

# The agent programming language Meta-APL

Thu Trang Doan, Natasha Alechina, and Brian Logan

University of Nottingham, Nottingham NG8 1BB, UK  
{ttd, nza, bs1}@cs.nott.ac.uk

**Abstract.** We describe a novel agent programming language, Meta-APL, and give its operational semantics. Meta-APL allows both agent programs and their associated deliberation strategy to be encoded in the same programming language. We define a notion of equivalence between programs written in different agent programming languages based on the notion of weak bisimulation equivalence. We show how to simulate (up to this notion of equivalence) programs written in other agent programming languages by programs of Meta-APL. This involves translating both the agent program and the deliberation strategy under which it is executed into Meta-APL.

## 1 Introduction

In this paper we sketch the agent programming language Meta-APL. Its distinguishing feature is that the agent’s deliberation strategy can be encoded as part of the agent program. Meta-APL is designed to form part of a platform for verifying multi-agent systems where the agents are implemented in different (BDI-based) agent programming languages. As part of the verification process, the agent programs will be translated into Meta-APL and verification tools designed for Meta-APL will be used to verify the system. Similar approaches have been proposed before, see for example [4]. The distinguishing feature of our approach is that Meta-APL is itself an agent programming language (rather than a special purpose library as in [4]) and that the deliberation strategy of the target agent can be expressed in the Meta-APL program along with the agent program itself. It is clearly necessary to encode the deliberation strategy correctly, as executing the same program under different strategies may give very different results. Other agent programming languages have also been used to program agent deliberation, see for example [1]. However Meta-APL is specifically designed with this capability in mind, for example it is possible to write rules in Meta-APL which match against the contents of the agent’s plan base.

This paper describes the first step towards this long term goal, concentrating on the design and the operational semantics of Meta-APL itself. We also define a notion of equivalence or bisimulation between programs written in different agent programming languages. We sketch how programs written in other agent programming languages plus their deliberation strategies can be translated in Meta-APL so that the resulting Meta-APL program is equivalent to the original program together with its deliberation strategy.

The remainder of this paper is organised as follows. In section 2 we introduce the syntax of Meta-APL. In section 3 we define its operational semantics. In section 4 we

define the notion of equivalence between programs and show how to simulate programs written in 3APL [2]. We conclude and outline directions for future work in section 5.

## 2 Syntax of Meta-APL

The language of beliefs and goals in Meta-APL is similar to that of other BDI-based agent programming languages, for example, 3APL [2], but we assume a propositional language for ease of presentation (to avoid extra notation to do with substitutions etc. The actual implementation has Prolog-like syntax for beliefs).

We assume that beliefs and goals are built using propositional atoms from a finite set  $Prop$ . Beliefs are either atoms  $p$ , or Horn clauses  $p : - q_1, \dots, q_n$ . A belief base is a finite set of beliefs. A belief query  $\phi$  is defined as  $\phi ::= p \mid \text{not } p \mid \phi_1 \text{ and } \phi_2 \mid \phi_1 \text{ or } \phi_2$ , where not is negation as failure. Goals are atoms or conjunctions of atoms. A goal base is a finite set of goals.

Before we introduce plans, we define the notion of a plan body. A plan body is a finite sequence of basic actions, test actions, meta-actions and sub-goals, where

- A basic action has the form of  $\#a$  where the symbol  $\#$  is used to indicate that this is a basic action, and  $a$  is the name of the basic action.
- A test action has the form of  $?\varphi$  where  $?$  is used to indicate that this is a test action, and  $\varphi$  is a belief query.
- A sub-goal has the form of  $!g$  where  $!$  is used to indicate that this is a sub-goal, and  $g$  is a goal.
- Meta-actions are actions that allow the agent to add and delete beliefs and goals, and to delete and execute applicable plans. The set of meta-actions is as follows:
  - $\text{add-bel}(d)$  is for adding a belief  $d$  into a belief base.
  - $\text{del-bel}(d)$  is for deleting a belief  $d$  from a belief base.
  - $\text{add-goal}(d)$  is for adding a goal  $d$  into a goal base.
  - $\text{del-goal}(d)$  is for deleting a goal  $d$  from a goal base.
  - $\text{del-plan}(i)$  is for deleting an applicable plan  $i$  from a plan base.
  - $\text{exec}(i)$  is for executing an applicable plan  $i$ .
  - $\text{step}(i)$  is for executing a single step of an applicable plan  $i$ .

A plan in Meta-APL represents the triggering conditions and execution state of a plan body. A plan is a tuple of the form  $(g, b, \pi, x, \pi')$  where:

- $g$  is a goal,
- $b$  is a context query (defined below),
- $\pi$  is an initial plan-body,
- $x$  is a flag for specifying the state of the plan which can have one of the following values:  $a$  (to say that it is active),  $ex$  (to say that it is executed),  $na$  (to say that it is not active).
- $\pi'$  is a partially executed plan-body.

A plan base is a finite set of plans.

Context queries are evaluated against the agent's belief and plan bases. In order to define context queries, we first need to introduce the notion of plan-body terms, goal

terms and flag terms. Informally, a plan-body term is a plan-body except that variables may occur where plan-bodies normally are. Given a set *Vars* of variables, the syntax of plan-body terms is as follows:

$$t_\pi ::= X \mid !s \mid ?\varphi \mid \#a \mid \text{ma}(d) \mid \text{mb}(t_\pi) \mid t_\pi; t'_\pi$$

where  $X \in \text{Vars}$ ,  $!s$  is a sub-goal,  $?\varphi$  is a test action,  $\#a$  is a basic action,  $\text{ma}(d)$  is a meta-action with  $\text{ma} \in \{\text{add-bel}, \text{del-bel}, \text{add-goal}, \text{del-goal}\}$ ,  $d$  is a belief, and  $\text{mb}(t_\pi)$  is also a meta-action with  $\text{mb} \in \{\text{del-plan}, \text{exec}, \text{step}\}$ .

A goal term is either a variable or a goal. A flag term is either a variable or a flag. Then, context queries are defined by the following syntax:

$$b ::= X \mid t_1 = t_2 \mid t_1 \neq t_2 \mid \varphi \mid p(t_g, b, t_\pi, t_x, t_{\pi'}) \mid \neg p(t_g, b, t_\pi, t_x, t_{\pi'}) \mid b \& b'$$

where  $X \in \text{Vars}$ ,  $t_1, t_2, t_\pi$  and  $t_{\pi'}$  are plan-body terms,  $\varphi$  is a belief query,  $t_g$  is a goal term, and  $t_x$  is a flag term. Evaluation of context queries will be defined in the next section when we define the operational semantics of Meta-APL. Informally, a belief query is evaluated against the belief base in a standard way,  $t_1 = t_2$  is used to check if two terms  $t_1$  and  $t_2$  are unifiable,  $p(t_g, b', t_\pi, t_s, t_{\pi'})$  means that there is a plan in the plan base that can unify with  $(t_g, b', t_\pi, t_s, t_{\pi'})$ ; and  $\neg p(t_g, b', t_\pi, t_s, t_{\pi'})$  means there is no applicable plan in the plan base which can unify with  $(t_g, b', t_\pi, t_s, t_{\pi'})$ .

A plan  $(g, b, \pi, x, \pi')$  is effectively identified by the first three components  $g, b$  and  $\pi$ . It can not happen that two applicable plans have the same first three components in a plan base at the same time. During the existence of the plan, its components  $g, b, \pi$  stay unchanged.

In Meta-APL, plans are generated by means of rules. A rule has the form:

$$g, b \rightarrow \pi.$$

where  $g$  is a goal and is optional,  $b$  is a context query where variables are not allowed to occur outside the scope of the atom  $p(\dots)$ , and  $\pi$  is a plan-body. When a rule has no goal, we define that its “hidden” goal is  $\top$ .

### 3 Operational semantics of Meta-APL

An agent program consists of an initial belief base, an initial goal base and a set of rules. The initial plan base is empty.

A configuration is a tuple of the form  $\langle \sigma, \gamma, \Pi, D \rangle$  where  $\sigma$  is a belief base,  $\gamma$  is a goal base,  $\Pi$  is a plan base and  $D$  is a phase indicator of the deliberation cycle which can have one of the following values to indicate in which phase the configuration is: `UpdatePercept`, `ApplyRule`, and `Exec`.

Informally, an agent runs by repeatedly performing a deliberation cycle. In the deliberation cycle, there are three main phases: updating percepts (`UpdatePercept`), matching and applying rules (`ApplyRule`), and executing executable intentions (`Exec`). At the beginning of a deliberation cycle, the agent first updates its percepts where the belief base of the agent is updated according to the percepts collected from the environment.

In this phase, the agent also updates its goal base by adding new goals which are received from outside and dropping goals which become achieved after the belief base is updated. In the next phase, the agent looks for applicable rules from the set of rules against the belief base, the goal base and the plan base. Then, an arbitrary applicable rule is applied to add a new applicable plan into the plan base. Notice that this newly added applicable plan may enable or disable the applicability of other rules. The agent then repeats looking for applicable rules and applying them, one by one, until no more are found. This is when the agent switches to the next phase where it executes executable intentions. In this phase, an executable applicable plan is selected and executed. After a plan is executed, the flag of the plan is set to be “ex”. The phase is continued until all executable plans are marked with “ex”. Then, the phase is changed to **Update-Percept** for starting a new deliberation cycle and all “ex” applicable plans are changed to “a”.

In the following, we discuss each phase of the deliberation cycle in more detail and define the operational semantics by describing transition rules which transform one configuration to another.

First we give a definition of evaluation of beliefs, goals and belief queries. For a belief base  $\sigma$  and a belief or goal  $d$ , we say that  $\sigma \models_{Pr} d$  iff  $d$  is propositionally entailed by  $\sigma$ . For a goal base  $\gamma$  and a goal  $d$ , we say that  $\gamma \models_g d$  iff  $d$  propositionally follows from one of the goals in  $\gamma$ . Finally, for belief query  $\phi$  and a belief base  $\sigma$ , we define  $\sigma \models_{naf} \phi$  as follows:

$$\begin{aligned} \sigma \models_{naf} p &\text{ iff } \sigma \models_{Pr} p. \\ \sigma \models_{naf} \neg p &\text{ iff } \sigma \not\models_{Pr} p. \\ \sigma \models_{naf} \phi_1 \text{ and } \phi_2 &\text{ iff } \sigma \models_{naf} \phi_1 \text{ and } \sigma \models_{naf} \phi_2. \\ \sigma \models_{naf} \phi_1 \text{ or } \phi_2 &\text{ iff } \sigma \models_{naf} \phi_1 \text{ or } \sigma \models_{naf} \phi_2. \end{aligned}$$

Next we define how to evaluate a context query, where variables can occur only within the scope of the literal  $p(\dots)$ , against a configuration  $\langle \sigma, \gamma, \Pi, D \rangle$ . We write  $t = g \mid \theta$  to say that two terms  $t$  and  $g$  are unifiable by the most general unifier (mgu)  $\theta$ . When two terms  $t$  and  $g$  fail to unify, we write  $t \neq g$ . We evaluate a context query against a configuration  $\langle \sigma, \gamma, \Pi, D \rangle$  inductively as follows:

- $\langle \sigma, \gamma, \Pi, D \rangle \models \varphi \mid \emptyset$  iff  $\sigma \models_{naf} \varphi$  where  $\emptyset$  is used to denote an empty substitution
- $\langle \sigma, \gamma, \Pi, D \rangle \models t_1 = t_2 \mid \theta$  iff the two terms  $t_1$  and  $t_2$  are unifiable by the mgu  $\theta$ , that is  $t_1 = t_2 \mid \theta$ .
- $\langle \sigma, \gamma, \Pi, D \rangle \models t_1 \neq t_2 \mid \emptyset$  iff the two terms  $t_1$  and  $t_2$  are not unifiable.
- $\langle \sigma, \gamma, \Pi, D \rangle \models p(t_g, b', t_\pi, t_s, t_{\pi'}) \mid \theta$  iff there exists an applicable plan  $i \in \Pi$  such that  $p(t_g, b', t_\pi, t_s, t_{\pi'}) = i \mid \theta$ .
- $\langle \sigma, \gamma, \Pi, D \rangle \models \neg p(t_g, b', t_\pi, t_s, t_{\pi'}) \mid \emptyset$  iff for all applicable plans  $i \in \Pi$ , we have that  $p(t_g, b', t_\pi, t_s, t_{\pi'}) \neq i$ .
- $\langle \sigma, \gamma, \Pi, D \rangle \models b_1 \& b_2 \mid \theta$  iff  $\langle \sigma, \gamma, \Pi, D \rangle \models b_1 \mid \theta$  and  $\langle \sigma, \gamma, \Pi, D \rangle \models b_2 \mid \theta$

Given a plan base  $\Pi$ , in order to determine which plans can be executed, we define the relation “ $>$ ” over plans in  $\Pi$  as follows. Given

$$\begin{aligned} i_1 &= (g_1, b_1, \pi_1, s_1, \pi'_1) \\ i_2 &= (g_2, b_2, \pi_2, s_2, \pi'_2) \end{aligned}$$

we say that  $i_1 > i_2$  iff  $p(g_2, b_2, \pi_2, -, -)$  occurs in  $b_1$  (is a subformula of  $b_1$ ). We use the notation underscore  $_$  as in Prolog for representing any value. This means  $i_1$  is created because of the existence of  $i_2$  and we shall call  $i_1$  to be the meta plan (of  $i_2$ ). In each deliberation cycle of the agent, we only execute the maximal meta applicable plan, which could indirectly lead to the execution of the lower meta plans through the help of the meta-actions for executing intentions.

Then, we define the set of active plans (those with the flag to be  $a$ ), the set of roots (the most meta applicable plan), the set of leafs (the least meta applicable plans or the object applicable plans), and the set of executable applicable plans (only roots which are not leafs are allowed to execute), respectively, with respect to a plan base  $\Pi$  as follows:

$$\begin{aligned} \text{active}(\Pi) &= \{(g, b, \pi, a, \pi') \in \Pi\} \\ \text{root}(\Pi) &= \{i \in \text{active}(\Pi) \mid \nexists i' \in \text{active}(\Pi) : i' > i\} \\ \text{leaf}(\Pi) &= \{i \in \text{active}(\Pi) \mid \nexists i' \in \text{active}(\Pi) : i > i'\} \\ \text{executable}(\Pi) &= \text{root}(\Pi) \setminus \text{leaf}(\Pi) \end{aligned}$$

Before defining transition rules for the operational semantics of Meta-APL, let us model an environment by two functions `env_percept` and `env_perform`. The effect of each function is as follows:

- `env_percept`( $\sigma, \gamma, e$ ) takes a belief base  $\sigma$ , a goal base  $\gamma$ , and the environment  $e$  as arguments. This function returns a pair of an updated belief base and an updated goal base. The current implementation assumes that percepts are atomic formulas, and the updated belief base is obtained by adding new beliefs and removing incorrect beliefs according to the percepts from the environment. Likewise, the updated goal base is obtained by adding new goals and removing achieved goals, also, according to the percepts from the environment.
- `env_action`( $\alpha, e$ ) takes a basic action and the environment as argument. This function performs the action on the environment. For the moment, we assume that this function returns the value *true* or *false* where it only returns *true* iff the basic action is supported by the environment.

At the beginning of a deliberation cycle, an agent always updates its belief base and goal base. The transition rule for the phase of updating percepts is as follows.

$$\frac{\text{env\_percept}(\sigma, \gamma) = (\sigma', \gamma')}{\langle \sigma, \gamma, \Pi, \text{UpdatePercept} \rangle \rightarrow \langle \sigma', \gamma', \Pi, \text{ApplyRule} \rangle} \quad (1)$$

The transition above expresses the phase `UpdatePercept` where the belief base and goal base are updated with the percepts. Apart from the belief base and the goal base being updated, the phase indicator also changes from `UpdatePercept` to `ApplyRule` so that the agent can start the next phase.

We say that a rule  $g, b \rightarrow \pi$  is applicable with respect to a configuration  $\langle \sigma, \gamma, \Pi, D \rangle$  by a substitution  $\theta$  iff the following conditions hold:

- $\gamma \models_g g$ ,

- $\langle \sigma, \gamma, II, D \rangle \models b \mid \theta$ ,
- There is no applicable plan  $(g, b\theta, \pi\theta, X, \pi') \in II$  where  $\pi' \neq \epsilon$ .

Let  $Applicable(\langle \sigma, \gamma, II, D \rangle)$  be the set of pairs of  $(\rho, \theta)$  where  $\rho$  is an applicable rule with respect to  $\langle \sigma, \gamma, II, D \rangle$  and  $\theta$  is the corresponding substitution. Moreover, the last condition is for avoiding the case when a rule may be fired more than once to produce the same applicable plan. The transitions rules for the phase **ApplyRule** are as follows:

$$\frac{\exists((g, b \rightarrow \pi), \theta) \in Applicable(\langle \sigma, \gamma, II, \mathbf{ApplyRule} \rangle)}{\langle \sigma, \gamma, II, \mathbf{ApplyRule} \rangle \rightarrow \langle \sigma, \gamma, II \cup \{(g, b\theta, \pi\theta, a, \pi\theta)\}, \mathbf{ApplyRule} \rangle} \quad (2)$$

The phase will change to the next one when there are no more applicable rules.

$$\frac{Applicable(\langle \sigma, \gamma, II, \mathbf{ApplyRule} \rangle) = \emptyset}{\langle \sigma, \gamma, II, \mathbf{ApplyRule} \rangle \rightarrow \langle \sigma, \gamma, II, \mathbf{Exec} \rangle} \quad (3)$$

Notice that programmers have the responsibility to make sure that the loop of applying applicable rules terminates. A neglectful design of rules can easily cause the phase **ApplyRule** to run forever, for example if one of the rules is  $p(G, B, II, a, X) \rightarrow \mathbf{exec}(G, B, II, a, X)$ .

In the phase of executing applicable plans, we execute the first step of every executable plan (those which are active, root and not leaf). Let us define transition rules corresponding to each type of the first step as follows.

The following rule is for executing a basic action.

$$\frac{i = (g, b, \pi, a, \# \alpha; \pi') \in executable(II) \text{ and } env\_action(\alpha) = true}{\langle \sigma, \gamma, II, \mathbf{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, ex, \pi')\}, \mathbf{Exec} \rangle} \quad (4)$$

When the basic action is not allowed (or supported) by the environment, the applicable plan is put into the inactive state as follows:

$$\frac{i = (g, b, \pi, a, \# \alpha; \pi') \in executable(II) \text{ and } env\_action(\alpha) = false}{\langle \sigma, \gamma, II, \mathbf{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, na, \# \alpha; \pi')\}, \mathbf{Exec} \rangle} \quad (5)$$

The test action simply checks if the belief query is true against the belief base.

$$\frac{i = (g, b, \pi, a, ? \varphi; \pi') \in executable(II) \text{ and } \sigma \models_{naf} \varphi}{\langle \sigma, \gamma, II, \mathbf{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, ex, \pi')\}, \mathbf{Exec} \rangle} \quad (6)$$

When it is not, the test action is not removed from the applicable plan.

$$\frac{i = (g, b, \pi, a, ? \varphi; \pi') \in executable(II) \text{ and } \sigma \not\models_{naf} \varphi}{\langle \sigma, \gamma, II, \mathbf{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, na, ? \varphi; \pi')\}, \mathbf{Exec} \rangle} \quad (7)$$

We execute a sub-goal by simply leaving it there and the programmer needs to define a suitable rule to process the subgoal.

$$\frac{i = (g, b, \pi, a, !h; \pi') \in executable(II)}{\langle \sigma, \gamma, II, \mathbf{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, ex, !h; \pi')\}, \mathbf{Exec} \rangle} \quad (8)$$

The following rule is for the case of the meta-action for deleting beliefs:

$$\frac{i = (g, b, \pi, a, \text{del-belief}(d); \pi') \in \text{executable}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma \setminus \{d\}, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, ex, \pi')\}, \text{Exec} \rangle} \quad (9)$$

The effect of the above rule is to remove the belief  $d$  from the belief base. Similarly, we have the following rules for the case for adding a new belief.

$$\frac{i = (g, b, \pi, a, \text{add-belief}(d); \pi') \in \text{executable}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma \cup \{d\}, \gamma', II \setminus \{i\} \cup \{(g, b, \pi, ex, \pi')\}, \text{Exec} \rangle} \quad (10)$$

where  $\gamma' = \gamma \setminus \{g \in \gamma \mid \sigma \cup \{d\} \models g\}$ . Besides the effect of adding new beliefs, we also remove achieved goals from the goal base. The following rule is for deleting a goal from the goal base:

$$\frac{i = (g, b, \pi, a, \text{del-goal}(d); \pi') \in \text{executable}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma \setminus \{d\}, II \setminus \{i\} \cup \{(g, b, \pi, ex, \pi')\}, \text{Exec} \rangle} \quad (11)$$

Similarly, we have the following rule for adding a new goal into the goal base:

$$\frac{i = (g, b, \pi, a, \text{add-goal}(d); \pi') \in \text{executable}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma', II \setminus \{i\} \cup \{(g, b, \pi, ex, \pi')\}, \text{Exec} \rangle} \quad (12)$$

where  $\gamma' = \gamma \cup \{d\}$  iff  $\sigma \not\models d$ ; otherwise,  $\gamma' = \gamma$ . This means  $d$  is added into the goal base only when it is not an achieved goal.

The next transition rule is for deleting an applicable plan.

$$\frac{i = (g, b, \pi, a, \text{del-plan}(i'); \pi') \in \text{executable}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i, i'\} \cup \{(g, b, \pi, ex, \pi')\}, \text{Exec} \rangle} \quad (13)$$

Then, we define the transition rules for the meta-actions for executing applicable plans. In principle, the “exec” meta-action makes similar effect as executing the first step of an intention.

For convenience, we also define a different transition rule, denoted as  $\rightarrow^i$  for executing an active applicable plan  $i$  which is not required to be a root, but has to be active (i.e. the flag is a). The definitions are the repetition of those above for the execution phase where we replace the condition  $i \in \text{executable}(II)$  by  $i \in \text{active}(II)$ . Then for every transition rule of the form  $\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma', \gamma', II \setminus \{i\} \cup \{i'\}, \text{Exec} \rangle$ , we also define:

$$\frac{i \in \text{active}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow^i \langle \sigma', \gamma', II \setminus \{i\} \cup \{i'\}, \text{Exec} \rangle} \quad (14)$$

Notice that the transition rules  $\rightarrow^i$  are not the operational semantics of Meta-APL but we use them as auxiliary transition rules for defining the operational semantics of the meta actions `exec` and `step` of Meta-APL. We shall define the transition rule for the meta-action  $\text{exec}(i')$  based on  $\rightarrow^{i'}$  as follows:

$$\frac{\begin{array}{l} i = (g, b, \pi, a, \text{exec}(i'); \pi') \in \text{executable}(II) \text{ and} \\ \langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow^{i'} \langle \sigma', \gamma', II \setminus \{i'\} \cup \{i''\}, \text{Exec} \rangle \end{array}}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma', \gamma, II \setminus \{i, i'\} \cup \{i'', (g, b, \pi, ex, \pi'')\}, \text{Exec} \rangle} \quad (15)$$

where  $\pi'' = \text{exec}(i'')$ ;  $\pi'$  if  $i''$  is not an empty plan; otherwise  $\pi'' = \pi'$ . The above rule means that if we have an executable applicable plan which starts with  $\text{exec}(i')$ , and  $i'$  can be executed by means of  $\rightarrow^{i'}$  to become  $i''$ , then  $\text{exec}(i')$  means to execute  $i'$  and to change to  $\text{exec}(i'')$ . The semantics  $\text{step}(i')$  is similar to  $\text{exec}(i')$  except that the action  $\text{step}(i')$  does not execute the remainder  $i''$  of  $i'$ . The transition rule for this meta-action is as follows:

$$\frac{i = (g, b, \pi, a, \text{step}(i'); \pi') \in \text{executable}(II) \text{ and } \langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow^{i'} \langle \sigma', \gamma', II \setminus \{i\} \cup \{i''\}, \text{Exec} \rangle}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i, i'\} \cup \{i'', (g, b, \pi, \text{ex}, \pi')\}, \text{Exec} \rangle} \quad (16)$$

In both cases of the transition rules for `exec` and `step`, if  $i' \notin II$  or is inactive, then  $i$  also becomes an inactive plan.

$$\frac{i = (g, b, \pi, a, \text{exec}(i'); \pi') \in \text{executable}(II) \text{ and } i' \notin \text{active}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, \text{na}, \text{exec}(i'); \pi')\}, \text{Exec} \rangle} \quad (17)$$

We also have:

$$\frac{i = (g, b, \pi, a, \text{step}(i'); \pi') \in \text{executable}(II) \text{ and } i' \notin \text{active}(II)}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma, II \setminus \{i\} \cup \{(g, b, \pi, \text{na}, \text{step}(i'); \pi')\}, \text{Exec} \rangle} \quad (18)$$

Notice that the new definition of the transition rules  $\rightarrow$  for the meta-actions `exec` and `step` also implicitly gives extra definitions of the transition rule  $\rightarrow^i$ . This helps us to define further transition rules for nested meta-actions `exec` and `step`.

Then, when there is no more executable plans, the phase turns back to `UpdatePercept` for a new deliberation cycle. All plans have the flag “ex” are also changed to “a” for further execution in the next deliberation cycle. We also have:

$$\frac{\text{executable}(II) = \emptyset}{\langle \sigma, \gamma, II, \text{Exec} \rangle \rightarrow \langle \sigma, \gamma, II', \text{UpdatePercept} \rangle} \quad (19)$$

where  $II' = (II \setminus \{(g, b, \pi, \text{ex}, \pi') \in II\}) \cup \{(g, b, \pi, a, \pi') \mid (g, b, \pi, \text{ex}, \pi') \in II \wedge \pi' \neq \epsilon\} \cup \{(g, b, \pi, \text{na}, \epsilon) \mid (g, b, \pi, \text{ex}, \epsilon) \in II\}$ .

In the above transition rule, when transiting from the phase `Exec` back to `UpdatePercept` in the deliberation cycle, all applicable plans in the plan base with the flag being `ex` are changed back to `a` so that they are ready for further execution, except those have an empty plan (denoted as  $\epsilon$ ) which are changed to the state `inactive`.

## 4 Simulating 3APL

In this section we show how to translate agent programs of 3APL into Meta-APL. Both languages share similar features, but have different deliberation cycles. We show how meta rules can be used to simulate deliberation cycle of other languages in Meta-APL.

First we define what we mean by simulating one program by another program. We use the concept of *weak bisimulation* [5]. We treat transitions other than basic actions as *internal* or  $\tau$  actions. By a *run of a program* we will mean a sequence



$s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \dots$  where  $s_i$  are agent's configurations and  $a_i$  are transitions of the agent's operational semantics which are either basic actions or other internal  $\tau$  transitions. Two runs  $r = s_0 \xrightarrow{a_0} s_1 \xrightarrow{a_1} s_2 \dots$  and  $r' = s'_0 \xrightarrow{a_0} s'_1 \xrightarrow{a_1} s'_2 \dots$  are *equivalent* if there is a symmetric relation  $R$  (later referred to as equivalence) between the configurations in  $r_1$  and  $r_2$  such that:

- $R(s_0, s'_0)$
- if  $R(s, s')$  then the agent's beliefs about environment in  $s$  and  $s'$  are the same
- if  $R(s, s')$  and  $s \xrightarrow{\tau^*} t_1 \xrightarrow{a} t_2$  in  $r$ , then  $s' \xrightarrow{\tau^*} t'_1 \xrightarrow{a} t'_2$  in  $r'$  and  $R(t_2, t'_2)$ , where  $\tau^*$  is a sequence of 0 or finitely many internal transitions, and  $a$  is the first basic action occurring after  $s$

Intuitively,  $R(s, s')$  means that in configurations  $s$  and  $s'$  the agent has the same beliefs and goals.

We say that a program  $p_1$  simulates another program  $p_2$  iff:

- For every run  $r_1$  of  $p_1$ , there is a run  $r_2$  of  $p_2$  such that  $r_1$  and  $r_2$  are equivalent.
- For every run  $r_2$  of  $p_2$ , there is a run  $r_1$  of  $p_1$  such that  $r_2$  and  $r_1$  are equivalent.

In this paper we show how to translate a program  $p_2$  of some agent programming language into a program  $p_1$  of Meta-APL so that  $p_1$  simulates  $p_2$ .

#### 4.1 3APL

In this paper, we refer to 3APL as the agent programming language which was presented in [2]. Since the version of Meta-APL introduced in this paper for simplicity only allows propositional beliefs and goals, we show how to simulate propositional 3APL programs, but the extension to full 3APL beliefs and goals is straightforward. Moreover, we also slightly modify 3APL in order to omit the plan constructs `if-then-else` and `while-do`.

An agent in 3APL can have three types of rules:

- PG (plan generation) rule:  $g \leftarrow b \mid \pi$
- GR (goal revision) rule:  $g \leftarrow b \mid g'$
- PR (plan revision) rule:  $\pi \leftarrow b \mid \pi'$

Where  $g$  and  $g'$  are goals,  $b$  is a belief query,  $\pi$  and  $\pi'$  are plans. A plan  $\pi$  is a sequence of basic actions, test actions and abstract plans. Moreover, test actions are of the form  $B(\varphi)$  only. Branching and looping constructs (`if-then-else` and `while-do`) in a plan are not allowed as they can be translated into rules in 3APL using abstract plans. Note that omitting `if-then-else` and `while-do` in plans does not reduce the expressiveness of 3APL as they can be represented by abstract plans by using PG and PR rules. For example, the following PG rule of 3APL:

$$g \leftarrow b \mid \pi; \text{if } b' \text{ then } \pi_1 \text{ else } \pi_2 \text{ end-if}; \pi'$$

is translated into an abstract plan by one PG rule and two PR rules as follows:

$$\begin{aligned} g &\leftarrow b \mid \pi; \text{abs}; \pi'. \\ \text{abs} &\leftarrow b' \mid \pi_1. \\ \text{abs} &\leftarrow \neg b' \mid \pi_2. \end{aligned}$$

Similarly, the `while-do` construct

$$g \leftarrow b \mid \pi; \text{while } b' \text{ do } \pi_1 \text{ end-while}; \pi'$$

is translated into abstract plans as follows

$$\begin{aligned} g &\leftarrow b \mid \pi; \text{abs}; \pi'. \\ \text{abs} &\leftarrow b' \mid \pi_1; \text{abs}. \\ \text{abs} &\leftarrow \neg b' \mid \epsilon. \end{aligned}$$

In order to prove the correctness of the simulation, we need to show the equivalence between runs in 3APL and those by the simulation. Firstly, we specify a deliberation cycle of 3APL. A basic deliberation cycle of 3APL is as follows:

1. Apply applicable rules.
2. Execute plans.

However, we have not defined in detail each of the above two stages. In the stage of applying applicable rules, two extreme approaches are either to apply only one rule, or to apply all the rules until no more are applicable (more precisely, compute a set of applicable rules, if it is not empty, choose a rule and apply it, recompute the set of applicable rules, etc. until the set of applicable rules is empty). In the stage of plan execution, a common approach is to pick a plan and to execute one step. If plans for executing a step are selected randomly, we call this an interleaved execution method. Otherwise, if a selected plan is executed completely before any other plan's steps are selected, we called this a non-interleaved execution method.

In order to demonstrate how to simulate 3APL by means of Meta-APL, we choose the following approach to define how the rule application is done:

1. Apply all applicable PG rules (in any order).
2. Continuously pick an applicable PR rule and apply until no more are applicable.

The two different methods of plan executions are considered in the next sections.

## 4.2 Simulating interleaved deliberation cycle of 3APL

Firstly, we translate a plan  $\pi$  in 3APL into Meta-APL by translating each element of  $\pi$ . Let us denote the translated plan as  $tr(\pi)$ .

We translate each type of rules into the corresponding ones in Meta-APL as follows:

- A PG rule  $g \leftarrow b \mid \pi$  is translated into:  $g, b \rightarrow tr(\pi)$ .
- A GR rule  $g \leftarrow b \mid g'$  is translated into the following rule:

$$g, b \ \& \ \neg p(g, -, \text{del-goal}(g); X, a, -) \rightarrow \text{del-goal}(g); \text{add-goal}(g').$$

The above rule is for replacing a goal  $g$  in the goal base with another goal  $g'$ . Comparing to the original rule in 3APL, the guard of the translated rule has an extra condition  $\neg p(\dots)$  which is for preventing from applying the same rule and other rules for revising the same goal  $g$  once this rule has been fired.

– A PR rule  $\pi \leftarrow b \mid \pi'$  is translated into the following rule:

$$\begin{aligned} & p(G, B, P, a, tr(\pi)) \ \& \ b \ \& \\ & \neg p(G', p(G, B, P, a, tr(\pi)) \ \& \ B', P', a, del-plan(G, B, P, a, tr(\pi)); X) \\ & \rightarrow del-plan(G, B, P, a, tr(\pi)); tr(\pi'). \end{aligned}$$

The above rule is for revising a plan  $tr(\pi)$  in the plan base with another plan  $tr(\pi')$ . Similar to the case of GR rules, the guard of the translated rule also has an extra condition  $\neg p(\dots)$  which is for preventing from applying the same rule and other rules for revising the same plan  $tr(\pi)$  once this rule has been fired.

Then, we implement the interleaved deliberation cycle. Below are the rules for implementing the interleaved deliberation cycle of 3APL:

$$\begin{aligned} & \neg p(G1, B1, P1, a, step(G2, B2, P2, a, Y)) \ \& \ p(G3, B3, P3, a, X) \\ & \rightarrow step(G3, B3, P3, a, X). \end{aligned}$$

$$\begin{aligned} & p(G1, B1, P1, a, step(G2, B2, P2, a, X)) \\ & \ \& \ p(G3, B3, P3, a, del-plan(G2, B2, P2, a, X); X') \\ & \rightarrow del-plan(G1, B1, P1, a, step(G2, B2, P2, a, X)). \end{aligned}$$

$$\begin{aligned} & p(G1, B1, P1, a, step(G2, B2, P2, a, X)) \\ & \ \& \ p(G3, B3, P3, a, del-plan(G2, B2, P2, a, X); X') \\ & \rightarrow step(G3, B3, P3, a, del-plan(G2, B2, P2, a, X); X'). \end{aligned}$$

$$\begin{aligned} & p(G1, B1, P1, a, step(G2, B2, P2, a, X)) \\ & \ \& \ X \neq del-plan(G3, B3, P3, a, Y); Y' \\ & \ \& \ \neg p(G4, B4, P4, a, del-plan(G2, B2, P2, a, X); X') \\ & \ \& \ p(G5, B5, P5, a, del-plan(G6, B6, P6, a, Z); Z') \\ & \rightarrow step(G5, B5, P5, a, del-plan(G6, B6, P6, a, Z); Z'). \end{aligned}$$

$$\begin{aligned} & p(G1, B1, P1, a, step(G2, B2, P2, a, X)) \\ & \ \& \ X \neq del-plan(G3, B3, P3, a, Y); Y' \\ & \ \& \ \neg p(G4, B4, P4, a, del-plan(G2, B2, P2, a, X); X') \\ & \ \& \ p(G5, B5, P5, a, exec(G6, B6, P6, a, Z)) \\ & \rightarrow step(G5, B5, P5, a, exec(G6, B6, P6, a, Z)). \end{aligned}$$

We also have the following rule for executing the rules which translate GR rules:

$$p(G1, B1, P1, a, del-goal(G); Z) \rightarrow exec(G1, B1, P1, a, del-goal(G); Z).$$

The first rule selects an arbitrary plan in the plan base for execution but of only one step. Then, once applied, the rule is not applicable any more in that cycle, hence, we



Notice that the above rules are in the abbreviated form.

We sketch here the proof that one cycle in 3APL corresponds to one or more cycles in the simulation by Meta-APL (hence the runs of the two programs are equivalent). There are two cases:

1. Consider a cycle in 3APL, if there are no PR and GR rules applicable, at the end of the cycle, a plan is selected for execution. The corresponding run in Meta-APL also contains only a single cycle, as no rules into which PR and GR rules are translated are applicable, there is no plan starting with `del-plan` or `exec`, hence, a plan which is selected for execution of one step is not blocked. The corresponding cycle is chosen by selecting the corresponding plan in the case of 3APL.
2. Consider a cycle in 3APL where some PR or GR rules are applicable, at the stage of applying PR and GR rules, it is repeated until no more PR and GR rules are applicable. We define sequences of PR rule application which are sequences of plans  $\pi_0, \dots, \pi_k$  where  $\pi_i$  is obtained by applying some PR rule to revise  $\pi_{i-1}$  for all  $i \geq 1$ . We also define a sequence of GR rule application as a sequence of goals  $g_0, \dots, g_m$  where  $g_i$  is replaced by applying some GR rule for all  $i \geq 0$ . Then, let  $n + 1$  be the length of the longest sequence among sequences of PR rule applications, then we construct a run in Meta-APL containing  $n + m + 1$  cycles where the starting and ending configurations are equivalent to configurations in 3APL before and after the cycle, respectively. In the first cycle, we apply all translated PR rules which are applicable by following the order of PR rule application in 3APL (ignore those which is not applicable yet). Of course, since some PR rule is applied, any plan which is selected for execution is blocked. We repeat this again and again and after  $n$  cycles, we must reach a configuration where no more PR rules are applicable. Then, the next cycles are for applying GR rules in the order of the corresponding to the sequence of GR rule application. The final cycle is just for selecting the corresponding plan in 3APL for execution (and it is not blocked as no more PR and GR rules are applicable).

The reverse direction can be shown similarly.

### 4.3 Simulating non-interleaved deliberation cycle of 3APL

In this section, we simulate the non-interleaved deliberation cycle of 3APL. In this deliberation cycle, we keep executing a plan and any plans which revise this plan until it becomes empty. The implementation of the non-interleaved deliberation cycle is quite similar to the case of interleaved deliberation cycle as in the previous section except that we interchange the use of the meta-actions `step` and `exec`.

In particular, we keep the function `tr` to translate plans, the translation of PG, GR and PR rules from 3APL to Meta-APL unchanged. The only difference comparing to the case of the interleaved deliberation cycle of 3APL is the implementation of the non-interleaved deliberation cycle. The rules to implement the non-interleaved deliberation cycle of 3APL are as follows:

$$\neg p(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, Y)) \ \& \ p(G3, B3, P3, a, X)$$

$\rightarrow \text{exec}(G3, B3, P3, a, X)$ .

$p(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, X))$   
 $\& p(G3, B3, P3, a, \text{del-plan}(G2, B2, P2, a, X); X')$   
 $\rightarrow \text{del-plan}(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, X))$ .

$p(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, X))$   
 $\& p(G3, B3, P3, a, \text{del-plan}(G2, B2, P2, a, X); X')$   
 $\rightarrow \text{exec}(G3, B3, P3, a, \text{del-plan}(G2, B2, P2, a, X); X')$ .

$p(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, X))$   
 $\& X \neq \text{del-plan}(G3, B3, P3, a, Y); Y'$   
 $\& \neg p(G4, B4, P4, a, \text{del-plan}(G2, B2, P2, a, X); X')$   
 $\& p(G5, B5, P5, a, \text{del-plan}(G6, B6, P6, a, Z); Z')$   
 $\rightarrow \text{step}(G5, B5, P5, a, \text{del-plan}(G6, B6, P6, a, Z); Z')$ .

$p(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, X))$   
 $\& X \neq \text{del-plan}(G3, B3, P3, a, Y); Y'$   
 $\& \neg p(G4, B4, P4, a, \text{del-plan}(G2, B2, P2, a, X); X')$   
 $\& p(G5, B5, P5, a, \text{step}(G6, B6, P6, a, \text{add-goal}(G); Z))$   
 $\rightarrow \text{step}(G5, B5, P5, a, \text{step}(\text{add-goal}(G); Z))$ .

$p(G1, B1, P1, a, \text{exec}(G2, B2, P2, a, X))$   
 $\& X \neq \text{del-plan}(G3, B3, P3, a, Y); Y'$   
 $\& \neg p(G4, B4, P4, a, \text{del-plan}(G2, B2, P2, a, X); X')$   
 $\& p(G5, B5, P5, a, \text{step}(G6, B6, P6, a, \text{del-goal}(G); Z))$   
 $\rightarrow \text{step}(G5, B5, P5, a, \text{step}(G6, B6, P6, a, \text{del-goal}(G); Z))$ .

Similar to the implementation of the interleaved deliberation cycle of 3APL, the first rule is also to select a plan to execute by using the meta action `exec`. Since `exec` is used, this selection of the plan to execute is still kept in the next deliberation cycle if it does not become empty. We also have the next two rules is for changing the selection of a plan to its parents when it is revised by some rule. Finally, the last three rules are also for blocking the selected plan from being executed if some plans or goals are revised. Notice that we have more than one rule comparing to the implementation of the interleaved deliberation cycle of 3APL.

In order to execute the rules which translate GR rules, we have the following rules:

$p(G1, B1, P1, a, \text{del-goal}(G); Z) \rightarrow \text{step}(G1, B1, P1, a, \text{del-goal}(G); Z)$ .  
 $p(G1, B1, P1, a, \text{add-goal}(G); Z) \rightarrow \text{step}(G1, B1, P1, a, \text{add-goal}(G); Z)$ .

It is straightforward to prove that the translated program in Meta-APL simulates the non-interleaved deliberation cycle of 3APL. The proof is similar to that of the interleaved case in the previous section.

## 5 Conclusions and future work

We have introduced the syntax and operational semantics of Meta-APL. We have sketched how it can be used to simulate programs written in other agent programming languages together with their operational semantics. In our future work, we plan to develop automatic methods for producing provably equivalent translations of agent programs in Meta-APL and a set of tools for automatically verifying properties of agent systems implemented in Meta-APL.

## References

1. M. Dastani, F. de Boer, F. Dignum, and J.J. Meyer. Programming agent deliberation: an approach illustrated using the 3APL language. In *The Second International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS 2003, July 14-18, 2003, Melbourne, Victoria, Australia, Proceedings*, pages 97–104. ACM, 2003.
2. M. Dastani, F. Dignum, and J.J. Meyer. 3APL: A Programming Language for Cognitive Agents. *ERCIM News, European Research Consortium for Informatics and Mathematics, Special issue on Cognitive Systems*, 2000.
3. Mehdi Dastani, Birna van Riemsdijk, Frank Dignum, and John-Jules Ch. Meyer. A Programming Language for Cognitive Agents Goal Directed 3APL. In Mehdi Dastani, Jürgen Dix, and Amal El Fallah-Seghrouchni, editors, *Programming Multi-Agent Systems, First International Workshop, PROMAS 2003, Melbourne, Australia, July 15, 2003, Selected Revised and Invited Papers*, volume 3067 of *Lecture Notes in Computer Science*, pages 111–130. Springer, 2003.
4. Louise A. Dennis, Berndt Farwer, Rafael H. Bordini, and Michael Fisher. A flexible framework for verifying agent programs. In Lin Padgham, David C. Parkes, Jörg P. Müller, and Simon Parsons, editors, *7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Estoril, Portugal, May 12-16, 2008, Volume 3*, pages 1303–1306. IFAAMAS, 2008.
5. R.J. Van Glabbeek and W.P. Weijland. Branching time and abstraction in bisimulation semantics. *Journal of the ACM (JACM)*, 43(3):555–600, 1996.