Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

An Intuitionistic Version of Alternating-Time Temporal Logic

Laura Bozzelli[1] , **Andrea Capone**[1] , **Davide Catta**[2] , **Aniello Murano**[1]
[1]University of Naples Federico II
[2]Université Sorbonne Paris Nord
{laura.bozelli, andrea.capone}@unina.it, catta@lipn.univ-paris13.fr, aniello.murano@unina.it

Abstract

Multi-Agent Systems (MAS) are essential for modelling strategic interactions between multiple agents, often involving partial information. Managing this partial information is crucial for accurate decision-making and strategy optimization. However, partial information combined with perfect recall strategies renders verifying strategic properties undecidable. Intuitionism, a form of partial information which has not yet been explored in the context of MAS, introduces a novel perspective. In this paper, we propose Intuitionistic Alternating Time Temporal Logic (IATL), an extension of ATL that incorporates intuitionistic logic, providing a specialized representation of imperfect information. We define its syntax, semantics, and key structural properties. Additionally, we propose a PTIME-complete algorithm for IATL model checking, supported by benchmarks demonstrating its efficiency.

1 Introduction

Formal methods for strategic reasoning play a fundamental role in Multi-Agent System (MAS) synthesis, specification and verification (Alur, Henzinger, and Kupferman 2002; Kupferman, Vardi, and Wolper 2001; Mogavero et al. 2014; Pnueli and Rosner 1989). This success story originated from the breakthrough idea of using temporal logics for the specification of behaviours of reactive systems (Clarke and Emerson 1981; Emerson and Halpern 1986; Pnueli 1977). Temporal logics are traditionally interpreted over Kripke structures, modelling closed systems, and quantify the computations of the systems universally and existentially. The need to reason about MAS led to the development of formalisms that enable the specification of strategic behaviours of agents (Alur, Henzinger, and Kupferman 2002; Laroussinie and Markey 2015; Mogavero et al. 2012). One of the main developments along this line has been Alternating-time Temporal Logic (ATL) (Alur, Henzinger, and Kupferman 2002), a logical formalism for the specification and verification of open systems involving multiple autonomous agents which allows expressing strategic cooperation and competition among agents in order to achieve certain goals. Imperfect information plays a critical role in the formulation of strategies in MAS (see e.g. (Belardinelli et al. 2020; Berthon et al. 2021; Jamroga and Bulling 2011; Kupferman and Vardi 2000; Reif 1984)), as these concepts are closely intertwined. However, the expressiveness of the imperfect information setting poses a significant challenge: when strategies involve perfect recall (i.e., agents remember all of the past), the verification of strategic properties becomes undecidable (Dima and Tiplea 2011). To address this issue, research has focused on reductions and alternative frameworks that ensure decidability. Notable examples include ATL with memoryless or reduced-memory strategies (Ågotnes et al. 2015; Bulling and Jamroga 2014; Lomuscio and Raimondi 2006; van der Hoek and Wooldridge 2002; Maubert and Murano 2018; Čermák et al. 2018; Jamroga, Malvone, and Murano 2019).

Intuitionistic Logic. In modal and temporal extensions of classical propositional logic, the law of the excluded middle $\varphi \vee \neg\varphi$ is valid. From an information-based perspective, this means that these logics can only represent complete information: every formula φ is either true or false in a model. The assumption of complete information is, however, inadequate when it comes to representing the information available to real-world agents. To represent the development of imperfect or fallible information over time, it turns out that constructive logics are useful as base logics for temporal reasoning (Dummett 2000; Van Benthem 2009; van Dalen and Troelstra 1988). Intuitionistic logic (IL) (Bauer-Mangelberg, van Heijenoort, and Bauer-Mengelberg 1970; Brouwer 1975; Mancosu 1997; Moschovakis 2023) is a subsystem of classical logic which historically arose out of intuitionism school developed in the early 1900s whose main intent was to formulate a more constructive foundation for mathematics. In this setting, the notion of truth for a formula is procedural and depends on the ability to know or prove it. As a result, in IL, the law of the excluded middle $\varphi \vee \neg\varphi$ does not hold in general, i.e., it is not always possible to have knowledge (i.e., prove or verify) φ or its negation. There have been several successful attempts to create semantics for IL such as Beth's tableaux (Kleene 1957), topological and algebraic models (Rasiowa 1963), and Kripke models (Kripke 1965). The best known semantics is based on Kripke models (Kripke 1965) where the accessibility relation is a partial order over the set of states or worlds which models knowledge or information accumulation. Intuitively, a model describes a pro-

cess of investigation where the agents learn progressively and procedurally by moving from less informative states to more informative ones. Thus, the truth of a formula at a state s depends upon the states s' which are reachable from s (in epistemic logic terminology, the states s' represent the information set associated with s). Intuitionistic extensions of modal logics have been explored in (Fischer Servi 1977; Plotkin and Stirling 1986; Simpson 1994), where semantics is based on birelational Kripke models with two accessibility relations: the intuitionistic information partial order and the modal relation. Moreover, several studies have explored the development of intuitionistic versions of temporal frameworks such as Linear Temporal Logic (LTL) (Balbiani and Diéguez 2016; Balbiani et al. 2019) and Computation Tree Logic (CTL) (Catta, Malvone, and Murano 2023). Additionally, in the area of non-monotonic reasoning, IL has played an important role within the well-known Answering Set Programming paradigm (ASP) (Brewka, Eiter, and Truszczynski 2011) leading to temporal extensions of ASP (Cabalar and Vega 2007; Bozzelli and Pearce 2015) that are supported by intuitionistic temporal logics like the temporal logic of here and there (Balbiani and Diéguez 2016). These contributions highlight a growing interest in integrating IL into temporal contexts, suggesting that a similar approach for temporal strategic reasoning could provide valuable insights for modelling and reasoning in multi-agent systems. To the best of our knowledge, the capabilities of IL for strategic reasoning and the related decidability issues have never been thoroughly explored.

Contribution. In this paper, we introduce Intuitionistic ATL (IATL), an extension of ATL that incorporates propositional IL to model a specific form of imperfect information, providing a computationally acceptable approach for representing and reasoning about such information in multi-agent systems. One of the core features of IATL lies in its ability to model and reason about information refinement, the process by which agents transition from states of imperfect knowledge to states of more complete understanding. This feature is particularly relevant in dynamic and uncertain environments, where decisions and strategies depend on the gradual accumulation of information over time. By means of a partial order relation on the game structure which captures the incremental nature of knowledge acquisition, IATL allows for a dynamic representation of evolving knowledge, where agents can gradually discover new facts over time, unlike classical approaches that assume static or complete information. This framework offers a more flexible and realistic way to handle incomplete knowledge, where agents iteratively refine their understanding through observation, analysis, and strategic interaction. We formally define the syntax and semantics of IATL, and furthermore, we provide key properties that underpin the structure of the logic, ensuring that strategies and decision-making processes are not only well-defined but also computationally tractable. To this end, we present an algorithm for model checking IATL and prove its PTIME-completeness, making it computationally equivalent to ATL. We support our theoretical contributions with a set of benchmarks, evaluating the efficiency of our approach in the MAS scenario of a well-known epistemic puzzle, the muddy children one (Fagin et al. 1995), highlighting the practicality and effectiveness of IATL in handling imperfect information in a computationally feasible manner.

Comparison with imperfect information ATL. IATL exploits intuitionistic truth to represent states where a fact may be unknown, and yet sufficient information may still be available to ensure a winning strategy. In particular, unlike imperfect information ATL, the law of the excluded middle $\varphi \vee \neg\varphi$ does not hold in IATL. On other hand, unlike IATL, model checking imperfect information ATL is undecidable (Dima and Tiplea 2011). Thus, the two frameworks are expressively incomparable. Note that in imperfect information ATL, information which is not available to the agents can be explicitly modeled, which makes this framework, unlike IATL, suitable for information-flow security analysis.

2 Preliminaries

We fix a finite non-empty set AP of atomic propositions. For a word (or sequence) w over some alphabet, $|w|$ denotes the length of w (we set $|w| = \infty$ if w is infinite) and for each $0 \leq i < |w|$, $w(i)$ is the $(i+1)^{th}$ letter of w.

2.1 Intuitionistic Propositional Logic

We first recall *Intuitionistic Propositional Logic* (IPL for short) and its standard Kripke semantics. In IPL, the truth of a formula φ is understood as φ is provable. The set of IPL formulas φ over AP is inductively defined as follows:

$$\varphi ::= \bot \mid p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \rightarrow \varphi$$

where $\bot$ is the falsehood symbol and $p \in \mathsf{AP}$. Negation of φ is defined as $\neg\varphi \stackrel{\text{def}}{=} \varphi \rightarrow \bot$. In the Kripke semantics, IPL formulas are interpreted over *Intuitionistic Kripke structures* (IKS) which are tuples $\mathcal{K} = \langle S, S_I, \preceq, \mathrm{V} \rangle$, where S is a set of states or worlds, $S_I \subseteq S$ is the set of initial states, $\preceq$ is a partial order over S, and $\mathrm{V} : S \mapsto 2^{\mathsf{AP}}$ is a *propositional valuation* that assigns to each state s the set of propositions holding at s. The valuation V satisfies the *monotonicity condition*, that is: for all states $s, t \in S$, if $s \preceq t$ then $\mathrm{V}(s) \subseteq \mathrm{V}(t)$. Intuitively, states represent partial information and $s \preceq s'$ means that information increases in moving from s to s'. Regarding the meaning of atomic propositions, they represent assertions or facts, as in the classical setting. However, in the intuitionistic framework, the truth value of a proposition p at a state s may be undetermined. Specifically, if $p \in \mathrm{V}(s)$, then the truth value of p at s is true. If $p \notin \mathrm{V}(s)$, then the truth value of p at s is not necessarily false. According to the intuitionistic interpretation of negation, p is false at s iff for every refinement s' of s (that is, $s \preceq s'$), $p \notin \mathrm{V}(s)$.

Semantics of IPL. For a state s of $\mathcal{K}$ and a formula φ, the satisfaction relation $\mathcal{K}, s \models \varphi$ is inductively defined as follows (we omit the semantics of $\vee$ and $\wedge$ which is classical):

$$\begin{array}{ll} \mathcal{K}, s \not\models \bot & \\ \mathcal{K}, s \models p & \Leftrightarrow p \in \mathrm{V}(s) \\ \mathcal{K}, s \models \varphi_1 \rightarrow \varphi_2 & \Leftrightarrow \text{for all states } t \in S \text{ such that } s \preceq t: \\ & \quad \mathcal{K}, t \models \varphi_1 \text{ implies } \mathcal{K}, t \models \varphi_2 \end{array}$$

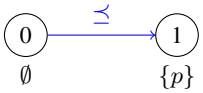

Figure 1: Counter model $\mathcal{K}$ for the excluded middle principle

Note that $\varphi_1 \to \varphi_2$ is checked at all the states greater or equal to the current state s. Moreover, $\mathcal{K}, s \models \neg\varphi$ if and only if for all states t such that $s \preceq t$, $\mathcal{K}, t \not\models \varphi$. Intuitionistic semantics has the feature that for any formula φ and states $s \preceq t$ of a Kripke model, if $\mathcal{K}, s \models \varphi$, then $\mathcal{K}, t \models \varphi$ holds as well; that is, *truth is monotone* (with respect to $\preceq$).
Due to intuitionistic semantics of implication, $\mathcal{K}, s \models \neg\varphi$ implies $\mathcal{K}, s \not\models \varphi$ but not the opposite. In other words, $\mathcal{K}, s \not\models \varphi$ just means that φ is not provable in s, but this does not imply that $\neg\varphi$ is provable in s. Indeed, IPL calculus has one axiom less than classical one, namely it does not contain the law of the excluded middle: $\varphi \vee \neg\varphi$. Figure 1 illustrates a *counterexample model* $\mathcal{K}$ for the formula $p \vee \neg p$. The formula is not satisfied at state 0. Indeed, $\mathcal{K}, 0 \not\models p$ because $p \notin \mathrm{V}(0)$ and $\mathcal{K}, 0 \not\models \neg p$ because $\mathcal{K}, 1 \models p$.

A formula φ is *satisfiable* (*valid*) if $\mathcal{K}, s \models \varphi$ for some (all) IKS $\mathcal{K}$ and state s. Due negation semantics, validity of a formula φ does not correspond to unsatisfiability of $\neg\varphi$. In fact, while the set of IPL satisfiable formulas equals the set SAT of classically satisfiable formulas, this correspondence does not extend to validity. As a matter of fact, checking validity in IPL is PSPACE-complete (Svejdar 2003).

2.2 Concurrent Game Structures

Concurrent Game Structures (CGS) (Alur, Henzinger, and Kupferman 2002) extend Kripke structures to a setting involving multiple agents. They can be viewed as multi-player games in which players perform concurrent actions, chosen strategically as a function of the history of the game.

Let Ag be a finite nonempty set of agents, and Act be a finite nonempty set of actions that can be made by agents. For a set $A \subseteq \mathsf{Ag}$ of agents, an *A-decision* d_A is an element in Act^A assigning to each agent $a \in A$ an action $d_A(a)$. Let $\mathsf{Dec} = \mathsf{Act}^{\mathsf{Ag}}$ be the set of *full decisions* of all the agents in Ag. A full decision d *extends* an A-decision d_A if $d(a) = d_A(a)$ for each agent $a \in A$.

Definition 1. *A concurrent game frame (CGF for short) $\mathcal{F}$ (over Ag and Act) is a tuple $\mathcal{F} = \langle S, S_I, \delta \rangle$, where S is a nonempty set of states or worlds, $S_I \subseteq S$ is the set of initial states, and $\delta : S \times \mathsf{Dec} \mapsto S \cup \{\dashv\}$ is a transition function that maps a state and a full decision either to a state or to the special symbol $\dashv$ ($\dashv$ is for 'undefined') such that for all states s, there exists $d \in \mathsf{Dec}$ so that $\delta(s, d) \neq \dashv$* (seriality).

A CGS $\mathcal{G}$ (over AP, Ag, and Act) is a tuple $\mathcal{G} = \langle S, S_I, \delta, V \rangle$ consisting of a CGF equipped with a propositional valuation $V : S \mapsto 2^{\mathsf{AP}}$.

Remark. In modelling *independent* agents, usually one assumes that at each state s, each agent a has a set $\mathsf{Act}_{a,s} \subseteq \mathsf{Act}$ of actions which are enabled at state s. This is reflected in the transition function δ by requiring that the set of full decisions d such that $\delta(s, d) \neq \dashv$ corresponds to $(\mathsf{Act}_{a,s})_{a \in \mathsf{Ag}}$.

Given a set $A \subseteq \mathsf{Ag}$ of agents, an A-decision d_A, and a state s, *d_A is available at s* if there is a full decision d which extends d_A such that $\delta(s, d) \neq \dashv$. Given two states s and t, t is a *successor* of s if $t = \delta(s, d)$ for some full decision. We say that *t is a successor of s consistent with an A-decision d_A* if $t = \delta(s, d)$ for some full decision d which extends d_A.

Strategies. Given a CGS, we consider *perfect recall* strategies where an agent takes a decision based on all available information up to the current round. A *play* is an infinite sequence of states $s_1 s_2 \ldots$ such that for all $i \geq 1$, s_{i+1} is a *successor* of s_i. An *history* ρ is a nonempty prefix of some play. Given a set $A \subseteq \mathsf{Ag}$ of agents, a *strategy for A* is a mapping f_A assigning to each history ρ (i.e., the history the agents saw so far) an A-decision available at the last state of ρ, denoted *last*(ρ). The *outcome* function $\mathsf{out}(s, f_A)$ for a state s and the strategy f_A returns the set of all the plays starting at s that can occur when agents A execute f_A from s on (with the opponent agents picking their actions arbitrarily). Formally, $\mathsf{out}(s, f_A)$ is the set of plays $\pi = s_1 s_2 \ldots$ such that $s_1 = s$ and for all $i \geq 1$, there is a full decision d that extends the A-decision $f_A(s_1 \ldots s_i)$ so that $s_{i+1} = \delta(s_i, d)$. We also consider *memoryless* A-strategies f_A where the A-decisions depend only on the state of the current round, that is $f_A(\rho) = f_A(\rho')$ for all histories ρ and ρ' such that *last*$(\rho) =$ *last*(ρ'). A memoryless A-strategy is represented as a map assigning to each state s an A-decision available at s.

3 Intuitionistic ATL

We now introduce Intuitionistic ATL (IATL) and provide a characterization of the birelational frames that ensure monotonicity of the satisfaction relation. We illustrate the novel framework by modelling the epistemic puzzle of muddy children (Fagin et al. 1995). Moreover, in Subsection 3.2 we address expressiveness issues, and give denotational and memoryless characterizations of the IATL semantics.

The syntax of IATL coincides with that of standard ATL but with emphasis on the use of the implication connective. Formally, for the given set AP of atomic propositions and set Ag of agents, the syntax of *state formulas* φ and *path formulas* ψ over AP and Ag is inductively defined as follows:

$$\varphi ::= \bot \mid p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \to \varphi \mid \langle\langle A \rangle\rangle \psi \mid [\![A]\!] \psi$$
$$\psi ::= \bigcirc \varphi \mid \varphi \mathsf{U} \varphi \mid \varphi \mathsf{R} \varphi$$

where $p \in \mathsf{AP}$, $A \subseteq \mathsf{Ag}$ (subsets of Ag will be called coalitions), $\bigcirc$, U, and R are the standard "next", "until", and "release" temporal modalities, respectively, and $\langle\langle A \rangle\rangle$ and $[\![A]\!]$ are the existential and universal strategic quantifiers parametrized by a set A of agents. Formula $\langle\langle A \rangle\rangle \varphi$ expresses the property that the coalition A has a collective strategy to enforce property φ, while formula $[\![A]\!] \varphi$ requires that no strategy of A can prevent property φ. Formulas of IATL are all and only the state formulas. The size $|\varphi|$ of a formula φ is the number of distinct (path or state) subformulas of φ.

We also consider the existential fragment IATL$_\exists$ and the universal fragment IATL$_\forall$ of IATL, obtained by disallowing the universal strategic quantifiers for IATL$_\exists$ and the existential strategic quantifiers for IATL$_\forall$, respectively. Note that

under the semantics of ATL, the two fragments are equivalent since $[\![A]\!]$ can be expressed in terms of $\langle\langle A\rangle\rangle$, and vice versa. However, as we will see in Subsection 3.2, this does not hold for the intuitionistic semantics of IATL. Similarly, the dual modalities U and R are not interdefinable in IATL.

IATL formulas are interpreted over *birelational* CGS which extend CGS by a partial order over the set of states.

Definition 2. *A* Birelational *Concurrent Game Frame (BCGF) is a tuple $\mathcal{F} = \langle S, S_I, \delta, \preceq\rangle$ where $\langle S, S_I, \delta\rangle$ is a concurrent game frame and $\preceq$ is a partial order over S.*

We say that $\mathcal{F}$ is well-behaved *if the partial order $\preceq$ satisfies the following two additional conditions for each coalition A, where for a state s, $\mathsf{Dec}(s, A)$ is the set of A-decisions available at state s of $\mathcal{F}$:*

- (C_1) *for all states s and s': if $s \preceq s'$, then for each A-decision $d_A \in \mathsf{Dec}(s, A)$, there is an A-decision $d'_A \in \mathsf{Dec}(s', A)$ such that, contravariantly, for each full decision $d' \in \mathsf{Dec}(s', \mathsf{Ag})$ which extends d'_A, there is a full decision $d \in \mathsf{Dec}(s, \mathsf{Ag})$ which extends d_A so that $\delta(s, d) \preceq \delta(s', d')$;*
- (C_2) *the condition obtained from condition C_1 by replacing $s \preceq s'$ with $s' \preceq s$, and $\delta(s, d) \preceq \delta(s', d')$ with $\delta(s', d') \preceq \delta(s, d)$.*

A Birelational *CGS (BCGS for short) is a tuple $\mathcal{G} = \langle S, S_I, \delta, \preceq, V\rangle$ consisting of a BCGF $\langle S, S_I, \delta, \preceq\rangle$ and a valuation $V : S \mapsto 2^{AP}$ satisfying the* monotonicity condition*, that is: for all $s, t \in S$, if $s \preceq t$ then $V(s) \subseteq V(t)$.*

A BCGS is *well-behaved* if the underlying frame is well-behaved. Conditions C_1 and C_2 formalize the interplay between the partial order $\preceq$ (which models information accumulation of the agents) and the time passage induced by agent actions, regulating the dynamic of information change which is induced by time passage. Intuitively, in a well-behaved BCGS $\mathcal{G}$, condition C_1 ensures that the partial order $\preceq$ behaves like an *alternating-time simulation* over $\mathcal{G}$ (Alur et al. 1998), and similarly for condition C_2 over the inverse $\succeq$ of $\preceq$: if a state s' is at least as informative as s, then for each A-decision d_A available at s, there is an A-decision available at s' that produces a set of outcomes at least as informative as those produced at s (and vice versa). As we will see, the well-behaved property is the minimal requirement for ensuring *truth monotonicity* of IATL, a crucial semantics constraint in the intuitionistic setting: in moving from a state to a more informative one, the truth of the statements is preserved. It is worth noting that by (Alur et al. 1998), checking the well-behaved requirement for *finite* BCGS can be done in polynomial time. Moreover, the BCGS models in which each state has exactly one successor correspond to the linear frames introduced in (Balbiani et al. 2019) for intuitionistic LTL (ILTL). In particular, for the subclass of linear frames, the additional conditions C_1 and C_2 coincide and are equivalent to the property of *forward confluence* as defined in (Balbiani et al. 2019) (*expanding frames*). It is worth emphasizing that the variant of ILTL we consider is the one interpreted over expanding frames (Balbiani et al. 2019) as our framework only requires forward confluence.

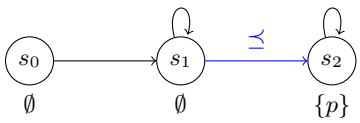

Figure 2: Well-behaved BCGS satisfying at state s_0 the formula $\neg\langle\langle\mathsf{Ag}\rangle\rangle \bigcirc p \wedge \neg\langle\langle\mathsf{Ag}\rangle\rangle \bigcirc \neg p$. For simplicity, we report only the minimal $\preceq$-edges which generate the partial order $\preceq$.

Semantics of IATL. The semantics of IATL differs from that of standard ATL by the intuitionistic interpretation of implication $\rightarrow$. Let $\mathcal{G}$ be a BCGS, $s \in S$, and π a play of $\mathcal{G}$. For a path formula ψ and a state formula φ, the satisfaction relations $\mathcal{G}, s \models \varphi$ and $\mathcal{G}, \pi \models \psi$ are inductively defined as follows, where we omit the base cases and the cases for $\vee$, $\wedge$, and the temporal modalities which are standard:

$$\begin{array}{ll} \mathcal{G}, s \models \varphi_1 \rightarrow \varphi_2 & \Leftrightarrow \text{for all states } t \in S \text{ such that } s \preceq t: \\ & \quad \mathcal{G}, t \models \varphi_1 \text{ implies } \mathcal{G}, t \models \varphi_2 \\ \mathcal{G}, s \models \langle\langle A\rangle\rangle \psi & \Leftrightarrow \text{there is an } A\text{-strategy } f_A \text{ such that} \\ & \quad \text{for each play } \pi \in \mathsf{out}(s, f_A), \mathcal{G}, \pi \models \psi \\ \mathcal{G}, s \models [\![A]\!] \psi & \Leftrightarrow \text{for each } A\text{-strategy } f_A \text{ there is a play} \\ & \quad \pi \in \mathsf{out}(s, f_A) \text{ such that } \mathcal{G}, \pi \models \psi \end{array}$$

We also consider the *memoryless* semantics $\models_m$ where strategic modalities only quantify over memoryless strategies. Moreover, we write $\mathcal{G}, s \models_{\text{ATL}} \varphi$ to mean that the CGS embedded into $\mathcal{G}$ satisfies φ at state s under the ATL semantics. We observe that unlike IPL versus propositional logic, IATL satisfiability *does not correspond* to ATL satisfiability.

Proposition 1. *Every IATL-formula φ which is satisfiable under the ATL-semantics is also IATL-satisfiable. However, there are IATL-satisfiable formulas which are unsatisfiable under the standard ATL semantics.*

Proof. Since CGS can be seen as well-behaved BCGS whose partial order $\preceq$ is the identity, the first part of Proposition 1 trivially holds. For the second part, we consider the formula $\varphi \stackrel{\text{def}}{=} \neg\langle\langle\mathsf{Ag}\rangle\rangle \bigcirc p \wedge \neg\langle\langle\mathsf{Ag}\rangle\rangle \bigcirc \neg p$. This formula is unsatisfiable under the ATL-semantics. On the other hand, the state s_0 of the well-behaved BCGS $\mathcal{G}$ illustrated in Figure 2 satisfies φ under the IATL-semantics (note that $\mathcal{G}, s_1 \not\models p$ and $\mathcal{G}, s_1 \not\models \neg p$). □

Generalization to multi-relational CGS. In this paper, for the ease of presentation, we consider a single partial order in BCGS. The framework can be extended by prescribing a partial order $\preceq_a$ for each agent a and by considering the versions $\rightarrow_A$ of the implication $\rightarrow$ parameterized by a set A of agents (the semantics of $\rightarrow_A$ is like the $\rightarrow$-semantics but we replace $\preceq$ with $\bigcap_{a \in A} \preceq_a$). All the results of this paper can be easily adapted to this extension (in particular, for Conditions C_1 and C_2, we replace $\preceq$ with $\bigcap_{a \in \mathsf{Ag}} \preceq_a$).

Monotonicity of truth of formulas. A BCGF $\mathcal{F} = \langle S, S_I, \delta, \preceq\rangle$ is *IATL-monotonic* if for all propositional valuations $\mathrm{V} : S \mapsto 2^{\mathsf{AP}}$ such that $(\mathcal{F}, \mathrm{V})$ is a BCGS (i.e., V satisfies the monotonicity condition), the following holds

for all states s, t and IATL formulas φ: if $\mathcal{G}, s \models \varphi$ and $s \preceq t$ then $\mathcal{G}, t \models \varphi$. The notions of IATL$_\exists$-monotonicity and IATL$_\forall$-monotonicity are similar. In the following, we show that conditions C_1 and C_2 in Definition 2 are necessary and sufficient conditions on the partial order $\preceq$ of a BCGF for ensuring IATL-monotonicity.

Let $\mathcal{F} = \langle S, S_I, \delta, \preceq \rangle$ be a BCGF. For two plays π and π' of $\mathcal{F}$, we write $\pi \preceq \pi'$ to mean that $\pi(i) \preceq \pi'(i)$ for all $i \geq 0$. Similarly for two histories ρ and ρ' of $\mathcal{G}$, we write $\rho \preceq \rho'$ to mean that $|\rho| = |\rho'|$ and $\rho(i) \leq \rho'(i)$ for all $0 \leq i < |\rho|$. Moreover, for a coalition A, an A-strategy f_A, a state s, and $k \geq 1$, $\mathsf{out}_k(s, f_A)$ is the set of histories of length k starting at state s which are consistent with strategy f_A.

Lemma 1. *Let $\mathcal{F}$ be a BCGF satisfying condition C_1, A a coalition, f_A an A-strategy, and $s \preceq s'$. Then, there exists an A-strategy f'_A such that for each $\pi' \in \mathsf{out}(s', f'_A)$, there is $\pi \in \mathsf{out}(s, f_A)$ so that $\pi \preceq \pi'$.*

Proof. Let $\mathcal{F} = \langle S, S_I, \delta, \preceq \rangle$. Since $\mathcal{F}$ satisfies condition C_1 in Definition 2, there must be a mapping Choice assigning to each triple (t, d_A, t') where $t, t' \in S$, $t \preceq t'$, and $d_A \in \mathsf{Dec}(t, A)$, an A-decision $\mathrm{d}'_A \in Dec(t', A)$ such that:

- for each full decision $\mathrm{d}' \in \mathsf{Dec}(s', \mathsf{Ag})$ which extends d'_A, there is a full decision $d \in \mathsf{Dec}(s, \mathsf{Ag})$ which extends d_A so that $\delta(s, d) \preceq \delta(s', \mathrm{d}')$.

By exploiting the mapping Choice, we prove the following.

Claim. There exist an A-strategy f'_A and for each $k \geq 1$, a mapping $H_k : \mathsf{out}_k(s', f'_A) \mapsto \mathsf{out}_k(s, f_A)$ such that:

- for each $\rho' \in \mathsf{out}_k(s', f'_A)$, $H_k(\rho') \preceq \rho'$;
- if $k > 1$: for each $\rho' \cdot t' \in \mathsf{out}_k(s', f'_A)$, there is $t \in S$ such that $H_k(\rho' \cdot t') = H_{k-1}(\rho') \cdot t$.

First, we show that Lemma 1 follows from the claim. Let f'_A and H_k, for $k \geq 1$, as in the claim. Pick an arbitrary play $\pi' \in \mathsf{out}(s', f'_A)$. We need to show that there is $\pi \in \mathsf{out}(s, f_A)$ such that $\pi \preceq \pi'$. For each $k \geq 1$, let π'_k be the prefix of the play π' of length k, and $\pi_k \stackrel{\text{def}}{=} H_k(\pi'_k)$. By the claim $\pi_k \in \mathsf{out}_k(s, f_A)$, $\pi_k \preceq \pi'_k$, and π_k is a prefix of π_{k+1}. Define $\pi \stackrel{\text{def}}{=} last(\pi_1)last(\pi_2)\ldots$. Evidently, $\pi \preceq \pi'$ and $\pi \in \mathsf{out}(s, f_A)$. Hence, the result follows.

We now prove the claim. We define the strategy f'_A by induction on the length $h \geq 1$ of the histories of $\mathcal{F}$. Since for any A-strategy f, the set of finite outcomes of f of length h is independent of the values assumed by the strategy over the histories of length equal or greater than h, we can assume that at the step h, the set $\mathsf{out}_h(s', f'_A)$ is already given, and there are functions $H_1, \ldots, H_h$ satisfying the claim. Note that for $h = 1$, since $\mathsf{out}_1(s', f') = \{s'\}$ for any A-strategy f', H_1 is uniquely determined by setting $H_1(s') = s$. Moreover, by hypothesis $s \preceq s'$. Hence, H_1 satisfies the claim.

Let ρ' be an history of length h. If $\rho' \notin \mathsf{out}_h(s', f'_A)$, we set $f'_A(\rho') \stackrel{\text{def}}{=} f_A(\rho')$. Otherwise, let $\rho = H_h(\rho')$. We have that $\rho \in \mathsf{out}_h(s, f_A)$ and $\rho \preceq \rho'$. Thus, we set:

$$f'_A(\rho') \stackrel{\text{def}}{=} \mathsf{Choice}(last(\rho), f_A(\rho), last(\rho')).$$

It remains to define the function H_{h+1}. Pick $\rho' \cdot t' \in \mathsf{out}_{h+1}(s', f'_A)$. By definition of the mapping Choice, there is $t \in S$ such that $t \preceq t'$ and $H_h(\rho') \cdot t \in \mathsf{out}_{h+1}(s, f_A)$. We define $H_{h+1}(\rho' \cdot t') \stackrel{\text{def}}{=} H_h(\rho') \cdot t$, and the result follows. □

For BCGF satisfying condition C_2 in Definition 2, we obtain a result similar to Lemma 1.

Lemma 2. *Let $\mathcal{F}$ be a BCGF satisfying condition C_2, A a coalition, f_A an A-strategy, and $s' \preceq s$. Then, there exists an A-strategy f'_A such that for each $\pi' \in \mathsf{out}(s', f'_A)$, there is $\pi \in \mathsf{out}(s, f_A)$ so that $\pi' \preceq \pi$.*

By exploiting Lemmata 1 and 2, we show that conditions C_1 and C_2 characterize the IATL-monotonic BCGF.

Theorem 1. *Let $\mathcal{F}$ be a BCGF. Then:*
1. *$\mathcal{F}$ satisfies condition C_1* iff *$\mathcal{F}$ is IATL$_\exists$-monotonic.*
2. *$\mathcal{F}$ satisfies condition C_2* iff *$\mathcal{F}$ is IATL$_\forall$-monotonic.*
3. *$\mathcal{F}$ is well-behaved* iff *$\mathcal{F}$ is IATL-monotonic.*

Proof. We prove Property (1). The proof of Property (2) is similar, and Property (3) easily follows from Properties (1) and (2). For the *left to right implication* of Property (1), assume that $\mathcal{F}$ satisfies condition C_1. Let $\mathcal{G} = \langle S, S_I, \delta, \preceq, \mathrm{V} \rangle$ be a BCGS whose underlying frame $\langle S, S_I, \delta, \preceq \rangle$ is $\mathcal{F}$, φ an IATL$_\exists$ formula, $s \preceq s'$, and $\mathcal{G}, s \models \varphi$. We prove by structural induction on φ that $\mathcal{G}, s' \models \varphi$. The cases where φ is an atomic proposition directly follows from the monotonicity condition on the propositional valuation V, while the cases where the root operator of φ is a Boolean connective directly follow from the induction hypothesis. Thus since φ is an IATL$_\exists$ formula, it remains to consider the case where φ is of the form $\langle\langle A \rangle\rangle \psi$ for some path formula ψ. Since $\mathcal{G}, s \models \langle\langle A \rangle\rangle \psi$, there exists an A-strategy f_A such that for each $\pi \in \mathsf{out}(s, f_A)$, $\mathcal{G}, \pi \models \psi$. Being $s \preceq s'$, by Lemma 1, there is an A-strategy f'_A such that for each $\pi' \in \mathsf{out}(s', f'_A)$, there is $\pi \in \mathsf{out}(s, f_A)$ so that $\pi(j) \preceq \pi'(j)$ for all $j \geq 0$. Assume that ψ is of the form $\varphi_1 \mathsf{U} \varphi_2$ (the cases where ψ is of the form $\bigcirc \varphi_1$ or $\varphi_1 \mathsf{R} \varphi_2$ are similar). Pick an arbitrary $\pi' \in \mathsf{out}(s', f'_A)$. We show that $\mathcal{G}, \pi' \models \varphi_1 \mathsf{U} \varphi_2$. Hence, the result follows. We know that there is $\pi \in \mathsf{out}(s, f_A)$ such that $\pi(j) \preceq \pi'(j)$ for all $j \geq 0$. Since $\mathcal{G}, \pi \models \varphi_1 \mathsf{U} \varphi_2$, there is $i \geq 0$ such that $\mathcal{G}, \pi(i) \models \varphi_2$ and $\mathcal{G}, \pi(k) \models \varphi_1$ for all $0 \leq k < i$. Thus, by the induction hypothesis, $\mathcal{G}, \pi'(i) \models \varphi_2$ and $\mathcal{G}, \pi'(k) \models \varphi_1$ for all $0 \leq k < i$. This means that $\mathcal{G}, \pi' \models \varphi_1 \mathsf{U} \varphi_2$ and we are done.

The *right to left implication* of Property (1) is proved by contrapositive. Assume that $\mathcal{F} = \langle S, S_I, \delta, \preceq \rangle$ does not satisfy condition C_1. We prove that $\mathcal{F}$ is not IATL$_\exists$-monotonic. By hypothesis there must be a coalition A, two states s, s' with $s \preceq s'$, and an A-decision $d_A \in \mathsf{Dec}(s, A)$ such that:

($\star$) let $\{t_1^{\text{bad}}, \ldots, t_n^{\text{bad}}\}$ be the set of successors of s consistent with the A-decision d_A. Then, for all A-decisions $\mathrm{d}'_A \in \mathsf{Dec}(s', A)$, there is a full decision $\mathrm{d}' \in \mathsf{Dec}(s', \mathsf{Ag})$ which extends d'_A such that $t_i^{\text{bad}} \not\preceq \delta(s', \mathrm{d}')$ for each $1 \leq i \leq n$.

Let $p \in \mathsf{AP}$ and $\mathrm{V} : S \mapsto 2^{\mathsf{AP}}$ be the valuation such that for all $t \in S$: $\mathrm{V}(t) = \{p\}$ if $t_i^{\text{bad}} \preceq t$ for some $1 \leq i \leq n$, and $\mathrm{V}(t) = \emptyset$ otherwise. We observe that V satisfies the monotonicity condition w.r.t. $\preceq$. Indeed, let $t \preceq t'$. If $\mathrm{V}(t) = \emptyset$, then $\mathrm{V}(t) \subseteq \mathrm{V}(t')$. Otherwise, $t_i^{\text{bad}} \preceq t$ for some

$1 \leq i \leq n$, and by transitivity of $\preceq$, $t_i^{\text{bad}} \preceq t'$ holds as well. Hence, $\mathrm{V}(t') = \{p\}$, and the result follows. Thus, the tuple $\mathcal{G} = \langle S, S_I, \delta, \preceq, \mathrm{V}\rangle$ is a BCGS with embedded frame $\mathcal{F}$.

By construction, $\mathrm{V}(t_i^{\text{bad}}) = \{p\}$ for all $1 \leq i \leq n$, and $\{t_1^{\text{bad}}, \ldots, t_n^{\text{bad}}\}$ is the set of successors of s consistent with the A-decision d_A. Hence, $\mathcal{G}, s \models \langle\langle A\rangle\rangle \bigcirc p$. Thus, since $s \preceq s'$, in order to show that $\mathcal{F}$ is not IATL$_\exists$-monotonic, we prove that $\mathcal{G}, s' \not\models \langle\langle A\rangle\rangle \bigcirc p$. By Condition ($\star$), for each A-decision $\mathrm{d}'_A \in \mathsf{Dec}(s', A)$, there is a full decision $\mathrm{d}' \in \mathsf{Dec}(s', \mathsf{Ag})$ which extends d'_A such that $\mathrm{V}(\delta(s', \mathrm{d}')) = \emptyset$. This means that $\mathcal{G}, s' \not\models \langle\langle A\rangle\rangle \bigcirc p$, and we are done. □

Well Behaved BCGS and Excluded Middle. We now observe that the addition of the law of the excluded middle causes the logic IATL to collapse to standard ATL. We say that a well-behaved BCGS is *Aristotelian* iff $\mathcal{G}, s \models p \vee \neg p$ for every $p \in \mathsf{AP}$ and state s of $\mathcal{G}$. It is easily seen that in an Aristotelian BCGS if $s \preceq s'$ then $\mathcal{V}(s) = \mathcal{V}(s')$. From this, and by condition C_1 and C_2 of Definition 2, we obtain that if $s \preceq s'$ then s and s' are alternating bisimilar (see (Alur, Henzinger, and Kupferman 2002) for a definition) and thus they satisfy the same (classical) ATL formulas. We can prove the following by induction on the structure of φ.

Theorem 2. *Let $\mathcal{G}$ be an Aristotelian BCGS and s one of its states. For every formula φ: $\mathcal{G}, s \models_{ATL} \varphi$ iff $\mathcal{G}, s \models \varphi$*

Comparison with imperfect information CGS (ICGS). ICGS (Reif 1984) and BCGS provide different and incomparable ways to represent partial information. In ICGS, states represent complete descriptions of facts and the partial knowledge of an agent *a* about these facts is modeled by an equivalence relation $\equiv_a$ over the states. Thus, in ICGS, a strategy of agent *a* must be *uniform* (Reif 1984), i.e., it must prescribe the same action in histories which are equivalent w.r.t. $\equiv_a$. Note that if $s \equiv_a s'$, in general, *there is no correlation* between the propositional labelings in s and s'. This lack of correlation is one of the technical motivations for undecidability of model checking ICGS against ATL (Dima and Tiplea 2011). In BCGS, the partial knowledge of an agent is directly modelled in the states and the partial order $\preceq$ models information refinement which does not allow for the rejection of facts (in contrast with belief revision): this justifies the monotonicity of propositional labeling in moving from a state to a more informative one. Since $\preceq$ is a partial order, two states are equivalent (i.e, $s \preceq s'$ and $s' \preceq s$) iff $s = s'$. Thus, in BCGS, a strategy is always uniform. In fact, for well-behaved BCGS, the notion of uniformity is replaced with that of strategic refinement. By Lemmata 1–2, if $s \preceq s'$, then for each A-strategy f_A starting at s (resp., A-strategy f'_A starting at s'), there is an A-strategy f'_A starting at s' (resp., A-strategy f_A starting at s) such that the strategy-tree rooted at s' induced by f'_A is a behavioural refinement of the strategy-tree rooted at s induced by f_A.

3.1 Example: Muddy Children Puzzle

We illustrate the IATL framework with the scenario of the well-known epistemic puzzle of the muddy children (Fagin et al. 1995). In this setting, there are $k \geq 1$ of n children with mud on their foreheads. They can only see the others, so they do not know their own status. Thus, they are uncertain whether the number of dirty children is k or $k-1$. Now their Father enters the scene and says: "at least one of you is dirty. Do you know whether your own forehead is dirty?" Children answer truthfully, and this is repeated round by round. The classical statement is that the k muddy children will learn that they are muddy answering "Yes" after the father repeats his question exactly k-times. There is a straightforward *constructive* proof of this fact by induction on k. When $k = 1$, the first question of the father is sufficient to solve the uncertainty, and the unique muddy child can answer "Yes". If $k > 1$, after the $(k-1)^{th}$-question the dirty children learn that none of the others has solved the uncertainty. By the induction hypothesis, they conclude that there cannot be only $k-1$ dirty children (*logical inference steps*). Thus, they will answer "Yes" at the next question of the father.

We model the considered setting with n children by a BCGS $\mathcal{G}_n = \langle S, S_I, \delta, \preceq, \mathrm{V}\rangle$ over $\mathsf{Ag} = \{F, 1, \ldots, n\}$, where the states specify partial knowledge, the transition function describes the temporal evolution of each single round, while the partial order captures the monotonic epistemic updating in moving from a round to the next one and the logical inference steps described above. In detail, S consists of the tuples (k, τ, ℓ, C) where (i) $1 \leq k \leq n$ is the number of muddy children, (ii) $\tau \in \{0, 1, 2\}$ marks the current step of a puzzle round, (iii) $\ell \in \{1, \ldots, k\}$ means that s is associated with the ℓ^{th} round of the scenario with k muddy children, (iv) $C \subseteq \{\mathit{know}\}$, where $C = \{\mathit{know}\}$ (resp., $C = \emptyset$) means that each muddy child knows (resp., does not know) its state, and (v) $C = \{\mathit{know}\}$ iff either $\ell = k$, or $\ell = k-1$ and $\tau = 2$. The initial states are of the form $(k, 0, 1, \emptyset)$ and the transition function δ is defined as:

- From states $s = (k, 0, \ell, C)$, s has the successor $(k, 1, \ell, C)$ reached when the father makes a question, and the same state s otherwise.
- States $s = (k, 1, \ell, C)$ have a unique successor $(k, 2, \ell, C')$ where $C' = \{\mathit{know}\}$ iff $\ell \geq k-1$.
- States $s = (k, 2, \ell, C)$ are sink states.

For the partial order, $\preceq$ is the transitive and reflexive closure of $\preceq'$ where $(k, \tau, \ell, C) \preceq' (k', \tau', \ell', C')$ if

- *either* $k = k'$, $\tau = \tau'$, $\ell \leq \ell'$, and $C \subseteq C'$ (i.e, the knowledge of the muddy children increases in moving from a round to the next one),
- *or* $k' = k-1$, $\tau' = \tau = 2$, $\ell' = \ell = k-1$, and $C = C' = \{\mathit{know}\}$ (*logical inference steps*).

Finally, $\mathsf{AP} = \{p_1, \ldots, p_n, \mathit{an}_1, \ldots, \mathit{an}_n, \mathit{know}, \mathit{yes}\}$ and

- $\mathrm{V}(k, 0, \ell, C) = \{p_k\}$, $\mathrm{V}(k, 1, \ell, C) = \emptyset$, and
- $\mathrm{V}(k, 2, \ell, C) = \{\mathit{an}_1, \ldots, \mathit{an}_\ell\} \cup D$, where $D = C$ if $\ell < k$ and $D = C \cup \{\mathit{yes}\}$ otherwise.

The number of states of $\mathcal{G}_n$ is quadratic in n. This is in contrast with classical formalizations of the problem, where the number of states typically grows exponentially in n. In Figure 3, we give a graphical representation of the BCGS $\mathcal{G}_n$ with $n = 3$. Note that the model is subdivided in three submodels, where each of them is associated with a distinct number $k \in \{1, 2, 3\}$ of muddy children. The submodel

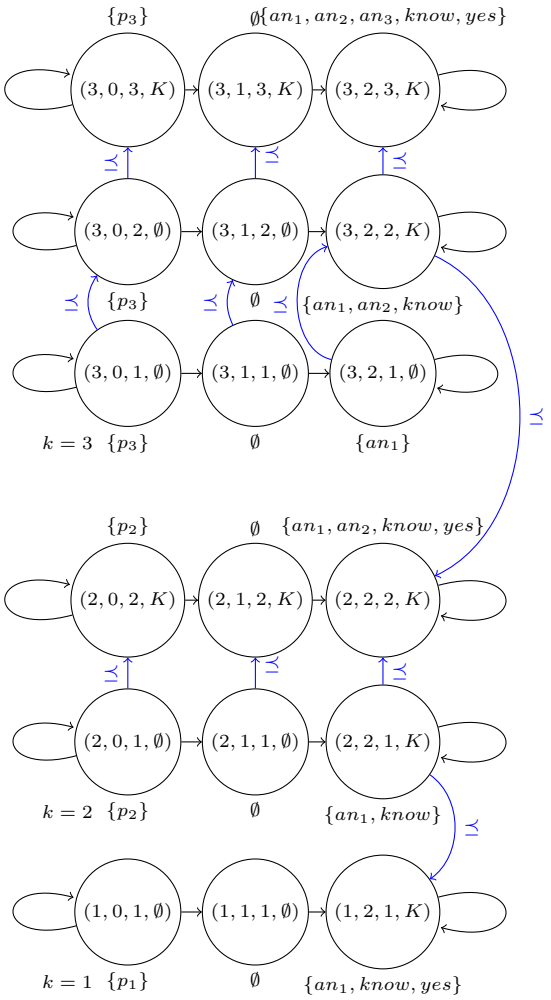

Figure 3: BCGS for Muddy Children problem with 3 children

for k muddy children consists of k rounds and the transition function describes the temporal evolution of each single round. Moreover, the k rounds are only connected by the edges of the partial order $\preceq$ (which are drawn in blue). The $\preceq$-edges connecting distinct submodels represent logical inference steps of the muddy children problem. In Figure 3, for simplicity, we report only the minimal $\preceq$-edges which generate the partial order $\preceq$. Moreover, for the states of the form $(k, \tau, \ell, \{know\})$, $\{know\}$ is replaced with the letter K.

Note that $\mathcal{G}_n$ is a well-behaved BCGS. Thus by truth monotonicity, the following IATL formula φ_n when asserted at the initial state $(k, 0, 1, \emptyset)$ with $k > 1$, expresses that the k muddy children will know their state exactly after the father announcement of the $(k-1)^{th}$ round, i.e., after the father repeats his question exactly $(k-1)$-times.

$$\varphi_n := \textstyle\bigwedge_{k=2}^{n} p_k \rightarrow \langle\!\langle \{F\} \rangle\!\rangle \top \mathsf{U} (an_1 \wedge (an_{k-1} \leftrightarrow know))$$

3.2 Properties of IATL

In this section, we first establish some expressiveness results about IATL. Then, for the class of *finitely-branching* BCGS (i.e., BCGS where each state has a finite number of successors), we provide least and greatest fix-point characterizations of the strategic quantifiers. Moreover, we show that for this class of BCGS, the perfect recall semantics and the memoryless semantics are equivalent. Note that these results are independent of the well-behaved assumption on CGS which characterizes truth monotonicity in IATL. However, without monotonicity, IATL would be not an intuitionistic logic, i.e., it would be not conservative over IPL.

For a formula φ and a BCGS $\mathcal{G}$ with set of states S, we denote by $[\![\varphi]\!]_{\mathcal{G}}$ the set of $\mathcal{G}$-states satisfying φ that is $[\![\varphi]\!]_{\mathcal{G}} = \{s \in S \mid \mathcal{G}, s \models \varphi\}$. We simply write $[\![\varphi]\!]$ if $\mathcal{G}$ is clear from the context. Two formulas φ_1 and φ_2 are *equivalent*, written $\varphi_1 \equiv \varphi_2$, if for each BCGS $\mathcal{G}$, $[\![\varphi_1]\!]_{\mathcal{G}} = [\![\varphi_2]\!]_{\mathcal{G}}$.

Expressiveness issues. We establish the following results, where IATL_{U} and IATL_{R} are the fragments of IATL obtained by disallowing the release modality for IATL_{U} and the until modality for IATL_{R}, respectively.

Theorem 3. *$\text{IATL}_{\exists}$ and $\text{IATL}_{\forall}$ are expressively incomparable, and IATL_{U} and IATL_{R} are expressively incomparable.*

Proof. The proof for $\text{IATL}_{\exists}$ and $\text{IATL}_{\forall}$ is by contraposition. Assume that either $\text{IATL}_{\exists}$ is subsumed by $\text{IATL}_{\forall}$ or $\text{IATL}_{\forall}$ is subsumed by $\text{IATL}_{\exists}$. We examine the first case (the second case being similar). Hence, for each $\text{IATL}_{\exists}$ formula, there exists an equivalent $\text{IATL}_{\forall}$ formula. Let us consider a BCGF $\mathcal{F}$ with partial order $\preceq$ which does not satisfy condition C_1 in Definition 2, but satisfies condition C_2 (obviously, such frames exist). By Theorem 1, $\mathcal{F}$ is not $\text{IATL}_{\exists}$-monotonic. Hence, there is a BCGS $\mathcal{G}$ whose frame is $\mathcal{F}$, two states s and s' with $s \preceq s'$, and an $\text{IATL}_{\exists}$ formula $\varphi_{\exists}$ such that $\mathcal{G}, s \models \varphi_{\exists}$ and $\mathcal{G}, s' \not\models \varphi_{\exists}$. By hypothesis, there exists an $\text{IATL}_{\forall}$ formula $\varphi_{\forall}$ such that $\varphi_{\forall} \equiv \varphi_{\exists}$. Hence, $\mathcal{G}, s \models \varphi_{\forall}$ and $\mathcal{G}, s' \not\models \varphi_{\forall}$. Since $\mathcal{F}$ satisfies C_2, by Theorem 1, $\mathcal{F}$ is $\text{IATL}_{\forall}$-monotonic. Hence, we deduce that $\mathcal{G}, s' \models \varphi_{\forall}$, and we reach a contradiction.

For the second part of the theorem, it is known (Balbiani et al. 2019) that modalities U and R are both necessary for ensuring the full expressivity of intuitionistic LTL (ILTL). Since the models of ILTL correspond to BCGS where each state has exactly one successor, and over these models, IATL coincides with ILTL, the result follows. □

Fix-point characterizations of strategic quantifiers. Let $\mathcal{G} = \langle S, S_I, \delta, \preceq, \mathrm{V} \rangle$ be a BCGS. Given $X \subseteq S$, we denote by $X^{\mathbf{c}}$ its complement, i.e., the set $S \setminus X$, and by $X^{\uparrow}$ the set of states $s \in S$ such that for each state s' with $s \preceq s'$, $s' \in X$. Note that $X^{\uparrow} \subseteq X$.

Moreover, for a coalition A, we write $\text{Pre}^{\exists}(A, X)$ for the set of states s such that for *some* A-decision d_A available at s, *all the successors* of s consistent with d_A are in X. Finally, we denote by $\text{Pre}^{\forall}(A, X)$, the set of states s such that for *all* A-decisions d_A available at s, *some successor* of s consistent with d_A is in X. The following two propositions easily follows from the semantics of IATL.

Proposition 2. *Given a BCGS $\mathcal{G}$ with set of states S and a coalition A, the following holds:*

- $[\![\varphi_1 \rightarrow \varphi_2]\!] = ([\![\varphi_1]\!]^{\mathbf{c}} \cup [\![\varphi_2]\!])^{\uparrow}$.

- $[\![\langle\!\langle A\rangle\!\rangle \bigcirc \varphi]\!] = \mathit{Pre}^{\exists}(A, [\![\varphi]\!])$.
- $[\![[\![A]\!] \bigcirc \varphi]\!] = \mathit{Pre}^{\forall}(A, [\![\varphi]\!])$.
- $(\mathit{Pre}^{\forall}(A, [\![X]\!]))^{\mathbf{c}} = \mathit{Pre}^{\exists}(A, [\![X]\!]^{\mathbf{c}})$ *for each* $X \subseteq S$.

Proposition 3. *The following equivalences hold:*
- $\langle\!\langle A\rangle\!\rangle(\varphi_1 \mathsf{U} \varphi_2) \equiv \varphi_2 \vee (\varphi_1 \wedge \langle\!\langle A\rangle\!\rangle \bigcirc \langle\!\langle A\rangle\!\rangle(\varphi_1 \mathsf{U} \varphi_2))$.
- $\langle\!\langle A\rangle\!\rangle(\varphi_1 \mathsf{R} \varphi_2) \equiv \varphi_2 \wedge (\varphi_1 \vee \langle\!\langle A\rangle\!\rangle \bigcirc \langle\!\langle A\rangle\!\rangle(\varphi_1 \mathsf{R} \varphi_2))$.
- $[\![A]\!](\varphi_1 \mathsf{U} \varphi_2) \equiv \varphi_2 \vee (\varphi_1 \wedge [\![A]\!] \bigcirc [\![A]\!](\varphi_1 \mathsf{U} \varphi_2))$.
- $[\![A]\!](\varphi_1 \mathsf{R} \varphi_2) \equiv \varphi_2 \wedge (\varphi_1 \vee [\![A]\!] \bigcirc [\![A]\!](\varphi_1 \mathsf{R} \varphi_2))$.

Let $\mathcal{G}$ be a BCGS with set of states S and g be a monotonic function $2^S \mapsto 2^S$. The *dual* $\widetilde{g}$ of g is the function $2^S \mapsto 2^S$ defined as follows for each $X \subseteq S$, $\widetilde{g}(X) \stackrel{\text{def}}{=} (g(X^{\mathbf{c}}))^{\mathbf{c}}$. Note that $\widetilde{g}$ is monotonic too. Recall that by Tarski theorem, both g and $\widetilde{g}$ have the least and the greatest fixpoints. Moreover, the following holds.

Proposition 4. *Let $\mathcal{G}$ be a BCGS with set of states S and $g : 2^S \mapsto 2^S$. If g is monotonic, then the greatest fixpoint (resp., least fixpoint) of g coincides with the complement of the least fixpoint (resp., greatest fixpoint) of $\widetilde{g}$.*

For all formulas φ_1 and φ_2 and BCGS $\mathcal{G}$ with set of states S, we consider the following monotonic functions $2^S \mapsto 2^S$ defined as follows for each $X \subseteq S$:
- $\langle\!\langle A\rangle\!\rangle \mathsf{U}_{\varphi_1,\varphi_2}(X) \stackrel{\text{def}}{=} [\![\varphi_2]\!] \cup ([\![\varphi_1]\!] \cap \mathrm{Pre}^{\exists}(A, X))$
- $\langle\!\langle A\rangle\!\rangle \mathsf{R}_{\varphi_1,\varphi_2}(X) \stackrel{\text{def}}{=} [\![\varphi_2]\!] \cap ([\![\varphi_1]\!] \cup \mathrm{Pre}^{\exists}(A, X))$
- $[\![A]\!] \mathsf{U}_{\varphi_1,\varphi_2}(X) \stackrel{\text{def}}{=} [\![\varphi_2]\!] \cup ([\![\varphi_1]\!] \cap \mathrm{Pre}^{\forall}(A, X))$
- $[\![A]\!] \mathsf{R}_{\varphi_1,\varphi_2}(X) \stackrel{\text{def}}{=} [\![\varphi_2]\!] \cap ([\![\varphi_1]\!] \cup \mathrm{Pre}^{\forall}(A, X))$

We deduce the following fix-point characterizations of the strategy quantifiers.

Theorem 4. *For all formulas φ_1 and φ_2 and every finitely-branching BCGS $\mathcal{G}$, the following holds:*
1. $[\![\langle\!\langle A\rangle\!\rangle(\varphi_1 \mathsf{U} \varphi_2)]\!]$ *is the least fix-point of* $\langle\!\langle A\rangle\!\rangle \mathsf{U}_{\varphi_1,\varphi_2}$;
2. $[\![\langle\!\langle A\rangle\!\rangle(\varphi_1 \mathsf{R} \varphi_2)]\!]$ *is the greatest fix-point of* $\langle\!\langle A\rangle\!\rangle \mathsf{R}_{\varphi_1,\varphi_2}$;
3. $[\![[\![A]\!](\varphi_1 \mathsf{U} \varphi_2)]\!]$ *is the least fix-point of* $[\![A]\!] \mathsf{U}_{\varphi_1,\varphi_2}$;
4. $[\![[\![A]\!](\varphi_1 \mathsf{R} \varphi_2)]\!]$ *is the greatest fix-point of* $[\![A]\!] \mathsf{R}_{\varphi_1,\varphi_2}$.

Proof. We prove Properties 1 and 4. The proofs for Properties 2 and 3 are similar but do not require the finitely-branching assumption.

Property 1. let $X = [\![\langle\!\langle A\rangle\!\rangle \varphi_1 \mathsf{U} \varphi_2]\!]$. By Propositions 2 and 3, it is clear that $\langle\!\langle A\rangle\!\rangle \mathsf{U}_{\varphi_1,\varphi_2}(X) \subseteq X$. Hence, X is a pre-fixed point of $\langle\!\langle A\rangle\!\rangle \mathsf{U}_{\varphi_1,\varphi_2}$. Pick an arbitrary pre-fixed point Y of $\langle\!\langle A\rangle\!\rangle \mathsf{U}_{\varphi_1,\varphi_2}$. Hence, $\langle\!\langle A\rangle\!\rangle \mathsf{U}_{\varphi_1,\varphi_2}(Y) \subseteq Y$. By the Tarski theorem, it suffices to show that $X \subseteq Y$. Fix $s \in X$. We prove that $s \in Y$. Let f_A be an A-strategy such that for each $\pi \in \mathsf{out}(s, f_A)$, $\mathcal{G}, \pi \models \varphi_1 \mathsf{U} \varphi_2$. Let Γ be the set of *finite* outcomes of f_A starting at state s. Evidently, Γ represents a tree whose root is s. We consider the tree Γ_{min} obtained from Γ by removing all the nodes (finite outcomes) which are children of finite outcomes whose last state is in $[\![\varphi_2]\!]$. We easily deduce that Γ_{min} is non-empty and the following holds:

- for each leaf ρ in Γ_{min}, $\mathit{last}(\rho) \in [\![\varphi_2]\!]$ and $\rho(i) \in [\![\varphi_1]\!]$ for all $0 \leq i < |\rho| - 1$;
- for each internal node ρ of Γ_{min}, the children of ρ in Γ are in Γ_{min} as well;
- there is no infinite branch in Γ_{min}.

Since $\mathcal{G}$ is finitely-branching, by König Lemma, Γ_{min} is *finite*. For each $\rho \in \Gamma_{min}$, let h_ρ be the height of the subtree of Γ_{min} rooted at ρ. We show by induction on h_ρ that $\mathit{last}(\rho) \in Y$. Hence, the result follows. For the base case ($h_\rho = 0$), ρ is a leaf, hence, $\mathit{last}(\rho) \in [\![\varphi_2]\!]$. Thus, since $[\![\varphi_2]\!] \subseteq Y$, we obtain that $\mathit{last}(\rho) \in Y$. Now, assume that $h_\rho > 0$. Since ρ is a prefix of a leaf, we have that $\mathit{last}(\rho) \in [\![\varphi_1]\!]$. Moreover, since the children of ρ in Γ are in Γ_{min} as well, by the induction hypothesis, we obtain that $\mathit{last}(\rho) \in \mathrm{Pre}^{\exists}(A, Y)$. Being $[\![\varphi_1]\!] \cap \mathrm{Pre}^{\exists}(A, Y) \subseteq Y$, we conclude that $\mathit{last}(\rho) \in Y$.

Property 4. Let $X = [\![[\![A]\!] \varphi_1 \mathsf{R} \varphi_2]\!]$. We introduce an additional temporal modality $\mathsf{U}^{\mathbf{c}}$ which is the dual of R:

$$\mathcal{G}, \pi \models \varphi_1 \mathsf{U}^{\mathbf{c}} \varphi_2 \quad \Leftrightarrow \text{there is } i \geq 0 \text{ s.t. } \mathcal{G}, \pi(i) \not\models \varphi_2 \text{ and } \mathcal{G}, \pi(k) \not\models \varphi_1 \text{ for all } 0 \leq k < i$$

By the semantics, $[\![[\![A]\!] \varphi_1 \mathsf{R} \varphi_2]\!] = [\![\langle\!\langle A\rangle\!\rangle \varphi_1 \mathsf{U}^{\mathbf{c}} \varphi_2]\!]^{\mathbf{c}}$. Let us consider the monotonic function $\langle\!\langle A\rangle\!\rangle \mathsf{U}^{\mathbf{c}}_{\varphi_1,\varphi_2}$ from $2^S \to 2^S$ defined as follows for each $X \subseteq S$:

$$\langle\!\langle A\rangle\!\rangle \mathsf{U}^{\mathbf{c}}_{\varphi_1,\varphi_2}(X) \stackrel{\text{def}}{=} [\![\varphi_2]\!]^{\mathbf{c}} \cup ([\![\varphi_1]\!]^{\mathbf{c}} \cap \mathrm{Pre}^{\exists}(A, X)).$$

By proceeding as in the proof of Property 1, we deduce that $[\![\langle\!\langle A\rangle\!\rangle \varphi_1 \mathsf{U}^{\mathbf{c}} \varphi_2]\!]$ is the least fixpoint of $\langle\!\langle A\rangle\!\rangle \mathsf{U}^{\mathbf{c}}_{\varphi_1,\varphi_2}$. By Proposition 2, $\langle\!\langle A\rangle\!\rangle \mathsf{U}^{\mathbf{c}}_{\varphi_1,\varphi_2}$ is the dual of the mapping $[\![A]\!] \mathsf{R}_{\varphi_1,\varphi_2}$. Thus, since $[\![[\![A]\!] \varphi_1 \mathsf{R} \varphi_2]\!] = [\![\langle\!\langle A\rangle\!\rangle \varphi_1 \mathsf{U}^{\mathbf{c}} \varphi_2]\!]^{\mathbf{c}}$, by Proposition 4, the result follows, i.e., $[\![[\![A]\!](\varphi_1 \mathsf{R} \varphi_2)]\!]$ is the greatest fix-point of $[\![A]\!] \mathsf{R}_{\varphi_1,\varphi_2}$. □

Equivalence with memoryless semantics. We can show that the perfect recall semantics and the memoryless semantics of IATL are equivalent over the class of finitely-branching BCGS. The result holds over arbitrary finitely-branching BCGS where the well-behaved requirement may not be fulfilled.

Theorem 5. *For any finitely-branching BCGS $\mathcal{G}$, state s of $\mathcal{G}$, and IATL formula φ: $\mathcal{G}, s \models \varphi \Leftrightarrow \mathcal{G}, s \models_m \varphi$.*

Proof. The proof is by structural induction on φ. We focus on the case where φ is of the form $\langle\!\langle A\rangle\!\rangle \varphi_1 \mathsf{U} \varphi_2$. The other cases are simpler or similar. The $\Leftarrow$-implication is trivial. Now, let us consider the $\Rightarrow$- implication. We say that an history ρ is good if $\mathcal{G}, \mathit{last}(\rho) \models_m \varphi_2$ and $\mathcal{G}, \rho(i) \models_m \varphi_1$ for all $0 \leq i < |\rho| - 1$. Assume that $\mathcal{G}, s \models \varphi$. Hence, by the induction hypothesis, there is an A-strategy f_A such that $\mathcal{G}, \pi \models_m \varphi_1 \mathsf{U} \varphi_2$ for all $\pi \in \mathsf{out}(s, f_A)$. As in the proof of Property 1 of Theorem 4, we consider the set Γ of the finite outcomes of f_A starting at state s (Γ corresponds to a tree rooted at s) and the tree Γ_{min} obtained from Γ by removing all the nodes (finite outcomes) which are children of finite outcomes whose last state is in $[\![\varphi_2]\!]$. We recall that:

- Γ_{min} is finite and each leaf ρ in Γ_{min} is good;
- for each internal node ρ of Γ_{min}, the children of ρ in Γ are in Γ_{min} as well.

Let $S(\Gamma_{min})$ be the set of states visited by the histories in Γ_{min}. For each $t \in S(\Gamma_{min})$, we choose any ρ in Γ_{min} leading to t such that the height of the subtree of Γ_{min} rooted at ρ is minimal. We denote by ρ_t the chosen history and by $\mathrm{rank}(t)$ the height of the subtree rooted at ρ_t. Let f_A^m be the memoryless A-strategy associating to each state $t \in S(\Gamma_{min})$, $f_A(\rho_t)$, and to each other state t some chosen A-decision available at t. In order to conclude the proof of Property (1), since $s \in S(\Gamma_{min})$, it suffices to show that for each $t \in S(\Gamma_{min})$ and for each $\pi \in \mathsf{out}(t, f_A^m)$, there is a non-empty prefix of π which is good. The proof is by induction on $\mathrm{rank}(t)$.

- $\mathrm{rank}(t) = 0$: by construction ρ_t is a leaf of Γ_{min}. Hence, $\mathcal{G}, t \models_m \varphi_2$, and in this case, the result holds.
- $\mathrm{rank}(t) > 0$: hence, ρ_t is an internal node of Γ_{min} and $\mathcal{G}, t \models_m \varphi_1$. Since the children of ρ_t in Γ are in Γ_{min} as well, by construction, $\pi(1) \in S(\Gamma_{min})$, and $\mathrm{rank}(\pi(1)) < \mathrm{rank}(t)$. By the induction hypothesis, the result follows. □

4 IATL Model Checking

We address the model-checking problem for IATL, that is, checking for given *finite* BCGS $\mathcal{G}$ and IATL formula φ, whether $\mathcal{G}, s \models \varphi$ for each initial state s. By applying the results of Section 3.2, we show that IATL model checking has the same complexity as standard ATL model checking.

Theorem 6. *IATL model checking is* PTIME*-complete and can be solved in time $O(|\mathcal{G}|^2 \cdot |\Phi|)$ for given finite BCGS $\mathcal{G}$ and formula Φ.*

Proof. PTIME-hardness follows from PTIME-hardness of ATL model checking (Alur, Henzinger, and Kupferman 2002) and the fact that checking that a finite CGS $\mathcal{G}$ is an ATL-model of a formula φ reduces to checking that the BCGS extension of $\mathcal{G}$ with the identity partial order is an IATL-model of φ. For the matching upper bound, we consider the model checking algorithm illustrated in Figure 4 which computes, for a given finite BCGS $\mathcal{G}$ and IATL formula Φ, the set of states $[\![\Phi]\!]$ which satisfy Φ. The algorithm computes the denotations $[\![\varphi]\!]$ of the subformulas φ of Φ by a bottom-up traversal of the parse tree of the formula Φ in accordance with Proposition 2 and Theorem 4. All cases are identical to the standard ATL model checking algorithm, except for the case of the intuitionistic implication $\rightarrow$, reported below. Each subformula is processed exactly once. Hence, the algorithm runs in time $O(|\mathcal{G}|^2 \cdot |\Phi|)$. □

We address the model-checking problem for IATL, that is, checking for given *finite* BCGS $\mathcal{G}$ and IATL formula φ, whether $\mathcal{G}, s \models \varphi$ for each initial state s. By applying the results of Section 3.2, we show that IATL model checking has the same complexity as standard ATL model checking.

Theorem 7. *IATL model checking is* PTIME*-complete and can be solved in time $O(|\mathcal{G}|^2 \cdot |\Phi|)$ for given finite BCGS $\mathcal{G}$ and formula Φ.*

Proof. PTIME-hardness follows from PTIME-hardness of ATL model checking (Alur, Henzinger, and Kupferman 2002) and the fact that checking that a finite CGS $\mathcal{G}$ is an ATL-model of a formula φ reduces to checking that the BCGS extension of $\mathcal{G}$ with the identity partial order is an IATL-model of φ.

for each $\varphi \in \mathrm{Sub}(\Phi)$

case $\varphi = p$: $[\![\varphi]\!] := \{s | p \in \mathrm{V}(s)\}$;
case $\varphi = \varphi_1 \vee \varphi_2$: $[\![\varphi]\!] := [\![\varphi_1]\!] \cup [\![\varphi_2]\!]$;
case $\varphi = \varphi_1 \wedge \varphi_2$: $[\![\varphi]\!] := [\![\varphi_1]\!] \cap [\![\varphi_2]\!]$;
case $\varphi = \varphi_1 \rightarrow \varphi_2$: $[\![\varphi]\!] := ([\![\varphi_1]\!]^{\mathbf{c}} \cup [\![\varphi_2]\!])^{\uparrow}$;
case $\varphi = \langle\!\langle A \rangle\!\rangle \bigcirc \varphi_1$: $[\![\varphi]\!] := \mathrm{Pre}^{\exists}(A, [\![\varphi_1]\!])$;
case $\varphi = [\![A]\!] \bigcirc \varphi_1$: $[\![\varphi]\!] := \mathrm{Pre}^{\forall}(A, [\![\varphi_1]\!])$;
case $\varphi = \langle\!\langle A \rangle\!\rangle \varphi_1 \mathsf{U} \varphi_2$:
  $Q_1 := \emptyset$; $Q_2 := [\![\varphi_1]\!]$; $Q_3 := [\![\varphi_2]\!]$;
  while $Q_3 \not\subseteq Q_1$ **do** $Q_1 := Q_1 \cup Q_3$;
  $Q_3 := \mathrm{Pre}^{\exists}(A, Q_1) \cap Q_2$;
  $[\![\varphi]\!] := Q_1$;
case $\varphi = [\![A]\!] \varphi_1 \mathsf{U} \varphi_2$: as in the previous case but we use $\mathrm{Pre}^{\forall}$ instead of $\mathrm{Pre}^{\exists}$;
case $\varphi = \langle\!\langle A \rangle\!\rangle \varphi_1 \mathsf{R} \varphi_2$:
  $Q_1 := S$; $Q_2 := [\![\varphi_1]\!]$; $Q_3 := [\![\varphi_2]\!]$;
  while $Q_1 \not\subseteq Q_3$ **do** $Q_1 := Q_1 \cap Q_3$;
  $Q_3 := \mathrm{Pre}^{\exists}(A, Q_1) \cup Q_2$;
  $[\![\varphi]\!] := Q_1$;
case $\varphi = [\![A]\!] \varphi_1 \mathsf{R} \varphi_2$:as in the previous case but we use $\mathrm{Pre}^{\forall}$ instead of $\mathrm{Pre}^{\exists}$;

return $[\![\Phi]\!]$;

Figure 4: Algorithm for IATL model checking

For the matching upper bound, we consider the model checking algorithm illustrated in Figure 4 which computes for a given finite BCGS $\mathcal{G}$ and IATL formula Φ, the set of states $[\![\Phi]\!]$ which satisfy Φ. The algorithm computes the denotations $[\![\varphi]\!]$ of the subformulas φ of Φ by a bottom-up traversal of the parse tree of the formula Φ in accordance with Proposition 2 and Theorem 4. In particular, the subformulas of Φ are ordered in a queue $\mathrm{Sub}(\Phi)$, where a subformula φ_1 precedes φ_2 iff $|\varphi_1| \leq |\varphi_2|$. Each while loop requires at most $O(|S|)$ iterations, and the time spent by a while loop is at most $O(|\mathcal{G}|)$, as the computation of preimages can be implemented in such a way that each transition of the model is visited exactly once during the various iterations of the loop. Note that the computation of the upward closure $X^{\uparrow}$ of a subset X of S requires quadratic time in $|X|$. Each subformula is processed exactly once. Hence, the algorithm runs in time $O(|\mathcal{G}|^2 \cdot |\Phi|)$. □

Implementation & Benchmark. To evaluate the performance of the IATL model checking algorithm, we extended the VITAMIN model checker (Ferrando and Malvone 2024) to support the specific semantics of intuitionistic ATL. The implementation adheres to privacy and copyright constraints while ensuring reproducibility of our results. The data representation follows the schema introduced in (Ferrando and Malvone 2024), and the model structure is based on the

	n_agents	n_states	t_mc (sec)
1	25	975	0.07944
2	50	3825	0.59195
3	75	8550	2.20910
4	100	15150	7.88473
5	125	23625	22.7658
6	150	33975	52.6759
7	175	46200	109.529
8	200	60300	211.067
9	250	94125	374.853
10	300	135400	742.362

Table 1: Benchmark for IATL model checking

setting illustrated in Figure 3. A model generation algorithm was developed to generalize over a variable number k of children, enabling scalability and adaptability of our approach. The experimental evaluation was conducted using the models $\mathcal{G}_n$ and the IATL formulas φ_n described in Section 3.1. All experiments were run on a desktop machine equipped with an Intel i7 12700K processor and 64GB of RAM, providing a consistent and controlled testing environment. Table 1 reports the average execution time (in seconds) across 10 independent runs for increasing model sizes, achieved by varying the number n of agents (children). The results show that our IATL model checking algorithm exhibits efficient scalability: as the number of agents and the corresponding state space grow, execution times remain stable without significant degradation. Our implementation of the model checking process, depicted in Figure 4, accepts both the model and the logical formula as inputs and systematically performs the checks by reproducing the testing procedure described. This ensures that the experiments can be reliably replicated under similar conditions, fostering transparency and facilitating the validation of our findings.

5 Conclusion

In this work, we have introduced Intuitionistic ATL (IATL), an extension of ATL that incorporates intuitionistic logic to offer a novel and expressive framework for strategic reasoning in multi-agent systems, which is suitable for modelling partial information and dynamic information refinement. Classical ATL typically assumes that a fact is either true or false in a state, thus encoding complete knowledge. In contrast, IATL exploits intuitionistic truth to represent states where a fact may be unknown, and yet sufficient information may still be available to ensure a winning strategy. This viewpoint captures realistic multi-agent scenarios (e.g., sensing, incremental knowledge in protocols) in which full knowledge is rarely at hand. In traditional temporal logics for strategic reasoning, imperfect information is usually modelled by the notion of *possible worlds*, between which an agent is unable to distinguish. Thus, unlike IATL, these logics are unable to capture dynamic scenarios where agents acquire new knowledge that allows agents to distinguish between worlds that they previously could not separate. As illustrated by the well-known muddy children puzzle, the IATL framework can model epistemic scenarios where agents must reason about their own and others' knowledge. This highlights IATL's potential for real-world applications in areas requiring reasoning under uncertainty, such as multi-agent planning, automated verification of distributed protocols, and strategic decision-making in AI. Importantly, our model-checking algorithm retains computational feasibility, offering a PTIME-complete solution that mirrors the complexity of classical ATL while delivering enhanced expressive power for imperfect information settings. We plan to get more insights on IATL by developing a complete axiomatization of IATL. This would allow to solve the satisfiability and validity problem. Another direction is to extend the intuitionistic approach beyond ATL to more expressive logics such as ATL* or Strategy Logic (Mogavero et al. 2014).

Acknowledgments

This work has been partially funded by the MUR PRIN project RIPER (No. 20203FFYLK), and the PNRR MUR projects FAIR (No. PE0000013-FAIR), INFANT (No. E23C24000390006) and APLAND (No. E63C24002020001).

References

Ågotnes, T.; Goranko, V.; Jamroga, W.; and Wooldridge, M. 2015. Knowledge and ability. In *Handbook of Epistemic Logic*. College Publications. 543–589.

Alur, R.; Henzinger, T.; Kupferman, O.; and Vardi, M. 1998. Alternating Refinement Relations. In *Proc. 9th CONCUR*, LNCS 1466, 163–178. Springer.

Alur, R.; Henzinger, T.; and Kupferman, O. 2002. Alternating-time temporal logic. *Journal of the ACM* 49(5):672–713.

Balbiani, P., and Diéguez, M. 2016. Temporal Here and There. In *Proc. 15th JELIA*, volume 10021 of *LNCS*, 81–96.

Balbiani, P.; Boudou, J.; Diéguez, M.; and Fernández-Duque, D. 2019. Intuitionistic linear temporal logics. *ACM Transactions on Computational Logic* 21(2):1–32.

Bauer-Mangelberg, S.; van Heijenoort, J.; and Bauer-Mengelberg, S. 1970. On the significance of the principle of excluded middle in mathematics, especially in function theory. *Journal of Symbolic Logic* 35(2):332–333.

Belardinelli, F.; Lomuscio, A.; Murano, A.; and Rubin, S. 2020. Verification of multi-agent systems with public actions against strategy logic. *Artificial Intelligence* 285:103302.

Berthon, R.; Maubert, B.; Murano, A.; Rubin, S.; and Vardi, M. 2021. Strategy logic with imperfect information. *ACM Transactions on Computational Logic* 22(1):1–51.

Bozzelli, L., and Pearce, D. 2015. On the complexity of temporal equilibrium logic. In *Proc. 30th LICS*, 645–656. IEEE Computer Society.

Brewka, G.; Eiter, T.; and Truszczynski, M. 2011. Answer set programming at a glance. *Communications of the ACM* 54(12):92–103.

Brouwer, L. 1975. Intuitionism and formalism. In *Philosophy and Foundations of Mathematics*. Elsevier. 123–138.

Bulling, N., and Jamroga, W. 2014. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Autonomous agents and multi-agent systems* 28:474–518.

Cabalar, P., and Vega, G. P. 2007. Temporal equilibrium logic: A first approach. In *Proc. 11th EUROCAST*, LNCS 4739, 241–248. Springer.

Catta, D.; Malvone, V.; and Murano, A. 2023. Reasoning about intuitionistic computation tree logic. In *Proc. 3rd AREA@ECAI 2023*, EPTCS 391, 42–48.

Čermák, P.; Lomuscio, A.; Mogavero, F.; and Murano, A. 2018. Practical verification of multi-agent systems against SLK specifications. *Information and Computation* 261:588–614.

Clarke, E., and Emerson, E. 1981. Design and synthesis of synchronization skeletons using branching time temporal logic. In *Workshop on Logics of Programs*, LNCS 131, 52–71. Springer.

Dima, C., and Tiplea, F. 2011. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR* abs/1102.4225.

Dummett, M. 2000. *Elements of intuitionism*, volume 39. Oxford University Press.

Emerson, E., and Halpern, J. 1986. "Sometimes" and "Not Never" revisited: on branching versus linear time temporal logic. *Journal of the ACM* 33(1):151–178.

Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning About Knowledge*. MIT Press.

Ferrando, A., and Malvone, V. 2024. VITAMIN: A compositional framework for model checking of multi-agent systems. *CoRR* abs/2403.02170.

Fischer Servi, G. 1977. On modal logic with an intuitionistic base. *Studia Logica* 36(3):141–149.

Jamroga, W., and Bulling, N. 2011. Comparing variants of strategic ability. In *Proc. 22nd IJCAI*, 252–257. IJCAI/AAAI.

Jamroga, W.; Malvone, V.; and Murano, A. 2019. Natural strategic ability. *Artificial Intelligence* 277.

Kleene, S. 1957. Semantic construction of intuitionistic logic. *The Journal of Symbolic Logic* 22(4):363–365.

Kripke, S. 1965. Semantical analysis of intuitionistic logic. I. In *Studies in Logic and the Foundations of Mathematics*, volume 40. Elsevier. 92–130.

Kupferman, O., and Vardi, M. 2000. Synthesis with incomplete informatio. *Advances in temporal logic* 109–127.

Kupferman, O.; Vardi, M.; and Wolper, P. 2001. Module checking. *Information and Computation* 164(2):322–344.

Laroussinie, F., and Markey, N. 2015. Augmenting ATL with strategy contexts. *Information and Computation* 245:98–123.

Lomuscio, A., and Raimondi, F. 2006. Model checking knowledge, strategies, and games in multi-agent systems. In *Proc. 5th AAMAS*, 161–168. ACM.

Mancosu, P., ed. 1997. *From Brouwer to Hilbert: The Debate on the Foundations of Mathematics in the 1920s*. Oxford, England: Oxford University Press USA.

Maubert, B., and Murano, A. 2018. Reasoning about knowledge and strategies under hierarchical information. In *Proc. 16th KR*, 530–540. AAAI Press.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. 2012. What makes ATL* decidable? A decidable fragment of strategy logic. In *Proc. 23rd CONCUR*, LNCS 7454, 193–208. Springer.

Mogavero, F.; Murano, A.; Perelli, G.; and Vardi, M. 2014. Reasoning about strategies: On the model-checking problem. *ACM Transactions on Computational Logic* 15(4):1–47.

Moschovakis, J. 2023. Intuitionistic Logic. In *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2023 edition.

Plotkin, G., and Stirling, C. 1986. A framework for intuitionistic modal logics. In *Proc. 1st TARK*, 399–406. Morgan Kaufmann.

Pnueli, A., and Rosner, R. 1989. On the synthesis of a reactive module. In *Proc. 16th POPL*, 179–190. ACM Press.

Pnueli, A. 1977. The Temporal Logic of Programs. In *Proc. 18th FOCS*, 46–57. IEEE Computer Society.

Rasiowa, H. 1963. *The mathematics of metamathematics*. Panstwowe Wydawn. Naukowe.

Reif, J. 1984. The complexity of two-player games of incomplete information. *Journal of Computer and System Sciences* 29(2):274–301.

Simpson, A. 1994. *The proof theory and semantics of intuitionistic modal logic*. Ph.D. Dissertation, University of Edinburgh, UK.

Svejdar, V. 2003. On the polynomial-space completeness of intuitionistic propositional logic. *Archive for Mathematical Logic* 42(7):711–716.

Van Benthem, J. 2009. The information in intuitionistic logic. *Synthese* 167(2):251–270.

van Dalen, D., and Troelstra, A. 1988. *Constructivism in Mathematics*. Amsterdam: North-Holland.

van der Hoek, W., and Wooldridge, M. 2002. Tractable multiagent planning for epistemic goals. In *Proc. 1st AAMAS*, 1167–1174. ACM.