

Fixpoint Approximation of Strategic Abilities under Imperfect Information

Wojciech Jamroga
Institute of Computer Science,
Polish Academy of Sciences
w.jamroga@ipipan.waw.pl

Michał Knapik
Institute of Computer Science,
Polish Academy of Sciences
michal.knapik@ipipan.waw.pl

Damian Kurpiewski
Institute of Computer Science,
Polish Academy of Sciences
damian.kurpiewski@ipipan.waw.pl

ABSTRACT

Model checking of strategic ability under imperfect information is known to be hard. In this paper, we propose translations of ATL_{ir} formulae that provide lower and upper bounds for their truth values, and are cheaper to verify than the original specifications. Most interestingly, the lower approximation is provided by a fixpoint expression that uses a nonstandard variant of the next-step ability operator. We show the correctness of the translations, establish their computational complexity, and validate the approach by experiments with a scalable scenario of Bridge play.

Keywords

strategic ability, alternating-time temporal logic, imperfect information, model checking, alternating mu-calculus

1. INTRODUCTION

There is a growing number of works that study the syntactic and semantic variants of the *strategic logic* ATL for agents with imperfect information [2]. The contributions are mainly theoretical, and include results concerning the conceptual soundness of a given semantics [29, 18, 1, 21, 10, 15, 2], meta-logical properties [16, 7], and the complexity of model checking [29, 20, 16, 30, 13, 5]. However, there is relatively little research on the *use* of the logics, in particular on practical algorithms for reasoning and/or verification in scenarios where agents have a limited view of the world.

This is somewhat easy to understand, since model checking of ATL variants with imperfect information has been proved Δ_2^P - to PSPACE -complete for agents playing memoryless strategies [29, 20, 5] and EXPTIME -complete to undecidable for agents with perfect recall of the past [13, 16]. Moreover, the imperfect information semantics of ATL does not admit alternation-free fixpoint characterizations [6, 11, 12], which makes incremental synthesis of strategies impossible, or at least difficult to achieve. Some early attempts at verification of imperfect information strategies made their way into the MCMAS model-checker [25, 28, 23, 24], but the issue was never at the heart of the tool. More dedicated attempts began to emerge only recently [26, 8, 17, 9]. Up until now, experimental results confirm that the initial in-

tuition was right: model checking of strategic modalities for imperfect information is hard, and dealing with it requires innovative algorithms and verification techniques.

In this paper, we propose that in some instances, instead of the exact model checking, it suffices to provide an upper and/or lower bound for the output. The intuition for the upper bound is straightforward: instead of checking existence of an imperfect information strategy, we can look for a perfect information strategy that obtains the same goal. If the latter is false, the former must be false too. Finding a reasonable lower bound is nontrivial, but we construct one by means of a fixpoint expression in alternating epistemic mu-calculus. We begin by showing that the straightforward fixpoint approach does not work. Then, we propose how it can be modified to obtain guaranteed lower bounds. To this end, we alter the next-step operator in such a way that traversing the appropriate epistemic neighborhood is seen as an atomic activity. We show the correctness of the translations, establish their computational complexity, and validate the approach by experiments with some scalable scenarios.

2. VERIFYING STRATEGIC ABILITY

In this section we provide an overview of the relevant variants of ATL . We refer the to [3, 31, 29, 6, 19] for details.

2.1 Models, Strategies, Outcomes

A *concurrent epistemic game structure* or *CEGS* is given by $M = \langle \text{Agt}, St, Props, V, Act, d, o, \{\sim_a \mid a \in \text{Agt}\} \rangle$ which includes a nonempty finite set of all agents $\text{Agt} = \{1, \dots, k\}$, a nonempty set of states St , a set of atomic propositions $Props$ and their valuation $V : Props \rightarrow 2^{St}$, and a nonempty finite set of (atomic) actions Act . Function $d : \text{Agt} \times St \rightarrow 2^{Act}$ defines nonempty sets of actions available to agents at each state, and o is a (deterministic) transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_k)$ to state q and a tuple of actions $\langle \alpha_1, \dots, \alpha_k \rangle$ that can be executed by Agt in q . We write $d_a(q)$ instead of $d(a, q)$. Every $\sim_a \subseteq St \times St$ is an epistemic equivalence relation. The CEGS is assumed to be *uniform*, in the sense that $q \sim_a q'$ implies $d_a(q) = d_a(q')$.

EXAMPLE 1. Consider a very simple voting scenario with two agents: the voter v and the coercer c . The voter casts a vote for a selected candidate $i \in \{1, \dots, n\}$ (action vote_i). Upon exit from the polling station, the voter can hand in a proof of how she voted to the coercer (action give) or refuse to hand in the proof (action ng). The proof may be a certified receipt from the election authorities, a picture of the ballot taken with a smartphone, etc. After that, the coercer can either punish the voter (pun) or not punish (np).

Appears in: Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017), S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.
Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.




Figure 1: A simple model of voting and coercion

The CEGS M_{vote} modeling the scenario for $n = 2$ is shown in Figure 1. Proposition $vote_i$ labels states where the voter has already voted for candidate i . Proposition pun indicates states where v has been punished. The indistinguishability relation for the coercer is depicted by dotted lines.

A strategy of agent $a \in \text{Agt}$ is a conditional plan that specifies what a is going to do in every possible situation. Formally, a perfect information memoryless strategy for a can be represented by a function $s_a : St \rightarrow Act$ satisfying $s_a(q) \in d_a(q)$ for each $q \in St$. An imperfect information memoryless strategy additionally satisfies that $s_a(q) = s_a(q')$ whenever $q \sim_a q'$. Following [29], we refer to the former as Ir-strategies, and to the latter as ir-strategies.

A collective x -strategy s_A , for $A \subseteq \text{Agt}$ and $x \in \{\text{Ir}, \text{ir}\}$, is a tuple of individual x -strategies, one per agent from A . The set of all such strategies is denoted by Σ_A^x . By $s_A|_a$ we denote the strategy of agent $a \in A$ selected from s_A .

Given two partial functions $f, f' : X \rightarrow Y$, we say that f' extends f (denoted $f \subseteq f'$) if, whenever $f(x)$ is defined, we have $f(x) = f'(x)$. A partial function $s'_a : St \rightarrow Act$ is called a partial x -strategy for a if s'_a is extended by some strategy $s_a \in \Sigma_a^x$. A collective partial x -strategy s_A is a tuple of partial x -strategies, one per agent from A .

A path $\lambda = q_0q_1q_2\dots$ is an infinite sequence of states such that there is a transition between each q_i, q_{i+1} . We use $\lambda[i]$ to denote the i th position on path λ (starting from $i = 0$). Function $out(q, s_A)$ returns the set of all paths that can result from the execution of strategy s_A from state q . We will sometimes write $out^{\text{Ir}}(q, s_A)$ instead of $out(q, s_A)$. Moreover, function $out^{\text{Ir}}(q, s_A) = \bigcup_{a \in A} \bigcup_{q \sim_a q'} out(q', s_A)$ collects all the outcome paths that start from states that are indistinguishable from q to at least one agent in A .

2.2 Alternating-Time Temporal Logic

We use a variant of ATL that explicitly distinguishes between perfect and imperfect information abilities. Formally,

the syntax is defined by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle A \rangle\rangle_x \mathbf{X}\varphi \mid \langle\langle A \rangle\rangle_x \mathbf{G}\varphi \mid \langle\langle A \rangle\rangle_x \varphi \mathbf{U} \varphi,$$

where $x \in \{\text{Ir}, \text{ir}\}$, $p \in \text{Props}$ and $A \subseteq \text{Agt}$. We read $\langle\langle A \rangle\rangle_{\text{ir}} \gamma$ as “ A can identify and execute a strategy that enforces γ ,” \mathbf{X} as “in the next state,” \mathbf{G} as “now and always in the future,” and \mathbf{U} as “until.” $\langle\langle A \rangle\rangle_{\text{Ir}} \gamma$ can be read as “ A might be able to bring about γ if allowed to make lucky guesses along the way.” We focus on the kind of ability expressed by $\langle\langle A \rangle\rangle_{\text{ir}}$. The other strategic modality (i.e., $\langle\langle A \rangle\rangle_{\text{Ir}}$) will prove useful when approximating $\langle\langle A \rangle\rangle_{\text{ir}}$.

The semantics of ATL can be defined as follows:

- $M, q \models p$ iff $q \in V(p)$,
- $M, q \models \neg\varphi$ iff $M, q \not\models \varphi$,
- $M, q \models \varphi \wedge \psi$ iff $M, q \models \varphi$ and $M, q \models \psi$,
- $M, q \models \langle\langle A \rangle\rangle_x \mathbf{X}\varphi$ iff there exists $s_A \in \Sigma_A^x$ such that for all $\lambda \in out^x(q, s_A)$ we have $M, \lambda[1] \models \varphi$,
- $M, q \models \langle\langle A \rangle\rangle_x \mathbf{G}\varphi$ iff there exists $s_A \in \Sigma_A^x$ such that for all $\lambda \in out^x(q, s_A)$ and $i \in \mathbb{N}$ we have $M, \lambda[i] \models \varphi$,
- $M, q \models \langle\langle A \rangle\rangle_x \psi \mathbf{U} \varphi$ iff there exists $s_A \in \Sigma_A^x$ such that for all $\lambda \in out^x(q, s_A)$ there is $i \in \mathbb{N}$ for which $M, \lambda[i] \models \varphi$ and $M, \lambda[j] \models \psi$ for all $0 \leq j < i$.

We will often write $\langle A \rangle \varphi$ instead of $\langle\langle A \rangle\rangle_{\text{ir}} \mathbf{X}\varphi$ to express one-step abilities under imperfect information. Additionally, we define “now or sometime in the future” as $\mathbf{F}\varphi \equiv \top \mathbf{U} \varphi$.

EXAMPLE 2. Consider model M_{vote} from Example 1. The following formula expresses that the coercer can ensure that the voter will eventually either have voted for candidate i (presumably chosen by the coercer for the voter to vote for) or be punished: $\langle\langle c \rangle\rangle_{\text{ir}} \mathbf{F}(\neg pun \rightarrow vote_i)$. We note that it holds in M_{vote}, q_0 for any $i = 1, 2$. A strategy for c that validates the property is $s_c(q_3) = np$, $s_c(q_4) = s_c(q_5) = s_c(q_6) = pun$ for $i = 1$, and symmetrically for $i = 2$.

Consequently, the formula $\langle\langle v \rangle\rangle_{\text{ir}} \mathbf{G}(\neg pun \wedge \neg vote_i)$ saying that the voter can avoid voting for candidate i and being punished, is false in M_{vote}, q_0 for all $i = 1, 2$.

We refer to the syntactic fragment containing only $\langle\langle A \rangle\rangle_{\text{ir}}$ modalities as \mathbf{ATL}_{ir} , and to the one containing only $\langle\langle A \rangle\rangle_{\text{Ir}}$ modalities as \mathbf{ATL}_{Ir} .

PROPOSITION 1 ([3, 29, 20]). Model checking \mathbf{ATL}_{Ir} is P-complete and can be done in time $O(|M| \cdot |\varphi|)$ where $|M|$ is the number of transitions in the model and $|\varphi|$ is the length of the formula.

Model checking \mathbf{ATL}_{ir} is Δ_2^P -complete wrt $|M|$ and $|\varphi|$.

REMARK 2. The semantics of $\langle\langle A \rangle\rangle_{\text{ir}} \gamma$ encodes the notion of “subjective” ability [29, 21]: the agents must have a successful strategy from all the states that they consider possible when the system is in state q . Then, they know that the strategy indeed obtains γ . The alternative notion of “objective” ability [7] requires a winning strategy from state q alone. We focus on the subjective interpretation, as it is more standard in ATL and more relevant in game solving (think of a card game, such as poker or bridge: the challenge is to find a strategy that wins for all possible hands of the opponents).

Note that if $[q]_{\sim_A} = \{q\}$ and γ contains no nested strategic modalities, then the subjective and objective semantics of $\langle\langle A \rangle\rangle_{\text{ir}} \gamma$ at q coincide. Moreover, model checking $\langle\langle A \rangle\rangle_{\text{ir}} p_1 \mathbf{U} p_2$

and $\langle\langle A \rangle\rangle_{\text{ir}} \mathbf{Gp}$ in M, q according to the objective semantics can be easily reduced to the subjective case by adding a spurious initial state q' , with transitions to all states in $[q]_{\sim_A^E}$, controlled by a “dummy” agent outside A [27].

2.3 Reasoning about Knowledge

Having indistinguishability relations in the models, we can interpret knowledge modalities K_a in the standard way:

- $M, q \models K_a \varphi$ iff $M, q' \models \varphi$ for all q such that $q \sim_a q'$.

The semantics of “everybody knows” (E_A) and *common knowledge* (C_A) is defined analogously by assuming the relation $\sim_A^E = \bigcup_{a \in A} \sim_a$ to aggregate individual uncertainty in A , and \sim_A^C to be the transitive closure of \sim_A^E . Additionally, we take \sim_\emptyset^E to be the minimal reflexive relation. We also use $[q]_{\mathcal{R}} = \{q' \mid q \mathcal{R} q'\}$ to denote the image of q wrt relation \mathcal{R} .

EXAMPLE 3. The following formulae hold in M_{vote}, q_0 for any $i = 1, 2$ by virtue of strategy s_c presented in Example 2:

- $\langle\langle c \rangle\rangle_{\text{ir}} \mathbf{F}((\neg K_c \text{vote}_i) \rightarrow \text{pun})$: The coercer has a strategy so that, eventually, the voter is punished unless the coercer has learnt that the voter voted as instructed;
- $\langle\langle c \rangle\rangle_{\text{ir}} \mathbf{G}((K_c \text{vote}_i) \rightarrow \neg \text{pun})$: Moreover, the coercer can guarantee that if he learns that the voter obeyed, then the voter will not be punished.

2.4 Alternating Epistemic Mu-Calculus

It is well known that the modalities in \mathbf{ATL}_{ir} have simple fixpoint characterizations [3], and hence \mathbf{ATL}_{ir} can be embedded in a variant of μ -calculus with $\langle\langle A \rangle\rangle_{\text{ir}} \mathbf{X}$ as the basic modality. At the same time, the analogous variant of μ -calculus for imperfect information has incomparable expressive power to \mathbf{ATL}_{ir} [6], which suggests that, under imperfect information, \mathbf{ATL} and fixpoint specifications provide different views of strategic ability.

Formally, *alternating epistemic μ -calculus* ($\mathbf{AE}\mu\mathbf{C}$) takes the next-time fragment of \mathbf{ATL}_{ir} , possibly with epistemic modalities, and adds the least fixpoint operator μ . The greatest fixpoint operator ν is defined as dual. Let \mathcal{Vars} be a set of second-order variables ranging over 2^{St} . The language of $\mathbf{AE}\mu\mathbf{C}$ is defined by the following grammar:

$$\varphi ::= p \mid Z \mid \neg \varphi \mid \varphi \vee \psi \mid \langle A \rangle \varphi \mid \mu Z(\varphi) \mid K_a,$$

where $p \in \text{Props}$, $Z \in \text{Vars}$, $a \in \text{Agt}$, $A \subseteq \text{Agt}$, and the formulae are Z -positive, i.e., each free occurrence of Z is in the scope of an even number of negations. We define $\nu Z(\varphi(Z)) \equiv \neg \mu Z(\neg \varphi(\neg Z))$. A formula of $\mathbf{AE}\mu\mathbf{C}$ is *alternation-free* if in its negation normal form it contains no occurrences of ν (resp. μ) on any syntactic path from an occurrence of μZ (resp. νZ) to a bound occurrence of Z .

The denotational semantics of $\mathbf{af}\text{-}\mathbf{AE}\mu\mathbf{C}$ (i.e., the alternation-free fragment of $\mathbf{AE}\mu\mathbf{C}$) assigns to each formula φ the set of states $[\varphi]_{\mathcal{V}}^M$ where φ is true under the valuation $\mathcal{V} \in \mathcal{Vals}$:

- $[p]_{\mathcal{V}}^M = V(p)$, $[Z]_{\mathcal{V}}^M = \mathcal{V}(Z)$,
- $[\neg \varphi]_{\mathcal{V}}^M = St \setminus [\varphi]_{\mathcal{V}}^M$,
- $[\varphi \vee \psi]_{\mathcal{V}}^M = [\varphi]_{\mathcal{V}}^M \cup [\psi]_{\mathcal{V}}^M$,
- $[\langle A \rangle \varphi]_{\mathcal{V}}^M = \{q \in St \mid \exists s_A \in \Sigma_A \forall \lambda \in \text{out}_M^{\text{tr}}(q, s_A) \lambda[1] \in [\varphi]_{\mathcal{V}}^M\}$,
- $[\mu Z(\varphi)]_{\mathcal{V}}^M = \bigcap \{Q \subseteq St \mid [\varphi]_{\mathcal{V}[Z:=Q]}^M \subseteq Q\}$,

- $[\langle K_a \varphi \rangle]_{\mathcal{V}}^M = \{q \in St \mid \forall q' (q' \sim_a q \text{ implies } q' \in [\varphi]_{\mathcal{V}}^M)\}$.

If φ contains no free variables, then its validity does not depend on \mathcal{V} , and we write $M, q \models \varphi$ instead of $q \in [\varphi]_{\mathcal{V}}^M$.

EXAMPLE 4. Consider the $\mathbf{AE}\mu\mathbf{C}$ formula $\mu Z.((\neg \text{pun} \rightarrow \text{vote}_i) \vee \langle c \rangle Z)$, i.e., the “naive” fixpoint translation of the formula $\langle\langle c \rangle\rangle_{\text{ir}} \mathbf{F}(\neg \text{pun} \rightarrow \text{vote}_i)$ from Example 2. The fixpoint computation produces the whole set of states St . Thus, in particular, $M_{\text{vote}}, q_0 \models \mu Z.((\neg \text{pun} \rightarrow \text{vote}_i) \vee \langle c \rangle Z)$.

PROPOSITION 3 ([6]). *Model checking $\mathbf{af}\text{-}\mathbf{AE}\mu\mathbf{C}$ with strategic modalities for up to 2 agents is \mathbf{P} -complete and can be done in time $O(|\sim| \cdot |\varphi|)$ where $|\sim|$ is the size of the largest equivalence class among \sim_1, \dots, \sim_k , and $|\varphi|$ is the length of the formula.*

For coalitions of size at least 3, the problem is between \mathbf{NP} and $\Delta_2^{\mathbf{P}}$ wrt $|\sim|$ and $|\varphi|$.

Thus, alternation-free alternating epistemic μ -calculus can be an attractive alternative to \mathbf{ATL}_{ir} from the complexity point of view. Unfortunately, formulae of \mathbf{ATL}_{ir} admit no universal translations to $\mathbf{af}\text{-}\mathbf{AE}\mu\mathbf{C}$. Formally, it was proved in [6, Proposition 6] that $\mathbf{af}\text{-}\mathbf{AE}\mu\mathbf{C}$ does not cover the expressive power of \mathbf{ATL}_{ir} . The proof uses formulae of type $\langle\langle a \rangle\rangle \mathbf{Fp}$, but it is easy to construct an analogous argument for $\langle\langle a \rangle\rangle \mathbf{Gp}$. In consequence, long-term strategic modalities of \mathbf{ATL}_{ir} do not have alternation-free fixpoint characterizations in terms of the next-step strategic modalities $\langle A \rangle$. A similar result was proved for \mathbf{ATL}_{IR} in [11, Theorem 11].

3. LOWER BOUNDS FOR ABILITIES

The complexity of $\mathbf{AE}\mu\mathbf{C}$ model checking seems more attractive than that of \mathbf{ATL}_{ir} . Unfortunately, the expressivity results cited in Section 2.4 imply that there is no simple fixpoint translation which captures *exactly* the meaning of \mathbf{ATL}_{ir} operators. It might be possible, however, to come up with a translation tr that provides a *lower bound* of the actual strategic abilities, i.e., such that $M, q \models tr(\langle\langle A \rangle\rangle_{\text{ir}} \gamma)$ implies $M, q \models \langle\langle A \rangle\rangle_{\text{ir}} \gamma$. In other words, a translation which can only reduce, but never enhance the abilities of the coalition.

We begin by investigating the “naive” fixpoint translation that mimics the one for \mathbf{ATL}_{ir} , and show that it works in some cases, but not in general. Then, we propose how to alter the semantics of the nexttime modality so that a general lower bound can be obtained. We focus first on reachability goals, expressed by formulae $\langle\langle A \rangle\rangle_{\text{ir}} \mathbf{F}\varphi$, and then move on to the other modalities.

3.1 Trying It Simple for Reachability Goals

We assume from now on that φ is a formula of \mathbf{ATL}_{ir} , M is a CEGS, and q is a state in M (unless explicitly stated otherwise). We start with the simplest translation, analogous to that of [3]: $tr_1(\langle\langle A \rangle\rangle_{\text{ir}} \mathbf{F}\varphi) = \mu Z.(\varphi \vee \langle A \rangle Z)$. Unfortunately, this translation provides neither a lower nor an upper bound. For the former, use model M_0 in Figure 2, and observe that $M_0, q_0 \models \mu Z.(\text{p} \vee \langle 1 \rangle Z)$ but $M_0, q_0 \not\models \langle\langle 1 \rangle\rangle_{\text{ir}} \mathbf{Fp}$. For the latter, take model M in [6, Figure 1], and observe that $M, q_0 \models \langle\langle 1 \rangle\rangle_{\text{ir}} \mathbf{Fp}$ but $M, q_0 \not\models \mu Z.(\text{p} \vee \langle 1 \rangle Z)$.

PROPOSITION 4. $M, q \models \mu Z.(\varphi \vee \langle A \rangle Z)$ does not imply $M, q \models \langle\langle A \rangle\rangle_{\text{ir}} \mathbf{F}\varphi$. The converse implication does not hold either.




Figure 2: CEGS M_0 : a counterexample for tr_1




Figure 3: M_1 : a counterexample for tr_2

Consider now a slightly stronger fixpoint specification: $tr_2(\langle\langle A \rangle\rangle_{ir}, \mathbf{F}\varphi) = \mu Z.(E_A\varphi \vee \langle A \rangle Z)$. This new translation works to an extent, as the following proposition shows.

PROPOSITION 5.

1. $M, q \models \mu Z.(E_{\emptyset}\varphi \vee \langle \emptyset \rangle Z)$ iff $M, q \models \langle\langle \emptyset \rangle\rangle_{ir}, \mathbf{F}\varphi$;
2. If $|A| = 1$, then $M, q \models \mu Z.(E_A\varphi \vee \langle A \rangle Z)$ implies $M, q \models \langle\langle A \rangle\rangle_{ir}, \mathbf{F}\varphi$, but the converse does not hold;¹
3. If $|A| > 1$, then $M, q \models \mu Z.(E_A\varphi \vee \langle A \rangle Z)$ does not imply $M, q \models \langle\langle A \rangle\rangle_{ir}, \mathbf{F}\varphi$, and vice versa.

PROOF. **Case 1:** follows from the fact that for the empty coalition the ir-reachability is equivalent to the IR-reachability, which in turn has a fixpoint characterization.

Case 2: Let us assume that $A = \{a\}$ for some $a \in \text{Agt}$. We define the sequence $\{F_j\}_{j \in \mathbb{N}}$ of **af-AE μ C** formulae s.t. $F_0 = K_a\varphi$ and $F_{j+1} = F_0 \vee \langle a \rangle F_j$, for all $j \geq 0$. From Kleene fixed-point theorem we have $\llbracket \mu Z.(K_a\varphi \vee \langle a \rangle Z) \rrbracket = \bigcup_{j=0}^{\infty} \llbracket F_j \rrbracket$, and $\{\llbracket F_j \rrbracket\}_{j \in \mathbb{N}}$ is a non-decreasing monotone sequence of subsets of St . Now, we prove that for each $j \in \mathbb{N}$ there exists a partial strategy s_a^j s.t. $\text{dom}(s_a^j) = \llbracket F_j \rrbracket$, $\forall q \in \text{dom}(s_a^j) \forall \lambda \in \text{out}^{ir}(q, s_a^j) \exists k \leq j \lambda[k] \models \varphi$, and $s_a^j \subseteq s_a^{j+1}$. The proof is by induction on j . We constructively build s_a^{j+1} from s_a^j for each $j \in \mathbb{N}$. The base case is trivial. For the inductive step, firstly observe that for each $j \in \mathbb{N}$ if $q \in \llbracket F_j \rrbracket$, then $[q]_{\sim_a} \subseteq \llbracket F_j \rrbracket$. As \sim_a is an equivalence relation, for each $q \in \llbracket F_{j+1} \rrbracket$ either $[q]_{\sim_a} \subseteq \llbracket F_j \rrbracket$ or $[q]_{\sim_a} \subseteq \llbracket F_{j+1} \rrbracket \setminus \llbracket F_j \rrbracket$. In the first case we put $s_a^{j+1}(q) = s_a^j(q)$. In the second case, we know that there exists a strategy s_a^q s.t. $\forall \lambda \in \text{out}^{ir}(q, s_a^q) \lambda[1] \in \llbracket F_j \rrbracket$. We thus put $s_a^{j+1}(q') = s_a^q(q')$ for all $q' \in [q]_{\sim_a}$, which concludes the inductive proof.

We finally define the partial strategy $s_a = \bigcup_{j \in \mathbb{N}} s_a^j$. For each $q \in St$ s.t. $M, q \models \mu Z.(K_a\varphi \vee \langle a \rangle Z)$, either $M, q \models \varphi$,

¹Note that, for $A = \{a\}$, $E_A\varphi$ is equivalent to $K_a\varphi$.




Figure 4: Lower bounds are not tight: (A) M_2 ; (B) M_3

or φ is reached along each path consistent with any extension of s_a to a full strategy.

For the converse implication, take model M in [6, Figure 1], and observe that $M, q_0 \models \langle\langle 1 \rangle\rangle_{ir}, \mathbf{F}p$ but $M, q_0 \not\models \mu Z.(K_{1p} \vee \langle 1 \rangle Z)$.

Case 3: Consider the CEGS M_1 presented in Figure 3. We assume that $d_1(q) = \{a, b\}$ and $d_2(q) = \{x, y\}$, for $q \in \{q_1, q_2, q_3, q_4\}$. In the remaining states the protocols allow only one action. For clarity, we omit from the figure the transitions leaving the states q_1, q_2, q_3 , and q_4 , leading to state *sink*. Assume now $\varphi \equiv p$. Note that $M, q_0 \models \mu Z.(E_{\{1,2\}}\varphi \vee \langle \{1,2\} \rangle Z)$ and $M, q_0 \not\models \langle\langle 1,2 \rangle\rangle_{ir}, \mathbf{F}\varphi$. For larger coalitions A , we extend the model with a sufficient number of spurious (idle) agents.

For the other direction, use the counterexample from Case 2, extended with appropriately many spurious agents. \square

As Propositions 4 and 5 show, translation tr_2 provides lower bounds for **ATL**_{ir} verification only in a limited number of instances. Also, the bound is rather loose, as the following example demonstrates.

EXAMPLE 5. Consider the single-agent CEGS M_2 presented in Figure 4A. The sole available strategy, in which agent 1 selects always action a , enforces eventually reaching p , i.e., $M_2, q_0 \models \langle\langle 1 \rangle\rangle_{ir}, \mathbf{F}p$. On the other hand, $M_2, q_0 \not\models \mu Z.(K_{1p} \vee \langle 1 \rangle Z)$. This is because the next-step operator in **ATL**_{ir} requires reaching p simultaneously from all the states indistinguishable from q_0 , whereas p is reached from q_0, q_1 in one and two steps, respectively.

3.2 Steadfast Next Step Operator

To obtain a tighter lower bound, and one that works universally, we introduce a new modality. $\langle A \rangle^\bullet$ can be seen as a semantic variant of the next-step ability operator $\langle A \rangle$ where: (i) agents in A look for a short-term strategy that succeeds from the “common knowledge” neighborhood of the initial state (rather than in the “everybody knows” neighborhood), and (ii) they are allowed to “steadfastly” pursue their goal in a variable number of steps within the indistinguishability class. In this section, we propose the semantics of $\langle A \rangle^\bullet$ and show how to revise the lower bound. Some additional insights are provided in Section 4.

We begin by defining the auxiliary function *Reach* so that $q \in \text{Reach}_M(s_A, Q, \varphi)$ collects all $q \in Q$ such that all the paths executing s_A from q eventually reach φ without leaving Q , except possibly for the last step:

$$\text{Reach}_M(s_A, Q, \varphi) = \{q \in Q \mid \forall \lambda \in \text{out}(q, s_A) \exists i. M, \lambda[i] \models \varphi \text{ and } \forall 0 \leq j < i. \lambda[j] \in Q\}.$$

The *steadfast next-step operator* $\langle A \rangle^\bullet$ is defined as follows:

- $M, q \models \langle A \rangle^\bullet \varphi$ iff there exists $s_A \in \Sigma_A^{\text{ir}}$ such that $\text{Reach}_M(s_A, [q]_{\sim_C^A}, \varphi) = [q]_{\sim_C^A}$.

Now we can propose our ultimate attempt at the lower bound for reachability goals: $\text{tr}_3(\langle\langle A \rangle\rangle_{\text{ir}} \mathbf{F}\varphi) = \mu Z.(E_A \varphi \vee \langle A \rangle^\bullet Z)$, with the following result.

PROPOSITION 6. *If $M, q \models \mu Z.(E_A \varphi \vee \langle A \rangle^\bullet Z)$, then $M, q \models \langle\langle A \rangle\rangle_{\text{ir}} \mathbf{F}\varphi$. The converse does not universally hold.*

PROOF. The proof is similar to the proof of Proposition 5. As previously, we define a sequence $\{F_j\}_{j \in \mathbb{N}}$ of **af-AE** $\mu\mathbf{C}$ formulae s.t. $F_0 = E_A \varphi$ and $F_{j+1} = F_0 \vee \langle A \rangle^\bullet F_j$, for all $j \geq 0$. We also use a sequence $\{H_j\}_{j \in \mathbb{N}}$ with $H_j = \langle A \rangle^\bullet F_j$. From Kleene fixed-point theorem we have $\llbracket \mu Z.(E_A \varphi \vee \langle A \rangle^\bullet Z) \rrbracket = \bigcup_{j=0}^{\infty} \llbracket F_j \rrbracket = \llbracket F_0 \rrbracket \cup \bigcup_{j=0}^{\infty} \llbracket H_j \rrbracket$. Observe that, as \sim_C^A is an equivalence relation, we have for each $q \in \text{St}$ and $j \in \mathbb{N}$ that if $[q]_{\sim_C^A} \cap \llbracket H_j \rrbracket \neq \emptyset$, then $[q]_{\sim_C^A} \subseteq \llbracket H_j \rrbracket$.

We prove that for each $j \in \mathbb{N}$ there exists a partial strategy s_A^j s.t. $\text{dom}(s_A^j) = \llbracket H_j \rrbracket$, $\forall q \in \text{dom}(s_A^j) \forall \lambda \in \text{out}^{\text{ir}}(q, s_A^j) \exists k \in \mathbb{N} \lambda[k] \models E_A \varphi$, and $s_A^j \subseteq s_A^{j+1}$. The proof is by induction on j . In the base case of $H_0 = \langle A \rangle^\bullet E_A \varphi$ observe that if $q \in \llbracket H_0 \rrbracket$ then there exists a partial strategy $s_A^{0,q}$ with $\text{dom}(s_A^{0,q}) = [q]_{\sim_C^A}$ s.t. every $\lambda \in \text{out}^{\text{ir}}(q, s_A^{0,q})$ stays in $[q]_{\sim_C^A}$ until it reaches a state where $E_A \varphi$ holds. We can now define $s_A^0 = \bigcup_{[q]_{\sim_C^A} \in \text{St}/\sim_C^A} s_A^{0,q}$ which is uniform, and reaches $E_A \varphi$ on all execution paths. For the inductive step, we divide the construction of s_A^{j+1} in two cases. Firstly, if $q \in \llbracket H_j \rrbracket$, then we put $s_A^{j+1}(q) = s_A^j(q)$. Secondly, let $q \in \llbracket H_{j+1} \rrbracket \setminus \llbracket H_j \rrbracket$. In this case there exists a partial strategy $s_A^{j+1,q}$ with $\text{dom}(s_A^{j+1,q}) = [q]_{\sim_C^A}$ s.t. each outcome $\lambda \in \text{out}^{\text{ir}}(q, s_A^{j+1,q})$ stays in $[q]_{\sim_C^A}$ until it reaches a state q' s.t. either $q' \models E_A \varphi$ or $q' \in \llbracket H_j \rrbracket$. In the latter, from the inductive assumption we know that following s_A^{j+1} always leads to reaching $E_A \varphi$ without leaving $\llbracket H_j \rrbracket$. We thus take $s_A^{j+1} = \bigcup_{[q]_{\sim_C^A} \in \text{St}/\sim_C^A} s_A^{j+1,q}$ which, again, is uniform, and reaches $E_A \varphi$ on all execution paths. This concludes the inductive part of the proof.

Finally, we build a partial strategy $s_A = \bigcup_{j \in \mathbb{N}} s_A^j$, whose any extension is s.t. for each $q \in \text{St}$, if $M, q \models \mu Z.(E_A \varphi \vee \langle A \rangle^\bullet Z)$, then a state in which $E_A \varphi$ holds is eventually reached along each outcome path $\lambda \in \text{out}^{\text{ir}}(q, s_A)$. This concludes the proof of the implication.

To see that the converse does not hold, consider model M_3 in Figure 4B. We have that $M_3, q_0 \models \langle\langle 1 \rangle\rangle_{\text{ir}} \mathbf{F}\mathbf{p}$, but $M_3, q_0 \not\models \mu Z.(K_1 \mathbf{p} \vee \langle 1 \rangle^\bullet Z)$. \square

Thus, tr_3 indeed provides a universal lower bound for reachability goals expressed in **ATL**_{ir}.

3.3 Lower Bounds for “Always” and “Until”

So far, we have concentrated on reachability goals. We now extend the main result to all the modalities of **ATL**_{ir}:

THEOREM 7.

1. If $M, q \models \nu Z.(C_A \varphi \wedge \langle A \rangle^\bullet Z)$, then $M, q \models \langle\langle A \rangle\rangle_{\text{ir}} \mathbf{G}\varphi$;
2. If $M, q \models \mu Z.(E_A \varphi \vee (C_A \psi \wedge \langle A \rangle^\bullet Z))$, then $M, q \models \langle\langle A \rangle\rangle_{\text{ir}} \psi \mathbf{U}\varphi$.

PROOF. **Case 1:** Let us define the sequence $\{G_j\}_{j \in \mathbb{N}}$ of formulae s.t. $G_0 = C_A \varphi$ and $G_{j+1} = G_0 \wedge \langle A \rangle^\bullet G_j$, for all $j \geq$

0. From Kleene fixed-point theorem, $\llbracket \nu Z.(C_A \varphi \wedge \langle A \rangle^\bullet Z) \rrbracket = \bigcap_{j=0}^{\infty} \llbracket G_j \rrbracket$. It suffices to prove that for each $j \in \mathbb{N}$ there exists a strategy s_A^j s.t. $\forall q \in \llbracket G_j \rrbracket \forall \lambda \in \text{out}^{\text{ir}}(q, s_A^j) \forall 0 \leq k \leq j \lambda[k] \models \varphi$. The proof is by induction on j , with the trivial base case. Assume that the inductive assumption holds for some $j \in \mathbb{N}$. From the definition of the steadfast next-step operator we can define for each equivalence class $[q]_{\sim_C^A} \in \llbracket G_{j+1} \rrbracket / \sim_C^A$ a partial strategy $s_A^{q,j+1}$ s.t. $\forall q' \in [q]_{\sim_C^A} \forall \lambda \in \text{out}^{\text{ir}}(q, s_A^{q,j+1}) \lambda[1] \in \llbracket G_j \rrbracket$. We now construct

$$s_A^{j+1} = \bigcup_{[q]_{\sim_C^A} \in \llbracket G_{j+1} \rrbracket / \sim_C^A} s_A^{q,j+1} \cup s_A^j \upharpoonright_{\llbracket C_A \varphi \rrbracket \setminus \llbracket G_j \rrbracket}.$$

Intuitively, s_A^j enforces that a path leaving each $q \in \llbracket G_{j+1} \rrbracket$ stays within $\llbracket C_A \varphi \rrbracket$ for at least j steps. Moreover, $s_A^j \subseteq s_A^{j+1}$ for all j . Thus, $s_A = \bigcup_{j \in \mathbb{N}} s_A^j$ enforces that a path leaving each $q \in \bigcup_{j \in \mathbb{N}} \llbracket G_j \rrbracket$ stays within $\llbracket C_A \varphi \rrbracket$ for infinitely many steps, which concludes the proof. Note that the correctness of the construction relies the fact that \sim_C^A is an equivalence relation.

Case 2: analogous to Proposition 6. \square

4. DISCUSSION & PROPERTIES

Theorem 7 shows that $\text{tr}_3(\varphi)$ provides a correct lower bound of the value of φ for all formulae of **ATL**_{ir}. In this section, we discuss the tightness of the approximation from the theoretical point of view. An empirical evaluation will be presented in Section 6.

4.1 Comparing tr_2 and tr_3 for Reachability Goals

Translation tr_3 updates tr_2 by replacing the standard next-step ability operator $\langle A \rangle$ with the “steadfast next-step ability” $\langle A \rangle^\bullet$. The difference between $\langle A \rangle \varphi$ and $\langle A \rangle^\bullet \varphi$ is twofold. First, $\langle A \rangle \varphi$ looks for a winning strategy in the “everybody knows” neighborhood of a given state (i.e., $[q]_{\sim_E^A}$), whereas $\langle A \rangle^\bullet \varphi$ looks at the “common knowledge” neighborhood (i.e., $[q]_{\sim_C^A}$). Secondly, $\langle A \rangle^\bullet$ allows to “zig-zag” across $[q]_{\sim_C^A}$ until a state satisfying φ is found.

Actually, the first change would suffice to provide a universally correct lower bound for **ATL**_{ir}. The second update makes it *more useful* in models where agents may not see the occurrence of some action, such as M_2 of Figure 4A. To see this formally, we show that tr_3 provides a strictly tighter approximation than tr_2 on singleton coalitions:

PROPOSITION 8. *For $A = \{a\}$, if $M, q \models \mu Z.(K_a \varphi \vee \langle a \rangle Z)$, then $M, q \models \mu Z.(K_a \varphi \vee \langle a \rangle^\bullet Z)$. The converse does not universally hold.*

PROOF. It suffices to observe that $M, q \models \langle a \rangle \varphi$ implies $M, q \models \langle a \rangle^\bullet \varphi$, for any $\varphi \in \mathbf{af-AE}\mu\mathbf{C}$. Note that this is true only for single-agent coalitions. For the converse, notice that in CEGS M_2 from Figure 4A we have $M_2, q_0 \models \mu Z.(K_1 \mathbf{p} \vee \langle 1 \rangle^\bullet Z)$ and $M_2, q_0 \not\models \mu Z.(K_1 \mathbf{p} \vee \langle 1 \rangle Z)$. \square

On the other hand, if agent a always sees whenever an action occurs, then tr_2 and tr_3 coincide for a ’s abilities. Formally, let us call CEGS M *lockstep for a* if, whenever there is a transition from q to q' in M , we have $q \not\sim_a q'$. The following is straightforward.

PROPOSITION 9. *If M is lockstep for a , then $M, q \models \langle a \rangle \varphi$ iff $M, q \models \langle a \rangle^\bullet \varphi$. In consequence, $M, q \models \text{tr}_2(\langle\langle a \rangle\rangle_{\text{ir}} \mathbf{F}\varphi)$ iff $M, q \models \text{tr}_3(\langle\langle a \rangle\rangle_{\text{ir}} \mathbf{F}\varphi)$.*

4.2 When is the Lower Bound Tight?

An interesting question is: what is the subclass of CEGS's for which tr_3 is tight, i.e., the answer given by the approximation is exact? We address the question only partially here. In fact, we characterize a subclass of CEGS's for which tr_3 is certainly *not* tight, by the necessary condition below.

Let $\gamma \equiv \mathbf{G}\psi$ or $\gamma \equiv \psi_1 \mathbf{U} \psi_2$ for some $\psi, \psi_1, \psi_2 \in \mathbf{ATL}_{ir}$. We say that strategy $s_A \in \Sigma_A^{ir}$ is *winning* for γ from q if it obtains γ for all paths in $out^{ir}(q, s_A)$. Moreover, for such s_A , let $RR(q, s_A, \gamma)$ be the *relevant reachable states* of s_A in the context of γ , defined as follows: $RR(q, s_A, \mathbf{G}\psi)$ is the set of states that occur anywhere in $out^{ir}(q, s_A)$; $RR(q, s_A, \psi_1 \mathbf{U} \psi_2)$ is the set of states that occur anywhere in $out^{ir}(q, s_A)$ before the first occurrence of ψ_2 .

PROPOSITION 10. *Let M be a CEGS, $q \in St_M$, and $\varphi \equiv \langle\langle A \rangle\rangle_{ir} \gamma$. Furthermore, suppose that φ and $tr_3(\varphi)$ are either both true or both false in M, q . Then:*

1. *either no strategy $s_A \in \Sigma_A^{ir}$ is winning for γ from q , or*
2. *there is a strategy $s_A \in \Sigma_A^{ir}$ which is winning for γ from every $q' \in RR(q, s_A, \gamma)$.*

Conversely, the approximation is *not* tight if there are winning strategies, but each of them reaches a intermediate state q' from which no winning substrategy can be computed. This can only happen if some states in $[q]_{\sim_A}^E$ are not reachable by s_A . In consequence, the agents in A forget relevant information that comes alone from the fact that they are executing s_A . We will use Proposition 10 in Section 6 to show that the few benchmarks existing in the literature are not amenable to our approximations.

5. APPROXIMATION SEMANTICS FOR \mathbf{ATL}_{ir}

Note that $M, q \models \langle\langle A \rangle\rangle_{ir} \gamma$ always implies $M, q \models E_A \langle\langle A \rangle\rangle_{ir} \gamma$. Based on this, and the lower bounds established in Theorem 7, we propose the *lower approximation* tr and the *upper approximation* TR for \mathbf{ATL}_{ir} as follows:

$$\begin{aligned} tr(p) &= p, & tr(\neg\varphi) &= \neg TR(\varphi), & tr(\varphi \wedge \psi) &= tr(\varphi) \wedge tr(\psi), \\ tr(\langle A \rangle \varphi) &= \langle A \rangle tr(\varphi), \\ tr(\langle\langle A \rangle\rangle_{ir} \mathbf{G}\varphi) &= \nu Z. (CAtr(\varphi) \wedge \langle A \rangle \bullet Z), \\ tr(\langle\langle A \rangle\rangle_{ir} \psi \mathbf{U} \varphi) &= \mu Z. (EAttr(\varphi) \vee (CAtr(\psi) \wedge \langle A \rangle \bullet Z)). \end{aligned}$$

$$\begin{aligned} TR(p) &= p, & TR(\neg\varphi) &= \neg tr(\varphi), \\ TR(\varphi \wedge \psi) &= TR(\varphi) \wedge TR(\psi), \\ TR(\langle A \rangle \varphi) &= E_A \langle\langle A \rangle\rangle_{ir} \mathbf{X} TR(\varphi), \\ TR(\langle\langle A \rangle\rangle_{ir} \mathbf{G}\varphi) &= E_A \langle\langle A \rangle\rangle_{ir} \mathbf{G} TR(\varphi), \\ TR(\langle\langle A \rangle\rangle_{ir} \psi \mathbf{U} \varphi) &= E_A \langle\langle A \rangle\rangle_{ir} TR(\psi) \mathbf{U} TR(\varphi). \end{aligned}$$

The following important results can be proved by straightforward induction on the structure of φ .

THEOREM 11. *For any \mathbf{ATL}_{ir} formula φ :*
 $M, q \models tr(\varphi) \Rightarrow M, q \models \varphi \Rightarrow M, q \models TR(\varphi)$.

THEOREM 12. *If φ includes only coalitions of size at most 1, then model checking $tr(\varphi)$ and $TR(\varphi)$ can be done in time $O(|M| \cdot |\varphi|)$. In the general case, the problem is between \mathbf{NP} and Δ_2^P wrt $\max_{A \in \varphi} (|\sim_A^C|)$ and $|\varphi|$.*

Thus, our approximations potentially offer computational advantage when we consider coalitions whose members have similar knowledge, and especially when verifying abilities of individual agents.

Approximation of abilities under perfect recall. In this paper, we focus on approximating abilities based on memoryless strategies. Approximations might be equally useful for \mathbf{ATL}_{ir} (i.e., the variant of \mathbf{ATL} using uniform perfect recall strategies); we simply begin with the problem that is easier in its exact form. The high intractability of \mathbf{ATL}_{ir} model checking suggests that a substantial extension will be needed to come up with satisfactory approximations.

We also observe that the benchmark in Section 6.2 is a *model of perfect recall*, i.e., the states explicitly encode the agents' memory of their past observations. In consequence, the memoryless and perfect recall semantics of \mathbf{ATL} coincide. The experimental results suggest that, for such models, verification of perfect recall abilities can be much improved by using the approximations proposed here.

k	#states	tgen	Lower approx.		Upper approx.		Match	Exact tg+tv
			tverif	result	tverif	result		
1	15	0.001	0.0001	True	0.00007	True	100%	0.006
2	225	0.02	0.002	True	0.001	True	100%	14.79
3	3375	0.50	0.14	True	0.03	True	100%	timeout
4	50625	14.39	22.78	True	0.77	True	100%	timeout

Figure 5: Experimental results for simple voting model (φ_1)

k	#states	tgen	Lower approx.		Upper approx.		Match	Exact tg+tv
			tverif	result	tverif	result		
1	15	0.001	0.00005	False	0.00003	False	100%	0.005
2	225	0.02	0.0005	False	0.0003	False	100%	0.02
3	3375	0.50	0.01	False	0.007	False	100%	0.04
4	50625	14.39	0.94	False	0.12	False	100%	0.12

Figure 6: Experimental results for simple voting model (φ_2)

6. EXPERIMENTAL EVALUATION

Theorem 11 and Proposition 12 validate the approximation semantics theoretically. In this section, we back up the theoretical results by looking at how well the approximations work in practice. We address two issues: the *performance* and the *accuracy* of the approximations.

6.1 Existing Benchmarks

The only publicly available tool that provides verification of \mathbf{ATL} with imperfect information is MCMAS [25, 28, 23, 24]. We note, however, that imperfect information strategies are not really at the heart of the model-checker, the focus being on verification of \mathbf{CTLK} and \mathbf{ATLK} with perfect information strategies. More dedicated attempts produced so far only experimental algorithms, with preliminary performance results reported in [26, 8, 17, 9, 27]. Because of that, there are few benchmarks for model checking \mathbf{ATL}_{ir} , and few experiments have actually been conducted.

The classes of models typically used to estimate the performance of \mathbf{ATL}_{ir} model checking are **TianJi** [28, 8] and **Castles** [26]. The properties to be verified are usually reachability properties, saying that Tian Ji can achieve a win over the king (in **TianJi**), or that a given coalition of workers can defeat another castle (for **Castles**). We observe that both **TianJi** and **Castles** *do not satisfy* the necessary condition in Proposition 10. This is because the states of the model do not encode some relevant information about the actions that have been already played by the coalition. Thus, even one

step before winning the game, the players take into account also some (possibly losing) states that couldn't be reached by the strategy that they are executing.

This means that the **AE μ C** approximations, proposed in this paper, are not useful for **TianJi** and **Castles**. It also means that the benchmarks arguably do not capture realistic scenarios. We usually do not want to assume agents to forget *their own actions* from a few steps back. In the remainder, we propose several new benchmarks that can be used to evaluate our approximation scheme.

Finally, we note that most experiments reported in the literature use very simple input formulae (no nested strategic modalities; singleton coalitions or groups of agents with identical indistinguishability relations). As the results show, verification of such formulae is complex enough – see the performance of exact model checking in the rest of this section.

6.2 Verifying the Simple Voting Scenario

For the first benchmark, we adapt the simple voting scenario from Example 1. The model consists of $k + 1$ agents (k voters v_1, \dots, v_k , and 1 coercer c). The module of voter v_i implements the transition structure from Figure 1, with three modifications. First, the voter can at any state execute the “idle” action *wait* (this is needed to ensure uniformity of the resulting CEGS). In consequence, synchronous voting as well as interleaving of votes is allowed. Secondly, in states q_3, \dots, q_6 , the coercer's action *np* (“no punishment”) leads to an additional final state (q'_7, \dots, q'_{10}), labeled accordingly. Thirdly, the old and new leaves in the structure (i.e., $q_7, \dots, q_{10}, q'_7, \dots, q'_{10}$) are labeled with an additional atomic proposition *finish_i*.

As specifications, we want to use the properties saying that: (i) the coercer can force the voter to vote for candidate 1 or else the voter is punished, and (ii) the voter can avoid voting for candidate 1 and being punished (cf. Example 2). Note, however, that the model used for the experiments is an unconstrained product of the voter modules. Thus, it includes also paths that were absent in the CEGS M_{vote} from Example 1 (in particular, ones where a voter executes *wait* all the time). To deal with this, we modify the specifications from Example 2 so that they discard such paths:

1. $\varphi_1 \equiv \langle\langle c \rangle\rangle_{\text{ir}} \mathbf{G}((\text{finish}_i \wedge \neg \text{pun}_i) \rightarrow \text{vote}_{i,1})$ which always holds in the voting scenario,
2. $\varphi_2 \equiv \langle\langle v_i \rangle\rangle_{\text{ir}} \mathbf{F}(\text{finish}_i \wedge \neg \text{pun}_i \wedge \neg \text{vote}_{i,1})$ which is always false.

The results of experiments for φ_1 are shown in Figure 5, and for φ_2 in Figure 6. The columns present the following information: parameter of the model (the number of voters k), size of the state space (#states), generation time for models (tgen), time and output of verification (tver, result) for model checking the lower approximation $tr(\varphi)$, and similarly for the upper approximation $TR(\varphi)$; the percentage of cases where the bounds have matched (match), and the total running time of the exact **ATL_{ir}** model checking for φ (tg+tv). The running times are given in seconds. *Timeout* indicates that the process did not terminate in 48 hours (!).

The computation of the lower and upper approximations was done with a straightforward implementation (in Python 3) of the fixpoint model checking algorithm for **AE μ C** and **ATL_{ir}**, respectively. We used the explicit representation of models, and the algorithms were not optimized in any way. The exact **ATL_{ir}** model checking was done with MCMAS

1.2.2 in such a way that the underlying CEGS of the ISPL code was isomorphic to the explicit models used to compute approximations. The subjective semantics of **ATL_{ir}** was obtained by using the option *-atlk 2* and setting the initial states as the starting indistinguishability class for the proponent. All the tests were conducted on a PC with an Intel Core i5-2500 CPU with dynamic clock speed of 3.30 GHz up to 3,60 GHz, 8 GB of RAM (two modules DDR3, 1600 MHz bus clock), and Windows 10 (64bit).

Discussion of results. Exact model checking with MCMAS performed well on the inputs where no winning strategy existed (formula φ_2), but was very bad at finding the existing winning strategy for formula φ_1 . In that case, our approximations offered huge speedup. Moreover, the approximations actually found the winning strategy in all the tested instances, thus producing fully conclusive output. This might be partly due to the fact that the scenario uses *perfect recall models*, i.e., ones encoding perfect memory of players explicitly in their local states.




Figure 7: Example 6-endplay in bridge:

6.3 Bridge Endplay

We use bridge play scenarios of a type often considered in bridge handbooks and magazines. The task is to find a winning strategy for the declarer, usually depicted at the South position (**S**), in the k -endplay of the game, see Figure 7 for an example. The deck consists of $4n$ cards in total (n in each suit),² and the initial state captures each player holding k cards in their hand, after having played $n - k$ cards. This way we obtain a family of models, parameterized by the possible values of (n, k) . A NoTrump contract is being played; the declarer wins if she takes more than $k/2$ tricks in the endplay.

The players' cards are played sequentially (clockwise). **S** plays first at the beginning of the game. Each next trick (i.e., the set of four played cards, one per player) is opened by the player who won the latest trick. The declarer handles her own cards and the ones of the dummy (**N**). The opponents (**W** and **E**) handle their own hands each. The cards of the dummy are visible to everybody; the other hands are only seen by their owners. Each player remembers the cards that have already been played, including the ones that were used up before the initial state of the k -endplay. That is, the local state of a player contains: the current hand of the player, the current hand of the dummy, the cards from the deck that were already used up in the previous tricks, the status

²In real bridge, $n = 13$.

of the current trick, i.e., the sequence of pairs $(player, card)$ for the cards already played within the trick (alternatively, the sequence of cards already played within the trick, plus who started the trick); and the current score (which team has won how many tricks so far).

We observe the following properties of the model. First, it is turn-based (with the “idle” action *wait* that players use when another player is laying down a card). Secondly, players have imperfect information, since they cannot infer (except for the last round) the hands of the other players. The missing information is relevant: anybody who has ever played bridge or poker knows how much the limited knowledge of the opponents’ hands decreases one’s chances of winning the game. Thirdly, this is a model of imperfect recall. The players do not remember in which order the cards have been played so far, and who had what cards;³ formally: the model is a DAG and not a tree as there are histories $h \not\approx_a h'$ such that $last(h) \sim_a last(h')$. Finally, the model is lockstep (everybody sees when a transition happens), and thus tr_2 and tr_3 coincide on singleton coalitions.

(n, k)	#states	tgen	Lower approx.		Upper approx.		Match	Exact tg+tv
			tverif	%true	tverif	%true		
(1, 1)	11	0.0005	0.0001	100%	7e-05	100%	100%	0.14
(2, 2)	310	0.017	0.002	60%	0.001	60%	100%	2.42 h*
(3, 3)	12626	0.92	0.16	70%	0.05	70%	100%	timeout
(4, 4)	534722	41.66	172.07	60%	2.61	60%	100%	timeout
(5, 5)*	2443467	2641.86	76 h	100%	1929	100%	100%	timeout

Figure 8: Experimental results: solving endplay in bridge

The results of the experiments for formula $\varphi \equiv \langle\langle \mathbf{S} \rangle\rangle_{ir} \mathbf{Fwin}$ are shown in Figure 8. The columns present the following information: parameters of the model (n, k) , size of the state space (#states), generation time for models (tgen), time and output of verification (tver, %true) for model checking the lower approximation $tr(\varphi)$, and similarly for the upper approximation $TR(\varphi)$; the percentage of cases where the bounds have matched (match), and the total running time of the exact \mathbf{ATL}_{ir} model checking for φ (tg+tv). The times are given in seconds, except where indicated.

The experiments were run in the same environment as for the voting scenario in Section 6.2. Again, we ran the experiments for up to 48h per instance. The results in each row are averaged over 20 randomly generated instances, except for (*) where only 1 instance was used.

Discussion of results. In the experiments, our approximations offered a dramatic speedup. Exact model checking of φ was infeasible except for the simplest models (hundreds of states), even with an optimized symbolic model checker like MCMAS. In contrast, the bounds were verified for models up to millions of states. Moreover, our approximations obtained an astonishing level of accuracy: the bounds matched in 100% of the analyzed instances, thus producing fully conclusive output. This was partly because we only considered endplays in relatively small decks. The gap grows for decks of more than 20 cards (we verified that by hand on selected instances from bridge literature).

³This reflects the capabilities of middle-level bridge players: they usually remember what has been played, but not in which order and by whom. Advanced players remember also who played what, and masters remember the whole history of the play.

(n, k)	#states	tgen	Lower approx.		Upper approx.		Match	Exact tg+tv
			tverif	%true	tverif	%true		
(1, 1)	19	0.001	0.0003	100%	0.0003	100%	100%	14.93 h*
(2, 2)	774	0.07	0.01	40%	0.02	50.00%	90%	timeout
(3, 3)	51865	6.71	29.31	65%	2.45	85%	80%	timeout

Figure 9: Experimental results for absent-minded declarer

(n, k)	#states	tgen	Lower approx.		Upper approx.		Match	Exact tg+tv
			tverif	%true	tverif	%true		
(1, 1)	19	0.002	0.0001	0%	0.0003	100%	0%	14.93 h*
(2, 2)	756	0.08	0.003	0%	0.03	95%	5%	timeout
(3, 3)	55688	9.99	0.09	0%	2.35	70%	30%	timeout

Figure 10: Absent-minded declarer, approximation tr_2

6.4 Bridge Endplay by Absentminded Declarer

In the bridge endplay models, the players always see when a move is made. Thus, for singleton coalitions, the steadfast next-time operator $\langle a \rangle^\bullet$ coincides with the standard next-time abilities expressed by $\langle a \rangle$. In order to better assess the performance, we have considered a variant of the scenario where the declarer is absentminded and does not see the cards being laid on the table until the end of each trick. Moreover, she can play her and the dummy’s cards at any moment, even in parallel with the opponents. This results in larger indistinguishability classes for \mathbf{S} , but also in a general increase of the number of states and transitions.

The results of the experiments are shown in Figure 9. Note that, for this class of models, the bounds do not match as tightly as before. Still, the approximation was conclusive in an overwhelming majority of instances. Moreover, it grossly outperformed the exact model checking which was (barely) possible only in the trivial case of $n = 1$.

The models are not turn-based, not lockstep, and not of perfect recall. Since they are not lockstep, approximations tr_2 and tr_3 do not have to coincide. In Figure 10, we present the experimental results obtained with tr_2 , which show that the improved approximation tr_3 provides tighter lower bounds also from the practical point of view.

7. CONCLUSIONS

Verification of strategic properties in scenarios with imperfect information is difficult, both theoretically and in practice. In this paper, we suggest that model checking of logics like \mathbf{ATL}_{ir} can be in some cases obtained by computing an under- and an overapproximation of the \mathbf{ATL}_{ir} specification, and comparing if the bounds match. In a way, our proposal is similar to the idea of may/must abstraction [14, 4, 22], only in our case the approximations are obtained by transforming formulae rather than models.

We propose such approximations, prove their correctness, and show that, for singleton coalitions, their values can be computed in polynomial time. We also propose novel benchmarks for experimental validation. Finally, we report very promising experimental results, in both performance and accuracy of the output. To our best knowledge, this is the first successful attempt at approximating strategic abilities under imperfect information by means of fixpoint methods.

Acknowledgements. The authors acknowledge the support of the National Centre for Research and Development (NCBR), Poland, under the PolLux project VoteVerif (POL-LUX-IV/1/2016).

REFERENCES

- [1] T. Ágotnes. A note on syntactic characterization of incomplete information in ATEL. In *Proceedings of Workshop on Knowledge and Games*, pages 34–42, 2004.
- [2] T. Ágotnes, V. Goranko, W. Jamroga, and M. Wooldridge. Knowledge and ability. In H. van Ditmarsch, J. Halpern, W. van der Hoek, and B. Kooi, editors, *Handbook of Epistemic Logic*, pages 543–589. College Publications, 2015.
- [3] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.
- [4] T. Ball and O. Kupferman. An abstraction-refinement framework for multi-agent systems. In *Proceedings of LICS*, pages 379–388, 2006.
- [5] N. Bulling, J. Dix, and W. Jamroga. Model checking logics of strategic ability: Complexity. In M. Dastani, K. Hindriks, and J.-J. Meyer, editors, *Specification and Verification of Multi-Agent Systems*, pages 125–159. Springer, 2010.
- [6] N. Bulling and W. Jamroga. Alternating epistemic mu-calculus. In *Proceedings of IJCAI-11*, pages 109–114, 2011.
- [7] N. Bulling and W. Jamroga. Comparing variants of strategic ability: How uncertainty and memory influence general properties of games. *Journal of Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
- [8] S. Busard, C. Pecheur, H. Qu, and F. Raimondi. Improving the model checking of strategies under partial observability and fairness constraints. In *Formal Methods and Software Engineering*, volume 8829 of *Lecture Notes in Computer Science*, pages 27–42. Springer, 2014.
- [9] S. Busard, C. Pecheur, H. Qu, and F. Raimondi. Reasoning about memoryless strategies under partial observability and unconditional fairness constraints. *Information and Computation*, 242:128–156, 2015.
- [10] C. Dima, C. Enea, and D. Guelev. Model-checking an alternating-time temporal logic with knowledge, imperfect information, perfect recall and communicating coalitions. In *Proceedings of GANDALF*, pages 103–117, 2010.
- [11] C. Dima, B. Maubert, and S. Pinchinat. The expressive power of epistemic μ -calculus. *CoRR*, abs/1407.5166, 2014.
- [12] C. Dima, B. Maubert, and S. Pinchinat. Relating paths in transition systems: The fall of the modal mu-calculus. In *Proceedings of MFCS*, volume 9234 of *Lecture Notes in Computer Science*, pages 179–191. Springer, 2015.
- [13] C. Dima and F. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
- [14] P. Godefroid and R. Jagadeesan. Automatic abstraction using generalized model checking. In *Proceedings of CAV*, volume 2404 of *Lecture Notes in Computer Science*, pages 137–150. Springer, 2002.
- [15] D. Guelev and C. Dima. Epistemic ATL with perfect recall, past and strategy contexts. In *Proceedings of CLIMA-XIII*, volume 7486 of *Lecture Notes in Computer Science*, pages 77–93. Springer, 2012.
- [16] D. Guelev, C. Dima, and C. Enea. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *Journal of Applied Non-Classical Logics*, 21(1):93–131, 2011.
- [17] X. Huang and R. van der Meyden. Symbolic model checking epistemic strategy logic. In *Proceedings of AAI*, pages 1426–1432, 2014.
- [18] W. Jamroga. Some remarks on alternating temporal epistemic logic. In B. Dunin-Keplicz and R. Verbrugge, editors, *Proceedings of Formal Approaches to Multi-Agent Systems (FAMAS 2003)*, pages 133–140, 2003.
- [19] W. Jamroga. *Logical Methods for Specification and Verification of Multi-Agent Systems*. ICS PAS Publishing House, 2015.
- [20] W. Jamroga and J. Dix. Model checking ATL_{ir} is indeed Δ_2^P -complete. In *Proceedings of EUMAS’06*, volume 223 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2006.
- [21] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 63(2–3):185–219, 2004.
- [22] A. Lomuscio and J. Michaliszyn. Verification of multi-agent systems via predicate abstraction against ATLK specifications. In *Proceedings of AAMAS*, pages 662–670, 2016.
- [23] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS : A model checker for the verification multi-agent systems. In *Proceedings of CAV*, volume 5643 of *Lecture Notes in Computer Science*, pages 682–688. Springer, 2009.
- [24] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: An open-source model checker for the verification of multi-agent systems. *International Journal on Software Tools for Technology Transfer*, 2015. Available online.
- [25] A. Lomuscio and F. Raimondi. Model checking knowledge, strategies, and games in multi-agent systems. In *Proceedings of AAMAS*, pages 161–168, 2006.
- [26] J. Pilecki, M. Bednarczyk, and W. Jamroga. Synthesis and verification of uniform strategies for multi-agent systems. In *Proceedings of CLIMA XV*, volume 8624 of *Lecture Notes in Computer Science*, pages 166–182. Springer, 2014.
- [27] J. Pilecki, M. Bednarczyk, and W. Jamroga. SMC: Synthesis of uniform strategies and verification of strategic ability for multi-agent systems. *Journal of Logic and Computation*, 2017. To appear.
- [28] F. Raimondi. *Model Checking Multi-Agent Systems*. PhD thesis, University College London, 2006.
- [29] P. Y. Schobbens. Alternating-time logic with imperfect recall. *Electronic Notes in Theoretical Computer Science*, 85(2):82–93, 2004.
- [30] W. van der Hoek, A. Lomuscio, and M. Wooldridge. On the complexity of practical ATL model checking. In *Proceedings of AAMAS’06*, pages 201–208. ACM, 2006.
- [31] W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In C. Castelfranchi and W. Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-02)*, pages 1167–1174. ACM Press, New York, 2002.