

Model-Checking for Ability-Based Logics with Constrained Plans

Stéphane Demri¹, Raul Fervari^{2,3}

¹Université Paris-Saclay, CNRS, ENS Paris-Saclay,

Laboratoire Méthodes Formelles, 91190, Gif-Sur-Yvette, France

²FAMAF, Universidad Nacional de Córdoba & CONICET, Argentina

³Guangdong Technion - Israel Institute of Technology, China

Abstract

We investigate the complexity of the model-checking problem for a family of modal logics capturing the notion of “knowing how”. We consider the most standard ability-based knowing how logic, for which we show that model-checking is PSpace-complete. By contrast, a multi-agent variant based on an uncertainty relation between plans in which uncertainty is encoded by a regular language, is shown to admit a PTime model-checking problem. We extend with budgets the above-mentioned ability-logics, as done for ATL-like logics. We show that for the former logic enriched with budgets, the complexity increases to at least ExpSpace-hardness, whereas for the latter, the PTime bound is preserved. Other variant logics are discussed along the paper.

Introduction

Knowing How Logics. The epistemic concept of “knowing how” has received considerable attention lately, as a new way of providing formal foundations to strategic reasoning in AI and automated planning (see e.g., (van Ditmarsch et al. 2015)). Most formalisations for this notion are based on simple combinations of standard knowledge modalities and abilities (Herzig and Troquard 2006; van der Hoek and Lomuscio 2003). However, arguably, such an approach does not lead to a proper characterisation of “knowing how”, as discussed in (Herzig 2015; Jamroga and Ågotnes 2007).

In (Wang 2018b), a novel proposal to reason about “knowing how” assertions was introduced. The logic (herein written \mathcal{L}_{kh}) includes a goal-direct knowing how modality $\text{Kh}(p, q)$, indicating that “the agent knows how to achieve the goal q , whenever the initial condition p holds”. As stated in (Wang 2018b), there is an explicit intention to keep the logical language neat and to introduce simple semantical structures, namely labelled directed graphs. The interpretation of the new modality is *ability-based*: the formula $\text{Kh}(p, q)$ holds whenever there is a plan σ (understood as a finite sequence of actions) such that from all the states satisfying p , the execution of σ leads only to states satisfying q ; and any partial execution of σ can be extended to a complete execution of σ (known as the strong executability condition, directly inspired from conformant planning (Smith

and Weld 1998)). Thus, “knowing how” is given by the abilities described by the graph. The simplicity of the logical language is partly reflected by the fact that formulae of the form $\text{Kh}(p, q)$ are global, no action symbol appears in formulae, a single agent is considered, and no “knowing that” modality is present. A complete axiomatisation is provided in (Wang 2018b) but more importantly, such a work has been a source of inspiration for many others. Some variants include: multiple agents, other classes of plans, or admit “knowing that” operators (see e.g. (Fervari et al. 2017; Li and Wang 2021b)). Other approaches, related to strategic games and coalitions, have been studied in (Naumov and Tao 2018c,a,b). Finally, the logic studied in (Areces et al. 2021) (called herein \mathcal{L}^U) is based on a notion of indistinguishability over plans. Arguably, such a proposal provides a more epistemic view of knowing how than other approaches.

Substantial progress has been already done related to philosophical motivations, axiom systems and combinations with other epistemic operators. However, much less contributions exist about the possibility to constraint plans (for instance, by adding budgets), or to perform the model-checking task (instead of checking theoremhood). These issues are addressed in this paper.

Our Motivations. The need to investigate knowing how logics with constrained plans is advocated in (Li 2017, Ch. 3). This is our first motivation to revisit results about numerical constraints (i.e. budgets) and/or regularity constraints about plans, over this family of logics. Moreover, it is more natural to investigate model-checking instead of the validity/satisfiability problem, not only because it is a standard problem in formal verification, but also because the power of adding constraints is best illustrated there. This is due to the fact that the available plans are part of the input, and a few clever tricks cannot be used anymore, such as guessing simply shaped plans witnessing the truth of a formula, see e.g. the completeness proofs in (Li 2017; Areces et al. 2021). Also, most of the works are dedicated to Hilbert-style axiomatisation and to the satisfiability problem, and relatively little was done on model-checking, apart from the works (Li, Yu, and Wang 2017; Areces et al. 2021; Li and Wang 2021a). Since the semantics of such logics involve complex conditions such as strong executability, PTime upper bounds seemed more difficult to

reach than what is known about the minimal modal logic K or alternating-time temporal logic ATL (known to admit a PTime model-checking problem, see e.g. (Alur, Henzinger, and Kupferman 2002)). An additional motivation to study the model-checking problem for ability-based logics rests on consequences for the satisfiability problem. It is known that when a logic has the (bounded) small model property, a decision procedure for checking satisfiability consists in guessing a small model and then perform a model-checking test. This is done in (Li 2017, Th. 3.4.8) to decide the satisfiability status of formulae in an extension of \mathcal{L}_{kh} , in which the satisfaction of Kh-formulae requires the satisfaction of constraints for intermediate states. Then, knowing the complexity of the model-checking problem becomes crucial to evaluate the complexity of the satisfiability problem. This is the case in (Areces et al. 2021, Sec. 3.3), where the satisfiability problem for an uncertainty-based knowing how logic is shown in NP by stating that model-checking is in PTime.

Our Contributions. We study the model-checking problem for known ability-based logics from the literature, typically from (Wang 2015) and (Li 2017, Ch. 3), but also for extensions by adding arithmetical and/or regular constraints on plans. Furthermore, we provide new relationships of this family of logics with formal languages and automata theory.

As a warm-up, we show that the model-checking problem for the logic of knowing how from (Wang 2015) (written \mathcal{L}_{kh}) is PSpace-complete (Thm. 1). This untractability result came a bit as a surprise, but witnesses the high expressive power of the Kh-like modalities despite its simplicity. A variant of the logic \mathcal{L}_{kh} designed in (Areces et al. 2021) and featuring a notion of indistinguishable plans, has been shown therein to admit a PTime model-checking problem. Here, we assume that each of those indistinguishability classes is specified by a regular language, thus, can be potentially infinite. We prove that the model-checking problem for this extension \mathcal{L}_{reg}^U can be also solved in PTime, using an algorithm based on reachability checks (Thm. 2).

We extend the logics \mathcal{L}_{kh} and \mathcal{L}_{reg}^U so that actions have costs, and executing a plan requires to stay within a certain budget. Adding budget-like constraints to plans is advocated in (Li and Wang 2017; Cao and Naumov 2017). The extensions we propose follow a standard pattern already used in ATL-like logics (see e.g. (Alechina et al. 2017, 2018; Bulling and Goranko 2022)), in energy games (see e.g. (Bouyer et al. 2008)) and in general for other transition systems (see e.g. (Cao and Naumov 2017)). We show that the PTime upper bound is preserved if budgets are added to \mathcal{L}_{reg}^U (written $\mathcal{L}_{reg}^U(\star)$, Thm. 3). This is done by taking advantage of algorithms for solving decision problems related to vector addition systems with states (VASS) (see e.g. (Karp and Miller 1969)), restricted to a single counter, and using Bellman-Ford algorithm for weighted directed graphs (see e.g. (Cormen et al. 2022)). By contrast, we show that adding budgets to \mathcal{L}_{kh} increases the complexity significantly. The model-checking problem for $\mathcal{L}_{kh}(\star)$ (with resources) is ExpSpace-hard (Lemma 7). The subproblem in which the cost of any action does not depend on the state it is triggered, is shown ExpSpace-complete (Thm. 4).

Preliminaries on a Simple Knowing How Logic

We present below the knowing how logic introduced in (Wang 2015, 2018b) (herein written \mathcal{L}_{kh}). Formulae of the form $\text{Kh}(\varphi, \psi)$ can be read as “when φ is the case, the agent knows how to make ψ true”. Since \mathcal{L}_{kh} has a central position in this paper, we dedicate a separate section for preliminary definitions, and for its model-checking problem. Notions introduced here, such as LTS, plans and strong executability, will be helpful in forthcoming sections.

Formulae and models. Let Prop be a countably infinite set of propositional symbols. **Formulae of the logic \mathcal{L}_{kh}** are defined by the grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \text{Kh}(\varphi, \varphi) \quad (p \in \text{Prop}).$$

Other Boolean connectives are defined as usual.

In (Wang 2015, 2018b), formulae are interpreted over *labelled transition systems* (LTSs): relational models in which each (basic) relation indicates the source and target of a particular type of action the agent can perform. Let Act be a countably infinite set of action symbols. A **labelled transition system (LTS)** is a tuple $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, V)$, where:

- S is a non-empty set of states (called the domain);
- $(R_a)_{a \in \text{Act}}$ is a collection of binary relations on S ;
- $V : S \rightarrow 2^{\text{Prop}}$ is a labelling function.

Given an LTS \mathcal{S} and s a state, the pair (\mathcal{S}, s) (parentheses usually dropped) is called a **pointed LTS**. Before defining the satisfaction relation for the logic \mathcal{L}_{kh} , we need to introduce the notion of linear plan, that is used all along the paper (for more complex plans, see (Li and Wang 2021b, Def. 18)).

Linear plans. Let Act^* be the set of finite sequences over Act. Elements of Act^* are called **plans**, with ε being the **empty plan**. Given $\sigma \in \text{Act}^*$, let $|\sigma|$ be the length of σ ($|\varepsilon| \stackrel{\text{def}}{=} 0$). For $0 \leq k \leq |\sigma|$, the plan σ_k is σ 's initial segment up to (and including) the k th position (with $\sigma_0 \stackrel{\text{def}}{=} \varepsilon$). For $0 < k \leq |\sigma|$, the action $\sigma[k]$ is the one in σ 's k th position.

Let $(R_a)_{a \in \text{Act}}$ be a family of binary relations with each $R_a \subseteq S \times S$. Define $R_\varepsilon \stackrel{\text{def}}{=} \{(s, s) \mid s \in S\}$ and, for all $\sigma \in \text{Act}^*$ and $a \in \text{Act}$, $R_{\sigma a} \stackrel{\text{def}}{=} \{(s, t) \in S \times S \mid \exists t' \in S \text{ s.t. } (s, t') \in R_\sigma \text{ and } (t', t) \in R_a\}$. Take a plan $\sigma \in \text{Act}^*$: for $s \in S$ define $R_\sigma(s) \stackrel{\text{def}}{=} \{t \in S \mid (s, t) \in R_\sigma\}$. For all $X \subseteq S$ and $\Pi \subseteq \text{Act}^*$, define $R_\Pi(X) \stackrel{\text{def}}{=} \bigcup_{s \in X, \sigma \in \Pi} R_\sigma(s)$.

Intuitively, (Wang 2015, 2018a,b) handle how to express that an agent knows how to achieve ψ given φ , when she has an appropriate plan that allows her to go from any state in which φ holds only to states in which ψ holds. In order to characterise what ‘appropriate’ means, one can impose restrictions on what qualifies as an adequate plan.

Let $(R_a)_{a \in \text{Act}}$ as before. A plan $\sigma \in \text{Act}^*$ is **strongly executable (SE)** at $s \in S$ iff for all $k \in [0, |\sigma| - 1]$ and $t \in R_{\sigma_k}(s)$, we have $R_{\sigma[k+1]}(t) \neq \emptyset$. We define the set $\text{SE}(\sigma) \stackrel{\text{def}}{=} \{s \in S \mid \sigma \text{ is SE at } s\}$. Interestingly, in the next section we show that for any state s , the set $\{\sigma \in \text{Act}^* \mid s \in \text{SE}(\sigma)\}$ is a regular language.




Figure 1: Example of LTS.

Satisfaction relation (\models). Let $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, V)$ be an LTS, $s \in S$ and φ be a formula of \mathcal{L}_{kh} , \models is defined as:

$$\begin{aligned} \mathcal{S}, s \models p &\stackrel{\text{def}}{\iff} p \in V(s), \\ \mathcal{S}, s \models \neg\varphi &\stackrel{\text{def}}{\iff} \mathcal{S}, s \not\models \varphi, \\ \mathcal{S}, s \models \varphi \vee \psi &\stackrel{\text{def}}{\iff} \mathcal{S}, s \models \varphi \text{ or } \mathcal{S}, s \models \psi, \\ \mathcal{S}, s \models \text{Kh}(\varphi, \psi) &\stackrel{\text{def}}{\iff} \text{there exists } \sigma \in \text{Act}^* \text{ such that} \\ &\quad (1) \llbracket \varphi \rrbracket^{\mathcal{S}} \subseteq \text{SE}(\sigma) \text{ and} \\ &\quad (2) R_{\sigma}(\llbracket \varphi \rrbracket^{\mathcal{S}}) \subseteq \llbracket \psi \rrbracket^{\mathcal{S}}, \end{aligned}$$

with $\llbracket \chi \rrbracket^{\mathcal{S}} \stackrel{\text{def}}{=} \{s \in S \mid \mathcal{S}, s \models \chi\}$. If a plan $\sigma \in \text{Act}^*$ satisfies conditions (1) and (2) above, we say that σ **witnesses** the satisfaction of $\mathcal{S}, s \models \text{Kh}(\varphi, \psi)$. In the LTS \mathcal{S} of Fig. 1, we have that $\mathcal{S}, s \models \text{Kh}(p, q)$ (via the single-action plan a) whereas $\mathcal{S}, s \not\models \text{Kh}(p, r)$, as the plan ab is not strongly executable at $t \in \llbracket p \rrbracket^{\mathcal{S}}$.

A complete Hilbert-style system for \mathcal{L}_{kh} can be found in (Wang 2015) (see also (Wang 2018b, Sec. 3)). Decidability of the satisfiability problem for \mathcal{L}_{kh} is shown in (Li 2017, Ch. 3) (the proof is done for a more general logic), but its complexity characterisation is open, as far as we know.

Model-checking Problem for \mathcal{L}_{kh}

An LTS $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, V)$ is **finite** iff S , Act and Prop are finite sets. The **model-checking problem for the logic \mathcal{L}_{kh}** , written $\text{MC}(\mathcal{L}_{kh})$, is defined as follows.

Input: a finite LTS \mathcal{S} and a formula φ over Prop , $s \in S$.

Question: $\mathcal{S}, s \models \varphi$?

Though the model-checking problem for many modal and temporal logics can be solved in PTime , see e.g. (Blackburn, de Rijke, and Venema 2001; Demri, Goranko, and Lange 2016), below, we show that $\text{MC}(\mathcal{L}_{kh})$ is PSPACE -complete, requiring more computational resources. This high complexity was probably not expected from the works in the literature. Indeed, the model-checking problem is either not considered explicitly at all, or shown in PTime for some variant knowing how logics (see e.g. (Li and Wang 2021a,b)) or shown PSPACE -hard but for some expressive logic, see e.g. (Li, Yu, and Wang 2017).

PSPACE -completeness of $\text{MC}(\mathcal{L}_{kh})$. To show PSPACE -hardness, we reduce the nonemptiness problem for intersection of deterministic finite-state automata that is a PSPACE -complete problem, see e.g. (Galil 1976; Kozen 1977).

Lemma 1. $\text{MC}(\mathcal{L}_{kh})$ is PSPACE -hard.

Proof. (sketch) Given $N \in \mathbb{N}$, for each $i \in [1, N]$, let $\mathcal{A}_i = (Q_i, \text{Act}, \delta_i, q_i, F_i)$ be a deterministic finite-state automaton accepting the regular language $L(\mathcal{A}_i) \subseteq \text{Act}^*$. We assume that each automaton \mathcal{A}_i is complete and the sets of locations are disjoint. We define an LTS $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, V)$

over the set $\{\text{init}, \text{fin}\} \subseteq \text{Prop}$, with $S \stackrel{\text{def}}{=} \bigsqcup Q_i$ such that $\mathcal{S}, q_1 \models \text{Kh}(\text{init}, \text{fin})$ iff $L(\mathcal{A}_1) \cap \dots \cap L(\mathcal{A}_N) \neq \emptyset$ (q_1 is arbitrary), as follows.

- For all $q, q' \in S$ and $a \in \text{Act}$, $(q, q') \in R_a \stackrel{\text{def}}{\iff}$ there is $i \in [1, N]$ such that $q, q' \in Q_i$ and $q \xrightarrow{a} q' \in \delta_i$.
- For all $i \in [1, N]$ and $q \in Q_i$, $\text{init} \in V(q) \stackrel{\text{def}}{\iff} q = q_i$, and $\text{fin} \in V(q) \stackrel{\text{def}}{\iff} q \in F_i$.

$\text{MC}(\mathcal{L}_{kh})$ is PSPACE -hard, since $\text{Kh}(\text{init}, \text{fin})$ and \mathcal{S} are computed in logspace in the size of the input automata. \square

In order to establish that $\text{MC}(\mathcal{L}_{kh})$ is in PSPACE , we show a small plan property based on the regular structure of the set of plans witnessing the satisfaction of $\mathcal{S}, s \models \text{Kh}(\varphi_1, \varphi_2)$. By regularity, we mean that the set of plans can be shown to be a regular language, and we can *effectively* compute a finite-state automaton accepting the language of witness plans. Hence, though the regularity property is instrumental to prove our PSPACE upper bound, we believe it is also interesting for its own sake to understand the expressive power of the modality Kh in \mathcal{L}_{kh} . Below, we state a few properties before getting to the PSPACE upper bound.

Given an LTS $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, V)$, let us characterise the plans $\sigma \in \text{Act}^*$ such that either (non1) $\llbracket \varphi_1 \rrbracket^{\mathcal{S}} \not\subseteq \text{SE}(\sigma)$ or (non2) $R_{\sigma}(\llbracket \varphi_1 \rrbracket^{\mathcal{S}}) \not\subseteq \llbracket \varphi_2 \rrbracket^{\mathcal{S}}$. Let $\mathcal{A}_{(t_1, t_2)} = (Q, \text{Act}, \delta, I, F)$ be the automaton defined as follows.

- $Q \stackrel{\text{def}}{=} S$, $I \stackrel{\text{def}}{=} \{t_1\}$, $F \stackrel{\text{def}}{=} \{t_2\}$.
- For all $t, t' \in Q$ and $a \in \text{Act}$, $t \xrightarrow{a} t' \in \delta \stackrel{\text{def}}{\iff} (t, t') \in R_a$.

It is worth observing that $t_2 \in R_{\sigma}(t_1)$ iff $\sigma \in L(\mathcal{A}_{(t_1, t_2)})$. So, the plans σ satisfying (non2) are exactly those in

$$\bigcup \{L(\mathcal{A}_{(t_1, t_2)}) \mid t_1 \in \llbracket \varphi_1 \rrbracket^{\mathcal{S}}, t_2 \in \llbracket \neg\varphi_2 \rrbracket^{\mathcal{S}}\},$$

defined from at most $|S|^2$ languages. We write $\bar{\mathcal{A}}_{(t_1, t_2)}$ to denote the powerset automaton built from $\mathcal{A}_{(t_1, t_2)}$ such that $L(\bar{\mathcal{A}}_{(t_1, t_2)}) = \text{Act}^* \setminus L(\mathcal{A}_{(t_1, t_2)})$ (complement language). Hence, σ satisfies (2) iff it belongs to the language below.

$$\bigcap \{L(\bar{\mathcal{A}}_{(t_1, t_2)}) \mid t_1 \in \llbracket \varphi_1 \rrbracket^{\mathcal{S}}, t_2 \in \llbracket \neg\varphi_2 \rrbracket^{\mathcal{S}}\}.$$

The locations of the automaton $\bar{\mathcal{A}}_{(t_1, t_2)}$ are subsets of S .

Let us handle now the condition (non1). We have (non1) iff there is $t_1 \in \llbracket \varphi_1 \rrbracket^{\mathcal{S}}$ such that $t_1 \notin \text{SE}(\sigma)$, i.e. for some $k \in [0, |\sigma| - 1]$ and $u \in R_{\sigma_k}(t_1)$, $R_{\sigma_{k+1}}(u) = \emptyset$. Equivalently, (non1) iff there are $X_1, X_2 \subseteq S$ such that for some $k \in [0, |\sigma| - 1]$, $X_1 = R_{\sigma_k}(t_1)$, $X_2 = R_{\sigma_{k+1}}(t_1)$ and there is $u \in X_1$ such that for no $u' \in X_2$, we have $(u, u') \in R_{\sigma_{k+1}}$. This characterisation serves as the basis to define the set $\{\sigma \mid s \in \text{SE}(\sigma)\}$ for some state $s \in S$. Let $\mathcal{A}_s^* = (Q, \text{Act}, \delta, I, F)$ be the finite-state automaton defined as follows.

- $Q \stackrel{\text{def}}{=} \mathcal{P}(S) \times \{\text{acc}, \text{rej}\}$, $I \stackrel{\text{def}}{=} \{(\{s\}, \text{acc})\}$, $F \stackrel{\text{def}}{=} \mathcal{P}(S) \times \{\text{acc}\}$ ('acc' stands for 'acceptance', 'rej' for 'rejection').
- For all $X \in \mathcal{P}(S)$ and $a \in \text{Act}$, $(X, \text{rej}) \xrightarrow{a} (X, \text{rej}) \in \delta$.
- For all $X, X' \in \mathcal{P}(S)$ and $a \in \text{Act}$, $(X, \text{acc}) \xrightarrow{a} (X', \text{acc}) \in \delta \stackrel{\text{def}}{\iff} (\star) R_a(X) = X'$ and $(\star\star)$ for all $t \in X$, there is $t' \in X'$ such that $(t, t') \in R_a$.

- For all $X, X' \in \mathcal{P}(S)$ and $a \in \text{Act}$, $(X, acc) \xrightarrow{a} (X', rej) \in \delta \stackrel{\text{def}}{\iff} (\star)$ and not $(\star\star)$.

Lemma 2. $\{\sigma \mid s \in \text{SE}(\sigma)\} = \text{L}(\mathcal{A}_s^*)$.

Thus, the plan σ satisfies (1) iff it belongs to the language

$$\bigcap \{\text{L}(\mathcal{A}_t^*) \mid t \in \llbracket \varphi_1 \rrbracket^S\},$$

defined from at most $|S|$ languages; moreover, each automaton \mathcal{A}_t^* has at most $2 \cdot 2^{\mathfrak{F}}$ locations. In conclusion, the set of plans witnessing \mathcal{S} , $s \Vdash \text{Kh}(\varphi_1, \varphi_2)$ is equal to

$$\bigcap \{\text{L}(\mathcal{A}_t^*) \mid t \in \llbracket \varphi_1 \rrbracket^S\} \cap \bigcap \{\text{L}(\bar{\mathcal{A}}_{(t_1, t_2)}) \mid t_1 \in \llbracket \varphi_1 \rrbracket^S, t_2 \in \llbracket \neg \varphi_2 \rrbracket^S\},$$

which amounts to build an automaton with at most $2^{\mathfrak{F}} \cdot |S|^2 \cdot 2^{2 \cdot \mathfrak{F}} \cdot |S|$ locations.

As we have provided a way of characterising conditions (1) and (2), we can state an interesting property about the plans witnessing the satisfaction of \mathcal{S} , $s \Vdash \text{Kh}(\varphi_1, \varphi_2)$. This property has been unnoticed so far, as far as we can judge.

Corollary 1. *Let S be an LTS, s be a state and $\text{Kh}(\varphi_1, \varphi_2)$ be a formula. Let L be set of plans witnessing \mathcal{S} , $s \Vdash \text{Kh}(\varphi_1, \varphi_2)$. Then, L is a regular language and if $L \neq \emptyset$, then it contains a plan of length at most $2^{3 \cdot \mathfrak{F}} \cdot |S|^3$.*

We state the key property to get the PSpace upper bound.

Lemma 3. *Checking whether \mathcal{S} , $s \Vdash \text{Kh}(p, q)$ (with $p, q \in \text{Prop}$) can be done in polynomial space.*

Proof. (sketch) To test whether \mathcal{S} , $s \Vdash \text{Kh}(p, q)$, we check on-the-fly the non-emptiness of the product automaton

$$\times \{\mathcal{A}_t^* \mid t \in \llbracket p \rrbracket^S\} \times \times \{\bar{\mathcal{A}}_{(t_1, t_2)} \mid t_1 \in \llbracket p \rrbracket^S, t_2 \in \llbracket \neg q \rrbracket^S\}.$$

There are at most $|S| + |S|^2$ automata and each location from those automata is a subset of S , possibly enriched with a flag either *acc* or *rej*. Polynomial space is enough to store two consecutive locations of the product automaton, and its transition relation can be also computed in polynomial space based on the LTS \mathcal{S} . Nonemptiness of finite-state automata can be checked in NLogSpace (the number of states of the product automaton is bounded by $2^{3 \cdot \mathfrak{F}} \cdot |S|^3$). So, we get a NPSpace decision procedure for checking nonemptiness of the product automaton. By Savitch's Theorem (Savitch 1970), we get the PSpace upper bound. \square

It is time to characterise the complexity of $\text{MC}(\mathcal{L}_{kh})$.

Theorem 1. $\text{MC}(\mathcal{L}_{kh})$ is PSpace-complete.

PSpace-hardness is a direct consequence of Lemma 1. In order to establish PSpace-easiness, we can design a standard labelling algorithm taking advantage of Lemma 3 (this idea can be adapted to all the logics studied in this paper).

Variants problems. Our approach using formal languages theory provides us with modular and standard tools in order to characterise the complexity of model-checking problems for several ability-based logics. This is the case for instance, of the logic EPDL, known to be PSpace-complete (Li, Yu, and Wang 2017, Th. 4.3), for which our results can be easily adapted. On the other hand, alternative constraints on

plans can be considered. For instance, a modality $\text{Kh}_{\mathcal{A}}$ (parametric on the finite-state automaton \mathcal{A}) such that a witness plan must necessarily belong to $\text{L}(\mathcal{A})$. Another option is to consider variant logics from the literature, such as the logic \mathcal{L}_{khm} from (Wang 2018b, Sec. 4) and (Li 2017, Ch. 3)). \mathcal{L}_{khm} contains a ternary modality $\text{Kh}^m(\varphi_1, \varphi_2, \varphi_3)$, and witness plans must satisfy φ_2 all along the traversed paths. It is not difficult to see that we can apply our techniques for all these logics (details are omitted due to lack of space).

Regular Classes in Uncertainty-Based Logics

It is argued in (Areces et al. 2021) that an epistemic notion of knowing how should be not only based on the given abilities, but also on some notion of indistinguishability/uncertainty between them. Therein, it is also argued that such an approach makes the framework closer to standard epistemic logics (van Ditmarsch et al. 2015). For instance, it makes easier to move to a multi-agent setting.

In the *uncertainty-based* setting, agents share the same set of *affordances* (provided by the actual environment, typically an LTS). Still, they have different *abilities* depending on which of these affordances are available for each of them, and how well they can tell these affordances apart.

In this section, we assume that the indistinguishability classes of plans are specified by regular languages and we prove that the model-checking problem for this extension \mathcal{L}_{reg}^U is also in PTime.

Uncertainty-Based Knowing How

Let Agt be a finite set of agent symbols. **Formulae of the logic \mathcal{L}^U** are defined by the grammar:

$$\varphi ::= p \mid \neg \varphi \mid \varphi \vee \varphi \mid \text{Kh}_a(\varphi, \psi) \quad (p \in \text{Prop}, a \in \text{Agt}).$$

A formula $\text{Kh}_a(\varphi, \psi)$ is read as “when φ is the case, the agent a knows how to make ψ true”.

Here, each agent's knowledge is not given simply by the abilities described by the LTS. Instead, the models are enriched with an **uncertainty relation** \sim_a for each agent a , which is an equivalence relation over a non-empty subset $\Pi_a \subseteq \text{Act}^*$ describing those plans that are indistinguishable from each other, from the agent's perspective. Below, each relation \sim_a is represented by its set of equivalence classes U_a (therefore, U_a satisfies a few simple properties recalled below). A **multi-agent uncertainty-based LTS** (LTS^U) is a tuple $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, V)$ where $(S, (R_a)_{a \in \text{Act}}, V)$ is an LTS and each U_a assigns to the agent a a non-empty collection of pairwise disjoint non-empty sets of plans with: (1) $U_a \neq \emptyset$, (2) $\emptyset \notin U_a$, and

$$(3) \Pi_1, \Pi_2 \in U_a \text{ with } \Pi_1 \neq \Pi_2 \text{ implies } \Pi_1 \cap \Pi_2 = \emptyset.$$

Intuitively, $\Pi_a \stackrel{\text{def}}{=} \bigcup_{\Pi \in U_a} \Pi$ is the set of plans that the agent a has at her disposal. Similarly, as in classical epistemic logic, $\sim_a \subseteq \Pi_a \times \Pi_a$ describes agent a 's indistinguishability. This relation is not defined over possible states of affairs, but rather over her available plans.

Again, in this setting we need to define the notion of being a “proper plan”. Let $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, V)$ be an LTS^U . For $\Pi \subseteq \text{Act}^*$ and $X \cup \{s\} \subseteq S$, we define $R_\Pi \stackrel{\text{def}}{=} \bigcup_{\sigma \in \Pi} R_\sigma$, $R_\Pi(s) \stackrel{\text{def}}{=} \bigcup_{\sigma \in \Pi} R_\sigma(s)$, and $R_\Pi(X) \stackrel{\text{def}}{=} \bigcup_{\sigma \in \Pi} R_\sigma(X)$.

$\bigcup_{t \in X} R_{\Pi}(t)$. A set of plans $\Pi \subseteq \text{Act}^*$ is **strongly executable** at $s \in S$ iff every plan $\sigma \in \Pi$ is strongly executable at s . Hence, $\text{SE}(\Pi) \stackrel{\text{def}}{=} \bigcap_{\sigma \in \Pi} \text{SE}(\sigma)$ is the set of the states in S where Π is strongly executable.

The notion of satisfiability for Kh_a -formulae is defined as:

$$\mathcal{S}, s \models \text{Kh}_a(\varphi, \psi) \stackrel{\text{def}}{\iff} \text{there exists } \Pi \in U_a \text{ such that} \\ (1) \llbracket \varphi \rrbracket^{\mathcal{S}} \subseteq \text{SE}(\Pi) \text{ and } (2) R_{\Pi}(\llbracket \varphi \rrbracket^{\mathcal{S}}) \subseteq \llbracket \psi \rrbracket^{\mathcal{S}}.$$

Let \mathcal{S} be the LTS from Fig. 1, enriched with $U_a = \{\{a, ab\}\}$ and $U_b = \{\{a\}, \{ab\}\}$, for some $a, b \in \text{Agt}$. With this setup, the agent a cannot distinguish between the plans a and ab (i.e., she is uncertain about whether these plans lead to the same outcome or not), whereas the agent b considers them different. Thus, $\mathcal{S}, s \models \neg \text{Kh}_a(p, q) \wedge \text{Kh}_b(p, q)$.

The satisfiability problem for \mathcal{L}^U is NP-complete, whereas $\text{MC}(\mathcal{L}^U)$ is in PTime (Areces et al. 2021). In an instance of $\text{MC}(\mathcal{L}^U)$, each U_a is defined such that Π_a is finite. Representing equivalence relations by their indistinguishability classes results handy for designing model-checking algorithms, as we can directly deal with the list of available plans. However, there is a limitation: we only consider a finite set of them. For instance, assuming that $L(a^*b) \subseteq \Pi_a$, we may wish to express that for all $\sigma \neq \sigma' \in L(a^*b)$, we have $\sigma \sim_a \sigma'$ (the action a is silent for the agent a). In a more concrete example, we can consider that an agent a is a server sending connection requests to another server. The indistinguishability class for a is given by aa^* , where action a stands for “sending a connection request through a secure channel”. Thus, the class establishes that after the first request is sent by agent a , all other requests are ignored. This type of class cannot be handled when U_a is a finite set of finite sets of plans. Instead, below, the elements of U_a are defined as finite-state automata \mathcal{A} such that $L(\mathcal{A})$ is understood as a (possibly infinite) equivalence class. This is a standard way to represent finitely a (potential) infinite set of words.

Generalisation with Regular Equivalence Classes

Strictly speaking, we generalise the problem $\text{MC}(\mathcal{L}^U)$ from (Areces et al. 2021), by requiring elements in U_a to be regular languages (instead of finite sets), but defined through finite-state automata. On the logical side, the models for the logic \mathcal{L}_{reg}^U defined below, are a subclass of the models for \mathcal{L}^U (because we assume regularity of each equivalence class). Hence, the introduction of the logic \mathcal{L}_{reg}^U serves at least two purposes: to generalise $\text{MC}(\mathcal{L}^U)$ and to introduce models in which the indistinguishability classes are regular languages, which might seem confusing at first glance.

An **LTS^U with regular constraints (reg-LTS^U)** is a tuple $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, V)$ where each U_a assigns to agent a a non-empty collection of finite-state automata $\emptyset \neq U_a = \{\mathcal{A}_1, \mathcal{A}_2, \dots\}$, such that each $\mathcal{A} \in U_a$ defines an equivalence class over a set $\Pi_a = \bigcup_{\mathcal{A} \in U_a} L(\mathcal{A}) \subseteq 2^{\text{Act}^*}$; i.e., for all $\mathcal{A}_j, \mathcal{A}_k \in U_a, j \neq k$ implies $L(\mathcal{A}_j) \cap L(\mathcal{A}_k) = \emptyset$.

The clause for Kh_a -formulae over reg-LTS^Us becomes:

$$\mathcal{S}, s \models \text{Kh}_a(\varphi, \psi) \stackrel{\text{def}}{\iff} \text{there is } \mathcal{A} \in U_a \text{ such that} \\ (1) \llbracket \varphi \rrbracket^{\mathcal{S}} \subseteq \text{SE}(L(\mathcal{A})) \text{ and } (2) R_{L(\mathcal{A})}(\llbracket \varphi \rrbracket^{\mathcal{S}}) \subseteq \llbracket \psi \rrbracket^{\mathcal{S}}.$$

Unlike (Areces et al. 2021), each equivalence class over Π_a is defined by a regular language. Interestingly, theoremhood in \mathcal{L}^U and \mathcal{L}_{reg}^U coincide (see (Areces et al. 2021, Th. 2)).

A reg-LTS^U $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, V)$ is **finite** iff $S, \text{Act}, \text{Prop}$ and the U_a 's are finite sets. The **model-checking problem for the logic \mathcal{L}_{reg}^U** , written $\text{MC}(\mathcal{L}_{reg}^U)$, is defined over finite reg-LTS^U.

In (Areces et al. 2021) it has been shown that the model checking problem for \mathcal{L}^U can be solved in PTime. However, herein the main difficulty is dealing with automata generating potentially infinite languages, and thus, equivalence classes with infinitely many plans. One can show that the method used in the previous section applies also here, but giving us (roughly) an ExpSpace upper bound. However, as we will show here, this bound is not optimal. We shall use a more fine-tuned approach to handle the conditions related to the modality Kh_a . With this at hand, we will show that $\text{MC}(\mathcal{L}_{reg}^U)$ belongs actually to PTime.

The semantics of $\text{Kh}_a(\varphi, \psi)$ requires to find some $\mathcal{A} \in U_a$, such that $L(\mathcal{A})$ satisfies two conditions. As U_a is finite, inspecting its members one by one, poses no difficulties. Again our method relies on characterising all the potential candidates for the satisfiability of a Kh_a -formula that fail to be a proper witness. For $\mathcal{A} \in U_a$, we proceed as follows.

1. Checking condition (2) $R_{L(\mathcal{A})}(\llbracket \varphi \rrbracket^{\mathcal{S}}) \subseteq \llbracket \psi \rrbracket^{\mathcal{S}}$ can be handled by checking that

$$L(\mathcal{A}) \cap \bigcup \{L(\mathcal{A}_{(t_1, t_2)}) \mid t_1 \in \llbracket \varphi \rrbracket^{\mathcal{S}}, t_2 \in \llbracket \neg \psi \rrbracket^{\mathcal{S}}\} = \emptyset,$$

with $\mathcal{A}_{(t_1, t_2)}$ defined as in the previous section. By the analysis provided there, and since nonemptiness of the intersection of two finite-state automata can be checked in PTime, we have a PTime procedure for checking (2).

2. Then, we will design a PTime algorithm for checking the strong executability of $L(\mathcal{A})$ at $\llbracket \varphi \rrbracket^{\mathcal{S}}$. The algorithm relies on constructing a directed graph as a product between \mathcal{A} and \mathcal{S} . Then, we proceed by looking for a state t that fails to execute a “productive” action, i.e., an action that leads to a final state of \mathcal{A} , and checking whether (q, t) is reachable, for some location q that is accepting in \mathcal{A} . For the sake of correctness, we assume that all the locations in \mathcal{A} are productive, i.e. for any location, there is a path in \mathcal{A} leading to a final location.

Let $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, V)$ be a reg-LTS^U and $\mathcal{A} = (Q, \text{Act}, \delta, I, F)$ be an automaton. We define a digraph $\mathcal{G} = (V, E)$ such that $V \stackrel{\text{def}}{=} Q \times S$ and for all $(q, t), (q', t') \in V$, we have $(q, t) \rightarrow (q', t') \in E \stackrel{\text{def}}{\iff}$ there is some $a \in \text{Act}$ such that $q \xrightarrow{a} q' \in \delta$ and $(t, t') \in R_a$. \mathcal{G} is therefore a (standard) product between \mathcal{A} and \mathcal{S} . Given $s \in S$, we define the procedure $\text{CheckSE}(\mathcal{S}, s, \mathcal{A})$ as follows.

1. For each $a \in \text{Act}$, for each $t \in S$ such that $R_a(t) = \emptyset$ and for each $q \in Q$ such that $\delta(q, a) \neq \emptyset$ do:
 - (a) if there is $q_0 \in I$ such that there is a path from (q_0, s) to (q, t) in \mathcal{G} , then: return True;
2. return False.

Lemma 4. $s \notin \text{SE}(L(\mathcal{A}))$ iff $\text{CheckSE}(\mathcal{S}, s, \mathcal{A})$ returns True.

Now it is time to establish the complexity bound.

Theorem 2. $MC(\mathcal{L}_{reg}^U)$ is in PTime.

Proof. (sketch) PTime is guaranteed if $\mathcal{S}, s \Vdash \text{Kh}_a(\varphi, \psi)$ can be checked in PTime. U_a contains a linear amount of finite-state automata and we showed that (2) can be checked in PTime. It remains to verify that so is the case for (1). By Lemma 4, (1) holds iff for all $s' \in \llbracket \varphi \rrbracket^S$, $\text{CheckSE}(\mathcal{S}, s', \mathcal{A})$ returns False. As $|\llbracket \varphi \rrbracket^S| \leq |S|$ and the for loop ranges over all S , Act and Q , checking (a) is performed $\mathcal{O}(|S|^2 \cdot |\text{Act}| \cdot |Q|)$ times, which is polynomial in the size of the instance $\mathcal{S}, s \Vdash \text{Kh}_a(\varphi, \psi)$. Moreover, (a) is an instance of the graph accessibility problem (GAP), known to be NLogSpace-complete, applied on the graph \mathcal{G} of quadratic size in the size of \mathcal{S} . Thus, (1) can be checked in PTime. \square

Ability-Based Logics with Budgets

The need to express budget-like constraints about plans has been advocated in (Li 2017, Sec. 3.1) and (Li and Wang 2017). Assuming that actions have costs, the execution of plans requires that the agent stays always within the budget. Adding resource reasoning is a well-known paradigm used in ATL-like logics, see e.g. (Alechina et al. 2017), in energy games, see e.g. (Chatterjee, Doyen, and Henzinger 2017), and in multi-agent systems, see e.g. (Cao and Naumov 2017). Herein, we study the complexity of adding resource reasoning in the logics \mathcal{L}_{kh} and \mathcal{L}_{reg}^U , respectively. It is particularly interesting to observe is that the worst-case complexity stays within PTime for the extension of \mathcal{L}_{reg}^U whereas it jumps to at least ExpSpace-hardness for \mathcal{L}_{kh} . We provide clues to understand this computational gap.

Given a set of states S and a finite set of actions Act, a **weight function** is a map $wf : S \times \text{Act} \rightarrow \mathbb{Z}^r$ for some $r \geq 0$; $wf(s, a)$ is understood as the cost of executing the action a at state s . Adding wf to an LTS shall be our standard way to enrich models with action costs. Given a computation $\lambda = s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \cdots \xrightarrow{a_K} s_K$, its weight is defined as $wf(\lambda) \stackrel{\text{def}}{=} \sum_{k=1}^K wf(s_{k-1}, a_k)$ (empty computations have zero cost).

PTime Upper Bound for Model-checking $\mathcal{L}_{reg}^U(\star)$

We write $\mathcal{L}_{reg}^U(r)$ to denote the ability-based logic \mathcal{L}_{reg}^U augmented with $r \geq 0$ resource types, and $\mathcal{L}_{reg}^U(\star)$ to denote the version with an arbitrary number of resource types. This is the version used in our model-checking problem.

For $r \geq 0$, the set of $\mathcal{L}_{reg}^U(r)$ **formulae** is defined below.

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \text{Kh}_a^{\vec{b}}(\varphi, \varphi) \quad (p \in \text{Prop}, a \in \text{Agt}, \vec{b} \in \mathbb{N}^r)$$

All the integers appearing in formulae and models are encoded with a binary representation. Formulae of the form $\text{Kh}_a^{\vec{b}}(\varphi, \psi)$ are read as “when φ is the case, the agent a knows how to make ψ true with budget \vec{b} ”. Models of the logic $\mathcal{L}_{reg}^U(r)$ are of the form $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, wf, V)$, where $wf : S \times \text{Act} \rightarrow \mathbb{Z}^r$ is a weight function. A plan $\sigma = a_1 \cdots a_K$ is \vec{b} -**compatible** at s ($\vec{b} \in \mathbb{N}^r$) $\stackrel{\text{def}}{\iff}$ for every computation $\lambda = s_0 \xrightarrow{a_1} s_1 \cdots \xrightarrow{a_K} s_K$ with $s_0 = s$, we have for all $L \in [1, K]$, $\vec{b} + wf(\lambda_{\leq L}) \geq \vec{0}$ (with $\lambda_{\leq L} \stackrel{\text{def}}{=} s_0 \xrightarrow{a_1} s_1 \cdots \xrightarrow{a_L} s_L$). The




Figure 2: A reg-LTS^U with budgets ($r = 1$).

plan σ is \vec{b} -**compatible** at a set $X \subseteq S \stackrel{\text{def}}{\iff}$ it is \vec{b} -compatible at all $s \in X$. \vec{b} is understood as the **initial budget**.

Given a model $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, (U_a)_{a \in \text{Agt}}, wf, V)$ and $s \in S$, we update the satisfiability clause for $\text{Kh}_a^{\vec{b}}$ -formulae:

$$\begin{aligned} \mathcal{S}, s \Vdash \text{Kh}_a^{\vec{b}}(\varphi, \psi) &\stackrel{\text{def}}{\iff} \text{there is } \mathcal{A} \in U_a \text{ such that} \\ (1) \llbracket \varphi \rrbracket^{\mathcal{M}} &\subseteq \text{SE}(L(\mathcal{A})), \quad (2) R_{L(\mathcal{A})}(\llbracket \varphi \rrbracket^{\mathcal{M}}) \subseteq \llbracket \psi \rrbracket^{\mathcal{M}}, \\ \text{and } (3) \text{ for all } \sigma &\in L(\mathcal{A}), \sigma \text{ is } \vec{b}\text{-compatible at } \llbracket \varphi \rrbracket^{\mathcal{M}}. \end{aligned}$$

Consider the model \mathcal{S} of Fig. 2 (we use standard notations for finite-state automata). Both plans a and ab lead to q states, but the cost of executing ab from s to u is -1 , whereas from s to t is -6 . Thus, $\mathcal{S}, s \Vdash \text{Kh}_b^5(p, q) \wedge \neg \text{Kh}_a^5(p, q)$.

In the budget-free logic \mathcal{L}_{reg}^U , the clause related to the satisfaction of $\text{Kh}_a(\varphi, \psi)$ uses exactly (1) and (2) above. We have shown that such conditions can be checked in PTime. To prove that $MC(\mathcal{L}_{reg}^U(\star))$ is in PTime too, we establish that (3) can be checked in PTime for a given \mathcal{A} . Since $|S|$ is less than the size of \mathcal{S} and $MC(\mathcal{L}_{reg}^U(\star))$ can be solved by a standard type of labelling algorithm, it is sufficient to show that given \mathcal{S} , $t \in S$ and a finite-state automaton \mathcal{A} , one can check in PTime that for all $\sigma \in L(\mathcal{A})$, σ is \vec{b} -compatible at t . This is the purpose of the rest of this subsection.

Interestingly, the models for $\mathcal{L}_{reg}^U(r)$ can be viewed as an extension of the vector addition systems with states (VASS) (Karp and Miller 1969), since in both models the transitions are labelled by tuples in \mathbb{Z}^r encoding an update function. Below, we recall a few standard definitions about VASS that are helpful in the sequel to characterise the complexity of $MC(\mathcal{L}_{reg}^U(\star))$.

A **vector addition system with states (VASS)** is a structure $\mathcal{V} = (Q, r, R)$, where Q is a finite set of **locations**, $r \in \mathbb{N}$ is its **dimension**, and R is a finite set of **transitions** in $Q \times \mathbb{Z}^r \times Q$. A **configuration** (resp. **pseudo-configuration**) in a VASS \mathcal{V} is a pair $(q, \vec{x}) \in Q \times \mathbb{N}^r$ (resp. in $Q \times \mathbb{Z}^r$). Given pseudo-configurations (q, \vec{x}) , (q', \vec{x}') and a transition $T = q \xrightarrow{\vec{u}} q'$, we write $(q, \vec{x}) \xrightarrow{T} (q', \vec{x}')$ whenever $\vec{x}' = \vec{u} + \vec{x}$. A **pseudo-run** is defined as a sequence $\rho = (q_0, \vec{x}_0) \xrightarrow{T_1} (q_1, \vec{x}_1) \xrightarrow{T_2} (q_2, \vec{x}_2) \cdots$ of pseudo-configurations, where (q_0, \vec{x}_0) is the **initial** pseudo-configuration. A **run** is a pseudo-run in which only configurations in $Q \times \mathbb{N}^r$ occur.

An r -VASS is a VASS with $r \geq 0$ counters. We begin by presenting a simple problem called **NONSAFE(VASS)**, strongly related to the non-satisfaction of the condition (3).

Input: a VASS \mathcal{V} and a configuration $(q_0, \vec{x}_0) \in Q \times \mathbb{N}^r$.

Question: is there a finite pseudo-run $\rho = (q_0, \vec{x}_0) \rightarrow (q_1, \vec{x}_1) \rightarrow \cdots \rightarrow (q_n, \vec{x}_n)$ such that $\vec{x}_n \notin \mathbb{N}^r$? (is it possible to reach a negative value from (q_0, \vec{x}_0) ?).

Lemma 5. *NONSAFE(VASS) is in PTime.*

Lemma 5 follows from the fact that one instance of NONSAFE(VASS) can be reduced to r instances of NONSAFE(1-VASS) (restriction to 1-VASS). The second step consists in showing that checking whether a negative value can be reached in a 1-VASS from a given initial configuration can be solved using Bellman-Ford algorithm working on weighted directed graphs.

Assume that $\mathcal{A} = (Q, \text{Act}, \delta, I, F)$ and all the locations in Q are productive, so any run reaching a given location from some initial location can be completed as an accepting run.

One more step is needed to establish that (3) can be solved in PTime. Namely, we build a VASS $\mathcal{V} = (Q', r, R')$ with $Q' \stackrel{\text{def}}{=} S \times Q$ such that not (3) iff there is $q_0 \in I$ and $t \in \llbracket \varphi \rrbracket^S$ such that $\mathcal{V}, ((t, q_0), \vec{b})$ is a positive instance of NONSAFE(VASS). It remains to define R' . We have $(s, q) \xrightarrow{\vec{u}} (s', q') \in R' \stackrel{\text{def}}{\iff}$ for some $a \in \text{Act}$, $(s, s') \in R_a$ and $q \xrightarrow{a} q' \in \delta$ (synchronisation on actions) with $\vec{u} = wf(s, a)$.

Lemma 6. *There are $t \in \llbracket \varphi_1 \rrbracket^S$ and $\sigma \in L(\mathcal{A})$ such that for some $j \in [1, |\sigma|]$, σ_j is not \vec{b} -compatible at t iff there are $t \in \llbracket \varphi_1 \rrbracket^S$ and $q_0 \in I$ such that $\mathcal{V}, ((t, q_0), \vec{b})$ is a positive instance of NONSAFE(VASS).*

Now, we are in position to state our best result as far as a PTime model-checking problem is concerned.

Theorem 3. *$\text{MC}(\mathcal{L}_{reg}^U(\star))$ is in PTime.*

Let us briefly provide the argument to get PTime. The proof of Thm. 3 uses a labelling algorithm as for Thm. 1 and PTime is guaranteed as soon as $\mathcal{M}, s \Vdash \text{Kh}_a^{\vec{b}}(p, q)$ can be checked in PTime. Now, $\mathcal{M}, s \Vdash \text{Kh}_a^{\vec{b}}(p, q)$ iff there is $\mathcal{A} \in U_a$ such that (1), (2) and (3) hold. There is a linear amount of automata in U_a and Thm. 2 guarantees that (1) and (2) can be checked in PTime. The remaining bit is to check that (3) can be done in PTime. Notice that when (1) and (2) hold true, the first statement of Lemma 6 is equivalent to (3) being false. Thus, checking (3) follows from the combination of Lemmas 5 and 6.

Extending the Logic \mathcal{L}_{kh} with Budgets

Let $\mathcal{L}_{kh}(r)$ be the ability-based logic \mathcal{L}_{kh} augmented with $r \geq 0$ resource types. The logic $\mathcal{L}_{kh}(\star)$ denotes the version in which the number of resource types is arbitrary.

Models of $\mathcal{L}_{kh}(r)$ are of the form $\mathcal{S} = (S, (R_a)_{a \in \text{Act}}, wf, V)$ where $wf : S \times \text{Act} \rightarrow \mathbb{Z}^r$ is a weight function. The relation \Vdash is updated as follows.

$\mathcal{S}, s \Vdash \text{Kh}^{\vec{b}}(\varphi, \psi) \stackrel{\text{def}}{\iff}$ there is a plan $\sigma \in \text{Act}^*$ such that (1) and (2) as for \mathcal{L}_{kh} , and (3) σ is \vec{b} -compatible at $\llbracket \varphi \rrbracket^{\mathcal{M}}$.

We already showed that the conditions (1) and (2) from \mathcal{L}_{kh} can be encoded by a finite-state automaton of exponential size, and that $\text{MC}(\mathcal{L}_{kh})$ is PSpace-complete. However, $\text{MC}(\mathcal{L}_{kh}(\star))$ witnesses at least an exponential blow-up, as stated below, partly due to condition (3) combined with (2).

Lemma 7. *$\text{MC}(\mathcal{L}_{kh}(\star))$ is ExpSpace-hard.*

The proof is by reduction from the control-state reachability problem for VASS, written **CREACH(VASS)**, known to be ExpSpace-complete, see e.g. (Lipton 1976; Rackoff 1978). CREACH(VASS) takes as inputs a VASS \mathcal{V} , a configuration (q_0, \vec{x}_0) , and a location q_f and asks whether there is a run from (q_0, \vec{x}_0) to a configuration with location q_f .

Cor. 1 states that small witness plans in \mathcal{L}_{kh} have length at most exponential in the joint size of the LTS and the formula. By contrast, since the proof of Lemma 7 uses a reduction from CREACH(VASS) for which witness runs can be of length doubly-exponential in the size of the input VASS (Lipton 1976), the witness plans in $\mathcal{L}_{kh}(\star)$ may have length doubly-exponential too (if not more). Up to now, no known upper bound exists for the length of witness plans in $\mathcal{L}_{kh}(\star)$ and the decidability status of $\text{MC}(\mathcal{L}_{kh}(\star))$ is open.

However, it is possible to characterise the complexity of a natural fragment of $\text{MC}(\mathcal{L}_{kh}(\star))$ by requiring a simple restriction on the weight function in LTS: we consider below the subproblem of $\text{MC}(\mathcal{L}_{kh}(\star))$ in which the action costs do not depend on the states the actions are triggered. Hence, to conclude, we assume that wf is of the form $\text{Act} \rightarrow \mathbb{Z}^r$.

Theorem 4. *$\text{MC}(\mathcal{L}_{kh}(\star))$ restricted to LTS with action costs independent of states is ExpSpace-complete. For $r \in \mathbb{N}$, its restriction to r resources is PSpace-complete.*

ExpSpace-hardness is inherited from the proof of Lemma 7. For ExpSpace-easiness, we show that testing $\mathcal{S}, s \Vdash \text{Kh}^{\vec{b}}(p, q)$ can be reduced to an instance of CREACH(VASS) that can be checked in nondeterministic exponential space (refined analysis required here). When r is fixed, we regain PSpace.

Concluding Remarks

We investigated the complexity of the model-checking problem for ability-based logics, possibly with plans constrained by budget-like requirements and/or by regularity constraints. First, we established that for the well-known knowing how logic introduced in (Wang 2015), the problem is PSpace-complete (Thm. 1). Interestingly, we have shown that the set of witness plans for a given Kh-formula is regular, a property used all along the paper. Then, we propose a generalisation of the uncertainty-based logic of (Areces et al. 2021), in which each equivalence class over sets of plans is defined by a regular language. We show that model-checking for this extension \mathcal{L}_{reg}^U is in PTime, using an algorithm based on graph accessibility (Thm. 2). The final part of the paper is devoted to add budget-like constraints to \mathcal{L}_{kh} and \mathcal{L}_{reg}^U , following a paradigm used in many formalisms. While model-checking for the extension $\mathcal{L}_{reg}^U(\star)$ is shown in PTime (Thm. 3), we prove that for $\mathcal{L}_{kh}(\star)$ is ExpSpace-hard (Lemma 7). Moreover, the restriction in which the action costs do not depend on states, a reasonable assumption in many contexts, is ExpSpace-complete (Thm. 4). In all cases, we rely on results from vector addition systems.

A puzzling question remains the decidability status of $\text{MC}(\mathcal{L}_{kh}(\star))$. Our investigations can be also broadened by going beyond sequential plans (see e.g. (Li and Wang 2021b)) or by investigating alternative knowing how modalities (see e.g. (Fervari et al. 2017)).

Acknowledgments

We thank the reviewers for their constructive comments and suggestions. Stéphane Demri is supported by Centre National de la Recherche Scientifique. Raul Fervari is supported by ANPCyT-PICT-2020-3780, CONICET PIP 11220200100812CO, and by the Laboratoire International Associé SINFIN.

References

- Alechina, N.; Bulling, N.; Demri, S.; and Logan, B. 2018. On the Complexity of Resource-Bounded Logics. *Theoretical Computer Science*, 750: 69–100.
- Alechina, N.; Bulling, N.; Logan, B.; and Nguyen, H. 2017. The virtues of idleness: A decidable fragment of resource agent logic. *Artificial Intelligence*, 245: 56–85.
- Alur, R.; Henzinger, T.; and Kupferman, O. 2002. Alternating-time temporal logic. *Journal of the ACM*, 49(5): 672–713.
- Areces, C.; Fervari, R.; Saravia, A.; and Velázquez-Quesada, F. 2021. Uncertainty-Based Semantics for Multi-Agent Knowing How Logics. In *TARK'21*, volume 335 of *EPTCS*, 23–37.
- Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*. Cambridge University Press.
- Bouyer, P.; Fahrenberg, U.; Larsen, K.; Markey, N.; and Srba, J. 2008. Infinite Runs in Weighted Timed Automata with Energy Constraints. In *FORMATS'08*, volume 5215 of *LNCS*, 33–47. Springer.
- Bulling, N.; and Goranko, V. 2022. Combining quantitative and qualitative reasoning in concurrent multi-player games. *Autonomous Agents and Multi-Agent Systems*, 36(2): 1–33.
- Cao, R.; and Naumov, P. 2017. Budget-Constrained Dynamics in Multiagent Systems. In *IJCAI'17*, 915–921. ijcai.org.
- Chatterjee, K.; Doyen, L.; and Henzinger, T. 2017. The Cost of Exactness in Quantitative Reachability. In *Models, Algorithms, Logics and Tools - Essays Dedicated to Kim Guldstrand Larsen on the Occasion of His 60th Birthday*, volume 10460 of *LNCS*, 367–381. Springer.
- Cormen, T.; Leiserson, C.; Rivest, R. L.; and Stein, C. 2022. *Introduction to Algorithms, 4th Edition*. MIT Press.
- Demri, S.; Goranko, V.; and Lange, M. 2016. *Temporal Logics in Computer Science*. Cambridge University Press.
- Fervari, R.; Herzig, A.; Li, Y.; and Wang, Y. 2017. Strategically knowing how. In *IJCAI'17*, 1031–1038. ijcai.org.
- Galil, Z. 1976. Hierarchies of complete problems. *Acta Informatica*, 6: 77–88.
- Herzig, A. 2015. Logics of knowledge and action: critical analysis and challenges. *Autonomous Agents and Multi Agent Systems*, 29(5): 719–753.
- Herzig, A.; and Troquard, N. 2006. Knowing how to play: uniform choices in logics of agency. In *AAMAS'06*, 209–216. ACM.
- Jamroga, W.; and Ågotnes, T. 2007. Constructive knowledge: what agents can achieve under imperfect information. *Journal of Applied Non Classical Logics*, 17(4): 423–475.
- Karp, R. M.; and Miller, R. E. 1969. Parallel Program Schemata. *JCSS*, 3(2): 147–195.
- Kozen, D. 1977. Lower Bounds for Natural Proof Systems. In *FOCS'77*, 254–266. IEEE Computer Society.
- Li, Y. 2017. *Knowing what to do: a logical approach to planning ad knowing how*. Ph.D. thesis, University of Groningen.
- Li, Y.; and Wang, Y. 2017. Achieving While Maintaining: - A Logic of Knowing How with Intermediate Constraints. In *ICLA'17*, volume 10119 of *LNCS*, 154–167. Springer.
- Li, Y.; and Wang, Y. 2021a. Knowing How to Plan. In *TARK'21*, volume 335 of *EPTCS*, 233–247.
- Li, Y.; and Wang, Y. 2021b. Planning-based knowing how: A unified approach. *Artificial Intelligence*, 296: 103487.
- Li, Y.; Yu, Q.; and Wang, Y. 2017. More for free: a dynamic epistemic framework for conformant planning over transition systems. *JLC*, 27(8): 2383–2410.
- Lipton, R. 1976. The Reachability Problem Requires Exponential Space. Technical Report 62, Department of Computer Science, Yale University.
- Naumov, P.; and Tao, J. 2018a. Second-Order Know-How Strategies. In *AAMAS'18*, 390–398.
- Naumov, P.; and Tao, J. 2018b. Strategic Coalitions With Perfect Recall. In *AAAI'18*, 4702–4709. AAAI Press.
- Naumov, P.; and Tao, J. 2018c. Together we know how to achieve: An epistemic logic of know-how. *Artificial Intelligence*, 262: 279–300.
- Rackoff, C. 1978. The covering and boundedness problems for vector addition systems. *Theoretical Computer Science*, 6(2): 223–231.
- Savitch, W. 1970. Relationships between nondeterministic and deterministic tape complexities. *JCSS*, 4(2): 177–192.
- Smith, D.; and Weld, D. 1998. Conformant Graphplan. In *AAAI'98*, 889–896. AAAI Press / The MIT Press.
- van der Hoek, W.; and Lomuscio, A. 2003. Ignore at your peril - towards a logic for ignorance. In *AAMAS'03*, 1148–1149. ACM.
- van Ditmarsch, H.; Halpern, J.; van der Hoek, W.; and Kooi, B., eds. 2015. *Handbook of Epistemic Logic*. College Pub.
- Wang, Y. 2015. A Logic of Knowing How. In *LORI'15*, volume 9394 of *LNCS*, 392–405. Springer.
- Wang, Y. 2018a. Beyond knowing that: a new generation of epistemic logics. In van Ditmarsch, H.; and Sandu, G., eds., *Jaakko Hintikka on knowledge and game theoretical semantics*, 499–533. Springer.
- Wang, Y. 2018b. A logic of goal-directed knowing how. *Synthese*, 195(10): 4419–4439.