.o.g. that all the state sets are pairwise disjoint. Observe that because each φ_i is a maximal state sub-formula, we have $I_{\varphi_i} = I_\varphi$, so that we do not need to narrow down automata \mathcal{A}_s^i and $\overline{\mathcal{A}}_s^i$. We define the ATA

$$\mathcal{A}_s^\varphi = (Q \cup \bigcup_{i, s'} Q_{s'}^i \cup \overline{Q}_{s'}^i, \delta', q_\iota, C'),$$

where the colours of states are left as they were in their original automaton, and δ is defined as follows. For states in $Q_{s'}^i$ (resp. $\overline{Q}_{s'}^i$), δ agrees with $\delta_{s'}^i$ (resp. $\overline{\delta}_{s'}^i$), and for $(q^\psi, s') \in Q$ and $a \in 2^{\text{AP}_\exists}$ we let $\delta'((q^\psi, s'), a)$ be the disjunction over $a' \in 2^{\text{max}(\psi)}$ of

$$\left(\delta((q^\psi, s'), a') \wedge \bigwedge_{\varphi_i \in a'} \delta_{s'}^i(q_{s'}^i, a) \wedge \bigwedge_{\varphi_i \notin a'} \overline{\delta}_{s'}^i(\overline{q}_{s'}^i, a) \right).$$

$\varphi = \exists^0 p. \varphi'$:

We build automaton $\mathcal{A}_s^{\varphi'}$ that works on $L_{\varphi'}$ -trees; because φ is hierarchical, we have that $\mathbf{o} \subseteq I_{\varphi'}$ and we can narrow down $\mathcal{A}_s^{\varphi'}$ to work on $L_{\mathbf{o}}$ -trees and obtain $\mathcal{A}_1 := \mathcal{A}_s^{\varphi'} \downarrow_{\mathbf{o}}$. By Theorem 6 we can nondeterminise it to get \mathcal{A}_2 , which by Theorem 5 we can project with respect to p , finally obtaining $\mathcal{A}_s^\varphi := \mathcal{A}_2 \downarrow_p$. \square

C. Proof of Theorem 4

We can now prove Theorem 4. Let S be a CKS, $s \in S$, and $\varphi \in \text{QCTL}_{i, \subseteq}^*$. By Lemma 1 one can build an ATA \mathcal{A}_s^φ such that for every labelled L_φ -tree t rooted in $s \downarrow_{I_\varphi}$, it holds that $t \in \mathcal{L}(\mathcal{A}_s^\varphi)$ iff $t \uparrow^{[n]} \bowtie t_S(s) \models \varphi$. Let τ be the full L_φ -tree rooted in $s \downarrow_\varphi$, and let $t = (\tau, \ell_\emptyset)$, where ℓ_\emptyset is the empty labelling. Clearly, $t \uparrow^{[n]} \bowtie t_S(s) = t_S(s)$, and because t is rooted in $s \downarrow_\varphi$, we have $t \in \mathcal{L}(\mathcal{A}_s^\varphi)$ iff $t_S(s) \models \varphi$. It only remains to build a simple deterministic tree automaton

\mathcal{A} over L_φ -trees such that $\mathcal{L}(\mathcal{A}) = \{t\}$, and check for emptiness of the alternating tree automaton $\mathcal{L}(\mathcal{A} \cap \mathcal{A}_s^\varphi)$. Because nondeterminisation makes the size of the automaton gain one exponential for each nested quantifier on propositions, the procedure is nonelementary, and hardness is inherited from the model-checking problem for QCTL [16].

V. MODEL-CHECKING HIERARCHICAL INSTANCES OF SL_{ii}

In this section we establish that the model-checking problem for SL_{ii} restricted to the class of hierarchical instances is decidable (Theorem 2).

We build upon the proof in [24] that establishes the decidability of the model-checking problem for ATL_{sc}^* by reduction to the model-checking problem for QCTL*. The main difference is that we reduce to the model-checking problem for QCTL* instead, using quantifiers on atomic propositions parameterised with observations that reflect the ones used in the SL_{ii} model-checking instance.

Let (Φ, \mathcal{G}) be a hierarchical instance of the SL_{ii} model-checking problem, where $\mathcal{G} = (\text{Ac}, V, E, \ell, v_\iota, \mathcal{O})$. We will first show how to define a CKS $\mathcal{S}_\mathcal{G}$ and a bijection $\rho \mapsto u_\rho$ between the set of finite plays ρ starting in a given position v and the set of nodes in $t_{\mathcal{S}_\mathcal{G}}(s_v)$.

Then, for every subformula φ of Φ and partial function $f : \text{Ag} \rightarrow \text{Var}$, we will define a QCTL* formula $(\varphi)^f$ (that will also depend on \mathcal{G}) such that the following holds:

Proposition 5. *Suppose that $\text{free}(\varphi) \cap \text{Ag} \subseteq \text{dom}(f)$, and $f(a) = x$ implies $\chi(a) = \chi(x)$ for all $a \in \text{dom}(f)$. Then*

$$\mathcal{G}, \chi, \rho \models \varphi \quad \text{if and only if} \quad t_{\mathcal{S}_\mathcal{G}}(s_\rho) \models (\varphi)^f.$$

Applying this to the sentence Φ , any assignment χ , and the empty function \emptyset , we get that

$$\mathcal{G}, \chi, v_\iota \models \Phi \quad \text{if and only if} \quad t_{\mathcal{S}_\mathcal{G}}(s_{v_\iota}) \models (\Phi)^\emptyset.$$

Constructing the CKS $\mathcal{S}_\mathcal{G}$. We will define $\mathcal{S}_\mathcal{G}$ so that (indistinguishable) nodes in its tree-unfolding correspond to (indistinguishable) finite plays in \mathcal{G} . The CKS will make use of atomic propositions $\text{AP}_v := \{p_v \mid v \in V\}$ (that we assume to be disjoint from AP). The idea is that p_v allows the QCTL* formula $(\Phi)^\emptyset$ to refer to the current position v in \mathcal{G} . Later we will see that $(\Phi)^\emptyset$ will also make use of atomic propositions $\text{AP}_c := \{p_c^x \mid c \in \text{Ac} \text{ and } x \in \text{Var}\}$ that we assume, again, are disjoint from $\text{AP} \cup \text{AP}_v$. This allows the formula to use p_c^x to refer to the actions c advised by strategies x .

Suppose $\text{Obs} = \{o_1, \dots, o_n\}$. For $i \in [n]$, define the local states $L_i := \{[v]_{o_i} \mid v \in V\}$ where $[v]_o$ is the equivalence class of v for relation \sim_o . Since we need to know the actual position of the CGS_{ii} to define the dynamics, we also let $L_{n+1} := V$.

Define the CKS $\mathcal{S}_\mathcal{G} := (S, R, s_\iota, \ell')$ where

- $S := \{s_v \mid v \in V\}$,
- $R := \{(s_v, s_{v'}) \mid \exists c \in \text{Ac}^{\text{Ag}} \text{ s.t. } E(v, c) = v'\} \subseteq S^2$,
- $s_\iota := s_{v_\iota}$,
- $\ell'(s_v) := \ell(v) \cup \{p_v\} \subseteq \text{AP} \cup \text{AP}_v$,

and $s_v := ([v]_{o_1}, \dots, [v]_{o_n}, v) \in \prod_{i \in [n+1]} L_i$.

We now show how to connect finite plays in \mathcal{G} with nodes in the tree unfolding of $\mathcal{S}_\mathcal{G}$. For every finite play $\rho = v_0 \dots v_k$, define the node $u_\rho := s_{v_0} \dots s_{v_k}$ in $t_{\mathcal{S}_\mathcal{G}}(s_{v_0})$ (which exists, by definition of $\mathcal{S}_\mathcal{G}$ and of tree unfoldings). Note that the mapping $\rho \mapsto u_\rho$ defines a bijection between the set of finite plays and the set of nodes in $t_{\mathcal{S}_\mathcal{G}}(s_\iota)$.

Constructing the QCTL* formulas $(\varphi)^f$. We now describe how to transform an SL_{ii} formula φ and a partial function $f : \text{Ag} \rightarrow \text{Var}$ into a QCTL* formula $(\varphi)^f$ (that will also depend on \mathcal{G}). Suppose that $\text{Ac} = \{c_1, \dots, c_l\}$, and define $(\varphi)^f$ by induction on φ . We begin with the simple cases: $(p)^f := p$; $(\neg\varphi)^f := \neg(\varphi)^f$; and $(\varphi_1 \vee \varphi_2)^f := (\varphi_1)^f \vee (\varphi_2)^f$.

We continue with the second-order quantifier case:

$$(\langle\langle x \rangle\rangle^o \varphi)^f := \exists \tilde{o} p_{c_1}^x \dots \exists \tilde{o} p_{c_l}^x. \varphi_{\text{str}}(x) \wedge (\varphi)^f$$

where $\tilde{o}_i := \{j \mid \mathcal{O}(o_i) \subseteq \mathcal{O}(o_j)\}$, and

$$\varphi_{\text{str}}(x) := \mathbf{AG} \bigvee_{c \in \text{Ac}} (p_c^x \wedge \bigwedge_{c' \neq c} \neg p_{c'}^x).$$

We describe this formula in words. For each possible action $c \in \text{Ac}$, an existential quantification on the atomic proposition p_c^x “chooses” for each finite play $\rho = v_0 \dots v_k$ of \mathcal{G} (or, equivalently, for each node u_ρ of the tree $t_{\mathcal{S}_\mathcal{G}}(s_{v_0})$) whether strategy x plays action c in ρ or not. Formula $\varphi_{\text{str}}(x)$ ensures that in each finite play, exactly one action is chosen for strategy x , and thus that atomic propositions p_c^x indeed characterise a strategy, call it σ_x .⁴

Moreover, a quantifier with concrete observation \tilde{o}_i receives information corresponding to observation o_i (observe that for all $i \in [n]$, $i \in \tilde{o}_i$) as well as information corresponding to coarser observations. Note that including all coarser observations does not increase the information accessible to the quantifier: indeed, one can show that two nodes are $\{i\}$ -indistinguishable if and only if they are \tilde{o}_i -indistinguishable. However, this definition of \tilde{o}_i allows us to obtain hierarchical formulas. Since quantification on propositions p_c^x is done uniformly with regards to concrete observation \tilde{o}_i , it follows that σ_x is an o_i -strategy.

Here are the remaining cases:

$$\begin{aligned} ((a, x)\varphi)^f &:= (\varphi)^{f[a \mapsto x]} \\ (\mathbf{X}\varphi_1)^f &:= \mathbf{A}(\psi_{\text{out}}(f) \rightarrow \mathbf{X}(\varphi_1)^f) \\ (\varphi_1 \mathbf{U} \varphi_2)^f &:= \mathbf{A}(\psi_{\text{out}}(f) \rightarrow (\varphi_1)^f \mathbf{U} (\varphi_2)^f) \end{aligned}$$

where

$$\psi_{\text{out}}(f) := \mathbf{G} \left(\bigwedge_{v \in V} \bigwedge_{c \in \text{Ac}^{\text{Ag}}} \left(p_v \wedge \bigwedge_{a \in \text{Ag}} p_{c_a}^{f(a)} \rightarrow \mathbf{X} p_{E(v, c)} \right) \right).$$

The formula $\psi_{\text{out}}(f)$ is used to select the unique path assuming that every player, say a , follows the strategy $\sigma_{f(a)}$. This completes the justification of Proposition 5.

⁴More precisely, if $\varphi_{\text{str}}(x)$ holds in node u_ρ , it ensures that propositions from AP_c define a partial strategy, defined on all nodes of the subtree rooted in u_ρ . This is enough because SL_{ii} can only talk about the future: when evaluating a formula in a finite play ρ , the definition of strategies on plays that do not start with ρ is irrelevant.

Preserving hierarchy. To complete the proof we show that $(\Phi)^\emptyset$ is a hierarchical QCTL_{ii}^* formula. This simply follows from the fact that Φ is hierarchical in \mathcal{G} and that for every two observations o_i and o_j in Obs such that $\mathcal{O}(o_i) \subseteq \mathcal{O}(o_j)$, by definition of \tilde{o}_k we have that $\tilde{o}_i \subseteq \tilde{o}_j$.

This completes the proof of Theorem 2.

VI. OUTLOOK

We introduced SL_{ii} , a logic for reasoning about strategic behaviour in multi-player games with imperfect information. The syntax mentions observations, and thus allows one to write formulas that talk about dynamic observations. We isolated the class of hierarchical formula/model pairs (Φ, \mathcal{G}) and proved that one can decide if $\mathcal{G} \models \Phi$. The proof reduces (hierarchical) instances to (hierarchical) formulas of QCTL_{ii}^* , a low-level logic that we introduced, and that serves as a natural bridge between SL_{ii} (that talks about players and strategies) and automata constructions.

We believe that QCTL_{ii}^* is of independent interest and deserves study in its own right. Indeed, it subsumes MSO with equal-level predicate, which is undecidable and of which we know no decidable fragment; yet its syntax and models make it possible to define a natural fragment (the hierarchical fragment) that has a simple definition, a decidable model-checking problem, and is suited for strategic reasoning.

Since one can alternate quantifiers in SL_{ii} , our decidability result goes beyond synthesis. As we showed, we can use it to decide if a game that yields hierarchical observation has a Nash equilibrium. A crude but easy analysis of our main decision procedure shows it is non-elementary.

This naturally leads to a number of avenues for future work: define and study the expressive power and computational complexity of fragments of SL_{ii} [38]; adapt the notion of hierarchical instances to allow for situations in which hierarchies can change infinitely often along a play [15]; and extend the logic to include epistemic operators for individual and common knowledge, as is done in [39], which are important for reasoning about distributed systems [23].

REFERENCES

- [1] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi, "Reasoning about strategies: On the model-checking problem," *TOCL*, vol. 15, no. 4, pp. 34:1–34:47, 2014.
- [2] N. Bulling and W. Jamroga, "Comparing variants of strategic ability: how uncertainty and memory influence general properties of games," in *AAMAS'14*, vol. 28, no. 3, 2014, pp. 474–518.
- [3] F. Laroussinie, N. Markey, and A. Sangnier, "ATLsc with partial observation," in *GandALF'15*, 2015, pp. 43–57.
- [4] C. Dima and F. L. Tiplea, "Model-checking ATL under imperfect information and perfect recall semantics is undecidable," *CoRR*, vol. abs/1102.4225, 2011.
- [5] A. Pnueli and R. Rosner, "Distributed reactive systems are hard to synthesize," in *FOCS'90*, 1990, pp. 746–757.
- [6] G. L. Peterson and J. H. Reif, "Multiple-person alternation," in *SFCS'79*, 1979, pp. 348–363.
- [7] O. Kupferman and M. Vardi, "Synthesizing distributed systems," in *LICS'01*, 2001, pp. 389–398.
- [8] G. Peterson, J. Reif, and S. Azhar, "Decision algorithms for multiplayer noncooperative games of incomplete information," *CAMWA*, vol. 43, no. 1, pp. 179–206, 2002.
- [9] B. Finkbeiner and S. Schewe, "Uniform distributed synthesis," in *LICS'05*, 2005, pp. 321–330.
- [10] S. Pinchinat and S. Riedweg, "A decidable class of problems for control under partial observation," *IPL*, vol. 95, no. 4, pp. 454–460, 2005.
- [11] S. Schewe and B. Finkbeiner, "Distributed synthesis for alternating-time logics," in *ATVA'07*, 2007, pp. 268–283.
- [12] K. Chatterjee and L. Doyen, "Games with a weak adversary," in *ICALP'14*, 2014, pp. 110–121.
- [13] R. Berthon, B. Maubert, and A. Murano, "Decidability results for ATL* with imperfect information and perfect recall," in *AAMAS'17*, 2017.
- [14] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin, "Verification of multi-agent systems with imperfect information and public actions," in *AAMAS'17*, 2017.
- [15] D. Berwanger, A. B. Mathew, and M. Van den Bogaard, "Hierarchical information patterns and distributed strategy synthesis," in *ATVA'15*, 2015, pp. 378–393.
- [16] F. Laroussinie and N. Markey, "Quantified CTL: expressiveness and complexity," *LMCS*, vol. 10, no. 4, 2014.
- [17] G. Peterson, J. Reif, and S. Azhar, "Lower bounds for multiplayer noncooperative games of incomplete information," *CAMWA*, vol. 41, no. 7, pp. 957–992, 2001.
- [18] J. Y. Halpern and M. Y. Vardi, "The complexity of reasoning about knowledge and time. i. lower bounds," *JCSS*, vol. 38, no. 1, pp. 195–237, 1989.
- [19] H. Läuchli and C. Savioz, "Monadic second order definable relations on the binary tree," *JSL*, vol. 52, no. 01, pp. 219–226, 1987.
- [20] P. Gastin, N. Sznajder, and M. Zeitoun, "Distributed synthesis for well-connected architectures," *FMSD*, vol. 34, no. 3, pp. 215–237, 2009.
- [21] B. Finkbeiner and S. Schewe, "Coordination logic," in *CSL'10*, 2010, pp. 305–319.
- [22] R. Alur, T. A. Henzinger, and O. Kupferman, "Alternating-time temporal logic," *JACM*, vol. 49, no. 5, pp. 672–713, 2002.
- [23] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning about knowledge*. MIT press Cambridge, 1995, vol. 4.
- [24] F. Laroussinie and N. Markey, "Augmenting ATL with strategy contexts," *JC*, vol. 245, pp. 98–123, 2015.
- [25] A. Sistla, "Theoretical Issues in the Design and Certification of Distributed Systems." Ph.D. dissertation, Harvard University, Cambridge, MA, USA, 1983.
- [26] E. A. Emerson and A. P. Sistla, "Deciding branching time logic," in *STOC'84*, 1984, pp. 14–24.
- [27] O. Kupferman, "Augmenting branching temporal logics with existential quantification over atomic propositions," in *CAV'95*, 1995, pp. 325–338.
- [28] O. Kupferman, P. Madhusudan, P. S. Thiagarajan, and M. Y. Vardi, "Open systems in reactive environments: Control and synthesis," in *CONCUR'00*, 2000, pp. 92–107.
- [29] T. French, "Decidability of quantified propositional branching time logics," in *AJCAI'01*, 2001, pp. 165–176.
- [30] C. C. Elgot and M. O. Rabin, "Decidability and undecidability of extensions of second (first) order theory of (generalized) successor," *JSL*, vol. 31, no. 2, pp. 169–181, 1966.
- [31] W. Thomas, "Infinite trees and automaton-definable relations over omega-words," *TCS*, vol. 103, no. 1, pp. 143–159, 1992.
- [32] O. Kupferman and M. Y. Vardi, "Church's problem revisited," *BSL*, pp. 245–263, 1999.
- [33] D. E. Muller and P. E. Schupp, "Alternating automata on infinite trees," *TCS*, vol. 54, pp. 267–276, 1987.
- [34] M. O. Rabin, "Decidability of second-order theories and automata on infinite trees," *TAMS*, vol. 141, pp. 1–35, 1969.
- [35] D. E. Muller and P. E. Schupp, "Simulating alternating tree automata by nondeterministic automata: New results and new proofs of the theorems of Rabin, McNaughton and Safra," *TCS*, vol. 141, no. 1&2, pp. 69–107, 1995.
- [36] O. Kupferman, M. Y. Vardi, and P. Wolper, "An automata-theoretic approach to branching-time model checking," *JACM*, vol. 47, no. 2, pp. 312–360, 2000.
- [37] M. Y. Vardi and P. Wolper, "Reasoning about infinite computations," *IC*, vol. 115, no. 1, pp. 1–37, 1994.
- [38] F. Mogavero, A. Murano, G. Perelli, and M. Y. Vardi, "What makes ATL* decidable? a decidable fragment of Strategy Logic," in *CONCUR'12*, 2012, pp. 193–208.
- [39] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano, "MCMAS-SLK: A model checker for the verification of strategy logic specifications," in *CAV'14*, 2014, pp. 525–532.