

*Autonomous agents coordination: Action languages meet CLP($\mathcal{F}\mathcal{D}$) and Linda**

AGOSTINO DOVIER

Dipartimento di Matematica e Informatica, Università di Udine, Udine, Italy
(e-mail: dovier@dimi.uniud.it)

ANDREA FORMISANO

Dipartimento di Matematica e Informatica, Università di Perugia, Perugia, Italy
(e-mail: formis@dmi.unipg.it)

ENRICO PONTELLI

Department of Computer Science, New Mexico State University, University Park, NM, USA
(e-mail: epontell@cs.nmsu.edu)

submitted 23 November 2010; revised 5 July 2011; accepted 31 August 2011

Abstract

The paper presents a knowledge representation formalism, in the form of a high-level *Action Description Language (ADL)* for multi-agent systems, where autonomous agents reason and act in a shared environment. Agents are autonomously pursuing individual goals, but are capable of interacting through a shared knowledge repository. In their interactions through shared portions of the world, the agents deal with problems of synchronization and concurrency; the action language allows the description of strategies to ensure a consistent global execution of the agents' autonomously derived plans. A distributed planning problem is formalized by providing the declarative specifications of the portion of the problem pertaining to a single agent. Each of these specifications is executable by a stand-alone CLP-based planner. The coordination among agents exploits a Linda infrastructure. The proposal is validated in a prototype implementation developed in SICStus Prolog.

KEYWORDS: Action languages, Multi-agent systems, Constraint logic programming

1 Introduction

Representing and reasoning in multi-agent domains are two of the most active research areas in *multi-agent system (MAS)* research. The literature in this area is extensive, and it provides a plethora of logics for representing and reasoning about various aspects of MAS domains (Fagin *et al.* 1995; van der Hoek *et al.* 2005; Gerbrandy 2006; Sauro *et al.* 2006; Spaan *et al.* 2006).

* Research partially funded by GNCS-INdAM projects, MUR-PRIN: *Innovative and multidisciplinary approaches for constraint and preference reasoning* project; NSF grants IIS-0812267 and HRD-0420407; and grants 2009.010.0336 and 2010.011.0403.

A large number of the logics proposed in the literature have been designed to specifically focus on particular aspects of the problem of modeling MAS, often justified by a specific application scenario. This makes them suitable to address specific subsets of the general features required to model real-world MAS domains. The task of generalizing some of these existing proposals to create a uniform and comprehensive framework for modeling several different aspects of MAS domains is an open problem. Although we do not dispute the possibility of extending several of these existing proposals in various directions, the task does not seem easy. Similarly, a variety of multi-agent programming platforms have been proposed, mostly in the style of multi-agent programming languages, like Jason (Bordini *et al.* 2007), ConGolog (De Giacomo *et al.* 2000), 3APL (Dastani *et al.* 2003), GOAL (de Boer *et al.* 2005), but with limited planning capabilities.

Our effort in this paper is focused on the development of a novel action language for multi-agent systems. The foundations of this effort can be found in the action language B^{MV} (Dovier *et al.* 2010); this is a flexible single-agent action language, which generalizes the action language B (Gelfond and Lifschitz 1998) with support for multi-valued fluents, non-Markovian domains, and constraint-based formulations—enabling, for example, the formulation of costs and preferences. B^{MV} has been implemented in $\text{CLP}(\mathcal{F}\mathcal{D})$.

In this work, we extend B^{MV} to support MAS domains. The perspective is that of a distributed environment, with agents pursuing individual goals but capable of interacting through shared knowledge and through collaborative actions. A first step in this direction has been described in the B^{MAP} language (Dovier *et al.* 2009), a multi-agent action language with capabilities for *centralized* planning. In this paper, we expand on this by moving B^{MAP} towards a truly distributed multi-agent platform. The language is extended with Communication primitives for modeling interactions among *Autonomous Agents*. We refer to this language simply as \mathcal{B}^{AAC} . Differently from B^{MAP} , agents in the framework proposed in this paper have private goals and are capable of developing independent plans. Agents' plans are composed in a distributed fashion, leading to replanning and/or introduction of communication activities to enable a consistent global execution.

The design of \mathcal{B}^{AAC} is validated in a prototype, available from <http://www.dimi.uniud.it/dovier/BAAC>, that uses $\text{CLP}(\mathcal{F}\mathcal{D})$ for the development of the individual plans of each agent and Linda for the coordination and interaction among them.

2 Syntax of the multiagent language \mathcal{B}^{AAC}

The signature of \mathcal{B}^{AAC} consists of:

- (1) A set \mathcal{G} of *agent* names, used to identify the agents in the system;
- (2) A set \mathcal{A} of *action* names;
- (3) A set \mathcal{F} of *fluent* names—i.e., predicates describing properties of objects in the world, and providing description of states of the world; such properties might be affected by the execution of actions; and
- (4) A set \mathcal{V} of values for the fluents in \mathcal{F} —we assume $\mathcal{V} = \mathbb{Z}$.

The behavior of each agent a is specified by an action description theory \mathcal{D}_a , composed of axioms of the forms described next.

Considering the action theory \mathcal{D}_a of an agent a , name and priority of the agent are specified by *agent declarations*:

$$\text{agent } a \text{ [priority } n \text{]} \quad (1)$$

where $n \in \mathbb{N}$. We adopt the convention that 0 denotes the highest priority—which also represents the default priority in the absence of a declaration. As we will see, priorities can be used to resolve possible conflicts among actions of different agents.

It is possible to specify which agents are known to the agent a , as follows:

$$\text{known_agents } a_1, a_2, \dots, a_k \quad (2)$$

Agent a can explicitly communicate with any of the agents a_i , as discussed below.

We assume the existence of a “global” set \mathcal{F} of fluents, and any agent a knows and can access only those fluents that are declared in \mathcal{D}_a by axioms of the form:

$$\text{fluent } f_1, \dots, f_h \text{ valued } \text{dom}_i \quad (3)$$

with $\{f_1, \dots, f_h\} \subseteq \mathcal{F}$, $h \geq 1$, and $\text{dom}_i \subset \mathcal{V}$ is a set of values representing the admissible values for each f_i (possibly represented as an interval $[v_1, v_2]$). These fluents describe the “local state” of the agent. We assume that the fluents accessed by multiple agents are defined consistently in each agent’s local theory.

Example 1

Let us specify a domain inspired by volleyball. There are two teams: *black* and *white*, with one player in each team; let us focus on the domain for the white team (Section 3.7 deals with the case that involves more players). We introduce fluents to model the positions of the players and of the ball, the possession of the ball, the score, and a numerical fluent `defense_time`. All players know the positions of all players. Since the teams are separated by the net, the x-coordinates of black and white players must differ. This can be stated by:

```
agent player(white,X) :- num(X).
known_agents player(black,X) :- num(X).
fluent x(player(white,X)) valued [B,E] :- num(X), net(NET),B is NET+1, linex(E).
fluent x(player(black,X)) valued [1,E] :- num(X), net(NET),E is NET-1.
fluent y(A) valued [1,MY] :- player(A), liney(MY).
fluent x(ball) valued [1,MX] :- linex(MX).
fluent y(ball) valued [1,MY] :- liney(MY).
fluent hasball(A) valued [0,1] :- agent(A).
fluent point(T) valued [0,1] :- team(T).
fluent defense_time valued [0,1].
team(black). team(white). num(1). linex(11). net(6). liney(5).
```

where `linex/liney` are the field sizes, and `net` is the x-coordinate of the net.

Fluents are used in *Fluent Expressions* (FE), which are defined as follows:

$$\text{FE} ::= n \mid f^t \mid \text{FE}_1 \oplus \text{FE}_2 \mid -(\text{FE}) \mid \text{abs}(\text{FE}) \mid \text{rei}(\text{C}) \quad (4)$$

where $n \in \mathcal{V}$, $f \in \mathcal{F}$, $t \in \{0, -1, -2, -3, \dots\}$, $\oplus \in \{+, -, *, /, \text{mod}\}$, and $r \in \mathbb{N}$. FE is referred to as a *timeless expression* if it contains no occurrences of f^t with $t \neq 0$. f can be used as a shorthand of f^0 . The notation f^t is an *annotated* fluent expression.

The expression refers to a relative time reference, indicating the value f had $-t$ steps in the past. The last alternative in (4), a *reified expression*, requires the notion of constraint C , introduced below. $\text{rei}(C)$ represents a Boolean value that reflects the truth value of C . A *Primitive Constraint* (PC) is formula $\text{FE}_1 \text{ op } \text{FE}_2$, where FE_1 and FE_2 are fluent expressions, and $\text{op} \in \{=, \neq, \geq, \leq, >, <\}$. A *constraint* C is a propositional combination of PCs. We will refer to the primitive constraints of the form $f = FE$, where $f \in \mathcal{F}$, as a *basic primitive constraint*. We accept the constraint $\text{pair}(\text{FE}_1, \text{FE}_3) = \text{pair}(\text{FE}_2, \text{FE}_4)$ as syntactic sugar of $\text{FE}_1 = \text{FE}_2$ and $\text{FE}_3 = \text{FE}_4$.

An axiom of the form $\text{action } x \text{ in } \mathcal{D}_a$ declares that the action $x \in \mathcal{A}$ is executable by the agent a . Observe that the same action name x can be used for different actions executable by different agents. This does not cause ambiguity, since each agent's knowledge is described by its own action theory. A special action, nop , is executable by every agent, and it does not change any of the fluents.

Example 2

The actions for each player A of Example 1 are:

- $A : \text{move}(d)$ one step in direction d , where d is one of the eight directions: north, north-east, east, ..., west, north-west.
- $A : \text{throw}(d, f)$ the ball in direction d (same eight directions as above) with strength f varying from 1 to a maximum throw power (5 in our example).

Moreover, the player of each team is in charge of checking if a point has been scored (in such case, he whistles). We write the actions as $\text{act}([A], \text{action_name})$ and state these axioms:

```
action act([A],move(D)) :-whiteplayer(A),direction(D).
action act([A],throw(D,F)):-whiteplayer(A),direction(D),power(F).
action act([player(white,1)],whistle).
```

where whiteplayer , power , and direction can be defined as follows:

```
whiteplayer(player(white,N)) :- agent(player(white,N)).
power(1). power(2). power(3). power(4). power(5).
direction(D) :- delta(D,-,-). delta(nw,-1,1). delta(n,0,1). delta(ne,1,1).
delta(w,-1,0). delta(e,1,0). delta(sw,-1,-1). delta(s,0,-1). delta(se,1,-1).
```

The executability of the actions is described by axioms of the form:

$$\text{executable } x \text{ if } C \tag{5}$$

where $x \in \mathcal{A}$ and C is a constraint. The axiom states that x is executable only if C is entailed by the state of the world. We assume that at least one executability axiom is present for each action; multiple executability axioms are treated disjunctively.

Example 3

In our working example, we can state executability as follows:

```
executable act([player(white,1)],whistle) if [S eq 0] :- build_sum(S).
executable act([A],move(D)) if [hasball(A) eq 0, defense_time gt 0,
Net lt x(A)+DX, x(A)+DX leq MX, 1 leq y(A)+DY, y(A)+DY leq MY] :-
action(act([A],move(D))), delta(D,DX,DY),
net(Net), linex(MX), liney(MY).
executable act([A],throw(D,F)) if
[hasball(A) gt 0,defense_time eq 0, 1 leq x(A)+DX*F, x(A)+DX*F leq MX,
1 leq y(A)+DY*F, y(A)+DY*F leq MY] :-
action(act([A],throw(D,F))), delta(D,DX,DY), linex(MX), liney(MY).
```

These axioms state that neither a player nor the ball can leave the field. `build_sum` is recursively defined to return the expression: `defense_time + hasball(A1) + ... + hasball(An)`, where A_1, \dots, A_n are the players (i.e., `player(white,1)` and `player(black,1)`). The operators $=, \neq, \leq, <$, etc. are concretely represented by `eq`, `neq`, `leq`, `lt`, respectively.

The effects of an action execution are modeled by *dynamic causal laws*:

$$x \text{ causes Eff if } Prec \quad (6)$$

where $x \in \mathcal{A}$, $Prec$ is a constraint, and Eff is a conjunction of basic primitive constraints. The axiom asserts that if $Prec$ is true with respect to the current state, then Eff must hold after the execution of x .

Since agents share fluents, their actions may interfere and cause inconsistencies. A *conflict* happens when the effects of different concurrent actions are incompatible and would lead to an inconsistent state; note that we allow only consistent states to exist during the evolution of the world. A procedure has to be applied to resolve a conflict and determine a consistent subset of the conflicting actions (see Section 3.3).

Example 4

Let us describe the effects of the actions in the volleyball domain. When the ball is thrown with force f in direction d , it reaches a destination cell whose distance is as follows: *a*) if d is either north or south, then $\Delta X = 0, \Delta Y = f$; *b*) if d is east or west, then $\Delta X = f, \Delta Y = 0$; *c*) if d is any other direction, $\Delta X = f, \Delta Y = f$. An additional effect is to set the fluent `defense_time` (to 1 in our example).

```
act([A],throw(D,F)) causes hasball(A) eq 0 :- action(act([A],throw(D,F))).
act([A],throw(D,F)) causes defense_time eq 1 :- action(act([A],throw(D,F))).
act([A],throw(D,F)) causes pair(x(ball),y(ball)) eq pair(x(A)-1+ F*DX,y(A)-1+ F*DY) :-
  action(act([A],throw(D,F))), delta(D,DX,DY).
act([A],throw(D,F), causes hasball(B) eq 1
  if [pair(x(B),y(B)) eq pair(x(A)+F*DX, y(A)+F*DY)] :-
    action(act([A],throw(D,F))), player(B), neq(A,B),delta(D,DX,DY).
act([A],throw(D,F)) causes point(black) eq 1 if [x(A)+F*DX eq Net] :-
  action(act([A],throw(D,F))), delta(D,DX,_), net(Net).
```

The effects of the other two actions `move` and `whistle` can be stated by:

```
act([player(white,1)],whistle) causes point(white) eq 1 if [x(ball) lt NET] :-
  net(NET).
act([player(white,1)],whistle) causes point(black) eq 1 if [NET lt x(ball)] :-
  net(NET).
act([A],move(D)) causes pair(x(A),y(A)) eq pair(x(A)-1+DX,y(A)-1+DY) :-
  action(act([A],move(D))), delta(D,DX,DY).
act([A],move(D)) causes defense_time eq defense_time-1- 1 :- action(act([A],move(D))).
act([A],move(D)) causes hasball(A) eq 1
  if [pair(x(ball),y(ball)) eq pair(x(A)+DX,y(A)+DY)] :-
    action(act([A],move(D))), delta(D,DX,DY).
```

In the presence of a conflict (i.e., two agents executing actions that assign a distinct value to the same fluent), at least two perspectives can be followed, by assigning either a passive or an active role to the conflicting agents. In the first case, a supervising entity is in charge of resolving the conflict, and all the agents will comply with the supervisor's decisions. Alternatively, the agents themselves are in charge of reaching an agreement, possibly through negotiation. In the latter case,

the following declarations allow one to specify in the action theories some basic reaction policies the agents might apply:

$$\text{action } x \text{ [OPT]} \tag{7}$$

with *OPT* defined as:

```

OPT ::= on_conflict OC [OPT]
      | on_failure OF [OPT]
OC ::= retry_after T [provided C]
      | forego [provided C]
OF ::= retry_after T [if C]
      | replan [if C] [add_goal C]
      | fail [if C]

```

where *T* is a number of steps and *C* is a constraint. Notice that one can also specify policies to be adopted whenever a failure occurs in executing an action.

We remark here the difference between *conflict* and *failure*. A conflict occurs whenever concurrent actions performed by different agents try to make inconsistent modifications to the state of the world. A failure occurs whenever an action *x* cannot be executed as planned by an agent *a*. This might happen, for instance, because after the detection of a conflict involving *x*, the outcome of the conflict resolution phase requires *x* to be inhibited. In this case, the agent *a* might have to reconsider its plan. Hence, reacting to a failure is a “local” activity the agent might perform after the state transition has been completed. In axioms of the form (7), one can specify different reactions to a conflict (resp., a failure) of the same action. Alternatives will be considered in their order of appearance.

Example 5

Let us assume that the agents *a* and *b* have priority 0, while agent *c* has lower priority 2. Let us also assume that the current state is such that actions *act_a*, *act_b*, and *act_c* are all executable (respectively, by agents *a*, *b*, and *c*), where their effects on fluent *f* are of setting it to 1, 2, and 3, respectively. This indicates a situation of conflict, since the effects of the concurrent execution of the three actions are inconsistent. Assume that the following options have been defined:

```

action act_a on_conflict retry_after 2
action act_b on_conflict forego
action act_c on_failure retry_after 3

```

and that the plan of agent *a* (resp., *b*, *c*) requires the execution of action *act_a* (resp., *act_b*, *act_c*) in the current state. One possible conflict resolution is to focus the priority of the agents. This causes *act_c* to be removed from the execution list. Thus, agent *c* fails in executing *act_c* and will retry the same action after three steps.

Some policy must now be chosen to resolve the conflict between *a* and *b*. The first possibility is that agents have passive roles in conflict resolution, and a supervisor selects, according to some criteria, a consistent subset of the actions/agents. For example, if *a* is selected (e.g., by lexicographic order), then the state will be modified by setting *f* = 1, declaring *act_a* successful, while agent *b* will fail.

An alternative is to allow the agents a and b to directly resolve the conflict, using their `on_conflict` options. This causes a to retry the execution of `act_a` after two time steps and b to forego the execution of `act_b`. Both of them will get a failure message, because neither `act_a` nor `act_b` is executed.

Apart from the possible communications occurring among agents during the conflict resolution phase, other forms of “planned” communication can be modeled in an action theory. An axiom of this form

$$\text{request } C_1 \text{ if } C_2 \quad (8)$$

describes a special static causal law that allows an agent to broadcast a request, whenever a certain condition (C_2) is encountered. By executing this action, an agent asks if there is another agent that can make the constraint C_1 true. Only an agent knowing all of the fluents occurring in C_1 is allowed to answer such a request.

Instead of broadcasting a help request, an agent a can send such a message directly to another agent by providing its name:¹

$$\text{request } C_1 \text{ to_agent } a' \text{ if } C_2 \quad (9)$$

The following communication primitive subsumes the previous ones:

$$\text{request } C_1 [\text{to_agent } a'] \text{ if } C_2 [\text{offering } C_3] \quad (10)$$

If the last option is used, the requesting agent also provides a “reward” by promising to ensure C_3 in case of acceptance of the proposal. Axioms of this type allow us to model negotiations and other forms of bargaining and transactions.

In turn, agents may declare their willingness to accept requests and serve other agents using statements of the form

$$\text{help } Agent_List [\text{if } C] \quad (11)$$

where *Agent_List* is either a list of agent names a_1, \dots, a_k —denoting that the agent in question can serve requests coming from the agents a_1, \dots, a_k —or the keyword *all*—denoting the fact that the agent can accept requests coming from any source. The optional condition allows the agent to select which requests to consider depending on properties of the current state of the world.

Example 6

Let us consider a domain with three agents: a guitar maker, a joiner that provides wooden parts of guitars (bodies and necks), and a seller that sells strings and pickups. We assume that the maker has plenty of money (so we do not take into account what it spends), that the seller wants to be paid for its materials, and that necks and bodies can be obtained for free (e.g., the joiner has a fixed salary paid by the maker). The income of the seller is modeled by changes to the value of the fluent `seller_account`. In Figure 1, we report an action description that models the agent `guitar_maker`—analogous theories can be formulated for the other two agents. Observe that two point-to-point interactions are modeled—namely, the one

¹ Any request sent to a nonexistent agent will never receive an answer.

```

agent guitar_maker.
action make_guitar.
executable make_guitar if neck > 0 and strings ≥ 6 and body > 0 and pickup > 0.

% actions for making two different kinds of guitars:
make_guitar causes guitars=guitars-1+1 and neck=neck-1-1 and body=body-1-1
and strings = strings-1 - 6 and pickup = pickup-1 - 2
if pickup ≥ 2.
make_guitar causes guitars=guitars-1+1 and neck=neck-1-1 and strings = strings-1 - 6
and body=body-1-1 and pickup=pickup-1-1
if pickup < 2.

% interaction with joiner:
request neck > 0 to_agent joiner if neck = 0.
request body > 0 to_agent joiner if body = 0.

% interaction with seller:
request strings > 5 to_agent seller if strings < 6
offering seller_account = seller_account-1 + 8.
request pickup > 0 to_agent seller if pickup = 0
offering seller_account = seller_account-1 + 60.

% the goal is to make 10 guitars:
goal guitars = 10.

% initially the maker owns some material:
initially guitars = 2 and body = 3 and neck = 5 and pickup = 6 and strings = 24.

```

Fig. 1. An action description in \mathcal{B}^{AAC} for a guitar maker agent.

between the `guitar_maker` and the `joiner`, to obtain necks and bodies, and the one between the `guitar_maker` and the `seller`, to buy strings (\$8) and pickups (\$60). Two kind of guitars can be made, differing in the number of pickups.

Various forms of *global constraint* can be exploited to impose control knowledge and maintenance goals. These constraints represent properties that must always persist in the world where the agents act. Some examples:

- *FC holds_at n*. This constraint is satisfied if the fluent constraint *FC* holds at the *n*th time step.
- *always FC*. This constraint imposes the condition that the fluent constraint *FC* holds in all the states of the evolution of the world.

Semantics of these constraints is reported in Section 3.1.

An *action domain description* consists of a collection \mathcal{D}_a of axioms of the form described so far, for each agent $a \in \mathcal{G}$. Moreover, it includes, for each agent a , a collection \mathcal{O}_a of goal axioms (objectives), of the form `goal C`, where *C* is a constraint, and a collection \mathcal{I}_a of initial state axioms of the form: `initially C`, where *C* is a constraint involving only timeless expressions. For the sake of simplicity, we assume that all the sets \mathcal{I}_a are drawn from a consistent global initial state description \mathcal{I} , i.e., $\mathcal{I}_a \subseteq \mathcal{I}$. A specific instance of a planning problem is a triple

$$\langle \langle \mathcal{D}_a \rangle_{a \in \mathcal{G}}, \langle \mathcal{I}_a \rangle_{a \in \mathcal{G}}, \langle \mathcal{O}_a \rangle_{a \in \mathcal{G}} \rangle.$$

3 System behavior

The behavior of \mathcal{B}^{AAC} can be split into two parts: the semantics of the action description language, parametric on the supervisor selection strategy, and these

strategies that can be programmed. We present the former in Section 3.1, the latter in Sections 3.2–3.5. Finally, some implementation notes are reported in Section 3.6.

3.1 Semantics of \mathcal{B}^{AAC}

The semantics of the action language are described by a transition function that operates on states. A state s is identified by a total function $v : \mathcal{F} \rightarrow \mathcal{V}$. We assume a given horizon \mathbf{N} , within which the planning activities of all agents have to be completed.

Let $\vec{v} = \langle v_0, \dots, v_i \rangle$ be a state sequence, with $0 \leq i \leq \mathbf{N}$. Given \vec{v} , $j \in \{0, \dots, i\}$, and a fluent expression φ , we define the concept of *value* of φ in \vec{v} at time j , denoted by $\vec{v}(j, \varphi)$, as follows:

$$\begin{aligned} \vec{v}(j, x) &= x \quad \text{if } x \in \mathcal{V} \\ \vec{v}(j, f^t) &= v_{j+t}(f) \quad \text{if } f \in \mathcal{F}, \text{ and } 0 \leq j+t \\ \vec{v}(j, f^t) &= v_0(f) \quad \text{if } f \in \mathcal{F}, \text{ and } j+t < 0 \\ \vec{v}(j, \text{abs}(\varphi)) &= |\vec{v}(j, \varphi)| \\ \vec{v}(j, \neg(\varphi)) &= -(\vec{v}(j, \varphi)) \\ \vec{v}(j, \varphi_1 \oplus \varphi_2) &= \vec{v}(j, \varphi_1) \oplus \vec{v}(j, \varphi_2) \\ \vec{v}(j, \text{rei}(C)) &= 1 \quad \text{if } \vec{v} \models_j C \\ \vec{v}(j, \text{rei}(C)) &= 0 \quad \text{if } \vec{v} \not\models_j C \end{aligned}$$

where $\oplus \in \{+, -, *, /, \text{mod}\}$. The last two cases specify the semantics of reification that relies on the notion of *satisfaction*, which, in turn, is defined by structural induction on constraints, as follows. Given a primitive constraint $\varphi_1 \text{ op } \varphi_2$ and a state sequence \vec{v} , the notion of satisfaction at time j is defined as: $\vec{v} \models_j \varphi_1 \text{ op } \varphi_2$ iff $\vec{v}(j, \varphi_1) \text{ op } \vec{v}(j, \varphi_2)$. The notion \models_j is generalized to the case of propositional combinations of fluent constraints in the usual manner. For the case of *pair*, we have that $\vec{v} \models_j \text{pair}(E_1, E_3) = \text{pair}(E_2, E_4)$ if and only if $\vec{v} \models_j E_1 = E_2 \wedge E_3 = E_4$.

We recall that a *timeless fluent* is a fluent expression of the form f^0 (and f).

Given a constraint C and a state sequence $\vec{v} = \langle v_0, \dots, v_i \rangle$, let $\text{fluents}(C)$ be the set of timeless fluents occurring in C . A function $\sigma : \text{fluents}(C) \rightarrow \mathcal{V}$ is a \vec{v} -*solution* of C if $\langle v_0, \dots, v_i, \sigma \rangle \models_{i+1} C$. Let us observe that this definition makes use of a slight abuse of notation, since σ is potentially not a complete state (some fluents may have not been assigned a value by σ). Nevertheless, the choice of fluents in $\text{fluents}(C)$ guarantees the possibility of correctly evaluating C . In other words, σ can be seen as a partial state contributing (with \vec{v}) to the satisfaction of C at time $i+1$. Let us see how to complete this state using inertia: if σ is a \vec{v} -solution of a constraint C , $\text{ine}(\sigma, \vec{v})$ is defined as follows:

$$\text{ine}(\sigma, \vec{v})(f) = \begin{cases} \sigma(f), & \text{if } f \in \text{fluents}(C), \\ v_i(f), & \text{otherwise.} \end{cases}$$

Fluents not appearing in C are considered inertial (namely they maintain their previous values) and therefore the state is completed using the function ine .

An action x is *executable* by agent a in a state sequence $\vec{v} = \langle v_0, \dots, v_i \rangle$ if there is at least an axiom *executable* x if C in \mathcal{D}_a and it holds that $\vec{v} \models_i C$. If there is more than one executability condition, it is sufficient for one of them to apply.

Let us denote with $Dyn(x)$ the set of dynamic causal law axioms for action x . The *desired effect* of executing x in state sequence $\vec{v} = \langle v_0, \dots, v_i \rangle$, denoted by $DEff(x, \vec{v})$, is a constraint defined as follows:

$$DEff(x, \vec{v}) = \bigwedge \{ Eff \mid x \text{ causes } Eff \text{ if } Prec \in Dyn(x), \vec{v} \models_i Prec \}.$$

Request accomplishment actions can be used in the construction of this set.

Given a state sequence $\vec{v} = \langle v_0, \dots, v_i \rangle$, a state v_{i+1} , and a set of actions X , a triple $\langle \vec{v}, X, v_{i+1} \rangle$ is a *valid state transition* if:

- for all $x \in X$, the action x is executable in \vec{v} by some agent a , and
- $v_{i+1} = \text{ine}(\sigma, \vec{v})$, where σ is a \vec{v} -solution of the constraint $\bigwedge_{x \in X} DEff(x, \vec{v})$.

Observe that if $X = \emptyset$, then $\langle \vec{v}, \emptyset, v_i \rangle$ will be a valid state transition.

Let $\vec{v} = \langle v_0, \dots, v_N \rangle$ be a sequence of states, $\langle (\mathcal{D}_a)_{a \in \mathcal{A}}, (\mathcal{J}_a)_{a \in \mathcal{A}}, (\mathcal{O}_a)_{a \in \mathcal{A}} \rangle$ an instance of a planning problem, and X_1, \dots, X_N be sets of actions. We say that $\langle v_0, X_1, v_1, \dots, X_N, v_N \rangle$ is a *valid trajectory* if:

- for each agent a and for each axiom of the form *initially* C in \mathcal{J}_a , we have that $\vec{v} \models_0 C$,
- for all $i \in \{0, \dots, N - 1\}$, $\langle v_0, \dots, v_i, X_{i+1}, v_{i+1} \rangle$ is a valid state transition.

A valid trajectory is *successful* for an agent a if, for each axiom of the form goal C in \mathcal{O}_a , it holds that $\vec{v} \models_N C$.

At each time step i , each agent might propose a set of actions for execution— we assume that all the proposed actions are executable in the state sequence $\vec{v}_i = \langle v_0, \dots, v_i \rangle$. Let Y_{i+1} be this set of actions. The supervisor selects a subset $X_{i+1} \subseteq Y_{i+1}$ such that the constraint $Eff(X_{i+1}, \vec{v}_i)$, defined as:

$$Eff(X_{i+1}, \vec{v}_i) = \bigwedge_{x \in X_{i+1}} DEff(x, \vec{v}_i)$$

is satisfiable with respect to \vec{v} —i.e., there exists a complete state v_{i+1} such that $\langle \vec{v}_i, X_{i+1}, v_{i+1} \rangle$ is a valid state transition. It is the job of the supervisor to determine the subset X_{i+1} given Y_{i+1} and \vec{v}_i —as a maximal consistent subset, using agent priorities or other approaches, as discussed in Section 3.3. If an agent cannot find a plan at the time step i , it will ask for a nop and try again the next step.

Let us complete the semantics of the language by dealing with request and help laws. A request of the agent a

$$\text{request } C_1 \text{ to_agent } a' \text{ if } C_2$$

is *executable* in a state sequence $\vec{v} = \langle v_0, \dots, v_i \rangle$ if it holds that $\vec{v} \models_i C_2$. If the request above is executable, it can be *accomplished* in the successive state v_{i+1} if there is an axiom

$$\text{help } \dots a \dots \text{ if } C_3$$

in $\mathcal{D}_{a'}$ and $\langle v_0, \dots, v_i, v_{i+1} \rangle \models_{i+1} C_3$. The semantics of the help law is that of enabling a request accomplishment (after a request demand) and it can be viewed as the execution of an ordinary action by agent a' .² We can view this as if a' had an additional action y defined in $\mathcal{D}_{a'}$ as:

$$\begin{array}{l} \text{executable } y \text{ if } C_3 \wedge (C_2)^{-1} \\ y \text{ causes } C_1 \text{ if true} \end{array}$$

Observe that, as happens for executability laws, multiple help preconditions are considered disjunctively. If the request includes also the option offering C_4 , then the action y will cause $C_1 \wedge C_4$ as effect.

Let us add some comments on agents' requests for action execution. Each agent wishes to execute some actions and to ask for some requests. After the supervisor has decided which actions will be executed, each agent retrieves the relevant requests and analyzes them in order to possibly fulfill them in the next time step (see below for further details). These requests behave like an action y , as stated above.

Two global constraints are allowed by the language \mathcal{B}^{AAC} . Their effect is to filter out sequences of states that do not fulfill those constraints:

- C holds_at i imposes that any valid trajectory $\langle v_0, X_1, v_1, \dots, X_N, v_N \rangle$ must satisfy $\langle v_0, v_1, \dots, v_N \rangle \models_i C$
- always C imposes that any valid trajectory $\langle v_0, X_1, v_1, \dots, X_N, v_N \rangle$ must satisfy $\langle v_0, v_1, \dots, v_N \rangle \models_i C$ for all $i \in \{0, \dots, N\}$.

The supervisor is in charge of checking if these constraints can be satisfied while selecting X_i as mentioned before. If the fluents involved in the constraints are all known to an agent a , the set of actions proposed by a are such that they will guarantee the property if all of them (and only them) are selected for application.

Each agent a , at each time step i , selects a set of actions Y_{i+1}^a it wishes to execute. For doing that, a looks for a sequence of (sets of) actions to achieve its local goal, given the current state sequence $\langle v_0, \dots, v_i \rangle$. The set of actions Y_{i+1}^a are those to be executed at the current time step. If the new state v_{i+1} communicated by the supervisor is different from the state it expected after the application of all the actions in the set Y_{i+1}^a (due either to the fact that some of these actions are not selected, or that other agents have executed actions that have unexpectedly changed some values), it will need to replan. Let us observe that, although globally the supervisor views a valid trajectory, locally this is not true (some state transitions are not justified by the actions of agent a alone). However, in looking for a plan (and in replanning), it reasons on an "internal" valid trajectory from the current time to the future.

Let us focus on the problem of reacting to requests. Suppose that an agent a' , at time i in a state sequence $\langle v_0, v_1, \dots, v_i \rangle$, receives the requests r_1, \dots, r_h , where r_j is of the form

$$\text{request } C_1^j \text{ to_agent } a' \text{ if } C_2^j$$

² We hypothetically assume that a' has access to all fluents of a .

and, moreover, assume that these requests are ordered (e.g., by the priorities of the requesting agent a_j). For $j = 1, \dots, h$, if $\mathcal{D}_{a'}$ contains an axiom

$$\text{help} \cdots a_j \cdots \text{ if } C_3^j$$

such that $\langle v_0, v_1, \dots, v_i \rangle \models_i C_3^j$, the agent a' adds temporarily to its theory the constraint

$$C_1^j \text{ holds_at } i + 1 \tag{12}$$

and looks for a plan in the enlarged theory. If such a plan exists, the constraint (12) is definitely stored in $\mathcal{D}_{a'}$, otherwise the request is ignored. In both cases, a' proceeds with next request ($j := j + 1$). At the end, some (possibly none) of the h constraints C_1^j, \dots, C_h^j will be fulfilled by a plan and the set of actions $Y_{a'}^{i+1}$ of the next step of this plan are passed to the supervisor.

Let us focus now on how the agent a deals with the options related to a failure (this is also developed in Section 3.4). Let us assume an action x submitted for execution at time i has not been selected by the supervisor, and, therefore, a failure signal is returned to the agent a . The current sequence of states is $\vec{v} = \langle v_0, v_1, \dots, v_{i+1} \rangle$.

Let us analyze what happens in the three options:

- **fail if C_1** : if $\vec{v} \models_{i+1} C_1$, then agent a declares its failure. From this point onwards, the agent will not generate any actions, nor interact with other agents.
- **replan if C_1 add_goal C_2** : if $\vec{v} \models_{i+1} C_1$, then goal C_2 is added in \mathcal{D}_a (and then the agent a starts replanning)
- **retry_after T if C_1** : if $\vec{v} \models_{i+1} C_1$, then for $T - 1$ time steps the agent a requires only nop to the supervisor, at time step $T + i$ the action x is required again.

If the **if** option is missing, the condition will be assumed to be satisfied. If the **add_goal** option is missing, no new goal will be added.

3.2 Concurrent plan execution

The agents are autonomous and develop their activities independently, except for the execution of the actions/plans. In executing their plans, the agents must take into account the effects of concurrent actions.

We developed the basic communication mechanism among agents by exploiting a *tuple space*, whose access and manipulation follows the blackboard principles introduced in the Linda model (Carriero and Gelernter 1989). Linda is a popular model for coordination and communication among processes; Linda offers coordination via a shared memory, commonly referred to as a *blackboard* or *tuple-space*. All the information are stored in the blackboard in the form of *tuples*—the shared blackboard provides atomic access and associative memory behavior (in retrieving and removing tuples). The SICStus Prolog implementation of Linda allows the definition of a server process, in charge of managing the blackboard, and client processes, that can add tuples (using the `out` operation), read tuples (using the `rd` operation), and remove tuples (using the `in` operation).

Most of the interactions among concurrent agents, especially those interactions aimed at resolving conflicts, are managed by a specific process, the *supervisor*, that also provides a global time to all agents, enabling them to execute their actions synchronously. The supervisor process stores the initial state and the changes caused by the successful executions of actions. It synchronizes the actions execution, and controls the coordination and the arbitration in case of conflicts. It also sends a success or a failure signal to each agent at each action execution attempt, together with the list of changes to its local state.

Let us describe how the execution of concurrent plans proceeds. As mentioned, each action description includes a set of constraints describing a portion of the initial state.

- (1) At the beginning, the supervisor acquires the specification $\mathcal{I} = \bigcup_{a \in \mathcal{G}} \mathcal{I}_a$ of the initial state.
- (2) At each time step, the supervisor starts a new state transition:
 - Each agent sends to the supervisor a request to perform an action—i.e., the next action of its locally computed plan—by specifying its effects on the (local) state.
 - The supervisor collects all these requests and starts an analysis, aimed at determining the subsets of actions/agents that conflict (if any). A conflict occurs whenever agents require incompatible assignments of values to the same fluents. The transition takes place once all conflicts have been resolved and a subset of compatible actions has been identified by means of some policy (see below). These actions are enabled while the remaining ones are inhibited.
 - All the enabled actions are executed, producing changes to the global state.
 - These changes are then sent back to all agents, to achieve the corresponding updates of each agent's local state. All agents are also notified about the outcome of the procedure. In particular, those agents whose actions have been inhibited receive a failure message.
- (3) The computation stops when the time \mathbf{N} is reached.

Observe that, after each step of the local plan execution, each agent needs to check if the reached state still supports its successive planned actions. If not, the agent has to reason locally and revise its plan, i.e., initiate a replanning phase. This is due to the fact that the reached state might be different from the expected one. This may occur in two cases:

- (1) The proposed action was inhibited, so the agent actually executed a nop; this case occurs when the agent receives a failure message from the supervisor.
- (2) The interaction was successful, i.e., the planned action was executed, but the effects of the actions performed by other agents affected fluents in its local state, preventing the successful continuation of the remaining part of the local plan. For instance, the agent a may have assumed that the fluent g maintained its value by inertia, but another agent, say b , changed such value. There is no direct conflict between the actions of a and b , but agent a has to verify that the rest

of its plan is still applicable (e.g., the next action in a 's plan may have lost its executability condition).

3.3 Conflict resolution

A conflict resolution procedure is invoked by the supervisor whenever it determines the presence of a set of incompatible actions. Different policies can be adopted in this phase and different roles can be played by the supervisor.

First of all, the supervisor exploits the priorities of the agents to attempt a resolution of the conflict, by inhibiting the actions issued by low-priority agents. If this does not suffice, further options are applied. We describe here some of the easiest viable possibilities, that we have already implemented in our prototype. The architecture of the system is modular (see Section 3.6), and can be easily extended to include more complex policies and protocols.

The two approaches we implemented so far differ by assigning the active role in resolving the conflict either (a) to the supervisor or (b) to the conflicting agents. In the *first case*, the supervisor has an *active role*—it acts as a referee and decides, without any further interaction with the agents, which actions have to be inhibited. In the current prototype, the arbitration strategy is limited to:

- A random selection of a single action to be executed; or
- The computation of a maximal set of compatible actions to be executed. This computation is done by solving a CSP—which is dynamically generated using a CLP($\mathcal{F}\mathcal{D}$) encoding.

Note that, in this strategy, the `on_conflict` policies assigned to actions by axioms (7) are ignored. This “centralized” approach is relatively simple; it has also strong potential of facilitating the creation of optimal plans. On the other hand, the adoption of a centralized approach to conflict resolution might become a bottleneck in the system, since all conflicting agents must wait for supervisor's decisions.

In the *second case*, the supervisor simply notifies the set of conflicting agents about the inconsistency of their actions. The agents involved in the conflict are completely in charge of resolving it by means of a negotiation phase. The supervisor waits for a solution from the agents. In solving the conflict, each agent a makes use of one of the `on_conflict` directives (7) specified for its conflicting action x . The semantics of these directives are as follows (in all the cases [provided C] is an optional qualifier; if it is omitted it is interpreted as `provided true`):

- The option `on_conflict forego provided C` causes the agent a to “search” among the other conflicting agents for someone, say b , that can guarantee the condition C . In this case, b performs its action while the execution of a 's action fails, and a executes a `nop` in place of its action x . Different strategies can be implemented in order to perform such a “search for help.” A simple one is the round-robin policy described below, but many other alternatives are possible and should be considered in completing the prototype.

- The option `on_conflict retry_after T` provided C , will cause a to execute `nop` during the following T time steps and then it will try again to execute its action (if the preconditions still hold).
- If there is no applicable option (e.g., no option is defined or none of the agents accept to guarantee C), the action is inhibited and its execution fails.

The way in which agents negotiate and exploit the `on_conflict` options can rely on several protocols, of different complexity. For instance, one possibility might be to nominate a “leader” within each of the conflicting sets S of agents. The leader is in charge of coordinating the agents in S to resolve the conflict without interacting with the supervisor.

Another approach consists of letting each agent in S free to proceed and to find an agreement by sending proposals to other agents (possibly by adopting some order of execution, some priorities, etc.) and receiving their proposals/answers. In the current prototype, we implemented a round-robin policy. Let us assume that the state sequence already constructed is $\vec{v} = \langle v_0, \dots, v_i \rangle$ and let us assume that the agents in the list $A = \langle a_1, \dots, a_m \rangle$ aim at executing the set of actions $Y = \langle y_1, \dots, y_m \rangle$, respectively. Furthermore, let us assume that the execution of all actions in Y will introduce a constraint that does not have a \vec{v} -solution. There is a sorting of the agents, and they take turn in resolving the conflict. Suppose that at a certain round of the procedure the agent a_k is selected. a_k tries its next unexplored `on_conflict` OP provided C option for its action and checks if $\vec{v} \models_i C$.

- If $\vec{v} \models_i C$, then a_k will apply the OP option and a_k and y_k are removed from A and Y , respectively.
- Otherwise, the next agent is selected and the successive call to a_k will consider the next `on_conflict` option.

If there are no successive options for a_k , then a_k, y_k will be removed from A, Y and a failure for a_k will occur. After each step, if Y has a \vec{v} -solution, then the procedure will terminate and the actions in Y will be executed. Observe that this procedure always terminates with a solution to the conflict, since a finite number of `on_conflict` options are defined for each action.

This is a relatively rigid policy, and it represents a simple example of how to realize a terminating protocol for conflict resolution. Alternative solutions can be added to the prototype thanks to its modularity.

Once all conflicts have been addressed, the supervisor applies the enabled actions, and obtains the new global state. Each agent receives a communication containing the outcome of its action execution and the changes to its local state. Moreover, further information might be sent to the participating agents, depending on the outcome of the coordination procedure. For instance, when two agents agree on an `on_conflict` option, they “promise” to execute specific actions (e.g., the fact that one agent has to execute T consecutive `nop`).

3.4 Failure policies

Agents receive a failure message from the supervisor whenever their requested actions have been inhibited. In such a case, the original plan of the agent has to be revised to detect if the local goal can still be reached, possibly by replanning. Also, in this case, different approaches can be applied. For instance, one agent could avoid developing an entire plan at each step, but limit itself to produce a partial plan for the very next step. Alternatively, an agent could attempt to determine the “minimal” modifications to the existing plan in order to make it valid with respect to the new encountered state.³

In this replanning phase, the agent can exploit the `on_failure` options associated with the corresponding inhibited action. The intuitive semantics of these options can be described as follows:

- `retry_after T [if C]`: the agent first evaluates the constraint C ; if C holds, then it executes the action `nop T` times and then tries again the failed action (provided that its executability conditions still hold).
- `replan [if C1] [add_goal C2]`: the agent first evaluates C_1 ; if it holds, then in the following replanning phase the goal C_2 is added to the current local goal. The option `add_goal C2` is optional; if it is not present, then nothing is added to the goal, i.e., it is the same as `add_goal true`.
- `fail [if C1]`: this is analogous to `replan [if C1] add_goal false`. In this case, the agent declares that it is impossible to reach its goal. It quits and does not participate in the subsequent steps of the concurrent plan execution.
- If none of the above options is applicable, then the agent will proceed as if the option `replan if true` is present.

All the options declared for the inhibited action are considered in the given order, executing the first applicable one.

It might be the case that some global constraints (such as `holds_at` and `always`, cf., Section 2) involve fluents that are not known by any of the agents. Therefore, none of the agents can consider such constraints while planning. Consequently, these constraints have to be enforced while merging the individual plans. In doing this, the supervisor adopts the same strategies introduced to deal with conflicts and failures among actions, as described earlier. Namely, whenever a global constraint would be violated by the concurrent execution of actions (taken from different agents' plans) a conflict is generated and a conflict resolution procedure executed. Thus, some of the conflicting actions will be inhibited causing their failure.

3.5 Broadcasting and direct requests

Let us describe a simple protocol for implementing the point-to-point and broadcast communications among agents, following an explicit request of the form (10). In particular, let us assume that the current state is the i th one of the plan execution—hence, the supervisor is coordinating the transition to the $(i+1)$ th state by executing

³ At this time, the prototype includes only replanning from scratch at each step.

the $(i+1)$ th action of each local plan. The handling of requests is interleaved with the agent–supervisor interactions that realize plan execution; nevertheless, the supervisor does not intervene on the requests, and the requests and offers are directly exchanged among agents. We can sketch the main steps involved in a state transition, from the point of view of an agent a , as follows:

- (1) Agent a tries to execute its action and sends this information to the supervisor (Section 3.2).
- (2) Possibly after a coordination phase, a receives from the supervisor the outcome of its attempt to execute the action (failure or success, the changes in the state, etc.)
- (3) If the action execution is successful, before declaring the current transition completed, the agent a starts an interaction with the other agents to handle pending requests. All the communications associated to such interactions are realized using Linda's tuple-space (requests and offers are posted and retrieved by agents).
 - (3.a) Agent a fetches the collection H of all the requests still pending and generated until step i . For each request of help $h \in H$, originating from some agent b , agent a decides whether to accept h or not. Such a decision might involve planning activities, in order to determine if the requested condition can be achieved by a , possibly by modifying its original plan. In the positive case, a posts its offer into the tuple-space and waits for a rendez-vous with b .
 - (3.b) Agent a checks whether there are replies to the requests it previously posted. For each request for which replies are available, a collects the set of offers/agents that expressed their willingness to help a . By using some strategy, a selects one of the responding agents, say b . The policy for choosing the responding agent can be programmed (e.g., by exploiting priorities, agent's knowledge on other agents, random selection, trust criteria, and utility and optimality considerations). Once the choice has been made, a establishes a rendez-vous with the selected agent and
 - declares its availability to b ,
 - communicates the fulfillment of the request to the other agents.

The request and the obsolete offers are removed from the tuple space.

- (4) At that point in time, the transition can be considered completed for the agent a . By taking into account the information about the outcome of the coordination phase in solving conflicts (point (2)) and the agreement reached in handling requests (point (3)), a might need to modify its plan. If the replanning phase succeeds, then a will proceed with the execution of the next action in its local plan.

Note that we have provided separate descriptions for steps (3.a) and (3.b). In a concrete implementation, these two steps have to be executed in an interleaved manner, to avoid that a fixed order in sending requests and offers causes deadlocks or starvation. Furthermore, if an agent fails in executing an action, then it will skip step (3) and proceed with step (4) in order to re-plan its activity.

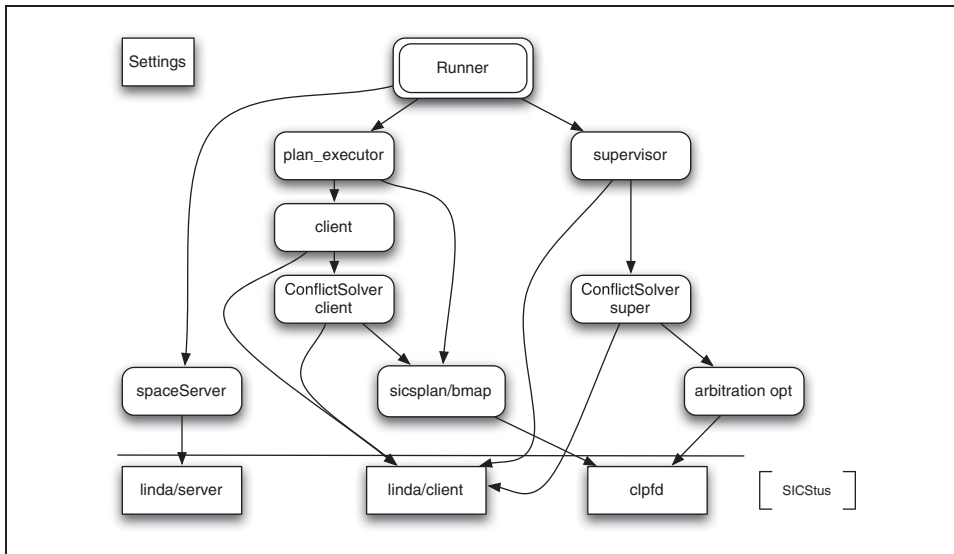


Fig. 2. The dependencies between modules in the system. The modules' names recall the corresponding Prolog-files names. The module `runner` is the starter of the application. The module `settings` specify user options (policies, strategies, etc.) and the sources files containing the action descriptions, since it is imported by all the others (we omitted drawing the corresponding arcs, as well as the nodes relative to less relevant SICStus libraries).

3.6 Implementation issues

A first prototype of the system has been implemented in SICStus Prolog, using the libraries `clpfd` for agents reasoning (by exploiting the interpreters for Action Description Languages described in Dovier *et al.* (2009, 2010)), and the libraries `system`, `linda/server`, and `linda/client` for handling process communication.

The system is structured in modules. Figure 2 displays the modules composing the Prolog prototype and their dependencies. The modules `spaceServer` (via `lindaServer`) and `lindaClient` implement the interfaces with the Linda tuple-space. These modules support all the communications among agents.

Each autonomous agent corresponds to an instance of the module `plan_executor`, which, in turn, relies on a planner (the module `sicsplan/bmap` in Fig. 2) for planning/replanning activities, and on `client` for interacting with other agents in the system. As explained previously, a large part of the coordination is guided by the module `supervisor`. Notice that both the `supervisor` and `client` act as Linda-clients. Conflict resolution functionalities are provided to the modules `client` and `supervisor` by the modules `ConflictSolver_client` and `ConflictSolver_super`, respectively. Finally, the `arbitration_opt` module implements the arbitration protocol(s). In the current code distribution, we provide an arbitration strategy that maximizes the number of actions performed at each step.

Let us remark that all the policies exploited in coordination, arbitration, and conflict handling can be customized by simply providing a different implementation of individual predicates exported by the corresponding modules. For instance, to

implement a conflict resolution strategy different from the round-robin described earlier, it suffices to add to the system a new implementation of the module `ConflictSolver_super` (and for `ConflictSolver_client`, if the specific strategy requires an active role of the conflicting agents). Similar extensions can be done for `arbitration_opt`.

The system execution is rooted in the server process runner—written either for Linux (`.sh`) or for Windows (`.bat`) platforms, in charge of generating the connection address that must be used by the client processes.

The file `settings.pl` describes the planning problem to be solved. In particular, the user must specify in this file, through Prolog facts, the number and the names of these files containing the action descriptions, a bound on the maximum length of the plan, and the selected strategies for conflict resolution and arbitration (default choices can be used).

As far as the reasoning/planning module is concerned, we slightly modified the interpreters of the B^{MV} and the B^{MAP} languages (Dovier *et al.* 2009, 2010) to accept the extended syntax presented here. However, the system is open to further extensions and different planners (even not necessarily based on Prolog technology) can be easily integrated thanks to the simple interface with the module `plan_executor`, which consists of a few Prolog predicates.

Currently, two planners have been integrated in the system: `sicsplan` is the constraint logic programming planner for the single-agent action language B^{MV} ; `bmap` is instead a constraint logic programming engine that supports centralized planning for multi-agent systems (capable, e.g., of collaborating in pursuing a common goal). Thus, the implementation allows each individual agent (according to the discussion from the previous sections) to be itself a complex system composed of multiple agents (operating in a cooperative fashion and planning in a centralized manner).

To accommodate for this perspective, the design of the supervisor has been modified. The framework allows each concurrent planner that executes a multiple-action step, to specify the desired granularity of the conflict resolution phase. This is done by specifying (for each step in a plan) a partition of the set of actions composing the step into those subsets of actions that have to be considered independently and as a whole.

For instance, in the next section, we describe a specification of a coordination problem between two multi-agent systems. Each multi-agent system develops a plan in a centralized manner. Each step of such plans consists of a set of, possibly complex, actions (instead of a single action, as happens for the planner `sicsplan`). The conflicts between the multi-agent plans occurring during the (i)th state transition are identified/resolved by considering a single action of each (i)th step proposed by each planner.

Let us make some considerations about the soundness of the implementation. Let us consider one step $i + 1$ in the construction of the trajectory. The state sequence already constructed is $\vec{v} = \langle v_0, \dots, v_i \rangle$. The agents propose some actions for execution; the overall set of all actions proposed by all agents is $Y_{i+1} = \{y_1, \dots, y_k\}$. Agents propose for execution actions that are executable in \vec{v} . At the implementation level,

the soundness property is guaranteed by the correctness of the `sicsplan/bmap` module—see Appendix (available online).

Let us denote with $C(y_j)$ the constraint that captures the effects of action y_j ; i.e., if the action y_j has dynamic causal laws y_j causes E_r if P_r for $r = 1, \dots, m$, then

$$C(y_j) \equiv \bigwedge_{r=1}^m P_r \rightarrow E_r.$$

Let $A(y_j)$ be a Boolean variable, intuitively denoting whether the supervisor has selected action y_j for execution at time $i + 1$.

The `arbitration_opt` implements an arbitration protocol $\Phi(\vec{v}, Y_{i+1})$ producing a substitution for $\{A(y_1), \dots, A(y_k)\}$ such that the constraint

$$\bigwedge_{j=1}^k \Phi(\vec{v}, Y_{i+1})(A(y_j)) \rightarrow C(y_j)$$

has a \vec{v} -solution σ .

For example, in the current code distribution, the protocol Φ is defined as a substitution that maximizes $\sum_{j=1}^k A(y_j)$.

From these definitions and from the properties of `sicsplan/bmap`, we have that $\langle \vec{v}, \{y_j \mid j \in \{1, \dots, k\}, \Phi(\vec{v}, Y_{i+1})(y_j) = 1\}, \text{ine}(\sigma, \vec{v}) \rangle$ is a valid state transition.

If the conflict resolution is left to the agents, then the protocol Φ is the outcome of the conflict resolution procedure, e.g., the round-robin analysis of the conflicting actions described in Section 3.3, which is currently implemented. It is immediately necessary to check that the round-robin procedure produces a protocol Φ that satisfies the properties shown above.

Due to the generality of the language for agent-based on-conflict resolution, the correctness of any conflict resolution procedure must be independently proved. Correctness is not an immediate consequence of the language itself but is dependent on the specific on-conflict declaration used in the specific procedure.

3.7 The volleyball domain

Let us describe a specification in \mathcal{B}^{AAC} of a coordination problem between two multi-agent systems—an extension of the domains described in Examples 1–4. There are two teams: *black* and *white* whose objective is to score a point, i.e., to throw the ball in the field of the other team (passing over the net) in such a way that no player of the other team can reach the ball before it touches the ground. Each team is modeled as a multi-agent system that elaborates its own plan in a centralized manner (thus, each step in the plan consists of a set of actions).

The playing field is discretized by fixing a `linex` \times `liney` rectangular grid that determines the positions where the players (and the ball) can move (see Fig. 3). The leftmost (rightmost) cells are those of the black (white) team, while the net ($x = 6$) separates the two subfields. There are p players per team ($p = 2$ in Fig. 3)—concretely, the fact `num(2)` is added to the theory. The allowable actions are: `move(d)`, `throw(d, f)`, and `whistle`. During the defense time, the players can move to catch

Time 0: ***** ***** * * * Y 0* * * * 0 * *X * ***** *****	Time 1: ***** ***** * * * Yo 0* * * * 0 * *Y * ***** *****	Time 2: ***** ***** * * * X 0 * * * * Y 0 * * * ***** *****	Time 3: ***** ***** * * * Y 0 * * * * Y 0 * * o * ***** *****	Time 4: ***** ***** * * * Y 0 * * * *Y * * Q * ***** *****
Time 5: ***** ***** * * * Y 0 * * * *Y * * o 0 * ***** *****	Time 6: ***** ***** * * * Y 0 * * * * 0 * * X * ***** *****	Time 7: ***** ***** * * * Y 0 * * * * 0 * * Y o * ***** *****	Time 8: ***** ***** * * *Y 0 * * * * 0 * *Y o * ***** *****	Time 9: ***** ***** * * *Y 0 * * * * 0 * *Y o * ***** *****

Fig. 3. A representation of an execution of the volleyball domain.

the ball and/or to re-position themselves on the court. When a player reaches the ball (s)he will have the ball and will throw the ball again. A team scores a point either if it throws the ball to a cell in the opposite subfield that is not reached by any player of the other team in the defense time, or if the opposite team throws the ball in the net. The captain (first player) of each team is in charge of checking if a point has been scored. In this case, (s)he whistles.

Each team (either **black** or **white**) is modeled as a centralized multi-agent system, which acts as a single agent in the interaction with the other team. Alternative options in modeling are also possible—for instance, one could model each single player as an independent agent that develops its own plan and interacts with all other players. The two teams have the goal of scoring a point: $goal(point(black) \text{ eq } 1)$. for blacks and $goal(point(white) \text{ eq } 1)$. for whites.

At the beginning of the execution, every team has a winning strategy, developed as a local plan; these are possibly revised after each play to accommodate for the new state of the world reached. An execution (as printed by the system) is reported in Figure 3, for a plan length of 9. The symbol 0 (respectively, Y) denotes the white (respectively, black) players, Q (resp. X) denotes a white player with the ball. The throw moves applied are:

```
[player(black,1)]:throw(ne,3)(time 1)[player(black,2)]:throw(se,3)(time 3)
[player(white,1)]:throw(w,5) (time 5)[player(black,1)]:throw(e,5) (time 7)
```

Let us observe that, although it would be, in principle, possible for the white team to reach the ball and throw it within the time allowed, it would be impossible to score a point. Therefore, players prefer to avoid to perform any move.

The complete description of the encoding of this domain is available at <http://www.dimi.uniud.it/dovier/BAAC>. The repository includes also additional domains—e.g., a domain inspired by games involving one ball and two-goals, as found in soccer. Although the encoding might seem similar to that of volleyball, the possibility of contact between two players makes this encoding more complex. Indeed, thanks to the fact that the net separates the two teams, in the volleyball domain rules like the following one suffice to avoid collisions:

```
always(pair(x(A),y(A)) neq pair(x(B),y(B))) :-
    A=player(black,N),B=player(black,M), num(N), num(M), N<M.
```

In a soccer world, this is not true because only the supervisor can be aware, in advance, of possible contacts between different team players originating from concurrent actions. This generates interesting concurrency problems, e.g., concerning the ball possession after a contact. A simple way to address this problem consists in assigning a fluent to each field cell, whose value can be -1 (free), 0 (resp., 1) if a white (resp. black) player is in the cell. The supervisor identifies a conflict when two opponent players move to the same cell, thus assigning to that fluent a different value. In this case, the supervisor arbitrarily enables one action, the other agent waits a turn to retry the action:

```
action act([A],move(D)) on_failure retry_after 1 on_conflict arbitrate :-
    agent(A), direction(D).
```

4 Conclusions and future work

In this paper, we illustrated the design of a high-level action description language for the description of multi-agent domains. The language enables the description of agents with individual goals operating in a shared environment. The agents can explicitly interact (by requesting help from other agents in achieving their own goals) and implicitly cooperate in resolving conflicts that may arise during execution of their individual plans. The main features of the framework we described in this paper have been realized into an implementation, based on SICStus Prolog. The implementation is fully distributed, and uses Linda to enable communication among agents. Such a prototype is currently being refined and extended with further features.

There have been many agent programming languages such as the BDI agent programming AgentSpeak (Rao 1996) (as implemented in Jason (Bordini *et al.* 2007)), JADE (Bellifemine *et al.* 2007) (and its extension Jadex (Pokahr *et al.* 2005)), ConGolog (De Giacomo *et al.* 2000), IMPACT (Subrahmanian *et al.* 2000), 3APL (Dastani *et al.* 2003), and GOAL (de Boer *et al.* 2005). A good comparison of many of these languages can be found in Mascardi *et al.* (2004). The emphasis of the effort presented in this paper is to expand our original work on constraint-based modeling of agents based on action languages. The generalization to a constraint-based multi-agent action language has been presented in Dovier *et al.* (2009). In this paper, we demonstrate a further extension to encompass distributed reasoning and distributed planning. Thus, the focus of the proposal remains on the level of creating an action language and demonstrating the suitability of constraint-based technology to support it. As such, we do not propose here a new agent programming language, rather we push an action language perspective and how action languages scale to multi-agent domains; our work could be used as the underlying formalism for the development of new agent programming languages. In this sense, our proposal is different than many of the MAS development platforms, which focus on programming languages for MAS and on complex protocols for advertising and interaction among agents (e.g., FIPA).

The choice of Linda came about for simplicity; we required the use of a CLP platform and SICStus provides support for both Linda and constraint handling—as a few other distributed communication platforms (e.g., OAA (Cheyer and Martin

2001)). In the long term, we envision mapping our agent design on an MAS infrastructure that enables discovery and addition of agents, handles network-wide distribution of agents, and maps the exchange of constraints to a standard agent communication language (e.g., FIPA-ACL/FIPA-SL (Hayzelden and Bourne 2001)). This will require a non-trivial engineering work, to map the reasoning with action languages (e.g., planning) to a platform that is not constraint-based—we are currently exploring the problem in the context of Jason (Bordini *et al.* 2007).

The work is an initial proposal that already shows strong potential and several avenues of research. The immediate goal in the improvement of the system consists of adding refined strategies and coordination mechanisms, involving, for instance, payoff, trust, etc. Then, we intend to evaluate the performance and quality of the system in several multi-agent domains (e.g., game playing scenarios, modeling of auctions, and other domains requiring distributed planning). We also plan to investigate strategies to enhance performance by exploiting features provided by the constraint solving libraries of SICStus (e.g., the use of the table constraint (Barták and Toropila 2008)).

We will investigate the use of future references in the fluent constraints (as fully supported in B^{MV})—we believe this feature may provide a more elegant approach to handle the requests among agents, and it is necessary to enable the expression of complex interactions among agents (e.g., to model forms of negotiation with temporal references). In particular, we view this platform as ideal to experiment with models of *negotiation* (e.g., as discussed in (Son *et al.* 2009)) and to deal with *commitments* (Mallya and Huhns 2003) (which often require temporal references).

We will also explore the implementation of different strategies associated to conflict resolution; in particular, we are interested in investigating how to capture the notion of “trust” among agents, as a dynamic property that changes depending on how reliable agents have been in providing services to other agents (e.g., accepting to provide a property but failing to make it happen). Also concerning trust evaluation, different approaches can be integrated in the system. For instance, a “controlling entity” (e.g., either the supervisor or a privileged/elected agent) could be in charge of assigning the “degree of trust” of each agent. Alternatively, each single agent could develop its own opinion on other agents’ reliability, depending on the behavior they manifested in past interactions.

Finally, work is needed to expand the framework to enable greater flexibility in several aspects, such as:

- Allow deadlines for requests—e.g., by allowing axioms of the form

$$\text{request } C_1 \text{ if } C_2 \text{ until } T$$
 indicating that the request is valid only if accomplished within T time steps.
- Allow constraint-based delays for requests:

$$\text{request } C_1 \text{ if } C_2 \text{ while } C_3$$
 indicating that the request is still valid while constraint C_3 is entailed.
- Allow dynamic changes in the agents’ knowledge about other agents (e.g., an action might make an agent aware of the existence of other agents), or about the world (e.g., an action might change the rights another agent has to access/modify some fluents).

Acknowledgement

The authors wish to thank the anonymous reviewers for their insightful comments.

References

- BARTÁK, R. AND TOROPILA, D. 2008. Reformulating constraint models for classical planning. In *Proc. of the International Florida AI Research Society Conference*, D. Wilson, Ed. AAAI Press, Menlo Park, CA, 525–530.
- BELLIFEMINE, F., CAIRE, G. AND GREENWOOD, D. 2007. *Developing Multi-Agent Systems with JADE*. John Wiley, Chichester, UK.
- BORDINI, R., HÜBNER, J. AND WOOLDRIDGE, M. 2007. *Programming Multi-Agent Systems in AgentSpeak Using Jason*. John Wiley, Chichester, UK.
- CARRIERO, N. AND GELERNTER, D. 1989. Coordination languages and their significance. *Communications of the ACM* 32, 4, 444–458.
- CHEYER, A. AND MARTIN, D. 2001. The open agent architecture. *Journal of Autonomous Agents and Multi-Agent Systems* 4, 1, 143–148.
- DASTANI, M., DIGNUM, F. AND MEYER, J.-J. 2003. 3APL: A programming language for cognitive agents. *ERCIM News* 53, 28–29.
- DE BOER, F., HINDRIKS, K., VAN DER HOEK, W. AND MEYER, J. 2005. A verification framework for agent programming with declarative goals. *Journal of Applied Logic*, 5, 277–302.
- DE GIACOMO, G., LESPÉRANCE, Y. AND LEVESQUE, H. 2000. ConGolog, a concurrent programming language based on the situation calculus. *Artificial Intelligence*, 121, 1–2, 109–169.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2009. Representing multi-agent planning in CLP. In *Proc. LPNMR*, Lecture Notes in Computer Science, E. Erdem, F. Lin and T. Schaub, Eds. vol. 5753. Springer, Berlin, 423–429.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2010. Multivalued action languages with constraints in CLP(FD). *Theory and Practice of Logic Programming* 10, 2, 167–235.
- FAGIN, R., HALPERN, J. Y., MOSES, Y. AND VARDI, M. Y. 1995. *Reasoning About Knowledge*. The MIT Press, Cambridge, MA, USA.
- GELFOND, M. AND LIFSCHITZ, V. 1998. Action languages. *Electronic Transactions on Artificial Intelligence* 2, 193–210.
- GERBRANDY, J. 2006. Logics of propositional control. In Nakashima et al. (2006), 193–200.
- HAYZELDEN, A. AND BOURNE, R. 2001. *Agent Technology for Communication Infrastructures*. John Wiley, Chichester, UK.
- MALLYA, A. AND HUHS, M. 2003. Commitments among agents. *IEEE Internet Computing* 7, 4, 90–93.
- MASCARDI, V., MARTELLI, M. AND STERLING, L. 2004. Logic-based specification languages for intelligent agents. *Theory and Practice of Logic Programming* 4, 4, 495–537.
- NAKASHIMA, H., WELLMAN, M. P., WEISS, G. AND STONE, P., Eds. 2006. In *Proc. of the International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, New York.
- POKAHR, A., BRAUBACH, L. AND LAMERSDORF, W. 2005. Jadex: A BDI reasoning engine. In *Multi-Agent Programming: Languages, Platforms and Applications*. R. H. Bordini, M. Dastani, J. Dix and A. El Fallah-Seghrouchni, Eds. Multiagent Systems, Artificial Societies, and Simulated Organizations Vol. 15, Springer, Berlin, 149–174.
- RAO, A. 1996. AgentSpeak: BDI agents speak out in a logical computable language. In *Proc. of the European Workshop on Modeling Autonomous Agents in a Multi-Agent World*, W. V.

- de Velde and J. W. Perram, Eds. Eindhoven, The Netherlands, January 22–25, Lecture Notes in Computer Science, vol. 1038, Springer, Berlin, 42–55.
- SAURO, L., GERBRANDY, J., VAN DER HOEK, W. AND WOOLDRIDGE, M. 2006. Reasoning about action and cooperation. See Nakashima *et al.* (2006), 185–192.
- SON, T., PONTELLI, E. AND SAKAMA, C. 2009. Logic programming for multiagent planning with negotiation. In *Proc. of the International Conference on Logic Programming*. P. M. Hill and D. S. Warren, Eds. Springer, Berlin, 99–114.
- SPAAN, M. T. J., GORDON, G. J. AND VLASSIS, N. A. 2006. Decentralized planning under uncertainty for teams of communicating agents. In *Proc. AAMAS*. H. Nakashima *et al.*, Eds. ACM Press, New York, 249–256.
- SUBRAHMANIAN, V. S., BONATTI, P., DIX, J., EITER, T., KRAUS, S., OZCAN, F. AND ROSS, R. 2000. *Heterogeneous Agent Systems: Theory and Implementation*. MIT Press, Cambridge, MA, USA.
- VAN DER HOEK, W., JAMROGA, W. AND WOOLDRIDGE, M. 2005. A logic for strategic reasoning. In *Proc. AAMAS*, F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh and M. Wooldridge, Eds. ACM Press, New York, 157–164.