PRACE IPI PAN ICS PAS REPORTS

Wojciech Jamroga
Michał Knapik

Underapproximating ATL with Imperfect Information and Imperfect Recall

Nr 1032

INSTYTUT PODSTAW INFORMATYKI POLSKIEJ AKADEMII NAUK
INSTITUTE OF COMPUTER SCIENCE POLISH ACADEMY OF SCIENCES
ul. Jana Kazimierza 5, 01-248 Warszawa http://www.ipipan.waw.pl

Wojciech Jamroga, Michał Knapik[1]

Underapproximating ATL with Imperfect Information and Imperfect Recall

Nr 1032

Warszawa, wrzesień 2015

[1]Michał Knapik is funded by DEC-2012/07/N/ST6/03426 NCN Preludium grant.

Abstract

We investigate the correspondence between model checking of af-AMC$_i$ and ATL$_{ir}$, on the example of reachability. We identify some of the reasons for the fact that these logics are of uncomparable expressivity. These observations form the basis for a novel method for underapproximating ATL$_{ir}$ by means of fixed-point calculations. We introduce a special version of the next-step operator, called Persistent Imperfect Next-Step Operator $\langle\cdot\rangle^F$ and show how it can be used to define a new version of reachability that carries to ATL$_{ir}$.

Keywords: ATL, model checking, approximation

Streszczenie

Dolna Aproksymacja Bezpamięciowego ATL
o Niepełnej Informacji

W pracy badane są związki pomiędzy weryfikacją modelową Bezpamięciowej Logiki Temporalnej Czasu Alternującego z Niepełną Informacją ATL$_{ir}$ i Epistemicznego Alternującego Mu-Rachunku af-AMC$_i$. Jak pokazano, naturalne uogólnienia pojęcia osiągalności z ATL$_{ir}$-a do af-AMC$_i$ nie przynoszą dobrych efektów: osiągalność w af-AMC$_i$ nie pociąga za sobą osiągalności w ATL$_{ir}$. Po zidentyfikowaniu części powodów dla których tak się dzieje, zaproponowano nową wersję operatora następnego kroku, który pozwala na przybliżanie osiągalności w ATL$_{ir}$ przy pomocy obliczeń stałopunktowych.

Słowa kluczowe: ATL, weryfikacja, aproksymacja

1 Introduction

Multi-agent systems consist of multiple entities called *agents*, often assumed to be intelligent, and able to interact with each other as well as with the environment. More and more practical problems are being modeled and solved under paradigms related to multi-agent systems. The examples of their industrial applications include space mission planning and air control logistics and production planning

Given the proliferation of multi-agent systems in the modern world, their automated verification, i.e., ensuring of the compliance with their specification, becomes an important task. Some model checking tools (e.g., MCK [5], MCMAS [6]) have been created or extended to accept multi-agent systems as their input models; some tools aim at the verification of programs specified in multi-agent programming languages (e.g., Agentspeak in Jason [2]).

In this paper we deal with the properties specified in Alternating-Time Temporal Logic [1] (ATL). In its basic version, this logic allows to specify strategic properties of groups of agents under assumption of the perfect knowledge about the state of affairs, i.e., an agent is able to observe the world in its entirety. One can argue that such a scenario is rather unrealistic, hence several versions of semantics have been introduced [7] to take into account the limited knowledge of the agents. This, however, typically leads to a greater complexity of the task of model checking. Model checking of ATL can be done in linear time in the size of the model and the formula.

In this paper we investigate a certain version of ATL, where the agents are assumed to be memoryless and of limited knowledge, denoted by ATL_{ir}. Model checking of ATL_{ir} is Δ_2^P–complete [7]. Moreover, there is no clear correspondence between the validity of the formulae specified in ATL_{ir} and their natural fixed-point counterparts specified in Alternating Epistemic Mu-Calculus [3] (AMC_i, for short). We analyse the reasons for this lack of correspondence and propose a novel method for underapproximating ATL_{ir} by means of fixed-point calculations. To this end we introduce a special version of the next-step operator, called Persistent Imperfect Next-Step Operator $\langle\cdot\rangle^F$ and show how it can be used to define a new version of reachability that carries to ATL_{ir}.

To the best knowledge of the authors this is the first successful attempt at underapproximating ATL_{ir} using fixed-point methods.

1.1 Related Work

Alternating-Time Temporal Logic and Alternating Mu-Calculus (AMC) are introduced in [1]. ATL deals with strategic abilities of coalitions of agents, while AMC combines the next-step operator of ATL with the operator of the least fixed-point μ. In [1] the authors show that AMC is strictly more expressive than the original ATL. However, the semantics of ATL is also considered in several flavours that rely upon the definition of strategy. In [7] a natural classification of these versions is presented, depending on (1) whether the agents make a decision about the next step based on the history of the visited states or based on the current state; (2) whether the agents can observe the entire state of affairs or only their local epistemic neighbourhood. In this paper we focus on the version of ATL with memoryless strategies and the agents that have a limited knowledge about the world, i.e., ATL_{ir} in the nomenclature of [7]. In [3] Alternating Epistemic Mu-Calculus, a logic built on top of the next-step operator of ATL_{ir}, using the least fixed-point operator, is presented and investigated; it is shown that the expressivity of ATL_{ir} is incomparable with the expressivity of af-AMC_i.

2 Preliminaries

Let us briefly recall the basic notions and properties of ATL_{ir} and af-AMC_i.

2.1 Imperfect Information Concurrent Game Structures

We interpret the presented logics in transition systems equipped with two relations: one modeling the transitions between the states and the other explicitly connecting the states that are indistinguishable from an agent's point of view.

Definition 1 (ICGS)**.** *We call an* Imperfect Information Concurrent Game Structure *a tuple* $\mathcal{M} = \langle \mathbb{A}\mathrm{gt}, \mathcal{Q}, \Pi, \pi, Acts, d, o, \{\sim_a \mid a \in \mathbb{A}\mathrm{gt}\}\rangle$, *where:*

- $\mathbb{A}\mathrm{gt} = \{1, \ldots, k\}$ *is a finite set of all the agents, for some* $k \in \mathbb{N}$,
- $\mathcal{Q}$ *is a finite set of states,*
- Π *is a set of atomic propositions,*
- $\pi \colon \Pi \to \mathcal{P}(\mathcal{Q})$ *is a labeling function,*
- *Acts is a finite set of atomic actions,*
- $d \colon \mathbb{A}\mathrm{gt} \times \mathcal{Q} \to \mathcal{P}(\mathcal{Q})$ *is a protocol function,*
- $o \colon \bigcup_{q\in\mathcal{Q}}\{(q, \alpha_1, \ldots, \alpha_k) \mid \forall_{1\leq a\leq k}\alpha_a \in d(a, q)\} \to \mathcal{Q}$ *is a transition function,*
- $\sim_a \subseteq \mathcal{Q} \times \mathcal{Q}$ *is an equivalence relation such that* $d(q, a) = d(q', a)$ *if* $q \sim_a q'$, *called the indistinguishability relation, for all* $a \in \mathbb{A}\mathrm{gt}$.

For any agent $a \in \mathbb{A}\mathrm{gt}$ and state $q \in \mathcal{Q}$ we write $d_a(q)$ instead of $d(a, q)$ and assume that $d_a(q) \neq \emptyset$. The set $d_a(q)$ consists of all the actions available to the agent a in the state q. The state $o(q, \alpha_1, \ldots, \alpha_k)$ is the *outcome* of simultaneously executing in the state q the actions $\alpha_1, \ldots, \alpha_k$ by the respective agents $1, \ldots, k$. We define the *coalition protocol function* d_A for a group $A \subseteq \mathbb{A}\mathrm{gt}$ by putting $d_A(q) = \prod_{a\in A} d_a(q)$, for each $q \in \mathcal{Q}$.

We call a *strategy* of $a \in \mathbb{A}\mathrm{gt}$ a function $s_a \colon \mathcal{Q} \to Acts$ satisfying $s_a(q) \in d(a, q)$ and $q \sim_a q' \implies s_a(q) = s_a(q')$, for all $q, q' \in \mathcal{Q}$. Intuitively, a strategy for an agent assigns to every state an action consistent with the protocol and the indistinguishable states are assigned the same actions. This type of strategy is also called *imperfect information, imperfect recall strategy* [7]. A set of strategies s_A, one per each agent from $A \subseteq \mathbb{A}\mathrm{gt}$, is called a *collective strategy for A*. By $s_A|_a$ we denote the strategy for agent $a \in A$ selected from s_A. The set of all the collective strategies for A is denoted by Σ_A. A partial function $s'_A \colon \mathcal{Q} \nrightarrow Acts$ is called *a partial strategy for A* if $s'_A \subseteq s_A$ for some strategy $s_A \in \Sigma_A$. If

s'_A and s''_A are partial strategies such that $s'_A \subseteq s''_A$, then we say that s''_A *extends* s'_A.

We call a *path* an infinite sequence of states $\lambda = q_0 q_1 q_2 \ldots$ such that for each $i \in \mathbb{N}$ there exist actions $\alpha_1, \ldots, \alpha_k \in Acts$ such that $q_{i+1} = o(q_i, \alpha_1, \ldots, \alpha_k)$. If $s_A \in \Sigma_A$ and for each $i \in \mathbb{N}$ there exist $\alpha_1, \ldots, \alpha_k \in Acts$ satisfying $\lambda_{i+1} = o(\lambda_i, \alpha_1, \ldots, \alpha_k)$ and $\alpha_a = s_A|_a(\lambda_i)$ for all $a \in A$, then λ is *consistent with* s_A. Intuitively, λ is consistent with s_A if it is the outcome of following the strategy s_A by the group A while all the remaining agents have free hand in selecting the actions. The set of all the paths starting from the state $q \in \mathcal{Q}$ and consistent with the strategy s_A is denoted by $out_{\mathcal{M}}(q, s_A)$. We define the *outcome of* s_A *in* q as $out^{ir}_{\mathcal{M}}(q, s_A) = \bigcup_{a \in A} \bigcup_{q' \sim_a q} out_{\mathcal{M}}(q', s_A)$ if $A \neq \emptyset$ and let $out^{ir}_{\mathcal{M}}(q, s_\emptyset) = out_{\mathcal{M}}(q, s_\emptyset)$.

Let us fix $\mathcal{M} = \langle \mathbb{A}\mathrm{gt}, \mathcal{Q}, \Pi, \pi, Acts, d, o, \{\sim_a \mid a \in \mathbb{A}\mathrm{gt}\}\rangle$.

2.2 ATL with Imperfect Information and Recall

Let us present the language of Alternating-Time Temporal Logic with Imperfect Information and Recall (ATL_{ir}, for short).

Definition 2. *The language of* ATL_{ir} *is defined by the following grammar:*

$$\phi ::= p \mid \neg\phi \mid \phi \vee \phi \mid \langle\langle A\rangle\rangle \bigcirc \phi \mid \langle\langle A\rangle\rangle \Box\phi \mid \langle\langle A\rangle\rangle \phi \mathcal{U} \phi,$$

where $p \in \Pi$ *and* $A \subseteq \mathbb{A}\mathrm{gt}$.

We read $\langle\langle A\rangle\rangle\xi$ as *"A can enforce* ξ*"*, $\bigcirc$ as *"in the next state"*, $\Box$ as *"now and always in the future"*, and $\mathcal{U}$ as *"until"*. To signify the fact that we use the semantics based on imperfect information and imperfect recall, we denote the satisfaction relation by $\models_{ir}$.

Definition 3. *Let* $q \in \mathcal{Q}$. *The semantics of* ATL_{ir} *is defined as follows:*

- $\mathcal{M}, q \models_{ir} p$ *iff* $q \in \pi(p)$,
- $\mathcal{M}, q \models_{ir} \neg\phi$ *iff* $\mathcal{M}, q \not\models_{ir} \phi$,
- $\mathcal{M}, q \models_{ir} \phi \vee \psi$ *iff* $\mathcal{M}, q \models_{ir} \phi$ *or* $\mathcal{M}, q \models_{ir} \psi$,

- *$\mathcal{M}, q \models_{ir} \langle\!\langle A \rangle\!\rangle \bigcirc \phi$ iff there exists a collective strategy $s_A \in \Sigma_A$ s.t. for all $\lambda \in out^{ir}_{\mathcal{M}}(q, s_A)$ we have $\mathcal{M}, \lambda_1 \models_{ir} \phi$,*
- *$\mathcal{M}, q \models_{ir} \langle\!\langle A \rangle\!\rangle \Box \phi$ iff there exists a collective strategy $s_A \in \Sigma_A$ s.t. for all $\lambda \in out^{ir}_{\mathcal{M}}(q, s_A)$ and every $i \in \mathbb{N}$ we have $\mathcal{M}, \lambda_i \models_{ir} \phi$,*
- *$\mathcal{M}, q \models_{ir} \langle\!\langle A \rangle\!\rangle \psi \mathcal{U} \phi$ iff there exists a collective strategy $s_A \in \Sigma_A$ s.t. for each $\lambda \in out^{ir}_{\mathcal{M}}(q, s_A)$ there exists $i \in \mathbb{N}$ for which $\mathcal{M}, \lambda_i \models_{ir} \phi$ and $\mathcal{M}, \lambda_j \models_{ir} \psi$ for all $0 \leq j < i$,*

where $p \in \Pi$, $\phi, \psi \in \mathrm{ATL}$, and $A \subseteq \mathbb{A}\mathrm{gt}$.

We also introduce the derived modality $\Diamond \phi \equiv \top \mathcal{U} \phi$, read as *"now or sometime in the future"*. It is easy to see that $\mathcal{M}, q \models_{ir} \langle\!\langle A \rangle\!\rangle \Diamond \phi$ iff there exists a collective strategy $s_A \in \Sigma_A$ such that following each path $\lambda \in out^{ir}_{\mathcal{M}}(q, s_A)$ leads to some state satisfying ϕ. In this case we say that ϕ is *ir–reachable from q*.

2.3 Alternating Epistemic Mu-Calculus

We now present the language of Alternating Epistemic Mu-Calculus (AMC_i, for short).

Definition 4. *Let $\mathcal{V}ars$ be a set of second-order variables ranging over $\mathcal{P}(\mathcal{Q})$. The language of AMC_i is defined by the following grammar:*

$$\phi ::= p \mid X \mid \neg\phi \mid \phi \vee \phi \mid \langle A \rangle \phi \mid \mu X(\phi) \mid \mathcal{K}_a,$$

where $p \in \Pi$, $X \in \mathcal{V}ars$, $a \in \mathbb{A}\mathrm{gt}$, $A \subseteq \mathbb{A}\mathrm{gt}$, and the formulae are X–positive, i.e., each free occurrence of X is in the scope of an even number of negations.

As usual, μ denotes the least fixed-point operator. We define the dual of μ as $\nu X(\phi(X)) \equiv \neg \mu X(\neg\phi(\neg X))$; ν corresponds to the greatest fixed-point operator. A formula of AMC_i is *alternation-free* if in its positive normal form it does not contain occurrences of ν (μ, resp.) on any syntactic path from an occurrence of μX (νX, resp.) to a bound occurrence of X (cf. [1]). Similarly to [3], we consider here only the alternation-free fragment of AMC_i, denoted by af-AMC_i.

We evaluate the formulae of af-AMC$_i$ with respect to the valuations of $\mathcal{V}ars$, i.e., functions $\mathcal{V}\colon \mathcal{V}ars \to \mathcal{P}(\mathcal{Q})$. We denote the set of all the valuations of $\mathcal{V}ars$ by $\mathcal{V}als$. If $X \in \mathcal{V}ars$, $Z \subseteq \mathcal{Q}$, and $\mathcal{V} \in \mathcal{V}als$, then by $\mathcal{V}[X := Z]$ we denote the valuation of $\mathcal{V}ars$ such that $\mathcal{V}[X := Z](Y) = \mathcal{V}(Y)$ for $Y \neq X$ and $\mathcal{V}[X := Z](X) = Z$.

We define the semantics of af-AMC$_i$ using the denotation function that assigns to each formula $\phi \in$ af-AMC$_i$ the set of states $[\![\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}}$ where ϕ is true under the valuation $\mathcal{V} \in \mathcal{V}als$.

Definition 5. *Let $q \in \mathcal{Q}$ and $\mathcal{V} \in \mathcal{V}als$. The denotation function for* af-AMC$_i$ *is defined as follows:*

- $[\![p]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \pi(p)$,
- $[\![X]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \mathcal{V}(X)$,
- $[\![\neg\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \mathcal{Q} \setminus [\![\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}}$,
- $[\![\phi \vee \psi]\!]^{\mathcal{M}}_{\mathcal{V}} = [\![\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} \cup [\![\psi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}}$,
- $[\![\langle A\rangle\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \{q \in \mathcal{Q} \mid \exists s_A \in \Sigma_A\ \forall\lambda \in out^{ir}_{\mathcal{M}}(q, s_A)\ \mathcal{M}, \lambda_1 \models_{ir} \phi\}$,
- $[\![\mu X(\phi)]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \bigcap\{Z \subseteq \mathcal{Q} \mid [\![\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}[X:=Z]} \subseteq Z\}$,
- $[\![\mathcal{K}_a\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \{q \in \mathcal{Q} \mid \forall_{q'}(q' \sim_a q \textit{ implies } q' \in [\![\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}})\}$,

where $\phi \in$ af-AMC$_i$*, $p \in \Pi$, $X \in \mathcal{V}ars$, $a \in \mathbb{A}\mathrm{gt}$, and $A \subseteq \mathbb{A}\mathrm{gt}$.*

We write $\mathcal{M}, q \models^{\mu,\mathcal{V}}_{ir} \phi$ iff $q \in [\![\phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}}$, for each $q \in \mathcal{Q}$ and $\phi \in$ af-AMC$_i$. If ϕ is a sentence, i.e., it contains no free variables, then its validity does not depend on the valuation $\mathcal{V}$ and we simply write $\mathcal{M}, q \models^{\mu}_{ir} \phi$ and $[\![\phi]\!]^{\mathcal{M}}_{\mu ir}$. We interpret $\mathcal{K}_a\phi$ as *"the agent a knows ϕ"*. For each $A \subseteq \mathbb{A}\mathrm{gt}$ we define the derived group knowledge modality *everybody knows* as $\mathcal{E}_A\phi \equiv \bigvee_{a\in A} \mathcal{K}_a\phi$. Recall that $\mathcal{M}, q \models^{\mu,\mathcal{V}}_{ir} \mathcal{E}_A\phi$ iff $\mathcal{M}, q' \models^{\mu,\mathcal{V}}_{ir} \phi$ for all $q' \sim^A_E q$, where $\sim^A_E = \bigcup_{a\in A} \sim_a$.

Model checking of ATL$_{ir}$ is Δ^P_2–complete [7]. Model checking of af-AMC$_i$ can be done in linear time with respect to the number of transitions for the formulae with coalitions that do not exceed 2; in general, model checking of af-AMC$_i$ is NP–hard and conjectured to be Δ^P_2–complete [3].

3 Underapproximating ATL_{ir}

Following [7], we denote the satisfaction in the vanilla ATL and AMC by $\models_{IR}$ and $\models^{\mu}$, respectively. Vanilla ATL can be embedded into AMC using the translation that assigns to any formula $\phi \in$ ATL the formula $\mathcal{TR}_{IR}(\phi) \in$ AMC defined recursively as follows [1]: $\mathcal{TR}_{IR}(p) = p$, $\mathcal{TR}_{IR}(\neg\phi) = \neg\mathcal{TR}_{IR}(\phi)$, $\mathcal{TR}_{IR}(\phi \vee \psi) = \mathcal{TR}_{IR}(\phi) \vee \mathcal{TR}_{IR}(\psi)$, $\mathcal{TR}_{IR}(\langle\langle A\rangle\rangle \bigcirc \phi) = \langle A\rangle\mathcal{TR}_{IR}(\phi)$, $\mathcal{TR}_{IR}(\langle\langle A\rangle\rangle\Box\phi) = \nu X(\mathcal{TR}_{IR}(\phi) \wedge \langle A\rangle X)$, $\mathcal{TR}_{IR}(\langle\langle A\rangle\rangle\psi\mathcal{U}\phi) = \mu X(\mathcal{TR}_{IR}(\phi) \vee (\mathcal{TR}_{IR}(\psi) \wedge \langle A\rangle X))$, where $p \in \Pi$, $\phi, \psi \in$ ATL, and $A \subseteq \mathbb{A}\mathrm{gt}$.

It is known [1] that $\mathcal{M}, q \models_{IR} \phi$ iff $\mathcal{M}, q \models^{\mu} \mathcal{TR}_{IR}(\phi)$, for each $\phi \in$ ATL and $q \in \mathcal{Q}$. Indeed, fixed-point computations are the usual way of verifying ATL. As shown in [3], this approach does not extend to ATL_{ir} and af-AMC_i. The fact that these logics are of incomparable expressivity means that model checking of af-AMC_i cannot be used, in general, for model checking of ATL_{ir}. Let us illustrate this on the example of reachability.

3.1 Reachability in ATL_{ir} versus af-AMC_i

It is not obvious how to define the reachability for af-AMC_i. A straightforward imitation of the fixed-point definition from ATL immediately proves problematic. Namely, if we introduce a new derived modality $\Diamond^{?}$ in such a way that $\langle\langle A\rangle\rangle\Diamond^{?}\phi \equiv \mu X(\phi \vee \langle A\rangle X)$, then it is easy to observe that $\mathcal{M}, q \models^{\mu}_{ir} \phi$ implies $\mathcal{M}, q \models^{\mu}_{ir} \langle\langle A\rangle\rangle\Diamond^{?}\phi$. This, however, is inconsistent with af-AMC_i, where $\mathcal{M}, q \models_{ir} \phi \implies \mathcal{M}, q \models_{ir} \langle\langle A\rangle\rangle\Diamond\phi$ does not hold. On the other hand, we have $\mathcal{M}, q \models_{ir} \mathcal{E}_A\phi \implies \mathcal{M}, q \models_{ir} \langle\langle A\rangle\rangle\Diamond\phi$, i.e., if the coalition A knows ϕ, then ϕ is bound to happen, as it is actually occuring. This suggests the following definition:

$$\langle A\rangle\Diamond^{\bullet}\phi \equiv \mu X(\mathcal{E}_A\phi \vee \langle A\rangle X), \tag{1}$$

for all $A \subseteq \mathbb{A}\mathrm{gt}$, $\phi \in$ af-AMC_i. However, as we show in the following proposition, this still does not capture the concept of *ir*–reachability.

Proposition 1. *Let $A \subseteq \mathbb{A}\mathrm{gt}$, $q \in \mathcal{Q}$, $\phi \in$ af-AMC_i, and $\mathcal{V} \in \mathcal{V}als$. The following conditions hold:*

1. $\mathcal{M}, q \models^{\mu}_{ir} \langle\!\langle \emptyset \rangle\!\rangle \Diamond^{\bullet} \phi$ *iff* $\mathcal{M}, q \models_{ir} \langle\!\langle \emptyset \rangle\!\rangle \Diamond \phi$,

2. *if* $|A| = 1$, *then* $\mathcal{M}, q \models^{\mu}_{ir} \langle\!\langle A \rangle\!\rangle \Diamond^{\bullet} \phi$ *implies* $\mathcal{M}, q \models_{ir} \langle\!\langle A \rangle\!\rangle \Diamond \phi$, *but the reverse implication is not true,*

3. *if* $|A| > 1$, *then* $\mathcal{M}, q \models^{\mu}_{ir} \langle\!\langle A \rangle\!\rangle \Diamond^{\bullet} \phi$ *does not imply* $\mathcal{M}, q \models_{ir} \langle\!\langle A \rangle\!\rangle \Diamond \phi$, *and the reverse implication is not true too.*

Proof. The first case follows from the fact that for the empty coalition of agents the ir–reachability is equivalent to the IR–reachability, which in turn has a fixed-point characterisation in AMC [1]. For the empty coalition, this characterisation carries over to af-AMC$_i$.

Let us move to the second case. We define the sequence $\{F_j\}_{j \in \mathbb{N}}$ of formulae of af-AMC$_i$ such that $F_0 = \mathcal{K}_a \phi$ and $F_{j+1} = F_0 \vee \langle a \rangle F_j$, for all $j \geq 0$. From Kleene fixed-point theorem we have $[\![\langle\!\langle a \rangle\!\rangle \Diamond \phi]\!]^{\mathcal{M}}_{\mu ir} = \bigcup_{j=0}^{\infty} [\![F_j]\!]^{\mathcal{M}}_{\mu ir}$, where $\{[\![F_j]\!]^{\mathcal{M}}_{\mu ir}\}_{j \in \mathbb{N}}$ is a non-decreasing monotone sequence of subsets of $\mathcal{Q}$. Now, we prove that for each $j \in \mathbb{N}$ there exists a partial strategy s_a^j such that $dom(s_a^j) = [\![F_j]\!]^{\mathcal{M}}_{\mu ir}$, $\forall q \in dom(s_a^j)\ \forall \lambda \in out^{ir}_{\mathcal{M}}(q, s_a^j)\ \exists k \leq j\ \lambda_k \models^{\mu}_{ir} \phi$, 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 $q \in [\![F_j]\!]^{\mathcal{M}}_{\mu ir}$, then $[q]_{\sim_a} \subseteq [\![F_j]\!]^{\mathcal{M}}_{\mu ir}$. Due to the fact that $\sim_a$ is an equivalence relation, for each $q \in F_{j+1}$ either $[q]_{\sim_a} \subseteq F_j$ or $[q]_{\sim_a} \subseteq F_{j+1} \setminus F_j$. 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 such that $\forall \lambda \in out^{ir}_{\mathcal{M}}(q, s_a^q)\ \lambda_1 \in [\![F_j]\!]^{\mathcal{M}}_{\mu ir}$. Moreover, due to the fact that $\sim_a$ is an equivalence relation, the set of such strategies is shared by the whole class $[q]_{\sim_a}$. We therefore put $s_a^{j+1}(q') = s_a^q(q')$ for all $q' \in [q]_{\sim_a}$. Finally, we can define the partial strategy $s_a = \bigcup_{j \in \mathbb{N}} s_a^j$. For each $q \in \mathcal{Q}$ such that $\mathcal{M}, q \models^{\mu}_{ir} \langle\!\langle A \rangle\!\rangle \Diamond^{\bullet} \phi$, either $\mathcal{M}, q \models^{\mu}_{ir} \phi$, or ϕ will be reached along each path consistent with any extension of s_a to a full strategy. The fact that the reverse implication is not true is shown in [3] (Proposition 4).

We now move to the last case and consider the ICGS $\mathcal{M}_1$ presented in Fig. 1. We assume that $d_1(s) = \{a, b\}$ and $d_2(s) = \{x, y\}$, for $s \in \{t, v, r, u\}$. In the remaining states the protols allow only one transition. For clarity, we omit from the figure the transitions leaving the states

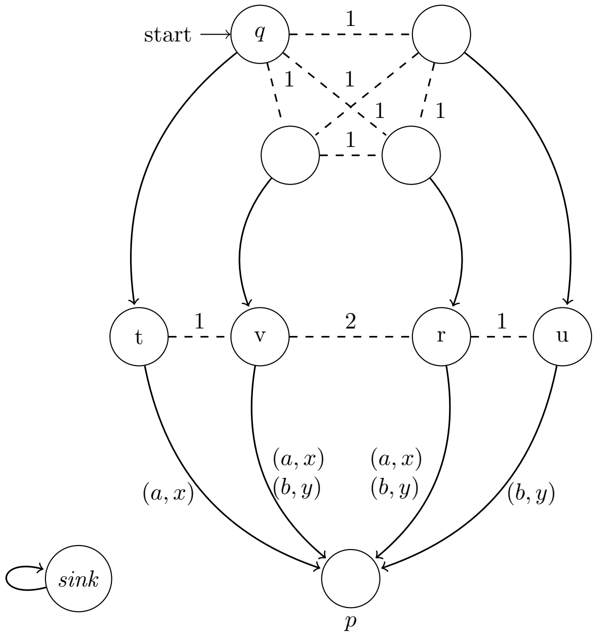

Figure 1: ICGS $\mathcal{M}_1$: a counterexample for reachability

t, v, r, and u, leading to *sink*. By hand-calculations one can observe that $\mathcal{M}, q \models^{\mu}_{ir} \langle\langle 1, 2\rangle\rangle \Diamond p$ and $\mathcal{M}, q \not\models_{ir} \langle\langle 1, 2\rangle\rangle \Diamond p$. This example can be easily extended with a larger number of (idle) agents. This concludes the proof of the case and the proposition. □

3.2 Altering the Next-Step Operator

From the practical point of view, the results obtained so far are rather discouraging. A fixed-point based logic such as af-AMC_i may enable more efficient procedures of model checking than ATL_{ir}. However, as shown in Proposition 1, the verification of the formulae of ATL_{ir} by their natural translation to af-AMC_i does not work in general.

In what follows, we propose a new version of the next-step operator, denoted by $\langle\cdot\rangle^F$. As we show, this construct retains the basic properties of the next-step operator of ATL_{ir} and the fixed-point reachability

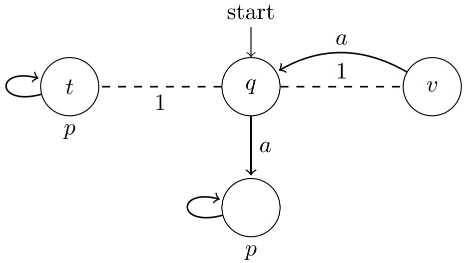

Figure 2: ICGS $\mathcal{M}_2$: the next-step operator limitations

defined using $\langle\cdot\rangle^F$ implies ir–reachability.

Let $A \subseteq \mathbb{A}\mathrm{gt}$. Following [4] we define the relation of the *common knowledge* as $\sim_C^A = (\sim_E^A)^\star$, where * is the operator of transitive closure. For our convenience, and with a slight notational abuse ($\sim_E^A$ is not an equivalence relation), here and in the remainder of this paper we denote $[q]_{\sim_E^A} = \{q' \in \mathcal{Q} \mid q' \sim_E^A q\}$.

Persistent Next-Step Operator Before we present the goal construct, we introduce an intermediate next-step operator $\langle\cdot\rangle^P$ whose aim is to alleviate the following limitation of af-AMC$_i$. Consider the single-agent ICGS presented in Fig. 2. Observe that the sole possible strategy, in which the agent 1 selects the action a, enforces eventually reaching p, i.e., $\mathcal{M}_2, q \models_{ir} \langle\langle 1\rangle\rangle \Diamond p$. On the other hand, $\mathcal{M}_2, q \not\models_{ir}^{\mu} \langle\langle 1\rangle\rangle \Diamond p$. The reason for the latter is that the next-step operator ATL$_{ir}$ requires reaching p in the succeeding state from *all* the states indistinguishable from q, whereas p is reached from t, q, v in zero, one, and two steps, respectively.

Let $s_A \in \Sigma_A$ be a strategy for $A \subseteq \mathbb{A}\mathrm{gt}$ and $Q \subseteq \mathcal{Q}$. We build from $\mathcal{M} = \langle \mathbb{A}\mathrm{gt}, \mathcal{Q}, \Pi, \pi, \mathit{Acts}, d, o, \{\sim_a \mid a \in \mathbb{A}\mathrm{gt}\}\rangle$ the restricted model $\mathcal{M}_{s_A}^Q = \langle \mathbb{A}\mathrm{gt}, \mathcal{Q}, \Pi', \pi', \mathit{Acts}, d', o', \{\sim_a \mid a \in \mathbb{A}\mathrm{gt}\}\rangle$ as follows. Firstly, in each state $q \in \mathcal{Q}$ and agent $a \in A$ we limit the set of actions allowed to a to the one selected by s_a, i.e., $d_a(q)' = \{s_a(q)\}$. Secondly, we introduce a new, fresh proposition $\bar{Q}$ with labeling $\pi'(\bar{Q}) = Q$. All the remaining components of $\mathcal{M}_{s_A}^Q$ are inherited from $\mathcal{M}$.

Let $q \in \mathcal{Q}$, $\mathcal{V} \in \mathcal{V}\mathit{als}$, and $\phi \in$ af-AMC$_i$. We define the auxilliary operator AF by giving its denotation function as follows:

$$[\![AF_Q^{s_A}\phi]\!]_{\mu ir,\mathcal{V}}^{\mathcal{M}} = \{q \in \mathcal{Q} \mid \mathcal{M}_{s_A}^{Q}, q \models \bar{Q} \wedge \langle\!\langle \emptyset \rangle\!\rangle \bar{Q}\mathcal{U}\phi\}. \quad (2)$$

Intuitively, $\mathcal{M}, q \models_{ir}^{\mu} AF_Q^{s_A}\phi$ iff $q \in Q$ and all the paths starting in q eventually reach ϕ without leaving Q (except for the last step, possibly).

Now we are ready to define the new operator, called the *persistent next-step operator* $\langle\cdot\rangle^P$. Its denotation function is as follows:

$$[\![\langle A\rangle^P\phi]\!]_{\mu ir,\mathcal{V}}^{\mathcal{M}} = \{q \in \mathcal{Q} \mid \exists s_A \in \Sigma_A \text{ s.t. } [\![AF_{[q]_{\sim_E^A}}^{s_A}\phi]\!]_{\mu ir,\mathcal{V}}^{\mathcal{M}} = [q]_{\sim_E^A}\}, \quad (3)$$

where $A \subseteq \mathbb{A}\mathrm{gt}$. Observe that $\mathcal{M}, q \models_{ir}^{\mu} \langle A\rangle^P\phi$ iff there exists a strategy $s_A \in \Sigma_A$ such that every outcome $\lambda \in out_{\mathcal{M}}^{ir}(q, s_A)$ eventually reaches ϕ without leaving $[q]_{\sim_E^A}$. We therefore have the following proposition.

Proposition 2. *Let* $A \subseteq \mathbb{A}\mathrm{gt}$*,* $q \in \mathcal{Q}$*, and* $\phi \in$ af-AMC$_i$*. If* $\mathcal{M}, q \models_{ir}^{\mu} \langle A\rangle^P\phi$*, then* $\mathcal{M}, q \models_{ir} \langle\!\langle A\rangle\!\rangle\Diamond\phi$*.*

However, the operator of the persistent next-step is only an intermediate step of our construction. Let us define the new *persistent reachability* operator $\Diamond^P$ in the usual way, by putting $\langle A\rangle\Diamond^P\phi \equiv \mu X(\mathcal{E}_A\phi \vee \langle A\rangle^P X)$. One can observe that $\mathcal{M}, q \models_{ir}^{\mu} \langle A\rangle\phi$ implies $\mathcal{M}, q \models_{ir}^{\mu} \langle A\rangle^P\phi$, hence we have the following proposition.

Proposition 3. *Let* $A \subseteq \mathbb{A}\mathrm{gt}$*,* $q \in \mathcal{Q}$*, and* $\phi \in$ af-AMC$_i$*;* $\mathcal{M}, q \models_{ir}^{\mu} \langle\!\langle A\rangle\!\rangle\Diamond^P\phi$ *does not imply* $\mathcal{M}, q \models_{ir} \langle\!\langle A\rangle\!\rangle\Diamond\phi$*.*

Indeed, in the ICGS presented in Fig. 1 we have $\mathcal{M}, q \models_{ir}^{\mu} \langle\!\langle A\rangle\!\rangle\Diamond^P p$ and $\mathcal{M}, q \not\models_{ir} \langle\!\langle A\rangle\!\rangle\Diamond p$. Recall that the method of unification of strategies for the case of single agent, presented in the proof of Proposition 1, was based on the fact that the relation of everybody knows $\sim_A$ is an equivalence relation if $|A| = 1$. We thus alter the definition given in Eq. 3 by putting the relation of common knowledge in place of everybody knows to obtain the next type of the next-step operator.

Persistent Imperfect Next-Step Operator We define the *persistent imperfect next-step operator* $\langle\cdot\rangle^F$ by giving its denotation function:

$$[\![\langle A\rangle^F \phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = \{q \in \mathcal{Q} \mid \exists s_A \in \Sigma_A \text{ s.t. } [\![AF^{s_A}_{[q]_{\sim^A_C}} \phi]\!]^{\mathcal{M}}_{\mu ir,\mathcal{V}} = [q]_{\sim^A_C}\}, \quad (4)$$

where $\phi \in$ af-AMC$_i$, $A \subseteq \mathbb{A}\text{gt}$, and $\mathcal{V} \in \mathcal{V}als$. Intuitively, $\mathcal{M}, q \models^{\mu}_{ir} \langle A\rangle^F \phi$ if the coalition A has the common knowledge that ϕ is happening either now or in the next step. As previously, let us define a new, derived modality $\Diamond^F$ as follows:

$$\langle A\rangle \Diamond^F \phi \equiv \mu X(\mathcal{E}_A \phi \vee \langle A\rangle^F X), \quad (5)$$

for all $A \subseteq \mathbb{A}\text{gt}$, $\phi \in$ af-AMC$_i$. As shown in the following theorem, this type of reachability implies *ir*–reachability.

Theorem 1. *Let* $A \subseteq \mathbb{A}\text{gt}$*,* $q \in \mathcal{Q}$*,* $\phi \in$ af-AMC$_i$*, and* $\mathcal{V} \in \mathcal{V}als$*. If* $\mathcal{M}, q \models^{\mu}_{ir} \langle\!\langle A\rangle\!\rangle \Diamond^F \phi$*, then* $\mathcal{M}, q \models_{ir} \langle\!\langle A\rangle\!\rangle \Diamond \phi$*.*

Proof. The proof is similar to the proof of Proposition 1. As previously, let us define the sequence $\{F_j\}_{j\in\mathbb{N}}$ of formulae of af-AMC$_i$ such that $F_0 = \mathcal{E}_A\phi$ and $F_{j+1} = F_0 \vee \langle A\rangle^F F_j$, for all $j \geq 0$. In our proof we use the derived sequence $\{H_j\}_{j\in\mathbb{N}}$ of formulae of af-AMC$_i$ such that $H_j = \langle A\rangle^F F_j$ for all $j \in \mathbb{N}$. From Kleene fixed-point theorem we have $[\![\langle\!\langle A\rangle\!\rangle \Diamond^F \phi]\!]^{\mathcal{M}}_{\mu ir} = [\![F_0]\!]^{\mathcal{M}}_{\mu ir} \cup \bigcup_{j=0}^{\infty} [\![H_j]\!]^{\mathcal{M}}_{\mu ir}$. Observe that due to the fact that $\sim^A_C$ is an equivalence relation, for each $q \in \mathcal{Q}$ and $j \in \mathbb{N}$, if $[q]_{\sim^A_C} \cap [\![H_j]\!]^{\mathcal{M}}_{\mu ir} \neq \emptyset$, then $[q]_{\sim^A_C} \subseteq [\![H_j]\!]^{\mathcal{M}}_{\mu ir}$.

We prove that for each $j \in \mathbb{N}$ there exists a partial strategy s^j_A such that $dom(s^j_A) = [\![H_j]\!]^{\mathcal{M}}_{\mu ir}$, $\forall q \in dom(s^j_A)\ \forall \lambda \in out^{ir}_{\mathcal{M}}(q, s^j_A)\ \exists k \in \mathbb{N}\ \lambda_k \models^{\mu}_{ir} \mathcal{E}_A\phi$, and $s^j_A \subseteq s^{j+1}_A$. The proof is by induction on j. In the base case of $H_0 = \langle A\rangle^F \mathcal{E}_A\phi$ observe that if $q \in [\![H_0]\!]^{\mathcal{M}}_{\mu ir}$ then there exists a partial strategy $s^{0,q}_A$ with $dom(s^{0,q}_A) = [q]_{\sim^A_C}$ such that every $\lambda \in out^{ir}_{\mathcal{M}}(q, s^{0,q}_A)$ stays in $[q]_{\sim^A_C}$ until it reaches a state where $\mathcal{E}_A\phi$ holds. As $\sim^A_C$ is an equivalence relation, we can now define $s^0_A = \bigcup_{[q]_{\sim^A_C} \in \mathcal{Q}/\sim^A_C} s^{0,q}_A$, where any choice of the representative from a given abstraction class is correct. For the inductive step, we divide the construction of s^{j+1}_A

in two cases. Firstly, if $q \in [\![H_j]\!]^{\mathcal{M}}_{\mu ir}$, then we put $s^{j+1}_A(q) = s^j_A(q)$. Secondly, let $q \in [\![H_{j+1}]\!]^{\mathcal{M}}_{\mu ir} \setminus [\![H_j]\!]^{\mathcal{M}}_{\mu ir}$. In this case there exists a partial strategy $s^{j+1,q}_A$ with $dom(s^{j+1,q}_A) = [q]_{\sim^A_C}$ such that each outcome $\lambda \in out^{ir}_{\mathcal{M}}(q, s^{j+1,q}_A)$ stays in $[q]_{\sim^A_C}$ until it reaches a state q' in which either $\mathcal{E}_A\phi$ holds or $q' \in [\![H_j]\!]^{\mathcal{M}}_{\mu ir}$. In the latter case from the inductive assumption we know that following s^{j+1}_A (i.e., s^j_A, by the previous case) always leads to reaching $\mathcal{E}_A\phi$ without leaving $[\![H_j]\!]^{\mathcal{M}}_{\mu ir}$. We therefore put $s^{j+1}_A = \bigcup_{[q]_{\sim^A_C} \in \mathcal{Q}/\sim^A_C} s^{j+1,q}_A$, where, again, any choice of the representative from an abstraction class is correct.

Finally, we can build a partial strategy $s_A = \bigcup_{j \in \mathbb{N}} s^j_A$. Any extension s'_A of this strategy is such that for each $q \in \mathcal{Q}$, if $\mathcal{M}, q \models^{\mu}_{ir} \langle\!\langle A \rangle\!\rangle \Diamond^F \phi$ then a state in which $\mathcal{E}_A\phi$ holds is eventually reached along each outcome $\lambda \in out^{ir}_{\mathcal{M}}(q, s'_A)$. This concludes the proof. □

4 Conclusion

In this paper we have investigated the correspondence between model checking of af-AMC$_i$ and ATL$_{ir}$ on the example of reachability. We have identified some of the reasons for the fact that these logics are of uncomparable expressivity. These observations form the basis for a novel method for underapproximating ATL$_{ir}$ by means of fixed-point calculations. To this end, we introduce a special version of the next-step operator, called Persistent Imperfect Next-Step Operator $\langle\cdot\rangle^F$ and show how it can be used to define a new version of reachability that carries to ATL$_{ir}$.

In the future work we plan to analyse the properties of the operator $\langle\cdot\rangle^F$ in more detail. In particular, the reasons for which the relation of common knowledge seems to appear in a natural way need to be investigated. Moreover, we conjecture that this operator can be employed to building a new strategic logic, based on the syntax of ATL, that underapproximates ATL$_{ir}$.

Acknowledgments Michał Knapik is funded by NCN Preludium 4 DEC-2012/07/N/ST6/03426 grant. Wojciech Jamroga acknowledges the support of the 7th Framework Programme of the European Union under the Marie Curie IEF project ReVINK (PIEF-GA-2012-626398).

References

[1] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

[2] R. Bordini, M. Fisher, C. Pardavila, and M. Wooldridge. Model checking agentspeak. In *Proc. of the 2nd Int. Conf. on Autonomous Agents and Multi-Agent Systems (AAMAS'03)*, pages 409–416. ACM, 2003.

[3] N. Bulling and W. Jamroga. Alternating epistemic mu-calculus. In *Proc. of the 22nd Int. Joint Conf. on Artificial Intelligence, IJCAI*, pages 109–114. IJCAI/AAAI, 2011.

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

[5] P. Gammie and R. Meyden. MCK: Model checking the logic of knowledge. In *Proc. of the 16th Int. Conf. on Computer Aided Verification (CAV'04)*, volume 3114 of *LNCS*, pages 479–483. Springer-Verlag, 2004.

[6] A. Lomuscio, H. Qu, and F. Raimondi. MCMAS: A model checker for the verification of multi-agent systems. In *Proc. of 21st International Conference on Computer Aided Verification (CAV'09)*, volume 5643 of *LNCS*, pages 682–688. Springer, 2009.

[7] P. Schobbens. Alternating-time logic with imperfect recall. *Electr. Notes Theor. Comput. Sci.*, 85(2):82–93, 2004.

Pracę zgłosił Wojciech Penczek.

Adres autora:	Wojciech Jamroga, Michał Knapik Instytut Podstaw Informatyki PAN ul. Jana Kazimierza 5 01-248 Warszawa, Polska
E-mail:	mknapik@ipipan.waw.pl w.jamroga@ipipan.waw.pl

Symbol klasyfikacji rzeczowej (ACM): D.2.4, F.3.1, F.4.1

Na prawach rękopisu
Printed as manuscript

Nakład 100 egzemplarzy. Papier kserograficzny klasy III. Oddano do druku we wrześniu 2015 r. Wydawnictwo IPI PAN. ISSN 0138-0648.