Evolving Multi-Agent Viewpoints - an architecture

Pierangelo Dell'Acqua[1], João Alexandre Leite[2], Luís Moniz Pereira[2]

[1] Department of Science and Technology, Campus Norrköping
Linköping University, Norrköping, Sweden, pier@itn.liu.se
[2] Centro de Inteligência Artificial - CENTRIA, Universidade Nova de Lisboa
2825-114 Caparica, Portugal, {jleite|lmp}@di.fct.unl.pt

Abstract. We present an approach to agents that can reason, react to the environment and are able to update their own knowledge as a result of new incoming information. Each agents' view of the social relationships among agents (itself and others) is represented by a graph, which in turn can be updated, allowing for the representation of such social evolution.

1 Introduction and motivation

Over recent years, the notion of agency has claimed a major role in defining the trends of modern research. Influencing a broad spectrum of disciplines such as Sociology, Psychology, among others, the agent paradigm virtually invaded every sub-field of Computer Science. Although commonly implemented by means of imperative languages, mainly for reasons of efficiency, the agent concept has recently increased its influence in the research and development of computational logic based systems. Since efficiency is not always the crucial issue, but clear specification and correctness is, *Logic Programming* and *Non-monotonic Reasoning* have been revived from the shade back into the spotlight. To this accrues the recent significant improvements in the efficiency of *Logic Programming* implementations for *Non-monotonic Reasoning* [5, 16, 18]. Besides allowing for a unified declarative and procedural semantics, eliminating the traditional wide gap between theory and practice, the use of several and quite powerful results in the field of non-monotonic extensions to *Logic Programming (LP)*, such as belief revision, inductive learning, argumentation, preferences, abduction, etc.[15, 17] can represent an important composite added value to the design of rational agents. These results, together with the improvement in efficiency, allow the referred mustering of *Logic Programming* and *Non-monotonic Reasoning* to accomplish a felicitious degree of combination between reactive and rational behaviours of agents, the *Holly Grail* of modern *Artificial Intelligence*, whilst preserving clear and precise specification enjoyed by declarative languages. This goal in mind, Kowalski and Sadri [10] advanced an agent architecture (KS-agents) based on an observe-think-act cycle. It was further developed by Dell'Acqua and Pereira [6] to allow the agents to dynamically update their own knowledge base (whether intentional or extensional) as well as their own goals. The updates richness provided to these agents is inherited from *Dynamic Logic Programming* [1, 2].
Dynamic Logic Programming (DLP) was introduced, following the eschewing of performing updates on a model basis, to envisage updates as applying to logic programming rules making up a theory [13, 14]. According to *DLP*, knowledge is conveyed by a set of theories (encoded as generalized logic programs) representing different states of the world. Different states may represent distinct

dimensions such as different time periods, different hierarchical instances, or even different domains. Consequently, the individual theories may contain mutually contradictory and overlapping information. The role of *DLP* is to take into account the mutual relationships extant between different states to precisely determine the declarative and the procedural semantics of the combined theory comprised of all individual theories and the way they relate.

Although *DLP* can represent several states in one evolving dimension or aspect of a system, no more than one such aspectual evolution can be encoded and combined simultaneously. This is so because *DLP* is defined only for linear sequences of states. In [15], the states represent different time instants. To overcome this drawback, Leite et al. introduced *Multi-dimensional Dynamic Logic Programming (MDLP)* [11], which generalizes *DLP* to allow for collections of states organized in arbitrary directed acyclic graphs (DAGs). Within this more general theory, one can encode simultaneously all representational dimensions, which can be particularly useful in the context of multi-agent systems.

In this paper, we formalize agents with such capabilities, generalizing the framework of [15] to allow for an arbitrary number of dimensions represented by an arbitrary DAG. We will show how this new theory is useful for an agent to represent and reason about its own and other agents' knowledge and its evolution in time. Each agent's view of the evolving social relationships among agents (itself and others) is represented by a DAG, itself in turn updatable, to capture the representation of social evolution.

The contribution of the paper is therefore twofold. On the one side, the paper presents an extension of the framework for multi-dimensional dynamic logic programming to incorporate integrity constraints and active rules, and to make the DAG of each agent updatable. On the other, the paper provides a means to incorporate the obtained framework into a multi-agent architecture. For simplicity, we have considered propositionalised generalized logic programs.

2 Preliminaries

2.1 Graphs

A **directed graph** $D = (V, E, \delta)$ is a composite notion of two sets $V = V_D$ of *vertices* and $E = E_D$ of (*directed*) *edges* and a mapping $\delta : E \to V \times V$. If $\delta(e) = (v, w)$ then v is called the *initial vertex* and w the *final vertex* of the edge e. A *directed edge sequence from* v_0 *to* v_n in a directed graph is a sequence of edges $e_1, e_2, \ldots, e_n$ such that $\delta(e_i) = (v_{i-1}, v_i)$ for $i = 1, \ldots, n$. A *directed path* is a directed edge sequence in which all the edges are distinct. A **directed acyclic graph** (DAG) is a directed graph D such that there are no directed edge sequences from v_i to v_i, for all vertices v_i of D. A **source** is a vertex with in-valency 0 (number of edges for which it is a final vertex) and a **sink** is a vertex with out-valency 0 (number of edges for which it is a final vertex). We say that $v_i < v_j$ if there is a directed path from v_i to v_j and that $v_i \leq v_j$ if $v_i < v_j$ or $i = j$.

For simplicity, we will omit the explicit representation of the mapping δ of a graph, and represent its edges $e \in E$ by their corresponding pairs of vertices (v, w) such that $(v, w) = \delta(e)$. Therefore, a graph D will be represented by the pair (V, E) where V is a set of vertices and E is a set of pairs of vertices. Next we introduce the notion of relevancy DAG.

Definition 1. *Let $D = (V, E)$ be an acyclic directed graph. Let v be a vertex in V. The* **relevancy DAG** *of D wrt. v is $D_v = (V_v, E_v)$, where:*

$$V_v = \{\ v_i \in V \mid \textit{there is a directed path from } v_i \textit{ to } v\}$$
$$E_v = \{\ (v_i, v_j) \in E \mid v_i \in V_v \textit{ and } v_j \in V_v\}$$

Intuitively the relevancy DAG of D wrt. v is the subgraph of D consisting of all vertices and edges contained in all directed paths to v.

2.2 Logic Programming framework

In order to represent negative information in logic programs, we need more general logic programs that allow default negation *not* A not only in premises of their clauses but also in their heads[1]. We call such programs generalized logic programs. It is convenient to syntactically represent generalized logic programs as propositional Horn theories. In particular, we represent default negation *not* A as a standard propositional variable (atom). Propositional variables whose names do not begin with "*not*" and do not contain the symbols ":" and "÷" are called **objective atoms**. Propositional variables of the form *not* A are called **default atoms**. Propositional variables of the form i:C (where C is defined below) are called **projects**. i:C denotes the intention (of some agent j) of updating the theory of agent i with C. Propositional variables of the form $j \div C$ are called **updates**. $j \div C$ denotes an update (proposed by j) of the current theory of agent i with C. We assume that updates cannot be negated (i.e., we disallow *not* $i{\div}C$). Instead projects can be negated. A negated project *not* i:C denotes the intention of the agent not to perform the project i:C.
Suppose that $\mathcal{K}$ is an arbitrary set of propositional variables consisting of objective atoms and projects such that $\textit{false} \notin \mathcal{K}$. By the propositional language $\mathcal{L}_{\mathcal{K}}$ generated by the set $\mathcal{K}$ we mean the language which consists of the following set of propositional variables:

$$\begin{aligned}\mathcal{L}_{\mathcal{K}} \;=\; \mathcal{K} &\cup \{\textit{not}\, A \mid \text{for every atom } A \in \mathcal{K}\} \\ &\cup \{\textit{not}\, i{:}C \mid \text{for every project } i{:}C \in \mathcal{K}\} \\ &\cup \{i \div C \mid \text{for every project } i{:}C \in \mathcal{K}\}.\end{aligned}$$

Definition 2. *A* **generalized rule** *in the language $\mathcal{L}_{\mathcal{K}}$ is a rule of the form:*

$$L_0 \leftarrow L_1 \wedge \ldots \wedge L_n \qquad (n \geq 0)$$

where every L_i $(0 \leq i \leq n)$ is an objective or default atom from $\mathcal{L}_{\mathcal{K}}$.

Definition 3. *An* **integrity constraint** *in the language $\mathcal{L}_{\mathcal{K}}$ is a rule of the form:*

$$\textit{false} \leftarrow L_1 \wedge \ldots \wedge L_n \wedge Z_1 \wedge \ldots \wedge Z_m \qquad (n \geq 0, m \geq 0)$$

where every L_i $(1 \leq i \leq n)$ is an objective or default atom, and every Z_j $(1 \leq j \leq m)$ is a project from $\mathcal{L}_{\mathcal{K}}$.

Integrity constraints are rules that enforce some condition over the state, and therefore take the form of denials.

[1] For further motivation and intuitive reading of logic programs with default negations in the heads see [1].

Definition 4. *A* **generalized logic program** P *in the language* $\mathcal{L}_{\mathcal{K}}$ *is a set of generalized rules and integrity constraints in the language* $\mathcal{L}_{\mathcal{K}}$.

Definition 5. *A* **query** Q *in the language* $\mathcal{L}_{\mathcal{K}}$ *is of the form:*

$$?- \ L_1 \wedge \ldots \wedge L_n \qquad (n \geq 1)$$

where every L_i $(1 \leq i \leq n)$ *is an objective or default atom from* $\mathcal{L}_{\mathcal{K}}$.

The following definition introduces rules that are evaluated bottom-up. To emphasize this aspect, we employ a different notation for them.

Definition 6. *An* **active rule** *in the language* $\mathcal{L}_{\mathcal{K}}$ *is a rule of the form:*

$$L_1 \wedge \ldots \wedge L_n \Rightarrow Z \qquad (n \geq 0)$$

where every L_i $(1 \leq i \leq n)$ *is an objective or default atoms, and* Z *is a project from* $\mathcal{L}_{\mathcal{K}}$.

Active rules are rules that modify the current state when executed. Active rules take the form:

$$action_name \wedge Preconditions \Rightarrow Z$$

where *action_name* is an abducible. If the *Preconditions* of the rule are satisfied, then the project (fluent) Z can be selected and executed. The head of an active rule must be a project that is either internal or external. An *internal project* operates on the state of the agent, e.g., if an agent gets an observation, then it updates its knowledge, or if some conditions are met, then it proves some goal, etc. *External projects* instead are performed on the environment, e.g., when an agent sends a message to another agent.
Given a set of vertices V, we assume that for every project $i{:}C$ in $\mathcal{K}$, C is either a generalized rule, an integrity constraint, a query, an active rule or an atom of the form $edge(u,v)$, $not\ edge(u,v)$, $modify_relation(j,u,v,x)$, or $not\ modify_edge(j,u,v,x)$, where $u,v \in V$. Thus, a project can only take one of the following form:

$$\begin{array}{ll} i{:}(L_0 \leftarrow L_1 \wedge \ldots \wedge L_n) & i{:}(L_1 \wedge \ldots \wedge L_n \Rightarrow Z) \\ i{:}(false \leftarrow L_1 \wedge \ldots \wedge L_n \wedge Z_1 \wedge \ldots \wedge Z_m) & i{:}(?{-}L_1 \wedge \ldots \wedge L_n) \\ i{:}modify_edge(j,u,v,x) & i{:}edge(u,v) \\ i{:}not\ modify_edge(j,u,v,x) & i{:}\ not\ edge(u,v) \end{array}$$

Note that the predicates *edge* and *modify_edge* can only occur inside projects or updates since those predicates do not belong to $\mathcal{L}_{\mathcal{K}}$. We assume that $i{:}edge(u,v)$ and $i{:}(?{-}L_1,\ldots,L_n)$ can only be internal projects, that is, only the agent itself can issue a project to update its own DAG (i.e., $edge(u,v)$) and goals (i.e., $?{-}L_1,\ldots,L_n$). The project $i{:}edge(u,v)$ issued by the agent i denotes the intention of i to modify its own DAG by establishing an edge between the vertices u and v in V. By issuing a project $i{:}modify_edge(j,u,v,x)$ the agent i expresses the intention to modify the DAG of another agent j by adding/deleting (depending on whether $x = a$ or $x = d$) an edge between u and v.

3 Agents and multi-agent systems

This section presents the notion of agent and multi-agent system. The initial theory of an agent is characterized by a multi-dimensional abductive logic program, inspired by [11], which expresses the agent's viewpoint on the relationships amongst a collection of agents, and encoded in a directed acyclic graph.

Definition 7. *Let $\mathcal{L}_{\mathcal{K}}$ be a propositional language and $\mathcal{M}$ a multi-agent system (defined below). A* **multi-dimensional abductive logic program** *for an agent α, Φ^{α}, is a tuple $\mathcal{T} = \langle \mathcal{P}_D, \mathcal{A}, \mathcal{R}_D, D \rangle$ where:*
- *$D = (V, E)$ is an acyclic directed graph such that $V = \{v \mid \text{for every } \Phi^v \in \mathcal{M}\} \cup \{\alpha'\}$;*
- *$\mathcal{P}_D = \{P_v \mid v \in V\}$ is a set of generalized logic programs in the language $\mathcal{L}_{\mathcal{K}}$, indexed by the vertices $v \in V$ of D;*
- *$\mathcal{A}$ is a set of object and default atoms;*
- *$\mathcal{R}_D = \{R_v \mid v \in V\}$ is a set of sets of active rules in the language $\mathcal{L}_{\mathcal{K}}$, indexed by the vertices $v \in V$ of D.*

We call the distinguished vertex $\alpha' \in V$ the **inspection point** *of agent α, and we call the atoms in $\mathcal{A}$* **abducibles**.

The initial theory of an agent α is determined (1) by a DAG D that represents the relation between the agents (represented by the vertices in V) in the multi-agent system; (2) by a set of generalized logic programs, one program P_v for each agent $v \in V$; (3) by a set of abducibles; and (4) by a set of sets of active rules R_v, for each agent $v \in V$.
Without loss of generality, we can assume that abducibles have no matching clauses in P. Abducibles can be thought of as hypotheses that can be used to extend the given logic program in order to provide an "explanation", or conditional answer, to a query. Explanations are required to satisfy the integrity constraints in P.

Definition 8. *Let $\mathcal{M}$ be a multi-agent system. An* **updating program** *U is a finite set of updates such that if an update $v \div C \in U$ then $\Phi^v \in \mathcal{M}$.*

The following definition introduces the notion of sequences of updating programs. Given the theory of an agent α (the initial theory[2] of the agent), a sequence of updating programs is defined as a set consisting of a number of updating programs containing the updates to be performed over the initial theory of α. Let $S = \{1, \ldots, m\}$ be a set of natural numbers. We call the elements $s \in S$ *states*.

Definition 9. *Let $S = \{1, \ldots, m\}$ be a set of states. An* **agent** *α* **at state** *m, $\Phi^{\alpha,m}$, is defined by a pair $(\mathcal{T}, \mathcal{U})$, where $\mathcal{T}$ is the multi-dimensional abductive logic program for the agent α, and $\mathcal{U} = \{U_s \mid s \in S\}$ is a sequence of updating programs.*

Definition 10. *A* **multi-agent system** *$\mathcal{M} = \{\Phi^{v_1,m}, \ldots, \Phi^{v_n,m}\}$ at state m is a set of agents $\Phi^{v_1}, \ldots, \Phi^{v_n}$ at state m.*

Within logic programs we refer to agents by using the corresponding superscript. For instance, if we want to update the theory of agent Φ^{v_i} with C, we write the project $v_i{:}C$.
Note that the definition of multi-agent system characterizes a static society of agents in the sense that it is not possible to add/remove agents from the system, and all the agents are at the same state. Note however that the relationships amongst the agents are dynamically changable.
The agent's semantics is given by a syntactical transformation (defined below), producing a generalized logic program for each agent, and apply to queries for it as well. Nevertheless, to increase readability, we will immediately present an

[2] Initial knowledge can be made into permanent background knowledge by including it in every update.

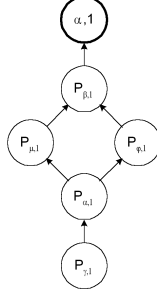

Fig. 1.

illustrative example that, though simple, allows one to glean some of the capabilities of the framework. To increase readability, we rely on natural language when explaining some details. The agent cycle, mentioned in the example and also defined in a section below, operates on the results of the transformations, but the illustrative example can be understood, on a first reading, without requiring detailed knowledge of the transformations.

3.1 Illustrative Example

Alfredo has a girlfriend and lives together with his mother and father. Being a conservative lad, Alfredo never does anything that contradicts those he considers to be higher authorities, in this case his mother, his father, and the judge who will appear later in the story. Furthermore, he considers the judge to be hierarchically superior to each of his parents. As for the relationship with his girlfriend, he hears her opinions but his own always prevail. Therefore Alfredo's view of the relationships between himself and other agents can be represented by the following DAG D, depicted in Fig.1 (where the inspection point of Alfredo, i.e. the vertex corresponding to Alfredo's overall semantics, is represented by a bold circle):

$$D = (V, E)$$
$$V = \{\alpha, \beta, \gamma, \mu, \varphi, \alpha'\}$$
$$E = \{(\gamma, \alpha), (\alpha, \mu), (\alpha, \varphi), (\mu, \beta), (\varphi, \beta), (\beta, \alpha')\}$$

where α is Alfredo, β is the judge, γ is the girlfriend, μ is the mother, φ is the father and α' is the inspection point of Alfredo.
Initially, Alfredo's programs P_{α_1} and R_{α_1} at state 1 contain the following rules:

$$r1 : girlfriend \leftarrow$$
$$r2 : move_out \Rightarrow \alpha : rent_apartment$$
$$r3 : get_married \wedge girlfriend \Rightarrow \gamma : propose$$
$$r4 : not\, happy \Rightarrow \varphi : (?{-}happy)$$
$$r5 : not\, happy \Rightarrow \mu : (?{-}happy)$$
$$r6 : modify_edge(\beta, u, v, a) \Rightarrow \alpha : edge(u, v)$$

stating that: -Alfredo has a girlfriend (r1); -if he decides to move out, he has to rent an apartment (r2); -if he decides to get married, provided he has a girlfriend, he has to propose (r3); -if he is not happy he will ask his parents how to be happy (r4, r5); -Alfredo must create a new edge between two agents (represented by u and v) in his DAG every time he proves the judge told him so (r6). Alfredo's set of abducibles, corresponding to the actions he can decide to perform to reach his goals, is:

$$A = \{move_out, get_married\} .$$

Initially, Alfredo's only goal is to be happy, this being represented at the agent cycle level (cf. section 4) together with *not false* (to preclude the possibility of violation of integrity constraints) and with the conjunction of active rules:

$$G = (? - happy \wedge not\, false \wedge r2 \wedge r3 \wedge r4 \wedge r5 \wedge r6)$$

During the first cycle (we will assume that there are always enough computational units to complete the proof procedure), the following two projects are selected:

$$\varphi : (? - happy) \qquad \mu : (? - happy)$$

Note that these correspond to the only two active rules (r4 and r5) whose premises are verified. Both projects are selected to be performed, producing 2 messages to agents φ and μ (further details on how communication can be combined into this framework the reader is referred to [7]).
In response to Alfredo's request to his parents, Alfredo's mother, as most latin mothers, tells him that if he wants to be happy he should move out, and that he should never move out without getting married. This correspond to the following observed updates:

$$\mu \div (happy \leftarrow move_out) \qquad \mu \div (false \leftarrow move_out, not\, get_married)$$

which produces the program P_{μ_2} at state 2 containing the following rules:

$$r7 : happy \leftarrow move_out \qquad r8 : false \leftarrow move_out, not\, get_married$$

Alfredo's father, on the other hand, not very happy with his own marriage and with the fact that he had never lived alone, tells Alfredo that, to be happy he should move out and live by himself, i.e. he will not be happy if he gets married now. This correspond to the following updates observed:

$$\varphi \div (happy \leftarrow move_out) \qquad \varphi \div (not\, happy \leftarrow get_married)$$

which produces the program P_{φ_2} at state 2 containing the following rules:

$$r9 : happy \leftarrow move_out \qquad r10 : not\, happy \leftarrow get_married$$

The situation after these updates is represented in Fig.2. After another cycle, the IFF proof procedure returns no projects because the goal is not reachable without producing a contradiction. From r7 and r8 one could abduce *move_out* to prove *happy* but, in order to satisfy r8, *get_married* would have to be abduced also, producing a contradiction via r10. This is indeed so because, according to the DAG, the rules of P_{φ_2} and P_{μ_2} cannot reject one another other since there is no path from one to the other.
At the beginning of another cycle, Alfredo receives from the judge the update:

$$\beta \div modify_edge(\beta, \varphi, \mu, a)$$

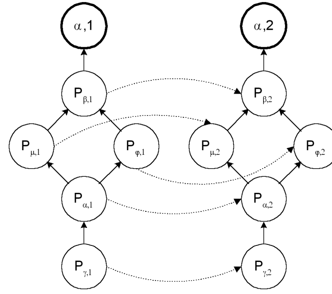

Fig. 2.

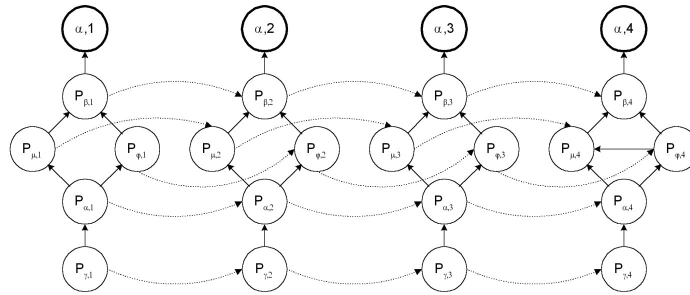

Fig. 3.

corresponding to the decision to give Alfredo's custody to his mother after his parents divorce. This produces the program P_{β_3} containing the rule:

$$r11 : modify_edge(\beta, \varphi, \mu, a) \leftarrow$$

After this update, the following three projects are selected for execution:

$$\varphi : (? - happy) \qquad \mu : (? - happy) \qquad \alpha : edge(\varphi, \mu)$$

which correspond to the evaluation of the active rules r4, r5 and r6.
In response to Alfredo's renewed request to his parents, given the precedence of opinion imposed by the judge, they nevertheless reply the same rules as before, producing the following two programs: P_{μ_4} containing the rules received from the mother:

$$r12 : happy \leftarrow move_out \qquad r13 : false \leftarrow move_out, not\, get_married$$

and P_{φ_4} containing the rules received from the father:

$$r14 : happy \leftarrow move_out \qquad r15 : not\, happy \leftarrow get_married$$

The update $\alpha \div edge(\varphi, \mu)$ produces a change in the graph, and the current situation is the one depicted in Fig.3.

After another cycle, the proof procedure returns the following two projects:

$$\alpha : rent_apartment \qquad \gamma : propose$$

These projects correspond to the evaluation of rules r2 and r3 after the abduction of $\{move_out, get_married\}$ to prove the goal $(?{-}happy)$. Note that now it is possible to achieve this goal without reaching a contradictory state since henceforth the advice from Alfredo's mother (r12 and r13) prevails over and rejects that of his father (r14 and r15). To end the story, Alfredo gets married, rents an apartment, moves out with his new wife and lives happily ever after.

3.2 Syntactical Transformation

The semantics of an agent α at state m, $\Phi^{\alpha,m}$, is established by a syntactical transformation (called *multi-dimensional dynamic program update* $\bigoplus$) that, given $\Phi^{\alpha,m} = (\mathcal{T},\mathcal{U})$, produces an *abductive logic program* [6] $\langle P, \mathcal{A}, R\rangle$, where P is a normal logic program (that is, default atoms can only occur in bodies of rules), $\mathcal{A}$ is a set of abducibles and R is a set of active rules. Default negation can then be removed from the bodies of rules in P via the *dual transformation* defined by Alferes et al [3]. The transformed program P' is a definite logic program. Consequently, a version of the IFF proof procedure proposed by Kowalski et al [9] and simplified by Dell'Acqua et al [6] to take into consideration propositional definite logic programs, can then be applied to $\langle P', \mathcal{A}, R\rangle$.
In the following, to keep notation simple we write rules as $A \leftarrow B \wedge not\, C$, rather than as $A \leftarrow B1 \wedge \ldots \wedge Bm \wedge not\, C1 \wedge \ldots \wedge not\, Cn$.
Suppose that $\langle \mathcal{P}_D, \mathcal{A}, \mathcal{R}_D, D\rangle$, with $D = (V, E)$, is the multi-dimensional abductive logic program of agent α, and $S = \{1, \ldots, m\}$ a set of states. Let the vertex $\alpha' \in V$ be the inspection point of α. We will extend the DAG D with two (source) vertices: an initial vertex sa for atoms and an initial vertex sp for projects. We will then extend D with a set of directed edges (sa_0, s') and (sp_s, s') connecting sa to all the sources s' in D at state 0, and connecting sp to all the sources s' in D at every state $s \in S$. Let $\mathcal{K}$ be an arbitrary set of propositional variables as described above, and $\overline{\mathcal{K}}$ the propositional language extending $\mathcal{K}$ to:

$$\overline{\mathcal{K}} = \mathcal{K} \cup \left\{ \begin{array}{l} A^-, A_{v_s}, A^-_{v_s}, A_{P_{v_s}}, A^-_{P_{v_s}}, reject(A_{v_s}), \\ reject(A^-_{v_s}) \mid A \in \mathcal{K},\ v \in V \cup \{sa, sp\},\ s \in S \end{array} \right\}$$
$$\cup \left\{ \begin{array}{l} rel_edge(u_s, v_s), rel_path(u_s, v_s), rel_vertex(u_s), \\ edge(u_s, v_s), edge(u_s, v_s)^- \mid u, v \in V \cup \{sa, sp\},\ s \in S \end{array} \right\}$$

Transformation of rules and updates

(RPR) Rewritten program rules: Let γ be a function defined as follows:

$$\gamma(v, s, F) = \begin{cases} A_{P_{v_s}} \leftarrow B \wedge C^- & \text{if } F \text{ is } A \leftarrow B \wedge not\, C \\ A^-_{P_{v_s}} \leftarrow B \wedge C^- & \text{if } F \text{ is } not\, A \leftarrow B \wedge not\, C \\ false_{P_{v_s}} \leftarrow B \wedge C^- \wedge & \text{if } F \text{ is } false \leftarrow B \wedge not\, C \wedge \\ \qquad \wedge Z1 \wedge Z2^- & \qquad \wedge Z1 \wedge not\, Z2 \\ B \wedge C^- \Rightarrow Z_{P_{v_s}} & \text{if } F \text{ is } B \wedge not\, C \Rightarrow Z \\ B \wedge C^- \Rightarrow Z^-_{P_{v_s}} & \text{if } F \text{ is } B \wedge not\, C \Rightarrow not\ Z \end{cases}$$

The rewritten rules are obtained from the original ones by replacing their heads A (respectively, the $not\, A$) by the atoms $A_{P_{v_s}}$ (respectively, $A^-_{P_{v_s}}$) and by replacing negative premises $not\, C$ by C^-.

(RU) Rewritten updates: Let σ be a function defined as follows:

$$\sigma(s, i \div C) = \gamma(i, s, C)$$

where $i \div C$ is not one of the cases below. Note that updates of the form $i \div (?{-}L_1 \wedge \ldots \wedge L_n)$ are not treated here. They will be treated at the agent cycle level (see Section 4). The following cases characterize the DAG rewritten updates:

$$\begin{array}{l} \sigma(s, i \div \mathit{edge}(u, v)) = \mathit{edge}(u_s, v_s) \\ \sigma(s, i \div \mathit{not\ edge}(u, v)) = \mathit{edge}(u_s, v_s)^- \\ \sigma(s, i \div \mathit{modify_edge}(j, u, v, x)) = \mathit{modify_edge}(j, u, v, x) \\ \sigma(s, i \div \mathit{not\ modify_edge}(j, u, v, x)) = \mathit{modify_edge}(j, u, v, x)^- \end{array}$$

(IR) Inheritance rules:

$$A_{v_s} \leftarrow A_{u_t} \wedge \mathit{not\,reject}(A_{u_t}) \wedge \mathit{rel_edge}(u_t, v_s)$$
$$A^-_{v_s} \leftarrow A^-_{u_t} \wedge \mathit{not\,reject}(A^-_{u_t}) \wedge \mathit{rel_edge}(u_t, v_s)$$

for all objective atoms $A \in \mathcal{K}$, for all $u, v \in V \cup \{sa\}$ and for all $s, t \in S$. The inheritance rules state that an atom A is true (resp., false) in a vertex v at state s if it is true (resp., false) in any ancestor vertex u at state t and it is not rejected, i.e., forced to be false.

(RR) Rejection Rules:

$$\mathit{reject}(A^-_{u_s}) \leftarrow A_{P_{v_t}} \wedge \mathit{rel_path}(u_s, v_t)$$
$$\mathit{reject}(A_{u_s}) \leftarrow A^-_{P_{v_t}} \wedge \mathit{rel_path}(u_s, v_t)$$

for all objective atoms $A \in \mathcal{K}$, for all $u, v \in V \cup \{sa\}$ and for all $s, t \in S$. The rejection rules say that if an atom A is true (respectively, false) in the program P_v at state t, then it rejects inheritance of any false (respectively, true) atoms of any ancestor u at any state s.

(UR) Update Rules:

$$A_{v_s} \leftarrow A_{P_{v_s}} \qquad A^-_{v_s} \leftarrow A^-_{P_{v_s}}$$

for all objective atoms $A \in \mathcal{K}$, for all $v \in V \cup \{sa\}$ and for all $s \in S$. The update rules state that an atom A must be true (respectively, false) in the vertex v at state s if it is true (respectively, false) in the program P_v at state s.

(DR) Default Rules:

$$A^-_{sa_1}$$

for all objective atoms $A \in \mathcal{K}$. Default rules describe the initial vertex sa for atoms (at state 1) by making all objective atoms initially false.

($\mathbf{CSR}_{v_s}$) Current State Rules:

$$A \leftarrow A_{v_s} \qquad A^- \leftarrow A^-_{v_s}$$

for every objective atoms $A \in \mathcal{K}$. Current state rules specify the current vertex v and state s in which the program is being evaluated and determine the values of the atoms A and A^-.

Transformation of projects

(PPR) Project Propagation rules:

$$Z_{u_s} \wedge not\, rejectP(Z_{u_s}) \wedge rel_edge(u_s, v_s) \Rightarrow Z_{v_s}$$
$$Z^-_{u_s} \wedge not\, rejectP(Z^-_{u_s}) \wedge rel_edge(u_s, v_s) \Rightarrow Z^-_{v_s}$$

for all projects $Z \in \mathcal{K}$, for all $u, v \in V \cup \{sp\}$ and for all $s \in S$. The project propagation rules propagate the validity of a project Z through the DAG at a given state s. In contrast to inheritance rules, the validity of projects is not propagated along states.

(PRR) Project Rejection Rules:

$$rejectP(Z^-_{u_s}) \leftarrow Z_{P_{v_s}} \wedge rel_path(u_s, v_s)$$
$$rejectP(Z_{u_s}) \leftarrow Z^-_{P_{v_s}} \wedge rel_path(u_s, v_s)$$

for all projects $Z \in \mathcal{K}$, for all $u, v \in V \cup \{sp\}$ and for all $s \in S$. The project rejection rules state that if a project Z is true (respectively, false) in the program P_v at state s, then it rejects propagation of any false (respectively, true) project of any ancestor u at state s.

(PUR) Project Update Rules:

$$Z_{P_{v_s}} \Rightarrow Z_{v_s} \qquad Z^-_{P_{v_s}} \Rightarrow Z^-_{v_s}$$

for all projects $Z \in \mathcal{K}$, for all $v \in V \cup \{sp\}$ and for all $s \in S$. The project update rules declare that a project Z must be true (respectively, false) in v at state s if it is true (respectively, false) in the program P_v at state s.

(PDR) Project Default Rules:

$$Z^-_{sp_s}$$

for all projects $Z \in \mathcal{K}$ and for all $s \in S$. Project default rules describe the initial vertex sp for projects at every state $s \in S$ by making all projects initially false at every state s.

(PAR$_{v_s}$) Project Adoption Rules:

$$Z_{v_s} \Rightarrow Z$$

for every projects $Z \in \mathcal{K}$. Project adoption rules specify the current vertex v and state s in which the project is being evaluated. Note that this rule only applies to positive projects. If a project is evaluated to true, then it will be selected at the agent cycle level and executed.

Transformation rules for edge

(ER) Edge Rules

$$edge(u_1, v_1)$$

for all $(u, v) \in E$. Edge rules describe the edges in graph D at state 1. Plus the following rules that characterize the edges of the initial vertices for atoms and projects (i.e., sa and sp):

$$edge(sa_1, u_1) \qquad edge(sp_s, u_s)$$

for all sources $u \in V$ and for all $s \in S$.

(EIR) Edge Inheritance rules:

$$
\begin{aligned}
edge(u_{s+1}, v_{s+1}) &\leftarrow edge(u_s, v_s) \wedge not\ edge(u_{s+1}, v_{s+1})^- \\
edge(u_{s+1}, v_{s+1})^- &\leftarrow edge(u_s, v_s)^- \wedge not\ edge(u_{s+1}, v_{s+1}) \\
edge(u_s, u_{s+1}) &\leftarrow
\end{aligned}
$$

for all $u, v \in V$ such that u is not the inspection point of the agent, and for all $s \in S$. Note that the rules EIR do not apply to edges containing the vertices sa and sp.

($\mathbf{RER}_{v_s}$) Current State Relevancy Edge Rules

$$
\begin{aligned}
rel_edge(X, v_s) &\leftarrow edge(X, v_s) \\
rel_edge(X, Y) &\leftarrow edge(X, Y) \wedge rel_path(Y, v_s)
\end{aligned}
$$

Current state relevancy edge rules define which edges are in the relevancy graph wrt. the vertex v at state s.

(RPR) Relevancy Path Rules

$$
\begin{aligned}
rel_path(X, Y) &\leftarrow rel_edge(X, Y) \\
rel_path(X, Z) &\leftarrow rel_edge(X, Y) \wedge rel_path(Y, Z).
\end{aligned}
$$

Relevancy path rules define the notion of relevancy path in a graph.

3.3 Multi-dimensional updates for agents

This section introduces the notion of multi-dimensional dynamic program update, a transformation that, given an agent $\Phi^{\alpha,m} = (\mathcal{T}, \mathcal{U})$, produces an abductive logic program (as defined in [6]). Note first that an updating program U for an agent α can contain updates of the form $\alpha{\div}(?{-}L_1 \wedge \ldots \wedge L_n)$. Such updates are intended to add $L_1 \wedge \ldots \wedge L_n$ to the current query of the agent α. Thus, those updates have to be treated differently. To achieve this, we introduce a function β defined over updating programs.

Definition 11. *Let U be an updating program and $\Phi^\alpha \in \mathcal{M}$ an agent.*

- *If U does not contain any update starting with "$\alpha{\div}(?{-}$", then* $\beta(\alpha, U) = (true, U)$.
- *(Otherwise) Suppose that U contains n updates starting with "$\alpha{\div}(?{-}$", say $\alpha{\div}(?{-}C_1), \ldots, \alpha{\div}(?{-}C_n)$ for n conjunctions $C_1, \ldots, C_n$ of objective and default atoms. Then* $\beta(\alpha, U) = (C_1 \wedge \ldots \wedge C_n, U - \{\alpha{\div}(?{-}C_1), \ldots, \alpha{\div}(?{-}C_n)\})$.

Given a set Q of rules, integrity constraints and active rules, we write Ω_1 and Ω_2 to indicate:

$$
\begin{aligned}
\Omega_1(Q) &= \{C \mid \text{for every rule and integrity constraint } C \text{ in } Q\} \\
\Omega_2(Q) &= \{R \mid \text{for every active rule } R \text{ in } Q\}
\end{aligned}
$$

Now we can present the notion of multi-dimensional dynamic program updates for agents.

Definition 12. *Let $S = \{1, \ldots, m\}$ a set of states, and $\Phi^{\alpha,m} = (\mathcal{T}, \mathcal{U})$ an agent α at state m, where $\mathcal{T} = \langle \mathcal{P}_D, \mathcal{A}, \mathcal{R}_D, D \rangle$ and $\mathcal{U} = \{U_s \mid s \in S\}$. The*

multi-dimensional dynamic program update at the state m **for agent** *α, written as $\bigoplus \Phi^{\alpha,m}$, is the tuple which consists of the following rules:*

$$
\begin{aligned}
\bigoplus \Phi^{\alpha,m} &= \langle P', \mathcal{A}, R' \rangle \text{ with} \\
P' &= \bigcup_{v \in V} \gamma(v, 1, P_v) \cup \bigcup_{s \in S} (\Omega_1(Q_s)) \cup \\
&\quad IR \cup RR \cup UR \cup DR \cup CSR(\alpha'_m) \cup PRR \cup \\
&\quad PDR \cup ER \cup EIR \cup RER(\alpha'_m) \cup RPR \cup \\
R' &= \bigcup_{v \in V} \gamma(v, 1, R_v) \cup \bigcup_{s \in S} (\Omega_2(Q_s)) \cup \\
&\quad PPR \cup PUR \cup PAR(\alpha'_m)
\end{aligned}
$$

where $\beta(\alpha, U_s) = (_, P_s)$, and $\sigma(s, P_s) = Q_s$, for every $s \in S$.

Note that P' is a normal logic program, that is, default atoms can only occur in bodies of clauses.

4 The agent cycle

Every agent can be thought of as a multi-dimensional abductive logic program equipped with a set of *inputs* represented as *updates*. The abducibles are (names of) *actions* to be executed as well as explanations of observations made. Updates can be used to solve the goals of the agent as well as to trigger new goals. The basic "engine" of the agent is the IFF proof procedure [9, 6], executed via the cycle represented in Fig.4. Assume that $\beta(i, P_s) = (L_1 \wedge \ldots \wedge L_n, U_s)$, $\sigma(s, U_s) = Q_s$ and $\Phi^{\alpha,s-1} = (\mathcal{T}, \{U_1, \ldots, U_{s-1}\})$. The cycle of an agent α

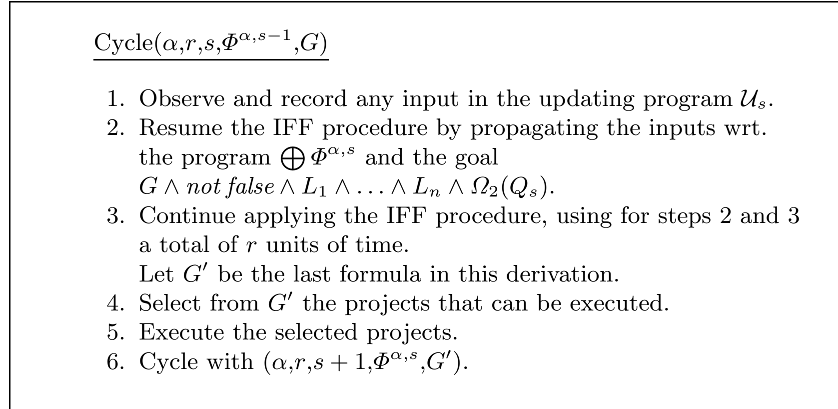

Fig. 4. *The agent cycle*

starts at state s by observing any inputs (projects from other agents) from the environment (step 1), and by recording them in the form of updates in the updating program U_s. Then, the proof procedure is applied for r units of time (steps 2 and 3) with respect to the abductive logic program $\bigoplus \Phi^{\alpha,s}$ obtained by updating the current one with the updating program U_s. The new goal is obtained by adding the goals $\textit{not false} \wedge L_1 \wedge \ldots \wedge L_n$ and the active rules

in $\Omega_2(Q_s)$ received from U_s to the current goal G. The amount of resources available in steps 2 and 3 is bounded by some amount r. By decreasing r the agent is more *reactive*, by increasing r the agent is more *rational*.
Afterwards (steps 4 and 5), the selectable projects are selected from G', the last formula of the derivation, and executed under the control of an abductive logic programming proof procedure such as ABDUAL in [3]. If the selected project takes the form $j : C$ (meaning that agent α intends to update the theory of agent j with C at state s), then (once executed) the update $\alpha \div C$ will be available (at the next cycle) in the updating program of j. Selected projects can be thought of as outputs into the environment, and observations as inputs from the environment. From every agent's viewpoint, the environment contains all other agents. Every disjunct in the formulae derived from the goal G represents an *intention*, i.e., a (possibly partial) plan executed in stages. A sensible action selection strategy may select actions from the same disjunct (intention) at different iterations of the cycle. Failure of a selected plan is obtained via logical simplification, after having propagated `false` into the selected disjunct.
Integrity constraints provide a mechanism for constraining explanations and plans, and action rules for allowing a condition-action type of behaviour.

5 Conclusions and Future Work

Throughout this paper we have extended the *Multi-dimensional Dynamic Logic Programming* framework to include integrity constraints and active rules, as well as to allow the associated acyclic directed graph to be updatable.
We have shown how to express, in a multi-agent system, each agent's viewpoint regarding its place in its perception of others. This is achieved by composing agents' view of one another in acyclic directed graphs, one for each agent, which can evolve over time by updating its edges.
Agents themselves have their knowledge expressed by abductive generalized logic programs with integrity constraints and active rules. They are kept active through an observe-think-act cycle, and communicate and evolve by realizing their projects to do so, in the form of updates acting on other agents or themselves. Projects can be knowledge rules, active rules, integrity constraints, or abductive queries.
With respect to future work, we are investigating into the representation and evolution of dynamic societies of agents. A first step towards this goal is to allow the set of vertices of the acyclic directed graph to be also updatable, representing the birth (and death) of the society's agents.

Acknowledgments We would like to thank José Júlio Alferes for all his comments. This work was partially supported by PRAXIS XXI 2/2.1/TIT/1593/95 and PRAXIS XXI BD/13514/97.

References

1. J. J. Alferes, J. A. Leite, L. M. Pereira, H. Przymusinska, and T. Przymusinski. *Dynamic logic programming.* In A. Cohn and L. Schubert (eds.), Proc. of KR'98, pages 98–109. Morgan Kaufmann, 1998.
2. J. J. Alferes, J. A. Leite, L. M. Pereira, H. Przymusinski and T. C. Przymusinski. *Dynamic Updates of Non-Monotonic Knowledge Bases.* The Journal of Logic Programming 45(1-3): 43-70, August 2000.

3. J. J. Alferes, L. M. Pereira and T. Swift. *Well-founded Abduction via Tabled Dual Programs*. Proc. of ICLP'99.
4. M. Bozzano, G. Delzanno, M. Martelli, V. Mascardi and F. Zini, *Logic Programming and Multi-Agent System: A Synergic Combination for Applications and Semantics*. In K. Apt, V. Marek, M. Truszczynski and D. S. Warren (eds.), The Logic Programming Paradigm - A 25-Year Perspective, pages 5-32, Springer 1999.
5. D. De Schreye, M. Hermenegildo and L. M. Pereira, *Paving the Roadmaps: Enabling and Integration Technologies*. Available from `http://www.compulog.org/net/Forum/ Supportdocs.html`
6. P. Dell'Acqua, L. M. Pereira, *Updating Agents*. In S. Rochefort, F. Sadri and F. Toni (eds.), Procs. of the ICLP'99 Workshop on Multi-Agent Systems in Logic (MASL'99), 1999.
7. P. Dell'Acqua, F. Sadri, and F. Toni. *Combining introspection and communication with rationality and reactivity in agents*. In J. Dix, F.L. Del Cerro, and U. Furbach (eds.), *Logics in Artificial Intelligence*, LNCS 1489, pages 17–32, Berlin, 1998. Springer-Verlag.
8. N. Jennings, K. Sycara and M. Wooldridge. *A Roadmap of Agent Research and Development*. In Autonomous Agents and Multi-Agent Systems, 1, 275-306, Kluwer, 1998.
9. T. H. Fung and R. Kowalski. *The IFF proof procedure for abductive logic programming. J. Logic Programming*, 33(2):151–165, 1997.
10. R. Kowalski and F. Sadri. *Towards a unified agent architecture that combines rationality with reactivity*. In D. Pedreschi and C. Zaniolo (eds), Procs. of LID'96, pages 137-149, Springer-Verlag, LNAI 1154, 1996.
11. J. A. Leite, J. J. Alferes and L. M. Pereira, *Multi-dimensional Dynamic Logic Programming*, In F. Sadri and K. Satoh (eds.), Procs. of the CL-2000 Workshop on Computational Logic in Multi-Agent Systems (CLIMA'00), pages 17-26, London, England, July 2000.
12. J. A. Leite. *Logic Program Updates*. M.Sc. Dissertation, Universidade Nova de Lisboa, 1997.
13. J. A. Leite and L. M. Pereira. *Generalizing updates: from models to programs*. In J.Dix, L.M. Pereira and T.C.Przymusinski (eds), Procs. of LPKR'97, pages 224-246, Springer, LNAI 1471, 1998.
14. J. A. Leite and L. M. Pereira. *Iterated Logic Program Updates*. In J. Jaffar (ed.), Procs. of the 1998 Joint International Conference and Symposium on Logic Programming, Manchester, England, pages 265-278. MIT Press, June 1998.
15. S. Rochefort, F. Sadri and F. Toni, editors, *Proceedings of the International Workshop on Multi-Agent Systems in Logic Programming*, Las Cruces, New Mexico, USA,1999. Available from `http://www.cs.sfu.ca/conf/MAS99`.
16. I. Niemelä and P. Simons. *Smodels - an implementation of the stable model and well-founded semantics for normal logic programs*. In Procs. of the 4th LPNMR'97, Springer, July 1997.
17. F. Sadri and F. Toni. *Computational Logic and Multiagent Systems: a Roadmap*, 1999. Available from `http://www.compulog.org`.
18. The XSB Group. *The XSB logic programming system, version 2.0*, 1999. Available from `http://www.cs.sunysb.edu/~sbprolog`.