

# Using Temporal Logic to integrate Goals and Qualitative Preferences into Agent Programming

Koen V. Hindriks<sup>1</sup> and M. Birna van Riemsdijk<sup>2</sup>

<sup>1</sup> EEMCS, Delft University of Technology, Delft, The Netherlands

<sup>2</sup> LMU, Munich, Germany

**Abstract.** The core component of a rational agent is its capability to make rational decisions or choices of action, in order to satisfy its design objectives. In agent programming languages for rational agents, such decisions or choices are derived from the agent's beliefs and goals. At any one time, an agent can typically choose from multiple actions, which may all lead to goal achievement. Existing approaches usually select one of those actions non-deterministically. In this paper, we propose the use of goals as hard constraints and qualitative preferences as soft constraints for choosing a most preferred action among the available ones. We use temporal logic for the representation of various kinds of goals and preferences, leading to a uniform framework for integrating goals and preferences into agent programming.

## 1 Introduction

The core component of a rational agent is its capability to make rational decisions or choices of action, in order to satisfy its design objectives. In agent programming languages for rational agents, such decisions or choices are derived from the agent's beliefs and goals. These goals are often achievement goals, i.e., goals that define states that are to be achieved. Typically, agent programs consist of a number of rules (usually called action selection or plan rules) that provide the agent with the means to choose actions that achieve its goals. Such rule sets, however, do not necessarily completely determine the choice of action, and, at any one time, an agent can typically choose from *multiple* actions, which may all lead to goal achievement. Existing approaches usually select one of those actions non-deterministically.

Recent research in both the planning and agent programming field [1, 9, 2, 11, 10] has shown that it can be useful to have additional mechanisms for further *constraining* the choice of action in an agent program, i.e., for reducing the non-determinism of the initial agent program. In [11], for example, we have proposed the use of maintenance goals, i.e., goals that define states that must remain true throughout the execution of the agent, for making a further selection among the actions available for execution. The idea is that an agent should avoid selecting actions which violate maintenance goals.

In this paper, we investigate such mechanisms for further constraining an agent's choice of action in a more general and uniform setting. We distinguish

between hard constraints which *must* be satisfied, and soft constraints or preferences, by means of which one can distinguish more preferred courses of action from less preferred ones. We propose a layered rational action selection architecture which specifies a rational action selection mechanism in which these types of constraints are integrated.

We show how the proposed architecture can be formalized in the context of the GOAL agent programming language [6, 12]. Where GOAL is based on pure propositional logic, we use in this paper temporal logic for the representation of goals and preferences. A key motivation for integrating temporal logic into agent programming languages is the potential that the additional expressiveness thus introduced provides for defining a uniform framework that naturally allows for the integration of hard and soft constraints including such concepts as achievement goals, maintenance goals, as well as (temporally extended) preferences. The work reported is inspired by work in planning where temporal logic is used to “guide” planners through the search space [1] and to select preferred plans [9, 2]. The purpose of this paper is to define a conceptual framework that provides a starting point for future research and unifies the various kinds of goals and their role in rational action selection. We are thus not primarily concerned with the practical implementability of the approach.

In Section 2, we discuss rational action selection in planning and agent programming, and informally present the rational action selection architecture. In Section 3, we present the extension of GOAL with temporal logic, which forms the first layer of our action selection architecture. Section 4 defines some technical preliminaries with respect to temporal logic, which are needed in Sections 5 and 6 in which we realize the second and third layer, respectively, of the action selection architecture in GOAL. In Section 7 we conclude the paper.

## 2 Rational Action Selection Architecture (RASA)

### 2.1 Planning and Programming

Two important approaches for realizing rational decision making agents are *planning* [8] and *agent programming* [4]. The classical AI planning problem is to search for a plan (sequence of actions) to get from the current state to a goal state, given a description of the available actions. That is, one seeks a *complete plan*, the execution of which will result in the agent achieving its goal (given certain assumptions on the environment). In agent programming, on the other hand, the behavior of the agent is specified by means of a program. This program is executed *step by step*, without first checking whether the executed actions will eventually lead to the agent reaching its goal. If the programmer has written an effective program, the execution will indeed lead to goal achievement, but whether this is really the case only becomes clear as the execution of the program progresses.<sup>3</sup>

---

<sup>3</sup> Using verification, one could prove that a program satisfies its goals. However, verification is a non-trivial undertaking and it might be difficult to prove correctness.

The main advantage of programming over planning is that it typically provides a more efficient solution, since it requires less search through possibly very large search spaces. The main advantage of planning over programming is that it is potentially more flexible since the technique can in principle be applied regardless of where the agent currently is and where it wants to go. Moreover, one knows before execution that the goal will be reached, and it puts less burden on the programmer.

What planning and agent programming approaches have in common, is that they typically allow for various alternative plans to be executed for one particular goal (or set of goals) as in general there are multiple ways of getting from a current state to a goal state. A planner might thus yield various alternative plans, and an agent program typically does not completely determine the behavior of the agent in each state as it allows multiple actions to be executed in a particular state. Planning problems and agent programs thus typically *underspecify* an agent's behavior.

## 2.2 An Architecture for Rational Action Selection

The underspecification of agent behavior can have benefits from a design point of view. If it doesn't matter whether the agent executes one plan or another for reaching a goal, one can argue that it is more natural to let the planning problem or agent program reflect this by leaving these choices open. However, recent research in both the planning and agent programming field [1, 9, 2, 11, 10] has shown that it can be useful to have additional mechanisms for finding plans to solve the planning problem or to further constrain the choice of action in an agent program. The idea is that *additional selection mechanisms can be introduced on top of the existing planning or programming mechanisms*, by means of which a further selection among plans can be made, e.g., based on additional constraints or preferences, in order to optimize agent behavior.

Regarding the kinds of additional constraints that may be applied to select among the possible plans, we distinguish between *hard constraints* which *must* be satisfied, and *soft constraints* or *preferences*, by means of which one can distinguish more preferred plans from less preferred ones (see also [9, 11]).

What one takes as hard constraints varies across approaches. In [9], and more generally in planning, the (achievement) goals themselves are considered to be hard constraints.<sup>4</sup> In [11], maintenance goals are taken as hard constraints, i.e., an agent may never violate a maintenance goal. Preferences can be used to distinguish between optimal and suboptimal plans, e.g., in terms of costs, but may also be used to distinguish plans with respect to which goals are reached (if goals are not considered as hard constraints and not all goals can be reached).

These considerations lead to the following layered rational action selection architecture (RASA).

---

<sup>4</sup> That is, in that approach the initial goals are the basis for plan generation, while at the same time being considered as hard constraints. The initial goals are thus not used as an *additional* action selection mechanism.

- **Layer 1:** The RASA *generates* options for action, i.e., *plans*, often on the basis of an agent’s beliefs and (achievement) goals.
- **Layer 2:** The RASA verifies whether the plans from layer one satisfy *hard constraints* and discards those that do not.
- **Layer 3:** The RASA selects from the options remaining after the application of layer two those plans that maximize satisfaction of *soft constraints*.

In this paper, we show how this architecture can be made concrete in the context of the GOAL agent programming language.

It may sometimes be more efficient to implement the architecture by combining various layers, i.e., generating the possible plans while taking into account the constraints (see, e.g., [2]). This is, however, not the focus of this paper.

### 2.3 Rational Action Selection in Agent Programming

The general RASA as described in Section 2.2 can be applied both to approaches focussing on planning, as well as to those focussing on agent programming. While most existing approaches that follow an architecture similar to the RASA are planning approaches [1, 9, 2], in this paper we are interested in action selection in agent programming languages. The main difference between applying the RASA in the context of planning and in the context of agent programming is that in planning one compares *complete plans* against the available constraints, while in agent programming the constraints need to be taken into account continuously *during* execution.

The way in which we propose to take into account constraints during execution is partly based on our previous work on the incorporation of maintenance goals in agent programming languages [11]. The idea is that the agent has a fixed, usually finite, *lookahead horizon*, i.e., the agent can lookahead a certain number of execution steps. The agent then evaluates the possible paths (plans) it can take (up to the given horizon) using its constraints, in order to determine which paths are the best. It then takes one step along one of these paths, and the process is repeated.

This approach is thus similar to planning in the sense that the agent looks ahead before executing. The difference is that it is not required that the agent has reached its goals within the number of steps specified by the horizon, i.e., it does not have complete plans at its disposal when selecting an action. This means that the agent has to take into account how paths *might* progress beyond its lookahead horizon when selecting an action. For example, a path on which a goal might still be satisfied, is probably better than one on which this goal will not be satisfied for certain.

The technical tool we use for the representation of goals and preferences is linear temporal logic (LTL) [7]. This idea is inspired by work on the representation of qualitative preferences in the context of planning [1, 9]. As will become clear in the sequel, the fact that the agent needs to take into account how paths might progress beyond the lookahead horizon, requires an adaptation of the way in which linear temporal logic is used in the context of planning.

### 3 RASA Layer 1: Temporalized GOAL

In this section, we define the first layer of the RASA architecture in the context of the GOAL agent programming language [6, 12]. We extend the original GOAL language by allowing temporal formulas as goals, rather than only propositional formulas.

#### 3.1 The General Idea

In the GOAL language, an agent selects actions on the basis of its beliefs and goals. A program consists of (1) a set of initial beliefs, collectively called the (initial) *belief base* of the agent, (2) a set of initial goals, called the (initial) *goal base*, (3) an *action specification* which consists of a specification of the pre- and post-conditions of *basic actions* of the agent, and (4) a program section which consists of a set of *conditional actions rules*.

In the original GOAL language, the belief base and goal base are sets of propositional formulas. The goals are interpreted as achievement goals. That is, if a propositional formula  $\phi$  is in the goal base, this informally means that the agent wants to reach a situation in which  $\phi$  is (believed to be) the case. If a basic action is executed, the agent’s beliefs change as specified in the pre- and postconditions of the action, and the achievement goals that are believed to be reached through the execution of the action are removed from the goal base.<sup>5</sup> A conditional action rule consists of a basic action and a condition on the agent’s beliefs and goals. Such a rule expresses that the basic action may be executed, if the condition holds. During execution, a GOAL agent selects non-deterministically any of its enabled conditional action rules, i.e., a conditional action rule of which the condition holds, and then executes its corresponding basic action. At any one time, typically *multiple* conditional action rules are enabled, i.e., a GOAL program underspecifies an agent’s behavior.

In [11], we have extended the GOAL language with maintenance goals. Just like achievement goals, maintenance goals were also expressed by propositional formulas. A maintenance goal  $\phi$  expresses that the agent wants that  $\phi$  holds continuously throughout the execution of the agent. Both achievement goals and maintenance goals express particular desired properties of the behavior of the agent. These properties can be expressed conveniently in LTL. LTL has computation traces (sequences of states) as models. An achievement goal for  $\phi$  can be represented by the LTL formula  $\Diamond\phi$ , specifying that  $\phi$  should hold eventually, i.e., in some state on the computation trace. A maintenance goal for  $\phi$  can be represented by  $\Box\phi$ , specifying that  $\phi$  should always hold. This idea can be generalized by realizing that in fact any LTL formula can be used for expressing goals (as also observed in [13]).

---

<sup>5</sup> The idea is that the agent’s beliefs also represent the environment, and that basic actions change this environment. However, the environment is not modeled in the formal specification of GOAL, and consequently basic actions update only the belief base.

In this paper, we make this idea concrete in the context of the GOAL language by using LTL for the representation of goals. This increases the expressiveness of the language by allowing the representation of all kinds of goals, and allows for the representation of goals in a uniform way. For example, the goal  $\phi U \phi'$ , which expresses that the agent wants to ensure that  $\phi$  while it is trying to achieve  $\phi'$ , can now be expressed easily.

In a more practical setting, we will have to restrict the allowed LTL formulas again in order to make it computationally tractable. What is first needed, however, is a conceptual framework that provides the right starting point and unifies the various kinds of goals and their role in rational action selection. Integrating beliefs, achievements and maintenance goals, as well as preferences into a single framework thus first requires a well-defined theory of how to do that. In this paper, we present the informal motivation of a particular architecture (see Section 2) and the formal framework which provides a proper understanding of how to technically combine all of these concepts and how they are involved in action selection.

### 3.2 Formalization

In the next definition, we define the LTL language that we use for representing goals. For reasons of simplicity we do not investigate and extend the beliefs of an agent in the programming language to temporal formulas. In particular, it is not clear how to determine whether a goal is reached, if the beliefs of the agent are temporal formulas.

The states of the traces on which the LTL formulas are evaluated are belief bases (see also [5]), i.e., goals express how the belief base of the agent should evolve.<sup>6</sup> We assume a language  $\mathcal{L}_0$  of propositional logic with typical element  $\phi$ , and the standard entailment relation  $\models$ . A belief base  $\Sigma \subseteq \mathcal{L}_0$  is a set of propositional formulas. Our LTL language contains the standard (temporal) operators of LTL. The difference between our LTL language and standard LTL is that the states of the traces are belief bases, rather than worlds as valuations of propositional atoms. The semantics of non-temporal propositional formulas is thus defined on belief bases. We specify that a propositional formula  $\phi$  holds in a belief base  $\Sigma$  if  $\Sigma \models \phi$ .

**Definition 1** (*linear temporal logic (LTL)*)

Let  $\phi \in \mathcal{L}_0$ . The set of LTL formulas  $\mathcal{L}_{LTL}$  with typical element  $\chi$  is defined as follows.

$$\chi ::= \top \mid \phi \mid \neg\chi \mid \chi_1 \wedge \chi_2 \mid \diamond\chi \mid \bigcirc\chi \mid \chi_1 U \chi_2$$

Let  $t^b = \Sigma_0, \Sigma_1, \dots$  be an infinite trace of belief bases and let  $i \in \mathbb{N}$  be a position in a trace. The semantics of LTL formulas is defined on infinite traces  $t^b$  as

---

<sup>6</sup> Since the idea is that the beliefs also represent the environment of the agent, goals also express desired properties of the environment.

follows.

$$\begin{aligned}
t^b, i &\models_{LTL} \top \\
t^b, i &\models_{LTL} \phi && \Leftrightarrow \Sigma_i \models \phi \\
t^b, i &\models_{LTL} \neg\chi && \Leftrightarrow t^b, i \not\models_{LTL} \chi \\
t^b, i &\models_{LTL} \chi_1 \wedge \chi_2 && \Leftrightarrow t^b, i \models_{LTL} \chi_1 \text{ and } t^b, i \models_{LTL} \chi_2 \\
t^b, i &\models_{LTL} \diamond\chi && \Leftrightarrow \exists k \geq i : t^b, k \models_{LTL} \chi \\
t^b, i &\models_{LTL} \bigcirc\chi && \Leftrightarrow t^b, i+1 \models_{LTL} \chi \\
t^b, i &\models_{LTL} \chi_1 U \chi_2 && \Leftrightarrow \exists k \geq i : t^b, k \models_{LTL} \chi_2 \text{ and } \forall i \leq l < k : t^b, l \models_{LTL} \chi_1
\end{aligned}$$

The always operator is defined in terms of the eventually operator as follows:  $\Box\chi \equiv \neg\diamond\neg\chi$ . The entailment relation  $\Omega \models_{LTL} \chi$  with  $\Omega \subseteq \mathcal{L}_{LTL}$  is the standard entailment relation of LTL, except for the base case of purely propositional formulas, since the semantics of the base case is different from standard LTL.

The mental state of a GOAL agent consists of those components that change during execution. That is, a mental state consists of a belief base and a goal base. The goal base is typically denoted by  $\Gamma$  and in temporalized GOAL this is a set of LTL formulas. Mental states should satisfy a number of rationality constraints.

**Definition 2** (*Mental States*)

A mental state of a GOAL agent, typically denoted by  $m$ , is a pair  $\langle \Sigma, \Gamma \rangle$  with  $\Sigma \subseteq \mathcal{L}_0$  and  $\Gamma \subseteq \mathcal{L}_{LTL}$  where  $\Sigma$  is the belief base, and  $\Gamma$  with typical element  $\chi$  is the goal base. Additionally, mental states need to satisfy the following *rationality constraints*:

- (i) The belief base is consistent:  $\Sigma \not\models \perp$ ,
- (ii) The goal base is consistent:  $\Gamma \not\models_{LTL} \perp$ ,
- (iii) The goal base does not contain goals that have already been achieved.

The third rationality constraints is implemented by means of the progression operator which will be introduced below (Definition 4).

A GOAL agent derives its choice of action from its beliefs and goals. In order to do so, a GOAL agent inspects its mental state by evaluating so-called *mental state conditions*. The syntax and semantics of these conditions is defined next.

**Definition 3** (*Mental State Conditions*)

Let  $\phi \in \mathcal{L}_0$ ,  $\chi \in \mathcal{L}_{LTL}$ . The language  $\mathcal{L}_M$  of mental state conditions, typically denoted by  $\psi$ , is defined as follows.

$$\psi ::= \mathbf{B}\phi \mid \mathbf{G}\chi \mid \neg\psi \mid \psi_1 \wedge \psi_2$$

The truth conditions of mental state conditions  $\psi$ , relative to a mental state  $m = \langle \Sigma, \Gamma \rangle$ , are defined as follows.

$$\begin{array}{lll}
m \models_m \mathbf{B}\phi & \text{iff} & \Sigma \models \phi \\
m \models_m \mathbf{G}\chi & \text{iff} & \Gamma \models_{LTL} \chi \\
m \models_m \neg\psi & \text{iff} & m \not\models_m \psi \\
m \models_m \psi_1 \wedge \psi_2 & \text{iff} & m \models_m \psi_1 \text{ and } m \models_m \psi_2
\end{array}$$

The semantics of  $\mathbf{B}\phi$  is as in original GOAL, and defines that  $\mathbf{B}\phi$  holds iff  $\phi$  follows from the belief base under a standard proposition logic entailment relation. The semantics of the  $\mathbf{G}$  operator is different from original GOAL, where the goal base is a set of propositional formulas. Here, we use the LTL entailment relation to generalize the operator to be able to express arbitrary types of goals.

Before we can move on to defining how the execution of a basic action changes the agent's mental state, we need to explain how the goal base is updated. In the original GOAL language, achievement goals that are believed to be achieved after the execution of a basic action are removed from the goal base. Checking whether an achievement goal is achieved is simple if these are represented as propositional formulas: an achievement goal  $\phi$  is achieved in a mental state if it follows from the belief base in that mental state. In temporalized GOAL, goals are temporal formulas. In order to be able to evaluate the achievement of temporal formulas in a mental state, we use a technique from [1] for "progressing" LTL formulas.

Progression of an LTL formula is a transformation of this formula, which should be performed at each execution step. The idea is that the transformation yields a new formula in which those "parts" of the formula which have already been achieved are set to  $\top$ , leaving an LTL formula which expresses what still has to be satisfied.

For example, if the agent has a goal  $\diamond\phi$  ( $\phi \in \mathcal{L}_0$ ) in a particular mental state and  $\phi$  is achieved in that mental state, i.e., follows from the belief base, then the progression of this formula is " $\top$ ", as the agent has produced an execution trace on which the formula holds. If  $\phi$  does not hold, then the progression is the formula " $\diamond\phi$ " itself, since it still needs to be satisfied. If a formula has progressed to a formula equivalent to  $\top$ , it means the agent has produced an execution trace on which the formula holds. Note that formulas of the form  $\Box\chi$  can never progress to  $\top$ , since it needs to be checked continuously whether  $\chi$  holds.

The progression operator can be defined inductively for general LTL formulas, as specified in the next definition. We adapt the definition of [1] slightly, as we want a formula which is reached in a mental state to be true already in that state, rather than one mental state later. For this, we define the progression of  $\bigcirc\chi'$  as  $Progress(\chi')$ , rather than as  $\chi'$ .

**Definition 4** (*progression of LTL formulas*)

Let  $\Sigma$  be a belief base, let  $\chi \in \mathcal{L}_{LTL}$  be an LTL formula, and let  $\phi \in \mathcal{L}_0$ . The progression of  $\chi$  in  $\Sigma$ ,  $Progress(\chi, \Sigma)$  is then defined as follows.

form of $\chi$ is	$Progress(\chi, \Sigma) =$
$\top$	$\top$
$\phi$	$\top$ if $\Sigma \models \phi$ , $\perp$ otherwise
$\neg\chi'$	$\neg Progress(\chi', \Sigma)$
$\chi_1 \wedge \chi_2$	$Progress(\chi_1, \Sigma) \wedge Progress(\chi_2, \Sigma)$
$\diamond\chi'$	$Progress(\chi', \Sigma) \vee \chi$
$\bigcirc\chi'$	$Progress(\chi')$
$\chi_1 U \chi_2$	$Progress(\chi_2, \Sigma) \vee (Progress(\chi_1, \Sigma) \wedge \chi)$
$\Box\chi'$	$Progress(\chi', \Sigma) \wedge \chi$

We lift the progression function of Definition 4 to sets of LTL formulas  $\Gamma \subseteq \mathcal{L}_{LTL}$  as follows:  $Progress(\Gamma, \Sigma) = \bigcup_{\chi \in \Gamma} Progress(\chi, \Sigma)$ .

The next definition specifies how the execution of a basic action changes an agent's mental state. In the formal definition of GOAL, we use a *transition function*  $\mathcal{T}$  to model the effects of basic actions for technical convenience, rather than a specification of pre- and postconditions. The function  $\mathcal{T}$  maps a basic action  $\mathbf{a}$  and a belief base  $\Sigma$  to an updated belief base  $\mathcal{T}(\mathbf{a}, \Sigma) = \Sigma'$ . The transition function is undefined if an action is not enabled in a mental state. The GOAL language also includes special actions for adding and removing goals from the goal base, but we do not discuss these actions here (see e.g. [6]). The change of the goal base is defined by means of the progression operator. At each step, all goals of the goal base are progressed.

**Definition 5** (*Mental State Transformer  $\mathcal{M}$* )

Let  $\mathbf{a}$  be a basic action,  $\phi \in \mathcal{L}_0$  and  $\mathcal{T}$  be a transition function for basic actions. Then the *mental state transformer function*  $\mathcal{M}$  is defined as a mapping from actions and mental states to updated mental states as follows:

$$\mathcal{M}(\mathbf{a}, \langle \Sigma, \Gamma \rangle) = \begin{cases} \langle \Sigma', Progress(\Gamma, \Sigma') \rangle & \text{if } \mathcal{T}(\mathbf{a}, \Sigma) = \Sigma' \\ \text{undefined} & \text{otherwise} \end{cases}$$

The specification of when a basic action may be executed, is done by means of conditional action rules. A conditional action rule  $c$  has the form **if**  $\psi$  **then**  $\mathbf{a}$ , with  $\mathbf{a}$  a basic action. This conditional action rule specifies that  $\mathbf{a}$  may be performed if the mental state condition  $\psi$  holds and the transition function is defined for  $\mathbf{a}$ . In that case we say that conditional action  $c$  is *enabled*. During execution, a GOAL agent selects non-deterministically any of its enabled conditional actions. This is expressed in the following transition rule, describing how an agent gets from one mental state to another.

**Definition 6** (*Conditional Action Semantics*)

Let  $m$  be a mental state, and  $c = \mathbf{if} \ \psi \ \mathbf{then} \ \mathbf{a}$  be a conditional action. The transition relation  $\xrightarrow{c}$  is the smallest relation induced by the following transition rule.

$$\frac{m \models \psi \quad \mathcal{M}(\mathbf{a}, m) \text{ is defined}}{m \xrightarrow{c} \mathcal{M}(\mathbf{a}, m)}$$

The execution of a GOAL agent results in a *computation trace*. We define a trace as a sequence of mental states, such that each mental state can be obtained from the previous by applying the transition rule of Definition 6. As GOAL agents are non-deterministic, the semantics of a GOAL agent is defined as the *set* of possible computations of the GOAL agent, where all computations start in the initial mental state of the agent.

**Definition 7** (*Agent Computation*) A computation trace, typically denoted by  $t$ , is an infinite sequence of mental states  $m_0, m_1, m_2, \dots$  such that for each  $i$

there is an action  $c_i$  and  $m_i \xrightarrow{c_i} m_{i+1}$  can be derived using the transition rule of Definition 6, or  $m_i \xrightarrow{f^i}$  and for all  $j > i$ ,  $m_j = m_i$ . The meaning  $R_{\mathcal{A}}(m_0)$  of a GOAL agent named  $\mathcal{A}$  with initial mental state  $m_0$  is the set of all computations starting in that state.

Observe that a computation is infinite by definition, even if the agent is not able to perform any action anymore from some point in time on. Also note that the concept of a computation trace is a general notion in program semantics that is not particular to GOAL. The notion of a computation trace can be defined for any agent programming language that is provided with a well-defined operational semantics.

The semantics  $R_{\mathcal{A}}(m_0)$  thus consists of all traces that may be generated by the agent program. These traces form the first layer of the RASA architecture.

## 4 Evaluating Temporal Formulas on Prefixes of Traces

In Sections 5 and 6, we will show how to define the second layer (hard constraints) and third layer (soft constraints) on top of the first layer as defined in the previous section. Both hard constraints and soft constraints will be represented using LTL. In standard LTL, formulas are evaluated on *infinite* traces. As explained in Section 2.3, however, in our setting an agent can lookahead a *finite* number of steps and we want to evaluate LTL formulas on such finite traces. We thus need to define the semantics of LTL formulas on such finite traces. This introduces several issues, one of which is the definition of the semantics of the next operator  $\bigcirc$ . It is not immediately clear what the semantics of a formula  $\bigcirc\phi$  should be if it is evaluated in the last state of a finite trace.

For explaining our approach, it is important to realize that the finite trace on which we evaluate LTL formulas is only a *prefix* of a trace which will be continued beyond the lookahead horizon. Intuitively, the truth of a formula  $\bigcirc\phi$  evaluated in the last state of such a finite prefix cannot be established. Its truth depends on how the trace continues beyond the finite prefix under consideration. A formula  $\diamond\phi$ , on the other hand, clearly holds on a finite prefix if  $\phi$  holds on some state of this prefix. Similarly, a formula  $\Box\phi$  is clearly false on a finite prefix if  $\neg\phi$  holds in some state of this prefix.

From these examples, we can see that the truth of an LTL formula evaluated on a finite prefix of a trace depends on how this trace progresses beyond the finite prefix. If a formula is false on any progression, it is also false on the finite prefix. Similarly, if it is true on any progression, it is true on the finite prefix. In all other cases, we cannot determine the truth of the formula. This intuition is reflected by a 3-valued semantics of LTL as presented in [3] in the context of monitoring. In the 3-valued semantics, the truth value of an LTL formula evaluated on a finite trace is  $\top$  (true),  $\perp$  (false), or “?” (unknown). We use (a slightly adapted version of) the definition of [3] for evaluating LTL formulas on finite prefixes. We define the semantics of LTL formulas over traces of mental states in terms of the semantics over belief bases as follows, where  $t^b$  is derived from  $t$  by keeping only the belief bases of each mental state:  $t, i \models_{LTL} \chi \Leftrightarrow t^b, i \models_{LTL} \chi$ .

**Definition 8** (*3-valued LTL semantics*)

Let  $M$  be a set of mental states. Finite traces over  $M$  are elements of  $M^*$  and infinite traces are elements of  $M^\omega$ . Both are typically denoted by  $t$ . Let  $t, t' \in M^* \cup M^\omega$  be (finite or infinite) traces. Concatenation of traces  $t, t'$  is defined as usual and simply denoted as  $tt'$ ; we stipulate that if  $t \in M^\omega$ , then  $tt' = t$ . The truth value of an  $LTL_3$  formula  $\varphi$  with respect to a trace  $t \in M^* \cup M^\omega$ , denoted by  $[t \models \chi]$ , is an element of  $\{\top, \perp, ?\}$  and defined as follows.

$$[t \models \chi] = \begin{cases} \top & \text{if } \forall t' \in M^\omega : tt' \models_{LTL} \chi, \\ \perp & \text{if } \forall t' \in M^\omega : tt' \not\models_{LTL} \chi, \\ ? & \text{otherwise.} \end{cases}$$

**Corollary 1.** *If  $t = \epsilon$ ,  $[t \models \varphi] = ?$  if  $\varphi \neq \perp$  and  $\varphi \neq \top$ .*

The only difference between our definition and the one of [3] is that in our definition  $t$  can also be infinite. This allows us to investigate our semantics also for an infinite lookahead horizon in a uniform way. In [3], a technique based on the construction of a finite state machine from an LTL formula and a finite trace is presented for determining the truth value of an  $LTL_3$  formula at runtime in implemented systems.

The need for a 3-valued semantics in our approach highlights an important difference with the use of LTL in planning approaches such as [9]. In a planning context, the plans under consideration are always complete. That is, it is not taken into account that these plans might progress beyond their final state. This means that, e.g., a formula  $\diamond\phi$  in those approaches is considered to be false if  $\phi$  does not hold in some state resulting from execution of the plan, and true otherwise. In our semantics, on the other hand, the truth value of the formula is unknown if  $\phi$  does not hold on some state of the finite prefix.

## 5 RASA Layer 2: Goals as Hard Constraints

Goals of an agent are temporal formulas. The semantics of agent programs provided in Section 3 accounts for the role such goals have in selecting actions using action selection rules. Action selection rules allow an agent to derive actions from its beliefs and goals in a *reactive* manner, but we argue that this layer in the architecture does not yet account for the full role that such goals can have in the action selection mechanism of a rational agent.

If an agent has the ability to *lookahead* a (finite) number of steps, it can also use its goals to avoid selecting those actions that prevent the realization of (some of) the agent's goals. In this section we define the second layer of the RASA which accounts for this role of goals. Goals thus viewed introduce additional constraints on action selection to the effect of excluding those actions of which the agent foresees that they will not satisfy its goals. Here we consider such constraints to be *hard* constraints, meaning that any foreseen violation of goals by performing an action will force a rational agent to choose an alternative action (if possible) or become inactive (in line with previous work, cf. [11]). That

is, actions will only be selected if for all that is known may still eventually satisfy the goals of the agent.

In this view of goals as hard constraints, achievement goals of the form  $\diamond\phi$  per se do not add any additional constraints on action selection since an agent with a finite lookahead horizon will not be able to conclude that an action will prohibit realizing such a goal. Achievement goals thus may be said to have no “selective force” beyond the rule-based mechanism of the action selection architecture. The role of such goals thus is mainly to guide this selection process and, in order to avoid wasting resources, to remove such goals as reasons for action when they have been achieved (see the *Progress* operator of Section 3.2). Other types of goals such as maintenance goals do have a selective force in this sense; for example, if  $\square\phi$  is a goal of the agent, an agent can conclude that it is no longer possible to satisfy this goal if  $\phi$  does not hold in some state within the lookahead horizon. An agent thus can use such goals to filter certain options for action generated by the rule-based layer one.<sup>7</sup> Here the 3-valued semantics for LTL discussed in the previous section is particularly useful since it allows us to evaluate goals on finite prefixes of traces. The fact that an achievement goal  $\diamond\phi$  does not have selective force, is reflected by the fact that its truth value in the 3-valued LTL semantics will be “?” (the “unknown” value).

The second layer of the RASA can be defined formally in the context of GOAL as follows. A *prefix* of a computation  $t$  is an initial finite sequence of  $t$  or  $t$  itself. A prefix of length  $n$  of a computation  $t$  is denoted by  $t^{(n)}$  with  $n \in \mathbb{N} \cup \{\infty\}$ , where  $t^{(\infty)}$  is defined as  $t$ .  $\mathbb{N}$  is the set of natural numbers including 0, and  $\infty$  is the first infinite ordinal. The lookahead horizon, i.e., the number of execution steps that an agent can lookahead, is denoted by  $h$ . The idea is that we take the set of possible execution traces  $R_{\mathcal{A}}$  that may be selected according to the first layer, and filter out those traces that do not satisfy one or more of the agent’s goals, when looking ahead  $h$  steps from some state on the trace.

We define a filter function  $\sigma_{\mathcal{A}}^h$  inductively on the set of traces  $R_{\mathcal{A}}$  and on the time point  $i$  on such a trace. An agent will use its goals to restrict selection of actions from the start, i.e. time 0, which explains why the base case of the inductive definition starts at  $-1$ . The base case defines the starting point of traces  $R_{\mathcal{A}}$  that need to be filtered. For technical reasons the filter function is defined as two different components, and in addition to  $\sigma$  we define a function  $\zeta$ . The idea is that the filter function  $\sigma^h(i)$  returns all (infinite) traces that satisfy the goals of the agent on a finite prefix of length  $h$  starting from state  $i$ , while the function  $\zeta^h(i)$  returns all (finite) prefixes of traces that satisfy the goals of an agent given a lookahead capability of  $h$  until the end of that prefix but no longer in any possible next state.

---

<sup>7</sup> One could argue that layer one should also make sure that such maintenance goals are not violated. However, trying to account for such general constraints on agent behavior in the rules in an ad hoc manner will often lead to less understandable programs, which is why we argue for the separation of concerns provided by the proposed RASA.

**Definition 9** (*Goal Filter Functions  $\sigma$  and  $\varsigma$* )

Let  $\mathcal{A}$  be some agent with meaning  $R_{\mathcal{A}}$  and let  $h \in \mathbb{N}$  be a horizon. Let  $t[i]$  be the tail of trace  $t$  starting in the  $i$ -th state of  $t$ . Then the *goal filter functions*  $\sigma_{\mathcal{A}}^h$  and  $\varsigma_{\mathcal{A}}^h$  are defined by simultaneous induction as follows:

$$\begin{aligned}\sigma_{\mathcal{A}}^h(-1) &= R_{\mathcal{A}}, \\ \sigma_{\mathcal{A}}^h(i) &= \{t \in \sigma_{\mathcal{A}}^h(i-1) \mid \forall \chi \in \Gamma_i^t : [t[i]^{(h)}] \models_{LTL} \chi \neq \perp\}\end{aligned}$$

$$\begin{aligned}\varsigma_{\mathcal{A}}^h(-1) &= \emptyset, \\ \varsigma_{\mathcal{A}}^h(i) &= \varsigma_{\mathcal{A}}^h(i-1) \cup \{t^{(i)} \mid t \in \sigma_{\mathcal{A}}^h(i-1), \exists \chi \in \Gamma_i^t : [t[i]^{(h)}] \models \chi = \perp\}, \text{ for } i \geq 0\end{aligned}$$

To avoid complicating the definition of the function  $\varsigma$  it is defined slightly too general and includes prefixes that do have continuations that satisfy all of an agent's goals. In order to eliminate these prefixes and keep only maximal prefixes of traces that cannot be extended given the agent's goals we additionally introduce a function *maxPrefix*. Let  $t \sqsubset t'$  denote that  $t$  is a strict prefix of  $t'$  and  $T$  a set of (in)finite traces. Then  $t \in \text{maxPrefix}(T)$  iff there is no  $t' \in T$  such that  $t \sqsubset t'$ . Using the definitions of the goal filter function(s), we can now provide a simple definition of the semantics of traces induced by the second layer in the action selection architecture. The meaning of an agent at this layer is denoted by  $H_{\mathcal{A}}$  and defined as the maximal elements of the limit of the filter functions.

**Definition 10** (*Semantics of GOAL Agent with Goals as Hard Constraints*)

The meaning of a GOAL agent  $H_{\mathcal{A}}$  that applies hard constraints in addition to its rule-based action selection is defined by:

$$H_{\mathcal{A}} = \text{maxPrefix} \left( \bigcap_{i=-1}^{\infty} \sigma_{\mathcal{A}}^h(i) \cup \bigcup_{i=-1}^{\infty} \varsigma_{\mathcal{A}}^h(i) \right)$$

It should be clear from the definition of the meaning of a GOAL agent that the semantics introduced clearly distinguished between the different layers of the RASA. Moreover, given the computational properties of  $LTL_3$  the semantics of the second layer can be realized computationally if the initial prefixes of length  $h$  of the traces induced by the first layer can be efficiently generated. In the next section we introduce the third layer to complete our formal picture of the informal RASA discussed in Section 2.

## 6 RASA Layer 3: Preferences as Soft Constraints

In the third layer of the RASA an agent aims to select those traces that maximize satisfaction of its preferences. Similarly to goals, preferences are expressed using LTL. Since LTL formulas express properties of traces, this allows an agent to express that it prefers one way of achieving a goal, i.e., one particular trace, over another, if multiple plans for realizing the goal are available. Such preferences are represented by a so-called *preference structure*. A preference structure consists of a sequence of LTL formulas.

**Definition 11** (*Preference Structure*)

A *preference structure*  $\Psi$  is a sequence  $(\chi_1, \dots, \chi_n)$  of LTL formulas.

This preference structure expresses that traces on which some  $\chi_i$  is satisfied are preferred over traces on which  $\chi_i$  is not satisfied, and the satisfaction of  $\chi_i$  is preferred over the satisfaction of  $\chi_j$  for  $j > i$ . That is, if  $\chi_1$  is satisfied on a trace  $t$  but not on another trace  $t'$ ,  $t$  is preferred over  $t'$ . If both  $t$  and  $t'$  do not satisfy  $\chi_1$ , but  $t$  satisfies  $\chi_2$  and  $t'$  does not,  $t$  is again preferred over  $t'$ , etc. This interpretation of the preference structure thus induces a lexicographic ordering on traces.

This ordering can formally be defined as follows. We use  $\Psi_t$  to denote the sequence  $(b_1, \dots, b_n)$ , where  $b_i = [t \models \chi_i]$  for  $1 \leq i \leq n$ , i.e.,  $b_i \in \{\top, \perp, ?\}$ ; we use  $\Psi_t^i$  to denote  $b_i$ . We now use the ordering  $\perp < ? < \top$  to define a lexicographic preference ordering on traces on the basis of a preferences structure. That is, if for a  $\chi_i$  of the preference structure we have  $[t \models \chi_i] = \top$ , this is better than when  $[t \models \chi_i] = ?$ , which is again better than when  $[t \models \chi_i] = \perp$ .

**Definition 12** (*Lexicographic Preference Ordering*)

Let  $\Psi = (\chi_1, \dots, \chi_n)$  be a preference structure and  $t, t' \in M^* \cup M^\omega$  be finite or infinite traces. Then we say that trace  $t$  is (lexicographically) preferred over  $t'$  with respect to  $\Psi$ , written  $t \prec_\Psi t'$ , iff:

$$\exists 1 \leq j \leq n : \forall 1 \leq i < j : (l_t^i = l_{t'}^i \text{ and } l_t^j < l_{t'}^j)$$

We also write  $t \preceq_\Psi t'$  if  $t \prec_\Psi t'$  or  $\psi_t = \psi_{t'}$ .

The main advantage of using a lexicographic preference order  $\prec_\Psi$  is that such a preference order on traces is a total preorder, which means that any two traces are either equally good with respect to the preference order, or one is better than the other. It is inspired by the work of [9]. There are, however, many differences with [9] and our approach. Instead of integrating preferences into the situation calculus as in [9] we propose a uniform framework for goals and preferences that is integrated into a programming framework for rational agents, and instead of compiling preferences into certain programs we have taken a more direct approach by defining a layered action selection architecture that is made precise by means of a formal semantics. Finally, we use three-valued LTL to evaluate goals and preferences on (prefixes of) traces.

Preferences are progressed or updated through time, just like goals. The idea is that a preference such as  $p \wedge \bigcirc q \wedge \diamond r$  has been satisfied when now  $p$  holds, in the next state  $q$  and possibly sometime thereafter  $r$  holds. Such a preference cannot be satisfied anymore when  $p$  is not true currently, and to express this we use the *Progression* operator of Section 3.2 to keep track of such facts. Since preferences are progressed during execution of the agent, we extend mental states of an agent with the agent's preference structure.

**Definition 13** (*Mental State with Preferences*)

Let  $m = \langle \Sigma, \Gamma \rangle$  be a mental state, and  $\Psi$  a preference structure. A mental state with preferences then simply is the tuple  $\langle \Sigma, \Gamma, \Psi \rangle$ .

**Definition 14** (*Progression of Preference structure*)

The progression of a preference structure  $\Psi = (\chi_1, \dots, \chi_n)$  from a mental state  $\langle \Sigma, \Gamma, \Psi \rangle$  to a new state  $\langle \Sigma', \Gamma', \text{Progress}(\Psi, \Sigma') \rangle$  is simply defined as the progression of each of the individual preference in the structure, i.e.

$$\text{Progress}(\Psi, \Sigma') = (\text{Progress}(\chi_1, \Sigma'), \dots, \text{Progress}(\chi_n, \Sigma'))$$

The semantics of conditional action execution of Definition 6 is changed accordingly as follows, where  $m = \langle \Sigma, \Gamma, \Psi \rangle$  and  $c = \mathbf{if} \ \psi \ \mathbf{then} \ \mathbf{a}$  is a conditional action:

$$\frac{m \models \psi \quad \mathcal{M}(\mathbf{a}, m) = \langle \Sigma', \Gamma' \rangle}{\langle \Sigma, \Gamma, \Psi \rangle \xrightarrow{c} \langle \Sigma', \Gamma', \text{Progress}(\Psi, \Sigma') \rangle}$$

The main difference between goals and preferences in our approach is that goals are used as hard constraints to guide action selection whereas preferences are used as soft constraints. That is, an agent would like to satisfy all of its preferences but if this is not possible it will simply choose to satisfy those that are most preferred, if any can be satisfied at all. The preference order introduced above can be used for this purpose. In the third layer of the RASA those traces are selected from the remaining ones (those that survived filtering by layer 2) that are maximal elements in this order. In order to make things precise, we introduce some notation again. We use  $\max(T, \prec_\psi)$  to denote the maximal elements of a set of traces  $T$  under the lexicographic preference order  $\prec_\psi$ , induced by the preference structure  $\Psi$ . The preference filter function of layer 3 in our action selection architecture can then be defined similarly to that of the goal filter function.

**Definition 15** (*Preference Filter Function  $\psi$* )

Let  $\mathcal{A}$  be some agent using layer 2 goal filtering with meaning  $H_{\mathcal{A}}$  and let  $h \in \mathbb{N}$  be a horizon. Then the *preference function*  $\psi_{\mathcal{A}}^h$  is defined by induction as follows:

$$\begin{aligned} \psi_{\mathcal{A}}^h(-1) &= H_{\mathcal{A}}, \\ \psi_{\mathcal{A}}^h(i) &= \{t \mid t^{(i+h)} \in \max(\{t[i]^{(h)} \mid t \in \psi_{\mathcal{A}}^h(i-1)\}, \prec_\psi)\} \end{aligned}$$

Similar to Section 5, the semantics of an agent that uses all three layers of the action selection architecture is defined as the limit of the preference filter function over the traces obtained from layer 2. The definition is somewhat simpler since an agent can always continue on a given trace since it only needs to select a maximum continuation of its past action performance but does not have to stop acting altogether due to a potential violation of a hard constraint. This is one of the main differences between layer 2 and 3.

**Definition 16** (*Semantics of GOAL with Preferences*)

The meaning of a GOAL agent  $P_{\mathcal{A}}$  that applies soft constraints with horizon  $h$  in addition to the action selection mechanisms of layer 1 and 2 is defined by:

$$P_{\mathcal{A}} = \bigcap_{i=-1}^{\infty} \psi_{\mathcal{A}}^h(i)$$

Definition 16 completes the specification of each of the three layers of the rational action selection architecture. The informal discussion of the architecture in which we distinguished three layers of action selection is mirrored in respectively Definitions 7, 10, and 16. The formal approach has shown that it is feasible to define a transparent and rich RASA into an agent programming language, that integrates beliefs, goals and preferences as tools for choosing the right action. In particular, the approach offers a uniform framework based on temporal logic to express goals and preferences.

## 7 Conclusion and Related Work

In this paper, we have proposed a layered rational action selection architecture which specifies a rational action selection mechanism in which hard and soft constraints are integrated. We have shown how the proposed architecture can be formalized in the context of the GOAL agent programming language. In order to obtain a uniform framework that naturally allows for the integration of hard and soft constraints including such concepts as achievement goals, maintenance goals, as well as (temporally extended) preferences, we have used linear temporal logic for the representation of goals and preferences.

The way in which we use temporal logic is inspired by work in planning where temporal logic is used to “guide” planners through the search space [1] and to select preferred plans [9, 2]. The main difference between the realization of a RASA in the context of planning and in the context of agent programming languages, is that in planning one compares *complete plans* against the available constraints, while in agent programming the constraints need to be taken into account continuously *during* execution. Technically, this has resulted in the use of 3-valued LTL for the realization of the RASA in GOAL.

In [13], a unifying approach for goals in agent systems is proposed, in which LTL is used for illustrating a unifying definition of the notion of goal. That paper, however, does not propose the use of LTL in an agent programming language, and does not propose a RASA in which various kinds of constraints are integrated as we do in this paper.

## References

1. F. Bacchus and F. Kabanza. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 166, 2000.
2. Jorge A. Baier, Fahiem Bacchus, and Sheila A. McIlraith. A heuristic search approach to planning with temporally extended preferences. In *IJCAI*, pages 1808–1815, 2007.
3. Andreas Bauer, Martin Leucker, and Christian Schallhart. Monitoring of real-time properties. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)*, volume 4337 of *LNCS*, pages 260–272. Springer, 2006.
4. Rafael H. Bordini, Mehdi Dastani, Jürgen Dix, and Amal El Fallah Seghrouchni. *Multi-Agent Programming: Languages, Platforms and Applications*. Springer, Berlin, 2005.

5. Mehdi Dastani, M. Birna van Riemsdijk, and John-Jules Ch Meyer. A grounded specification language for agent programs. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)*, pages 578–585, Honolulu, Hawaii, 2007. IFAAMAS.
6. Frank de Boer, Koen Hindriks, Wiebe van der Hoek, and John-Jules Meyer. A Verification Framework for Agent Programming with Declarative Goals. *Journal of Applied Logic*, 5:277–302, 2007.
7. E.A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, pages 996–1072. Elsevier, Amsterdam, 1990.
8. R.E. Fikes and N.J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
9. Christian Fritz and Sheila A. McIlraith. Decision-theoretic golog with qualitative preferences. In *KR*, pages 153–163, 2006.
10. Koen Hindriks. Modules as policy-based intentions: Modular agent programming in goal. In *Proceedings of the International Workshop on Programming Multi-Agent Systems (PROMAS'07)*, number 4908 in LNAI. Springer, 2008.
11. Koen Hindriks and Birna van Riemsdijk. Satisfying maintenance goals. In *Proceedings of the International Workshop on Declarative Agent Languages and Theory (DALT'07)*, number 4897 in LNAI. Springer, 2008.
12. Koen V. Hindriks, Frank S. de Boer, Wiebe van der Hoek, and John-Jules Ch. Meyer. Agent Programming with Declarative Goals. In *Proceedings of ATAL00*, volume 1986 of *LNCIS*, pages 228–243, 2000.
13. M. Birna van Riemsdijk, Mehdi Dastani, and Michael Winikoff. Goals in agent systems: A unifying framework. In *Proceedings of the seventh international joint conference on autonomous agents and multiagent systems (AAMAS'08)*, Estoril, 2008. To appear.