1 Introduction

Research in Cognitive Robotics [22, 23] is forcing the area of reasoning about actions to go through a "reality check". It has shown that, when one wants to equip an actual robot with the ability of reasoning about its actions, it is essential to take into account that the robot normally operates in an environment which it only partially knows, and that it must dynamically acquire new information through its sensors when needed [17]. Moreover, the basic reasoning task of projection is in general not sufficient to reason about sophisticated aspects of the robot's behavior, such as being always responsive to requests of other agents, or guaranteeing the successful termination of certain activities. Finally, reasoning about actions must be effective, i.e., reasonably efficient, in this generalized setting.

In this paper we propose a logical formalism for reasoning about actions which has the potentiality of meeting all the requirements above. Specifically, we define a variant of modal mu-calculus [24], a logic of programs that subsumes both propositional dynamic logics, such as standard PDL and ΔPDL [13], and branching time temporal logics such as CTL and CTL* [9]. Modal mu-calculus is used in the verification of concurrent systems [12, 20], and for this task several automated model checking techniques and systems have been developed [24, 3, 19].

We extend modal mu-calculus with an autoepistemic modal operator in order to represent and reason about the epistemic state of the robot. Autoepistemic operators have already been introduced for reasoning about actions e.g. in [14, 18]. Here, following [6, 7], we use a minimal knowledge modality in a way that strongly characterizes how the deliberative behavior of the robot is modeled: the robot may perform an action if it *knows* that the preconditions for that action hold, not simply if the preconditions are true. Similarly, the effects of an action of interest for the robot are only the ones the robot is aware of, i.e. the effects that change its epistemic state. This is obtained by specifying what a robot knows after an action, instead of specifying what is true after an action. In this approach, the robot follows the changes in the world through the changes in its epistemic state only. The minimal knowledge modality is also used to provide a natural formalization of *sensing actions*, i.e., actions that allow the robot to know whether a certain property holds in the current state of the world [17, 11, 2].

The special use of the minimal knowledge modality, that we require in the axioms specifying preconditions and effects of actions, forces a strong uniformity on the models of the dynamic system specification. This uniformity allows for representing all models by means of a single transition graph, whose nodes correspond to epistemic states of the robot, and transitions reflect how the robot's epistemic state changes by executing actions.

The proposed extension inherits from modal mu-calculus the ability of expressing very general dynamic and temporal properties. Moreover, by exploiting the possibility of representing all models of a dynamic system specification as a single graph, it becomes possible to adapt model checking techniques for the modal mu-calculus to our setting. Essentially, such model checking techniques are used to visit the graph in a suitable fashion, checking validity (instead of truth) of propositional formulae on single states, while traversing the graph.

2 Logical formalism

The technical background of our proposal is constituted by a logical formalism $\mathcal{L}$ that originates from a suitable integration of modal mu-calculus and autoepistemic description logics (see respectively [24, 9] and [8] for an introduction to these formalisms). The basic elements of $\mathcal{L}$ are a finite set of actions *Act*, a countable set of propositions *Prop*, and a countable set of propositional variables *Var*.

Formulae of the formalism are divided in two layers.

- *State description formulae*:

$$p ::= A \mid p_1 \wedge p_2 \mid \neg p$$

- *Dynamic formulae*:

$$\phi ::= \mathbf{k}p \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid [a]\phi \mid X \mid \mu X.\phi$$

  where $A \in Prop$, p is a state description formula, $a \in Act$ and $X \in Var$. The formula ϕ in $\mu X.\phi$ must be syntactically monotone in X, that is the variable X must be in the scope of an even number of negations.

We use the usual abbreviations $\vee, \supset$, tt, ff (the last two denoting tautology and contradiction respectively), and also the abbreviations $\langle a\rangle\phi \doteq \neg[a]\neg\phi$ and $\nu X.\phi \doteq \neg\mu X.\neg\phi[X/\neg X]$ where $[X/\neg X]$ denotes the syntactic substitution of X by $\neg X$.

We give the semantics of $\mathcal{L}$ by first fixing once and for all a countable-infinite set $\mathcal{S}$ of *state names* which will constitute the interpretation domain of $\mathcal{L}$. We assume to have a set of constants $Const \subseteq \mathcal{S}$ that are used to denote state names. We assume that *Const* contains the name s_{init}.

A *pre-interpretation* $\mathcal{I}$ is a function over $\mathcal{S}$ which assigns to each constant in *Const* the corresponding state name, i.e. $s^{\mathcal{I}} = s$; to each atomic proposition in *Prop* a subset of $\mathcal{S}$, i.e. $A^{\mathcal{I}} \subseteq \mathcal{S}$; and to each action $a \in Act$ a functional relation over $\mathcal{S}$, i.e. $a^{\mathcal{I}} \subseteq \mathcal{S} \times \mathcal{S}$, with the restriction that for every $s, s', s'' \in \mathcal{S}$ if $(s, s') \in a^{\mathcal{I}}$ and $(s, s'') \in a^{\mathcal{I}}$ then $s' = s''$. In addition, the union of the relations interpreting the actions is backward functional, i.e. for every $s, s', s'' \in \mathcal{S}$ if $(s', s) \in \cup_{a\in Act} a^{\mathcal{I}}$ and $(s'', s) \in \cup_{a\in Act} a^{\mathcal{I}}$ then $s' = s''$. Finally, we assume that, for every $s, s' \in \mathcal{S}$, for each $a \in Act$ and for each pre-interpretation $\mathcal{I}$, if $(s, s') \in a^{\mathcal{I}}$ then $s' \neq s_{init}$.

Pre-interpretations are extended to state description formulae as follows:

$$\begin{array}{lcl} (p_1 \wedge p_2)^{\mathcal{I}} & = & p_1^{\mathcal{I}} \cap p_2^{\mathcal{I}} \\ (\neg p)^{\mathcal{I}} & = & \mathcal{S} - p^{\mathcal{I}} \end{array}$$

A *valuation* ρ is a function from *Var* to a subset of $\mathcal{S}$ such that $\rho(X) \subseteq \mathcal{S}$ for every $X \in Var$. Given a valuation ρ and $\mathcal{E} \subseteq \mathcal{S}$,, we denote by $\rho[X/\mathcal{E}]$ the valuation obtained from ρ by changing to $\mathcal{E}$ the subset assigned to the variable X.

An *interpretation* $\mathcal{W}$ is a set of pre-interpretations over $\mathcal{S}$. We define interpretations of state formulae and actions respectively as:

$$\begin{array}{lcl} p^{\mathcal{W}} & = & \cap_{\mathcal{I}\in\mathcal{W}}\, p^{\mathcal{I}} \\ a^{\mathcal{W}} & = & \cap_{\mathcal{I}\in\mathcal{W}}\, a^{\mathcal{I}} \end{array}$$

Interpretations and valuations are used to interpret dynamic formulae as follows:

$$\begin{array}{lcl}
(\mathbf{k}p)^{\mathcal{W}}_{\rho} & = & p^{\mathcal{W}} \\
(\phi_1 \wedge \phi_2)^{\mathcal{W}}_{\rho} & = & (\phi_1)^{\mathcal{W}}_{\rho} \cap (\phi_2)^{\mathcal{W}}_{\rho} \\
(\neg\phi)^{\mathcal{W}}_{\rho} & = & \mathcal{S} - \phi^{\mathcal{W}}_{\rho} \\
([a]\phi)^{\mathcal{W}}_{\rho} & = & \{s \in \mathcal{S} \mid \forall s'.(s,s') \in a^{\mathcal{W}} \supset s' \in \phi^{\mathcal{W}}_{\rho}\} \\
X^{\mathcal{W}}_{\rho} & = & \rho(X) \\
(\mu X.\phi)^{\mathcal{W}}_{\rho} & = & \cap\{\mathcal{E} \subseteq \mathcal{S} \mid \phi^{\mathcal{W}}_{\rho[X/\mathcal{E}]} \subseteq \mathcal{E}\}
\end{array}$$

In particular we will be interested in closed formulae (formulae with no free variables). Such formulae are interpreted independently from the valuation, hence we will interpret them using an interpretation $\mathcal{W}$ alone: $\phi^{\mathcal{W}}$.

A *knowledge base* Σ is defined as a pair $\Sigma = (\mathcal{T}, \mathcal{A})$, where $\mathcal{T}$ is a finite set of state description formulae and (closed) dynamic formulae, and $\mathcal{A}$ is a finite set of assertions of the form $\psi(s)$ with ψ either a state description formula or a dynamic formula, and $s \in \mathit{Const}$.

An interpretation $\mathcal{W}$ *satisfies* a formula $\psi \in \mathcal{T}$ iff $\psi^{\mathcal{W}} = \mathcal{S}$. $\mathcal{W}$ satisfies an assertion $p(s) \in \mathcal{A}$ iff $s \in p^{\mathcal{W}}$. $\mathcal{W}$ satisfies a knowledge base $\Sigma = (\mathcal{T}, \mathcal{A})$ iff $\mathcal{W}$ satisfies every formula from $\mathcal{T}$ and every assertion from $\mathcal{A}$.

An interpretation $\mathcal{W}$ is a *model* for Σ iff $\mathcal{W}$ is a maximal set of interpretations satisfying Σ, i.e., for each interpretation $\mathcal{W}'$, if $\mathcal{W} \subset \mathcal{W}'$ then $\mathcal{W}'$ does not satisfy Σ. This corresponds to impose a "minimal knowledge" semantics on the epistemic states of the agent [8]. In fact, each interpretation can be viewed as a Kripke structure in which each pre-interpretation is a possible world, and each world is connected to all worlds in the structure: only structures satisfying Σ and having a maximal set of possible worlds are considered, which maximizes ignorance of the agent in its epistemic states.

Finally, Σ *logically implies* a formula or an assertion σ, written $\Sigma \models \sigma$, iff every model for Σ satisfies σ.

3 Dynamic system representation

In this section we present our framework for representing dynamic systems in the logic $\mathcal{L}$. The framework essentially follows the one presented in [6, 7].

The formalization of a dynamic system is constituted by the following elements.

- *Initial state description* is formed by a finite set of assertions of the form

  $$p(s_{init})$$

  where p is a state description formula and s_{init} is a constant in Const. In fact we may assume that $\mathit{Const} = \{s_{init}\}$.

- *Static axioms* (also known as *state constraints*) are a finite set of state description formulae p, which are assumed valid, defining invariance properties of states.

- *Precondition axioms* specify under which conditions an action can be executed. In our case such a condition depends on the epistemic state of the agent and not on what is true in the world. Precondition axioms are dynamic formulae of the the form:

  $$\mathbf{k}p \supset \langle a \rangle \mathbf{k}tt$$

- *Effect axioms* specify the effects of an ordinary (i.e., non-sensing) action when executed under certain conditions. Again, in our approach both effects and conditions concern the epistemic state of the agent. Effect axioms are dynamic formulae of the form:

  $$\mathbf{k}p_1 \supset [a]\mathbf{k}p_2$$

  No special treatment of the frame problem is considered here; we simply make use of frame axioms constituted by effect axioms of the form:

  $$\mathbf{k}p \supset [a]\mathbf{k}p$$

- *Sensing effect axioms* are effect axioms of a special form, which specify the outcome of a sensing action. Suppose a_f is a generic sensing action whose effect is to let the agent know the truth value of the property f, where f is any state formula. Also, suppose p is the precondition for the execution of a_f. Such a sensing action is represented in our framework by an usual action precondition axiom $\mathbf{k}p \supset \langle a_f \rangle \mathbf{k}tt$, plus the sensing effect axiom

  $$\mathbf{k}p \supset [a_f](\mathbf{k}f \vee \mathbf{k}\neg f)$$

  which formalizes the fact that, after the execution of a_f, the robot knows whether f holds or $\neg f$ holds.

We assume that the union of the static axioms and the initial state description is consistent. Observe that only propositional reasoning is needed in order to verify that the specification satisfies such an assumption.

Finally, for each sensing action a_f, we enforce a *frame axiom schema* of the form:

$$\mathbf{k}\varphi \supset [a_f]\mathbf{k}\varphi$$

which formalizes the fact that all the properties known by the robot before the execution of the sensing action are still known after executing it. Observe that, as a consequence of the frame axiom schema, if the robot already knows the truth-value of f then the sensing action a_f does not have any effect, in the sense that, if the robot knows f ($\neg f$), then after executing a_f the robot will still know f ($\neg f$). It is possible to show that the above axiom schema can be represented, without loss of generality, through a finite (linear) number of instances, by replacing φ in the schema with the initial state description and with each effect appearing in effect axioms.

Let Σ be the knowledge base describing the dynamic system as above. We are interested in verifying if the system satisfies a certain dynamic property. Formally, we are interested in logical inference of the form

$$\Sigma \models \phi(s_{init}) \tag{1}$$

where ϕ can be any dynamic formula. As we shall see later, in this way we can deal for instance with the *projection problem*, "given a sequence of actions, does a given state description formula hold in the resulting state?"; the *planning problem*, "is there a sequence of actions such that the goal (a state description formula) holds in the resulting state?"; but also very sophisticated dynamic properties such as *liveness*, *safeness*, etc. that are easily expressed using fixpoint formulae.

4 Reasoning technique

Let us now turn our attention to the problem of computing the logical implication (1).

First of all, an $\mathcal{L}$ knowledge base Σ corresponding to a dynamic system specification has in general many models. However, if sensing actions are not considered, all models of Σ are isomorphic up to renaming of states. Hence, it is possible to do reasoning using a single model, since it can be shown that all the properties that are expressible in the right-hand side of (1) are independent of such state names. On the other hand, the presence of sensing effect axioms in Σ causes the existence of models which structurally differ from each other, in general. This can be intuitively explained by the fact that when the robot uses its sensing capabilities to know whether a certain boolean property p holds, its epistemic state changes according to one of the two possible outcomes of the sensing action. Hence, two different models for Σ represent the two different possible epistemic states resulting from the execution of the sensing action.

We are going to show that wrt logical implication, all the models of Σ can be taken into account by means of a single graph, called the *transition graph* (TG) of Σ. Roughly speaking, the transition graph of Σ is a graph in which:

- each node corresponds to a state and is labeled with propositional formulae representing the properties which are known in such a state;
- each edge is labeled with an action name, and denotes the transition caused by the execution of the corresponding action.

The idea behind the notion of transition graphis to exploit the correspondence between the notions of robot's *knowledge* (about propositional properties) and propositional *validity* that arises from the particular structure of the axioms in the specification Σ. What the robot knows in the initial state is the set of propositional formulae which are valid in s_{init}, i.e., the set of propositional formulae which are logically implied by $p(s_{init})$. Moreover, what the robot knows after executing an ordinary action is the set of propositional formulae which are logically implied by the postconditions representing the effects of the action execution, while what the robot knows after executing a sensing action is the truth value of the sensed fluent, plus the knowledge of the robot before executing the sensing action. Finally, to verify whether an action can be executed (that is, whether the preconditions are known by the robot) we simply check for the validity of the action precondition.

Formally, $TG(\Sigma) = (S, L_S, L_A)$, where $S \subseteq \mathcal{S}$ is the set of states which includes s_{init}, L_S is a function assigning a finite set of propositional formulae to each state in S, and L_A is a partial function assigning a state to a pair formed by a state and an action. Given a dynamic system specification $\Sigma = (\mathcal{T}, \mathcal{A})$, where $\mathcal{T}$ is the set of static axioms (Γ_S), precondition axioms (Γ_P), effect axioms (Γ_E), and initial state description $\mathcal{A} = \{p(s_{init})\}$, the transition graph $TG(\Sigma)$ is computed by the algorithm shown in Fig. 1.

In order to understand the algorithm, let us not consider sensing action at first. The algorithm, starting from the initial state s_{init}, iteratively proceeds as follows. First, it finds an action a which can be executed in the current state, by identifying in the set Γ_P a precondition axiom for a whose left-hand side is logically implied by the current knowledge base. Then, it propagates the effects of the action a, which again is based on checking whether the left-hand side of each effect axiom for a in the set Γ_E

ALGORITHM TG
INPUT: $\Sigma = (\Gamma_S \cup \Gamma_P \cup \Gamma_E, \{p(s_{init})\})$
OUTPUT: $TG(\Sigma)$ or *inconsistent*
PROCEDURE CREATE_NEW_STATE(s, a)
begin
  s' = NEW state name;
  $Prop(s') = \{q | (\mathbf{k}p' \supset [a]\mathbf{k}q \in \Gamma_E) \wedge (\Gamma_S \cup L_S(s) \models p')\}$;
  if $Prop(s') \cup \Gamma_S \models ff$
  then return(*inconsistent*);
  if there exists $s'' \in S$ **such that**
    $(\Gamma_S \cup L_S(s'))$ and $(\Gamma_S \cup L_S(s''))$ are logically equiv.
  then $L_A(s, a) = s''$
  else begin
    $S_{active} = S_{active} \cup \{s'\}$;
    $S = S \cup \{s'\}$;
    $L_S(s') = Prop(s')$
  end
end
begin
  $S_{active} = \{s_{init}\}$;
  $S = \{s_{init}\}$;
  $L_S(s_{init}) = \{p(s_{init})\}$;
  repeat
    s = choose(S_{active});
    for each ordinary action a **do**
      if $(\mathbf{k}p \supset \langle a \rangle \mathbf{k}tt) \in \Gamma_P$ **and** $(\Gamma_S \cup L_S(s) \models p)$
      then CREATE_NEW_STATE(s, a);
    for each sensing action a_f **do**
      if $(\mathbf{k}p \supset \langle a_f \rangle \mathbf{k}tt) \in \Gamma_P$ **and**
        $(\Gamma_S \cup L_S(s) \models p)$ **and**
        $(\Gamma_S \cup L_S(s) \not\models f)$ **and**
        $(\Gamma_S \cup L_S(s) \not\models \neg f)$
      then begin
          CREATE_NEW_STATE(s, a_f^+);
          CREATE_NEW_STATE(s, a_f^-)
      end;
    $S_{active} = S_{active} - \{s\}$
  until $S_{active} = \emptyset$;
  return (S, L_S, L_A)
end;

Figure 1: Algorithm computing $TG(\Sigma)$

is logically implied by the properties holding in the current state. In this way, the set of properties corresponding to the effect of the execution of a in the current state is computed. A new state (or two new states, if a is a sensing action) is then generated, unless a state with the same properties has already been created. This step is repeated until all actions executable in the current state have been considered. Then, a new current state is chosen among those previously created and the main iteration proceeds.

The transition graph is unique, that is, every order of extraction of the states from the set S_{active} produces the same set of assertions, up to the renaming of states. Moreover, the algorithm terminates, that is, the condition $S_{active} = \emptyset$ is eventually reached, since the number of states generated is bounded to the number of different subsets of the set $\mathcal{E} = \{q \mid \mathbf{k}p' \supset [a]\mathbf{k}q \in \Gamma_E\}$, i.e., 2^n, where n is the number of axioms in Γ_E. Finally, the condition

$$(\Gamma_S \cup L_S(s')) \text{ and } (\Gamma_S \cup L_S(s'')) \text{ are logically equivalent}$$

can be verified by a propositional validity check, as well as the propositional logical implication

$$\Gamma_s \cup L_s(s) \models p$$

Now, let us also define the *extension of a dynamic formula in* $TG(\Sigma)$ wrt a valuation ρ as follows:

$$\begin{array}{lcl}
(\mathbf{k}p)_\rho^{TG(\Sigma)} & = & \{s \in S \mid \Gamma_s \cup L_s(s) \models p\} \\
(\phi_1 \wedge \phi_2)_\rho^{TG(\Sigma)} & = & (\phi_1)_\rho^{TG(\Sigma)} \cap (\phi_2)_\rho^{TG(\Sigma)} \\
(\neg\phi)_\rho^{TG(\Sigma)} & = & S - (\phi)_\rho^{TG(\Sigma)} \\
([a]\phi)_\rho^{TG(\Sigma)} & = & \{s \in S \mid \forall s'.(L_A(s,a) = s') \supset \\
 & & \qquad (s' \in \phi_\rho^{TG(\Sigma)})\} \\
X_\rho^{TG(\Sigma)} & = & \rho(X) \\
(\mu X.\phi)_\rho^{TG(\Sigma)} & = & \cap\{E \subseteq S \mid \phi_{\rho[X/E]}^{TG(\Sigma)} \subseteq E\}
\end{array}$$

In fact, we are interested in closed formulae ϕ, whose extension in $TG(\Sigma)$ is independent of the valuation: each such formula will be denoted simply by $\phi^{TG(\Sigma)}$.

We now prove that entailment with respect to a specification Σ can be reduced to reasoning in $TG(\Sigma)$. To this aim, we prove some preliminary properties of $TG(\Sigma)$.

Given a set of pre-interpretations $\mathcal{W}$ and a state $s \in \Delta$, we denote as $VALID_{\mathcal{W}}(s)$ the set of propositional formulae φ such that $s \in (\varphi)^{\mathcal{W}}$, and denote as $ACTIONS_{\mathcal{W}}(s)$ the set of actions a such that $(s, s') \in (a)^{\mathcal{W}}$ for some $s' \in \Delta$. Moreover, given a set of propositional formulae Ψ, we denote as $Cn(\Psi)$ the propositional deductive closure of Ψ, i.e., the set of propositional formulae which are propositionally entailed by Ψ.

Given a specification Σ, we denote as $\mathcal{M}_{TG(\Sigma)}$ the set of all the pre-interpretations $\mathcal{I}$ over Δ satisfying the following conditions:

1. for each $s \in S$, if $\varphi \in Cn(\Gamma_S \cup L_S(s))$ then $s \in (\varphi)^{\mathcal{I}}$;
2. for each $s, s' \in \Delta$ and for each action $a \in Act$, if $L_A(s,a) = s'$ then $(s, s') \in (a)^{\mathcal{I}}$

¿From the above definition, it immediately follows that, since $\mathcal{M}_{TG(\Sigma)}$ is the maximal set of pre-interpretations satisfying the above conditions, for

each $s \in S$, $VALID_{\mathcal{M}_{TG(\Sigma)}}(s) = Cn(\Gamma_S \cup L_S(s))$, and, for each $s, s' \in \Delta$ and for each action $a \in Act$, $(s, s') \in (a)^{\mathcal{M}_{TG(\Sigma)}}$ iff $L_A(s, a) = s'$.

The following property is an immediate consequence of the previous definition.

Lemma 1 *Let $s \in S$, $\phi \in \mathcal{L}$. Then, $s \in \phi^{TG(\Sigma)}$ iff $\mathcal{M}_{TG(\Sigma)} \models \phi(s)$.*

We now turn our attention to the models of a specification Σ.

First of all, due to the assumption that the set of propositional formulae composed of the union of the static axioms and the initial state description is consistent, it is immediate to verify that, if the specification Σ has no models, then the algorithm computing $TG(\Sigma)$ returns *inconsistent*. Therefore, from now on we assume that the specification Σ has at least one model.

Given a set of pre-interpretations $\mathcal{W}$ over Δ, we denote as $\Delta_0^{\mathcal{W}}$ the following subset of Δ:

$$\Delta_0^{\mathcal{W}} = \{s, s' \mid \exists a \in Act \text{ such that } (s, s') \in (a)^{\mathcal{W}}\}$$

Then, we denote as $\mathcal{W}_0$ the set of pre-interpretations over $\Delta_0^{\mathcal{W}}$ obtained from $\mathcal{W}$ by restricting each pre-interpretation in $\mathcal{W}$ to the elements of $\Delta_0^{\mathcal{W}}$.

Lemma 2 *Let Σ be a specification of a dynamic system, and let $\mathcal{M}$ be a model for Σ. Then, for each $\phi \in \mathcal{L}$, $s \in \Delta_0^{\mathcal{M}}$, $\mathcal{M} \models \phi(s)$ iff $\mathcal{M}_0 \models \phi(s)$.*

Proof. By induction on the number of nested fixpoints. Base case. If no fixpoints are present in ϕ, then we proceed by induction on the structure of ϕ, using the property that in each model $\mathcal{M}$ the pre-interpretation of $\phi(s)$ in $\mathcal{M}$ only depends on the state s and on the states which can be transitively reached from s through the union of the relations $(a)^{\mathcal{M}}$. Inductive case. Let us assume that the thesis holds for the formula ψ with k nested fixpoint constructs. We prove it for $\phi = \mu X.\psi$ with $k+1$. We recall that, by the Tarski-Knaster Theorem on fixpoints [26], $s \in (\mu X.\psi)_\rho^{\mathcal{W}}$ iff there exists an ordinal ξ such that $s \in (\mu_\xi X.\psi)_\rho^{\mathcal{W}}$, where $(\mu_\xi X.\psi)_\rho^{\mathcal{W}}$ is defined by transfinite induction as:

- $(\mu_0 X.\psi)_\rho^{\mathcal{W}} = \emptyset$
- $(\mu_{\xi+1} X.\psi)_\rho^{\mathcal{W}} = \psi_{\rho[X/(\mu_\xi X.\psi)_\rho^{\mathcal{W}}]}^{\mathcal{W}}$
- $(\mu_\lambda X.\psi)_\rho^{\mathcal{W}} = \bigcup_{\xi<\lambda} (\mu_\xi X.\psi)_\rho^{\mathcal{W}}$, if λ is a limit ordinal.

Hence, we proceed by transfinite induction on ordinals ξ. □

Let us now restrict our attention to specifications which do not contain sensing actions. Given any mapping $\eta : \Delta_0^{\mathcal{M}} \to S$ and a pre-interpretation $\mathcal{I}$ over $\Delta_0^{\mathcal{M}}$, we denote as $\eta(\mathcal{I})$ the pre-interpretation over S obtained as follows:

- for each $A \in Prop$

$$(A)^{\eta(\mathcal{I})} = \{s' \mid s' = \eta(s) \wedge s \in (A)^{\mathcal{I}}\}$$

- for each $a \in Act$

$$(a)^{\eta(\mathcal{I})} = \{(s'_1, s'_2) \mid s'_1 = \eta(s_1) \wedge s'_2 = \eta(s_2) \wedge (s_1, s_2) \in (a)^{\mathcal{I}}\}$$

Furthermore, given a set of pre-interpretations $\mathcal{W}$, we denote as $\eta(\mathcal{W}_0)$ the set of interpretations $\{\eta(\mathcal{I}) \mid \mathcal{I} \in \mathcal{W}_0\}$.

Lemma 3 *Let Σ be a specification of a dynamic system without sensing actions, and let $\mathcal{M}$ be a model for Σ. Then, there exists a mapping $\eta_{\mathcal{M}} : \Delta_0^{\mathcal{M}} \rightarrow S$ such that $\eta_{\mathcal{M}}(\mathcal{M}_0) = \mathcal{M}_{TG(\Sigma)}$.*

Proof. First, it is immediate to see that $VALID_{\mathcal{M}}(s_{init}) = Cn(\Gamma_S \cup L_S(\eta(s_{init}))$, since the union of relations interpreting actions is backward functional. Moreover, by definition of $TG(\Sigma)$, for each state $s \in \Delta$ there exist a state s' and an action a such that $(s, s') \in (a)^{\mathcal{M}}$ iff there exists a state s'' such that $L_A(\eta(s), a) = s''$ in $TG(\Sigma)$, and $VALID_{\mathcal{M}}(s') = Cn(\Gamma_S \cup L_S(\eta(s'')))$ in $TG(\Sigma)$. Therefore, since the union of relations interpreting actions is backward functional, we can constructively define the mapping $\eta_{\mathcal{M}}$ as follows:

1. $\eta_{\mathcal{M}}(s_{init}) = s_{init}$;
2. $\eta_{\mathcal{M}}(s') = s''$ for each state s' such that there exist a state s (for which $\eta_{\mathcal{M}}(s)$ has already been defined) and an action a such that $(s, s') \in (a)^{\mathcal{M}}$ and $L_A(\eta_{\mathcal{M}}(s), a) = s''$ in $TG(\Sigma)$.

For such a mapping $\eta_{\mathcal{M}}$ we have that $\eta_{\mathcal{M}}(\mathcal{M}_0) = \mathcal{M}_{TG(\Sigma)}$, hence the thesis follows. □

Lemma 4 *Let Σ be a specification of a dynamic system without sensing actions, and let $\mathcal{M}$ be a model for Σ. Then, for each $\phi \in \mathcal{L}$, and for each state $s \in \Delta_0^{\mathcal{M}}$, $\mathcal{M}_0 \models \phi(s)$ iff $\eta_{\mathcal{M}}(\mathcal{M}_0) \models \phi(s)$.*

Proof. Let $\mathcal{R}$ be the relation defined as follows:

$$\mathcal{R} = \{(s, t) \in \Delta_0^{\mathcal{M}} \times S \mid t = \eta_{\mathcal{M}}(s)\}$$

We show that, for each $\mathcal{I} \in \mathcal{M}_0$, $\mathcal{R}$ is a bisimulation between $\mathcal{I}$ and $\eta_{\mathcal{M}}(\mathcal{I})$. Indeed, it is easy to verify that $\mathcal{R}$ is closed under the rules that define a bisimilarity relation (see e.g. [25]), i.e., if $(s, t) \in \mathcal{R}$ then:

1. for each $A \in Prop$, $s \in (A)^{\mathcal{I}}$ iff $t \in (A)^{\eta_{\mathcal{M}}(\mathcal{I})}$;
2. for each $a \in Act$ and for each s' such that $(s, s') \in (a)^{\mathcal{I}}$, there exists t' such that $(t, t') \in (a)^{\eta_{\mathcal{M}}(\mathcal{I})}$ and $(s', t') \in \mathcal{R}$;
3. for each $a \in Act$ and for each t' such that $(t, t') \in (a)^{\eta_{\mathcal{M}}(\mathcal{I})}$, there exists s' such that $(s, s') \in (a)^{\mathcal{I}}$ and $(s', t') \in \mathcal{R}$.

Since $\mathcal{I}$ and $\eta_{\mathcal{M}}(\mathcal{I})$ are bisimilar, by a standard result in μ-calculus (see e.g. [25]) it follows that, for every $\phi \in \mathcal{L}$ and for every $s \in \Delta_0^{\mathcal{M}}$, $s \in (\phi)^{\mathcal{I}}$ iff $\eta_{\mathcal{M}}(s) \in (\phi)^{\eta_{\mathcal{M}}(\mathcal{I})}$. Hence, the thesis follows. □

Theorem 5 *Let Σ be a specification of a dynamic system without sensing actions, and let $\phi \in \mathcal{L}$. Then, $\Sigma \models \phi(s_{init})$ iff $s_{init} \in \phi^{TG(\Sigma)}$.*

Proof. ¿From Lemma 2, Lemma 3, and Lemma 4 we have that $\Sigma \models \phi(s_{init})$ iff $\mathcal{M}_{TG(\Sigma)} \models \phi(s_{init})$. Hence, by Lemma 1, the thesis follows. □

Observe that, being $TG(\Sigma)$ essentially a finite "transition system" whose nodes represent sets of valid propositional formulae, it is immediate to modify model checking algorithms for modal mu-calculus formulae for finite transition systems [24, 3, 19], to verify whether s_{init} is in the extension of a formula in $TG(\Sigma)$, and hence, by Theorem 5, to reason about actions in our setting.

Next, we focus on sensing. In order to deal with sensing, we replace each sensing action a_f by two special actions a_f^+ and a_f^-. We denote by $\Gamma_E^{\pm}$ the set of effect axioms Γ_E in which those for the sensing action a_f are replaced by:

$$\mathbf{k}tt \supset [a_f^+]\mathbf{k}f \qquad\qquad \mathbf{k}tt \supset [a_f^-]\mathbf{k}\neg f.$$

Moreover, we instantiate suitably the frame axiom schemas. We denote by $\Gamma_{IFR}^{\pm}$ the set of axioms

$$\mathbf{k}p \supset [a_f^+]\mathbf{k}\varphi \qquad\qquad \mathbf{k}p \supset [a_f^-]\mathbf{k}\varphi$$

obtained by:

1. instantiating the frame axiom schemas in Γ_{FR} for each propositional formula φ such that either $\varphi(init) \in \Gamma_I$, or $\mathbf{k}\varphi$ is in the postcondition of some effect axiom in Γ_E (i.e., φ such that $\mathbf{k}p \supset [a]\mathbf{k}\varphi$, or $\varphi, \neg\varphi$ such that $\mathbf{k}\top \supset [a_\varphi]\mathbf{k}\varphi \sqcup \mathbf{k}\neg\varphi$ in Γ_E). Notice that the size of $\Gamma_{IFR}^{\pm}$, i.e., the number of instances of the frame axiom schemas, is polynomial in the size of Σ;
2. replacing each sensing action a_f by the two special actions a_f^+ and a_f^-.

Finally, we add the axioms in the set $\Gamma_{IFR}^{\pm}$ to Γ_E.

Now, given $\phi \in \mathcal{L}$, we denote with $d(\phi)$ the formula obtained from the negation normal form of ϕ by replacing each occurrence of a subformula of the form $[a_f]\psi$, in which a_f is a sensing action symbol, with the formula $[a_f^+]\psi \wedge [a_f^-]\psi$, and replacing each occurrence of a subformula of the form $\langle a_f\rangle\psi$, in which a_f is a sensing action symbol, with the formula $\langle a_f^+\rangle\psi \wedge \langle a_f^-\rangle\psi$.

In order to extend the result stated in Theorem 5 to specifications with sensing actions, we prove the correctness of the encoding (both in the specification Σ and in the formula ϕ) that expresses each sensing action a_f in terms of a pair of deterministic actions a_f^+, a_f^-.

We denote as $D(\Sigma)$ the specification obtained from Σ by:

1. replacing each precondition axiom for a sensing action a_f

   $$\mathbf{k}p \supset \langle a_f\rangle\mathbf{k}tt$$

   with the following precondition axioms:

   $$\mathbf{k}p \supset \langle a_f^+\rangle\mathbf{k}tt$$
   $$\mathbf{k}p \supset \langle a_f^-\rangle\mathbf{k}tt$$

2. replacing each sensing effect axiom of the form

   $$\mathbf{k}p \supset [a_f](\mathbf{k}f \vee \mathbf{k}\neg f)$$

   with the following effect axioms:

   $$\mathbf{k}p \supset [a_f^+](\mathbf{k}f)$$
   $$\mathbf{k}p \supset [a_f^-](\mathbf{k}\neg f)$$

In order to prove a property analogous to Theorem 5 in the presence of sensing actions in the specification Σ, we follow the same lines of the above proof. The main difference lies in the fact that, in the presence of sensing actions, the model $\mathcal{M}_{TG(\Sigma)}$ cannot be considered as a representative of each model $\mathcal{M}$ for Σ (up to the renaming of states $\eta_{\mathcal{M}}$). In fact, it can be proven that $\mathcal{M}_{TG(\Sigma)}$ represents a sort of "union" of all models for Σ. In particular, it is possible to extend the mapping $\eta_{\mathcal{M}}$ to any model $\mathcal{M}$ of a specification Σ containing sensing actions, in a way such that each renamed model $\eta_{\mathcal{M}}(\mathcal{M}_0)$ corresponds to a portion of $\mathcal{M}_{TG(\Sigma)}$. Therefore, we can state the following.

Lemma 6 *Let Σ be a dynamic system specification, and let $\phi \in \mathcal{L}$. Then, $\Sigma \models \phi(s_{init})$ iff $D(\Sigma) \models d(\phi)(s_{init})$.*

Finally, considering Lemma 5 and Lemma 6, and considering that, by definition of $TG(\Sigma)$, we have that $TG(\Sigma) = TG(D(\Sigma))$, we get:

Theorem 7 *Let Σ be a specification of a dynamic system, and let ϕ be any closed dynamic formula in $\mathcal{L}$. Then, $\Sigma \models \phi(s_{init})$ iff $s_{init} \in d(\phi)^{TG(\Sigma)}$.*

5 Reasoning about actions in $\mathcal{L}$

We illustrate the how the formalism proposed can be used for various forms of reasoning about actions. Below, we informally say that a formula "holds" in a state if the formula is "known" in the robot's corresponding epistemic state.

Projection problem

We start by expressing the *projection problem*: "does a proposition p hold after the execution of a given sequence of actions, say a_1, a_2, a_3?" This can be checked by verifying the following logical implication:

$$\Sigma \models [\langle a_1 \rangle \langle a_2 \rangle \langle a_3 \rangle \mathbf{K} p](s_{init})$$

where $\langle a_1 \rangle \langle a_2 \rangle \langle a_3 \rangle \mathbf{K} p$ expresses that the sequence of actions a_1, a_2, a_3 can indeed be executed and that it leads to a state where p holds.

Planning

Let us now consider the *planning problem*: "is there a sequence of actions that leads to a state where a given goal p_{goal} holds?". This can be expressed by

$$\Sigma \models [\mu X.\, \mathbf{K} p_{goal} \vee \bigvee_{a \in Act} \langle a \rangle X](s_{init}) \tag{2}$$

The dynamic formula on the right-hand side denotes the following inductive property: either p_{goal} holds in the current state, or there is an action a that leads to a state from which there exists a sequence of actions that leads to a state where p_{goal} holds. Notice that, in the presence of sensing actions, the planning process has to return a *conditional plan* [17]. In fact, the presence of sensing actions implies that if property (2) holds, then in each model of Σ there exists a sequence of actions, leading to the goal, which is different in the different models.

It can be shown that in our setting a conditional plan can be effectively returned by visiting $TG(\Sigma)$ and introducing an if-then-else statement on the sensed condition right after each sensing action. Notice that our formalization guarantees that the existence of a plan can be inferred if and only if there exists a *constructive* (conditional) plan which achieves the goal. That is, unrealizable plans are discarded a priori.

The planning problem can be more sophisticated than what shown above. For example we may want to do *planning with archiving and maintenance goals*: "is there a sequence of actions which achieves a certain goal p_{agoal} while another goal p_{mgoal} is always satisfied?". This can be expressed by modifying the formula used above as follows:

$$\mu X.\, \mathbf{k}p_{mgoal} \wedge (\mathbf{k}p_{agoal} \vee \bigvee_{a \in Act} \langle a \rangle X)$$

expressing the fact that, inductively, either both p_{mgoal} and p_{agoal} hold in the current state, or p_{mgoal} holds and there is an action a leading to a state where there exists a sequence achieving p_{agoal} while maintaining p_{mgoal}.

Safeness, invariance, and liveness

Next we consider *safeness properties*. These in general are properties that express that "something bad can never happen". For example, "it is not possible to reach a state from which there exists no plan to get the batteries charged"; in other words, in any reachable state the robot can formulate a plan to charge its battery. In $\mathcal{L}$, the existence of a plan to charge the batteries can be expressed, as shown above, by: $\phi_{pcb} \doteq \mu X.\mathbf{k}BttrChrgd \vee \bigvee_{a \in Act} \langle a \rangle X$. The fact that this can always be done (a safeness property) is expressed as

$$\nu X.\, \phi_{pcb} \wedge \bigwedge_{a \in Act} [a] X.$$

Invariance properties can be expressed in an analogous way, since they can be seen as safeness properties: the bad thing is the violation of the invariant. *Liveness properties*, that in general express that "something good is eventually achieved", can also be captured. For example, "a given job eventually comes to an end" can be expressed as

$$\mu X.\, \mathbf{k}JobEnded \vee (\bigvee_{a \in Act} \langle a \rangle \mathbf{k}tt) \wedge (\bigwedge_{a \in Act} [a] X)$$

Liveness and safeness conditions can be used together to express complex properties as "whenever a job is started, the job is also terminated":

$$\nu X.\, [startjob]\psi \wedge \bigwedge_{a \in Act} [a] X$$

where

$$\psi = \mu Y.\, (\bigvee_{a \in Act} \langle a \rangle \mathbf{k}tt) \wedge (\bigwedge_{a \in Act \wedge a \neq endjob} [a] X)$$

Observe the use of ψ to express the well-foundedness of all sequences of actions not including *endjob*.

Programs

Finally, we introduce a notion of robot program in order to enforce a control flow on actions. Robot programs are not part of the basic action theory specifying the general behavior of the robot; instead, they are used on top of the action theory to introduce a notion of control on the robot actions. This way to proceed mirrors the one used in developing GOLOG [16].

We consider a simple programming language that allows for building nondeterministic while-programs:

$$\begin{array}{rcl} \delta & ::= & \texttt{nop} \mid a \mid \delta_1; \delta_2 \mid \delta_1 | \delta_2 \mid \texttt{if}\ \phi\ \texttt{then}\ \delta_1\ \texttt{else}\ \delta_2 \mid \\ & & \texttt{while}\ \phi\ \texttt{do}\ \delta \end{array}$$

where `nop` is a special instruction that does nothing, a is the command requiring the execution of the action a, ";" is the sequential composition, "|" is nondeterministic choice, and $\texttt{if} \cdot \texttt{then} \cdot \texttt{else} \cdot$ and $\texttt{while} \cdot \texttt{do} \cdot$ are the classical if-then-else and while constructs. The semantics of the various constructs is the usual one (see e.g. [21]), except for atomic actions, whose semantics is given by the basic action theory.

As in the case of GOLOG [16], formally programs are not part of the formalism $\mathcal{L}$. They are used to define suitable macros that are translated into $\mathcal{L}$ dynamic formulae.

We illustrate this approach by showing how to express the property "there exists a terminating execution of program δ that terminates in a state where ϕ holds", which corresponds to the expression $\exists s'.DO(\delta, s, s') \wedge \Phi(s)$ used in GOLOG computations [16]. We can capture the property by defining a $\mathcal{L}$ dynamic formula $afterS(\delta, \phi)$ by induction on the structure of the program as follows (we define $\langle \texttt{nop} \rangle \phi = [\texttt{nop}]\phi = \phi$):

$$\begin{array}{l} afterS(\texttt{nop}, \phi) = \phi \\ afterS(a, \phi) = \langle a \rangle \phi \\ afterS(\delta_1; \delta_2, \phi) = \\ \quad afterS(\delta_1, afterS(\delta_2, \phi)) \\ afterS(\delta_1 | \delta_2, \phi) = \\ \quad afterS(\delta_1, \phi) \ \vee \ afterS(\delta_2, \phi) \\ afterS(\texttt{if}\ \phi'\ \texttt{then}\ \delta_1\ \texttt{else}\ \delta_2, \phi) = \\ \quad \phi' \wedge afterS(\delta_1, \phi) \ \vee \ \neg\phi' \wedge afterS(\delta_1, \phi) \\ afterS(\texttt{while}\ \phi'\ \texttt{do}\ \delta, \phi) = \\ \quad \mu X. \neg\phi' \wedge \phi \ \vee \ \phi' \wedge afterS(\delta, X) \end{array}$$

Notice that the formula $afterS(\delta, \phi)$ is particularly meaningful if we assume that, at the various choice points of the program, the robot can do the choice, choosing the execution that eventually leads to termination in a state where ϕ holds (exactly as assumed by GOLOG computations).

The expressive abilities of $\mathcal{L}$ allow for the formalization of a wide variety of program properties. As a further example, the property "all executions of program δ terminate in states where ϕ holds", can be expressed as the $\mathcal{L}$ formula $afterA(\delta, \phi)$ defined as $afterS(\delta, \phi)$ except that the disjunction in the fourth equation is replaced by a conjunction.[1] Hence, the only difference between the definitions of $afterA(\delta, \phi)$ and $afterS(\delta, \phi)$ is in the treatment of the choice construct: in the case of $afterA(\delta, \phi)$ we require that, independently of the choices made, the program terminates in a state satisfying ϕ,

[1] Observe that $\langle a \rangle \phi \equiv \langle a \rangle \mathbf{k} tt \wedge [a]\phi$, since actions are assumed to be deterministic.

while in the case of $afterA(\delta, \phi)$ only one such choice has to do so. That is, $afterA(\delta, \phi)$ is especially meaningful if the robot has no control on the choice points of the program, so we require that the program "does the right thing" independently of the choices made.[2] We also observe that typical *total correctness* conditions, usually written as $[\phi_1]\delta[\phi_2]$, are expressible by $\phi_1 \supset afterA(\delta, \phi_2)$. Instead, *partial correctness* conditions (correctness for terminating executions only), usually written as $\{\phi_1\}\delta\{\phi_2\}$, are expressible by $\phi_1 \supset afterAw(\delta, \phi_2)$, where $afterAw(\delta, \phi)$ is the formula obtained from $afterA(\delta, \phi)$ replacing $\langle a \rangle \phi$ in the first equation by $[a]\phi$, and the least fixpoint μ in the last equation by the greatest fixpoint ν.

6 Conclusions

In this paper we have shown a way to combine model checking for a very expressive logic of programs with propositional inference, in order to exploit model checking techniques and systems for sophisticated forms of reasoning about actions, including planning and reasoning about program executions. In particular, we have applied such techniques in a framework which enables for both representing rich dynamic systems (it allows for dealing with sensing actions and incomplete information) and efficiently verifying complex properties of such systems (expressed in the language of the modal mu-calculus).

The work presented is related to several proposals in reasoning about actions. We already mentioned the connection with GOLOG [16]. Moreover, in [15] a "validity/provability based GOLOG" has been developed, which shares, in fact, some of the ideas behind our transition graph construction.

There are also some similarities with $\mathcal{A}$-like action languages (see e.g. [1, 10]; indeed, the semantics of $\mathcal{A}$ languages is based on a single transition function, and this allows for building a single transition graph. States in such graph are characterized by the formulae that are *true* (vs. *valid*), while the initial state is replaced by a set of possible initial states. Notably, model checking techniques could be adopted in that setting as well [5].

Model checking is the basic reasoning technique used in [4], where a process algebra is introduced to specify the behavior of the dynamic system, and a suitable variant of modal mu-calculus is adopted as verification formalism. Interestingly, programs (processes) in that work have a somewhat different role, since they are used for specifying basic behavior of the robot and are not considered in the verification formalism.

References

[1] C. Baral and M. Gelfond. Representing concurrent actions in extended logic programming. In *Proc. of IJCAI'93*, pages 866–871, 1993.

[2] C. Baral and T. Son. Approximate reasoning about actions in presence of sensing and incomplete information. In *Proc. of ILPS-97*, 1997.

[3] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Information and Computation*, 98:142–170, 1992.

[2] Notice that $afterS(\delta, \phi)$ is expressible in PDL (leaving aside the **k** operator), while $afterA(\delta, \phi)$ is not.

[4] X. J. Chen and G. De Giacomo. Reasoning about nondeterministic and concurrent actions: A process algebra approach. *Artificial Intelligence*, 107:63–98, 1999.

[5] A. Cimatti, M. Roveri, and P. Traverso. Automatic OBDD-based generation of universal plans in non-deterministic domains. In *Proc. of AAAI'98*, pages 875–881, 1998.

[6] G. De Giacomo, L. Iocchi, D. Nardi, and R. Rosati. Moving a robot: the KR&R approach at work. In *Proc. of KR'96*, pages 198–209, 1996.

[7] G. De Giacomo, L. Iocchi, D. Nardi, and R. Rosati. Planning with sensing for a mobile robot. In *Proceedings of the Fourth European Conference on Planning (ECP'97)*, 1997.

[8] F. M. Donini, D. Nardi, and R. Rosati. Autoepistemic description logics. In *Proc. of IJCAI'97*, pages 136–141, 1997.

[9] E. A. Emerson. Automated temporal reasoning about reactive systems. In *Logics for Concurrency: Structure versus Automata*, number 1043 in Lecture Notes in Computer Science, pages 41–101. Springer-Verlag, 1996.

[10] E. Giunchiglia and V. Lifschitz. An action language based on causal explanations: preliminary report. In *Proc. of AAAI'98*, 1998.

[11] K. Golden and D. Weld. Representing sensing actions: the middle ground revisited. In *Proc. of KR'96*, pages 174–185, 1996.

[12] C. Hoare. *Communicating Sequential Processes.* Prentice Hall Int., London, 1985.

[13] D. Kozen and J. Tiuryn. Logics of programs. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science – Formal Models and Semantics*, pages 789–840. Elsevier, 1990.

[14] G. Lakemeyer and H. J. Levesque. AOL: a logic of acting, sensing, knowing, and only knowing. In *Proc. of KR'98*, pages 316–327. Morgan Kaufmann, Los Altos, 1998.

[15] Y. Lesperance and D. Tremaine. A procedural approach to belief update for agent programming. Unpublished Manuscript, Department of Computer Science York University, 1998.

[16] H. Levesque, R. Reiter, Y. Lesperance, F. Lin, and R. Scherl. GOLOG: A logic programming language for dynamic domains. *Journal of Logic Programming*, 31:59–84, 1997.

[17] H. J. Levesque. What is planning in presence of sensing? In *Proc. of AAAI'96*, pages 1139–1149. AAAI Press/The MIT Press, 1996.

[18] J. Lobo, G. Mendez, and S. R. Taylor. Adding knowledge to the action description language A. In *Proc. of AAAI'97*, pages 454–459, 1997.

[19] K. L. McMillan. *Symbolic Model Checking.* Kluwer Academic Publishers, 1993.

[20] M. Milner. *Communication and Concurrency.* Prentice-Hall, 1989.

[21] H. R. Nielson and F. Nielson. *Semantics with Applications.* Wiley, 1992.

[22] R. Reiter. *Knowledge in Action: Logical Foundation for Describing and Implementing Dynamical Systems.* 1998. In preparation.

[23] M. Shanahan. *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. The MIT Press, 1997.

[24] C. Stirling. Modal and temporal logics for processes. In *Logics for Concurrency: Structure versus Automata*, number 1043 in Lecture Notes in Computer Science, pages 149–237. Springer-Verlag, 1996.

[25] C. Stirling. Modal and temporal logics for processes. In F. Moller and G. Birtwistle, editors, *Logics for Concurrency: Structure versus Automata*, volume 1043 of *LNCS*, pages 149–237. Springer-Verlag, 1996.

[26] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.