

MINERVA - A Dynamic Logic Programming Agent Architecture

João Alexandre Leite, José Júlio Alferes, and Luís Moniz Pereira

Centro de Inteligência Artificial (CENTRIA)
Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
{jleite | jja | lmp}@di.fct.unl.pt

Abstract. The agent paradigm, commonly implemented by means of imperative languages mainly for reasons of efficiency, 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 brought back into the spotlight. To this accrues the recent significant improvements in the efficiency of *Logic Programming* implementations for *Non-monotonic Reasoning*.

This paper presents an overall description of *MINERVA*, an agent architecture and system designed with the intention of providing a common agent framework based on the unique strengths of *Logic Programming*, to allow for the combination of several non-monotonic knowledge representation and reasoning mechanisms developed in recent years. In [10], the semantics of the multi-dimensional structure and combination of the evolving societal knowledge of agents is described and discussed in detail.

1 Introduction

The *Logic Programming (LP)* paradigm provides a well-defined, general, integrative, encompassing, and rigorous framework for systematically studying computation, be it syntax, semantics, procedures, or attending implementations, environments, tools, and standards. *LP* approaches problems, and provides solutions, at a sufficient level of abstraction so that they generalize from problem domain to problem domain. This is afforded by the nature of its very foundation in logic, both in substance and method, and constitutes one of its major assets. To this accrues the recent significant improvements in the efficiency of *Logic Programming* implementations for *Non-monotonic Reasoning* [11,14,16]. 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 *LP*, such as belief revision, inductive learning, argumentation, preferences, abduction, etc.[12,13] 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 felicitous combination between reactive and rational behaviours of agents, the *Holy Grail* of modern *Artificial Intelligence*, whilst preserving clear and precise specification enjoyed by declarative languages.

Until recently, *LP* could be seen as a good representation language for static knowledge. If we are to move to a more open and dynamic environment, typical of the agency paradigm, we need to consider ways of representing and integrating knowledge from different sources which may evolve in time. To solve this limitation, the authors, with others, first introduced *Dynamic Logic Programming (DLP)* [1]. There, they studied and defined the declarative and operational semantics of sequences of logic programs (or dynamic logic programs). Each program in the sequence contains knowledge about some given state, where different states may, for example, represent different time periods or different sets of priorities. In [9] the *DLP* paradigm was then generalized in order to allow, not only for sequences of logic programs, but for collections of programs structured by acyclic digraphs (DAGs). By dint of such generalization, *Multi-dimensional Dynamic Logic Programming (MDLP)* affords extra expressiveness, thereby enlarging the latitude of logic programming applications unifiable under a single framework. The generality and flexibility provided by a *DAG* ensures a wide scope and variety of new possibilities. By virtue of the newly added characteristics of multiplicity and composition, *MDLP* provides a “societal” viewpoint in Logic Programming, important in these web and agent days, for combining knowledge in general. In [10] these new possibilities are explored.

Moreover, an agent not only comprises knowledge about each state, but also knowledge about the transitions between states. The latter can represent the agent’s knowledge about the environment’s evolution, and its own behaviour and evolution. Since logic programs describe knowledge states, it’s only fit that logic programs be utilized to describe transitions of knowledge states as well, by associating with each state a set of transition rules to obtain the next one. In [3], the authors, with others, introduce the language *LUPS* – “Language for dynamic updates” – designed for specifying changes to logic programs. Given an initial knowledge base (as a logic program) *LUPS* provides a way for sequentially updating it. The declarative meaning of a sequence of sets of update commands in *LUPS* is defined by the semantics of the dynamic logic program generated by those commands.

Based on the strengths of *MDLP* as a framework capable of simultaneously representing several aspects of a system in a dynamic fashion, and on those of *LUPS* as a powerful language to specify the evolution of such representations via state transitions, we have launched into the design of an agent architecture, *MINERVA*. Named after the Goddess of Wisdom, this architecture is conceived with the intention of providing, on a sound theoretical basis, a common agent framework based on the strengths of Logic Programming, so as to allow the combination of several non-monotonic knowledge representation and reasoning mechanisms developed in recent years. Rational agents, in our opinion, will require an admixture of any number of those reasoning mechanisms to carrying out their tasks. To this end, a *MINERVA* agent hinges on a modular design, whereby a common knowledge base, containing all knowledge shared by more than one sub-agent, is concurrently manipulated by specialized sub-agents.

Its description here is by no means exhaustive, a task that would be impossible to accomplish within this article alone. Rather, it aims at providing an overview highlighting the role of computational logic in its several components.

2 Dynamic Logic Programming

In this section we briefly review Dynamic Logic Programming [1], its generalization *MDLP* [9], and the update command language *LUPS* [3].

The idea of dynamic updates is simple and quite fundamental. Suppose that we are given a set of generalized logic program modules P_s (possibly with negation in rule heads), structured by a DAG (or simply index by a sequence number, in [1]). Each program P_s contains some knowledge that is supposed to be true at the state s . Different states may represent different time periods or different sets of priorities or perhaps even different viewpoints. Consequently, the individual program modules may contain mutually contradictory as well as overlapping information. The role of the dynamic program update is to use the mutual relationships existing between different states (and specified by the relation in the DAG) to precisely determine, at any given state s , the *declarative* as well as the *procedural* semantics of the combined program, composed of all modules. The declarative semantics at some state is determined by the stable models of the program that consists in all those rules that are “valid” in that state. Intuitively a rule is “valid” in a state if either it belong to the state or it belong to some previous state in the DAG and is not rejected (i.e. it is inherited by a form of inertia). A rule from a less priority state is rejected if there is another conflicting rule in a more priority state with a true body. A transformational semantics, that directly provides a means for *MDLP* implementation, has also been defined¹. For details on Dynamic Logic Programming and *MDLP* the reader is referred to [1,9]. The paper [10], contains a more detailed review of this subject, as well as an explanation of how *MDLP* can be employed to model the combination of inter- and intra-agent societal viewpoints.

LUPS [3] is a logic programming command language for specifying updates. It can be viewed as a language that declaratively specifies how to construct a sequence of logic programs. A sentence U in *LUPS* is a set of simultaneous update commands (or actions) that, given a pre-existing sequence of logic programs (or a *MDLP*), whose semantics corresponds to our knowledge at a given state (or program), produces a new *MDLP* with one more program, corresponding to the knowledge that results from the previous sequence after performing all the simultaneous commands. Different possibilities on how to connect the newly produced program with the previously extant ones give rise to different ways of combining the knowledge of the various agents. In [10] these possibilities are examined, and it is shown how to attain: equal role among agents, time prevailing and hierarchy prevailing representations. In this paper, if not otherwise stated, simply assume that the new program is concatenated immediately after the program where the command is evaluated.

A program in *LUPS* is a sequence of such sentences, and its semantics is defined by means of a dynamic logic program generated by the sequence of commands. In [3], a translation of a *LUPS* program into a generalized logic program is presented, where stable models exactly correspond to the semantics of the original *LUPS* program. This translation directly provides an implementation of *LUPS* (available at the URL given above).

¹ The implementation can be found in `centria.di.fct.unl.pt/~jja/updates/`

LUPS update commands specify assertions or retractions to the current program (a predefined one in the \mathcal{MDLP} , usually the one resulting from the last update performed). In *LUPS* a simple assertion is represented as the command:

$$\mathbf{assert} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (1)$$

meaning that if L_{k+1}, \dots, L_m holds in the current program, the rule $L \leftarrow L_1, \dots, L_k$ is added to the new program (and persists by inertia, until possibly retracted or overridden by some future update command). To represent rules and facts that do not persist by inertia i.e., that are one-state events, *LUPS* includes the modified form of assertion:

$$\mathbf{assert \ event} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (2)$$

The retraction of rules is performed with the two update commands:

$$\mathbf{retract} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (3)$$

$$\mathbf{retract \ event} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (4)$$

meaning that, subject to precondition L_{k+1}, \dots, L_m (verified at the current program) rule $L \leftarrow L_1, \dots, L_k$ is either retracted from its successor state onwards, or just temporarily retracted in the successor state (if governed by **event**).

Normally assertions represent newly incoming information. Although its effects remain by inertia (until contravened or retracted), the **assert** command itself does not persist. However, some update commands may desirably persist in the successive consecutive updates. This is especially the case of laws which, subject to some preconditions, are always valid, or of rules describing the effects of an action. In the former case, the update command must be added to all sets of updates, to guarantee that the rule remains indeed valid. In the latter case, the specification of the effects must be added to all sets of updates, to guarantee that, whenever the action takes place, its effects are enforced. To specify such persistent update commands, *LUPS* introduces:

$$\mathbf{always} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (5)$$

$$\mathbf{always \ event} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (6)$$

$$\mathbf{cancel} \ L \leftarrow L_1, \dots, L_k \ \mathbf{when} \ L_{k+1}, \dots, L_m \quad (7)$$

The first two statements signify that, in addition to any new set of arriving update commands, the persistent update command keeps executing along with them too. The first case without, and the second case with the **event** keyword. The third statement cancels execution of this persistent update, once the conditions for cancellation are met.

3 Overall Architecture

In the *MINERVA* agent architecture, an agent consists of several specialized, possibly concurrent, sub-agents performing various tasks while reading and manipulating a common knowledge base. A schematic view of the *MINERVA* architecture is depicted in Fig. 1. The common knowledge base contains knowledge about the self and

the agent community, and is conceptually divided into the following components: *Capabilities*, *Intentions*, *Goals*, *Plans*, *Reactions*, *Object Knowledge Base* and *Internal Behaviour Rules*. There is also an internal clock. Although conceptually divided in such components, all these modules will share a common representation mechanism based on *MDLP* and *LUPS*, the former to represent knowledge at each state and *LUPS* to represent the state transitions i.e., the common part of the agent's behaviour. Every agent is composed of specialized function related sub-agents, that execute their various specialized tasks. Examples of such sub-agents are those implementing the reactive, planning, scheduling, belief revision, goal management, learning, dialogue management, information gathering, preference evaluation, strategy, and diagnosis functionalities. These sub-agents contain a *LUPS* program encoding its behaviour, interfacing with the *Common Knowledge Base*. Whilst some of those sub-agent's functionalities are fully specifiable in *LUPS*, others will require private specialized procedures where *LUPS* serves as an interface language.

In all *LUPS* commands, both in the common knowledge base and in the sub-agents, by default the **when**-clauses are evaluated in the *Object Knowledge Base*, and the rules are added to the *Object Knowledge Base* also. More precisely, the **when**-condition of a sub-agent α_n *LUPS* command is evaluated at a special sink program of the *Object Knowledge Base*². If the condition is verified, the rule is asserted (resp. retracted), by default, at the state $\alpha_{n,c+1}$ i.e., the sub-agent's state corresponding to the next time state. In effect, at each time state c , each sub-agent's *LUPS* interpreter evaluates its commands at state α' (dynamically and therefore differently at each time state) and produces its corresponding next state, $c + 1$. Whenever some literal L in a **when**-condition is to be evaluated in a program different from the default one, we denote that by $L@ProgramName$. This @ notation will also be used to denote addition of rules to programs other than the default one.

4 Common Knowledge Base

Object knowledge base. The *Object Knowledge Base* is the main component containing knowledge about the world i.e., knowledge at the object level and info about the society where the agent is situated.

It is represented by an evolving *MDLP*. In it there is a sequence of nodes for each sub-agent of the agent α , representing its evolution in time. There is also a sequence of nodes for every other agent in the multi-agent system, representing α 's view of the evolution of those agents in time. As will be seen in the following Section, at each time state each sub-agent manipulates its corresponding state node. There is a *Dialoguer* sub-agent dealing with the interactions with other agents, and manipulating the state nodes corresponding to the agents it communicates with. Several methods are possible for combining these several sequences to form the *MDLP*'s DAG. Each corresponds to a different way of combining inter- and intra-agent social viewpoints. In [10] we show how this can be done.

² This special program is a sink of the DAG, and its semantics corresponds to the semantics of all the *Object Knowledge Base*, cf. [10].




Fig. 1. The MINERVA agent architecture

Capabilities. This block contains a description of the actions, and their effects, capable of being performed by the agent. *LUPS*, by allowing to declaratively specify both knowledge states and their transitions, can be used as a powerful representation language for specifying actions and their effects. Its diversity of update commands can model from the simplest condition-effect action to parallel actions and their coupled indirect effects. Typically, for every *action* ω , there will be one (or more) *LUPS* commands, where the action name ω appears in the ‘when’ clause. For example, in

$$\text{always } L \leftarrow L_1, \dots, L_k \text{ when } L_{k+1}, \dots, L_m, \omega$$

we have, intuitively, that ω is an action whose preconditions are L_{k+1}, \dots, L_m and whose effect is an update that, according to its type, can model different kinds of actions, all in one unified framework. Examples of kinds of actions are:

- actions of the form “ ω causes F if F_1, \dots, F_k ”, where F, F_1, \dots, F_k are fluents (such as in the language \mathcal{A} of [6]) translates into the update command

$$\text{always } F \text{ when } F_1, \dots, F_k, \omega$$

- actions whose epistemic effect is a rule update of the form “ ω updates with $L \leftarrow L_1, \dots, L_k$ if L_{k+1}, \dots, L_m ” translates into the update command

$$\text{always } L \leftarrow L_1, \dots, L_k \text{ when } L, \dots, L_m, \omega$$

- actions that, when performed in parallel, have different outcomes, of the form “ ω_a or ω_b cause L_1 if L_{k+1}, \dots, L_m ” and “ ω_a and ω_b in parallel cause L_2 if L_{k+1}, \dots, L_m ” translate into the three update commands:

$$\text{always } L_1 \text{ when } L_{k+1}, \dots, L_m, \omega_a, \text{not } \omega_b$$

$$\text{always } L_1 \text{ when } L_{k+1}, \dots, L_m, \text{not } \omega_a, \omega_b$$

$$\text{always } L_2 \text{ when } L_{k+1}, \dots, L_m, \omega_a, \omega_b$$

- actions with non-deterministic effects of the form “ α cause $L_1 \vee L_2$ if L_{k+1}, \dots, L_m ” translates into the update commands (where β_a, β_b are new auxiliary predicates and I is a unique identifier of an occurrence of α):

$$\text{always } L_1 \leftarrow \beta_a(I) \text{ when } L_{k+1}, \dots, L_m, \alpha(I)$$

$$\text{always } L_2 \leftarrow \beta_b(I) \text{ when } L_{k+1}, \dots, L_m, \alpha(I)$$

$$\text{always } \beta_a(I) \leftarrow \text{not } \beta_b(I) \text{ when } L_{k+1}, \dots, L_m, \alpha(I)$$

$$\text{always } \beta_b(I) \leftarrow \text{not } \beta_a(I) \text{ when } L_{k+1}, \dots, L_m, \alpha(I)$$

In this representation of the non-deterministic effects of an action α , we create two auxiliary actions (β_a, β_b) with deterministic effects, and make the effect of α be the non-deterministic choice between actions β_a or β_b .

Internal behaviour rules. The *Internal Behaviour Rules* are *LUPS* commands that specify the agent’s reactive internal epistemic state transitions. They are of the form:

$$\text{assert } L \leftarrow L_1, \dots, L_k \text{ when } L_{k+1}, \dots, L_m$$

whose effect is the immediate asserting of rule $L \leftarrow L_1, \dots, L_k$ in the *Object Knowledge Base* if L_{k+1}, \dots, L_m holds now. These commands will be executed by one of the sub-agents (the *Reactor*), and they are quite common since they may be used by more than one of the sub-agents. E.g. the *Planner* sub-agent must be aware of these rules when planning for goals. An example of internal behaviour rules would be:

$$\begin{aligned} \text{assert } \textit{jail}(X) \leftarrow \textit{abortion}(X) \text{ when } \textit{government}(\textit{republican}) \\ \text{assert } \textit{not } \textit{jail}(X) \leftarrow \textit{abortion}(X) \text{ when } \textit{government}(\textit{democrat}) \end{aligned}$$

stating that whenever the government is republican, an abortion implies going to jail, and whenever the government is democrat, an abortion does not imply going to jail.

Goals. The *Goals* structure is a DLP where at each state the program contains facts of the form $\textit{goal}(\textit{Goal}, \textit{Time}, \textit{Agent}, \textit{Priority})$ representing the goals that are to be achieved by the agent. *Goal* is a conjunction of literals, *Time* refers to a time state, *Agent* represents the origin of the goal and *Priority* contains the priority of the goal. *Goals* is a dynamic structure in the sense that new goals can be asserted onto this structure by any sub-agent, some of them originating in another one. They will be manipulated by a *Goal Manager* sub-agent that can assert, retract and change a goal’s priority.

Plans. As shown in [2], *LUPS*, together with abduction, can be used to represent and solve planning problems. There, the notion of *action update*, U^ω , is defined as a set of update commands of the form $U^\omega = \{\text{assert event } \omega\}$ where ω is an action name. Intuitively, each command of the form “assert event ω ” represents the performing of action ω . Note that performing an action itself is something that does not persist by inertia. Thus, according to the description of *LUPS*, the assertion must be of an event. By asserting event ω , the effect of the action will be enforced if the preconditions are met. Each *action update* represents a set of simultaneous actions. Plans are *LUPS* programs composed of action updates. An example of a plan to achieve a \textit{Goal}_1 at time T would be:

$$\begin{aligned} U_{T-3} &= \{\text{assert event } \omega_1; \text{assert event } \omega_2\} \\ U_{T-1} &= \{\text{assert event } \omega_4; \text{assert event } \omega_5\} \end{aligned}$$

meaning that in order to achieve \textit{Goal}_1 at time T , the agent must execute actions ω_1 and ω_2 at time $T - 3$ and actions ω_4 and ω_5 at time $T - 1$. Also, associated with each plan, there is a set of preconditions (*Conditions*) that must be met at the beginning of the plan. Each plan (*Plan*) for goal (*Goal*) is generated by a planner sub-agent, which asserts it into the *Common Knowledge Base* in the form $\textit{plan}(\textit{Goal}, \textit{Plan}, \textit{Conditions})$ where it is kept for future reuse.

By permitting plans to be represented in the *LUPS* language, we allow the use of more complex planners, namely conditional planners. Conditional action updates are to be represented by *LUPS* commands of the form:

$$\text{assert event } \omega \text{ when } L_{k+1}, \dots, L_m$$

An example of such a conditional plan, for $Goal_2$ would be:

$$U_{T-3} = \{\text{assert event } \omega_3 \text{ when } L; \text{assert event } \omega_6 \text{ when } \textit{not} L\}$$

$$U_{T-2} = \{\text{assert event } \omega_1; \text{assert event } \omega_2\}$$

meaning that in order to achieve $Goal_2$ at time T , the agent must execute, at time $T - 3$, the command **assert event** ω_3 if L holds at time $T - 3$, otherwise, if $\textit{not} L$ holds, it must execute the command **assert event** ω_6 . Then, it should execute the commands **assert event** ω_1 and **assert event** ω_2 at time $T - 2$. These conditions, in the **when** statement of these commands, can also be employed to include the verification of the success of previous actions.

The plans stored in the common knowledge are going to be processed, together with the reactions, by a specialized *Scheduler* sub-agent that will avail itself of them to determine the intentions of the agent.

Reactions. The *Reactions* block at the *Common Knowledge Base* is represented by a *very simple MDLP* whose object rules are just facts denoting actions ω or negation of actions names $\textit{not} \omega$. The *MDLP* contains a sequence of nodes for every sub-agent capable of reacting, and a set of edges representing a possible hierarchy of reactions. Those sub-agents will contain a set of *LUPS* commands of the form

$$\text{assert } \omega \text{ when } L_{k+1}, \dots, L_m$$

Whenever L_{k+1}, \dots, L_m is true, action ω should be reactively executed. This corresponds to the assertion of ω in the corresponding node of the *Reactions* module. These, together with the plans, will be processed by the *Scheduler* sub-agent to transform them into intentions. The richness of the *LUPS* language enables us to express action blockage i.e., under some conditions, preventing some action ω to be executed by the agent, possibly proposed by another sub-agent. This can be expressed by the command:

$$\text{assert } \textit{not} \omega \text{ when } L_{k+1}, \dots, L_m$$

whose effect would be the assertion of $\textit{not} \omega$. This assertion would block any ω proposed by a lower ranked reactive sub-agent. By admitting the buffering of reactions i.e., they are first proposed and then scheduled, this architecture allows for some conceptual control over them, namely by preventing some undesirable reactions being performed.

Intentions. *Intentions* are the actions the agent has committed to. A *Scheduler* sub-agent compiles the plans and reactions to be executed into the intentions, according to the resources of the agent. Control over the level of reactivity and deliberation can be achieved by this process. The *Intentions* structure is a *DLP* where the object language rules are simply facts of the form:

$$\textit{intention} (Action, Conditions, Time)$$

where $Action \in K_\omega$, $Conditions$ is a conjunction of literals from L , and $Time \in T$. They can be interpreted as: execute $Action$, at time T , if $Conditions$ is true. It is left

to the Actuator sub-agent to execute the intentions. For example, to achieve $Goal_2$ and $Goal_1$, the two plans previously exemplified can be compiled into the set of intentions (where c is the current time state):

$$\begin{aligned} &intention(\omega_3, L, c + 1); \quad intention(\omega_1, \neg, c + 3); \quad intention(\omega_4, \neg, c + 5) \\ &intention(\omega_6, not\ L, c + 1); \quad intention(\omega_2, \neg, c + 3); \quad intention(\omega_5, \neg, c + 5) \end{aligned}$$

Although typically intentions are internal commitments, nothing prevents us from having a specialized sub-agent manage (and possibly retract from) the set of intentions.

5 Sub-agents

The sub-agents constitute the labour part of the *MINERVA* architecture. Their tasks reside in the evaluation and manipulation of the *Common Knowledge Base* and, in some cases, in the interface with the environment in which the agent is situated. They differ essentially in their specialities, inasmuch as there are planners, reactors, schedulers, learners, dialoguers, etc. Conceiving the architecture based on the notion of, to some extent, independent sub-agents attain a degree of modularity, appropriate in what concerns the adaptability of the architecture to different situations.

Each of these sub-agents contains a *LUPS* program encoding its behaviour, interfacing with the *Common Knowledge Base* and possibly with private specialized procedures. Conceptually, each of these sub-agents also comprises a *LUPS* meta-interpreter that executes its *LUPS* program and produces a sequence of states in the structures of the *Common Knowledge Base*. The collection of all such sequences of states, produced by all the sub-agents will come to constitute the states of the *Object Knowledge Base*.

To permit private procedure calls, we extend *LUPS* by allowing those to be called in the **when** statement of a command:

$$\mathbf{assert}\ X \leftarrow L_1, \dots, L_k \mathbf{\ when}\ L_{k+1}, \dots, L_m, ProcCall(L_{m+1}, \dots, L_n, X)$$

In this case, if L_{k+1}, \dots, L_m is true at state α' , then $ProcCall(L_{m+1}, \dots, L_n, X)$ is executed and, if successful, X is instantiated and rule $X \leftarrow L_1, \dots, L_k$ is asserted into the sub-agent's state corresponding to the next time state. These procedure calls can read from the *Common Knowledge Base*, but cannot change it. All changes to the *Common Knowledge Base* are performed via *LUPS* commands, ensuring a clear and precise semantics of the overall system.

It is important to stress that *LUPS* constitutes a language to specify state transitions, namely by determining the construction of the programs that define each state, whose meaning is assigned by the semantics of *MDLP*. Some of the tasks to be performed by the sub-agents need to resort to procedure calls, but others are fully specifiable with *LUPS* commands that simply consult the *Common Knowledge Base*.

Sensor. The *Sensor* sub-agent receives input from the environment through procedure calls of the form $SensorProc(Rule)$. The information, in the form of logic program rules, is asserted in the object knowledge base. This behaviour is accomplished with *LUPS* commands of the form:

$$\mathbf{assert}\ Rule \mathbf{\ when}\ SensorProc(Rule)$$

The **when** statement can also be used as a filter, possibly specifying conditions for the acceptance of the input. Mark that since there is no explicit reference as where to execute the command, the assertions are, by default, performed at the *Sensor*'s node at the *Object Knowledge Base*.

Dialoguer. The Dialoguer sub-agent is similar to the *Sensor* one in what concerns receiving input from the environment. It differs in that it processes incoming messages from the other agents in the multi-agent system ($\beta, \gamma, \mu, \dots$), and the *LUPS* commands are then executed by asserting the information into the corresponding agent's node. Moreover, it is this sub-agent's charge, according to the message, to generate new goals, issue replies, etc. Examples of *LUPS* commands typical of this sub-agent would be:

```

assert goal(Goal, Time, Agent, _)@goals
           when MsgFrom(Agent, Goal, Time, Rule), cooperative(Agent).
assert Rule@Agent when MsgFrom(Agent, Goal, Time, Rule).
assert msgTo(Agent, Goal, plan(Goal, Plan, Cond))@reactions
           when goal(Goal, Time, Agent, _)@goals, Agent  $\neq$   $\alpha$ ,
                plan(Goal, Plan, Cond)@plans.
    
```

These commands specify a behaviour for the acceptance of goals and beliefs from another agent, and for dealing with the request. The first command insists that a goal should be asserted into the *Goals* of the *Common Knowledge Base* whenever a message from a cooperative agent requests it. The second asserts the rules communicated by another agent into its corresponding node of the *Object Knowledge Base*. The third issues a communication action, by asserting it into the *Reactions*, to be scheduled, whenever there is a request from another agent to solve a goal, and there is a plan for such a goal.

Actuator. The Actuator sub-agent interfaces with the environment by executing actions on it. At each cycle, the *Actuator* extracts the actions to be executed from the *Intentions* and performs them on the outside world. If successful, it event asserts the action name, ω , in the *Object Knowledge Base*. At each cycle, the *Actuator* executes *LUPS* commands of the form:

```

assert event  $\omega$  when intention (Action, Cond, Time) @intentions,
                    Current(Time), Cond, ActionProc( $\omega$ )
    
```

corresponding to the execution of the *Intentions* scheduled for the current time state.

Effector. The Effector, at each cycle, evaluates and executes the *LUPS* commands in the *Capabilities* and *Common behaviour Rules*. Note that, although the *Capabilities* and *Common Behaviour Rules* constitute the *LUPS* program of the *Effector* sub-agent, they belong to the *Common Knowledge Base* because they may be accessed by other sub-agents, namely by the Planner. The *Capabilities* require the prior successful execution of an action. This does not happen with the *Common Behaviour Rules*. These specify state transitions solely based on the internal state of the agent.

Reactor. The *Reactor* will contain reactive rules that, if executed, produce an action to be performed. They are of the form:

$$\text{assert event } \omega @ \text{reactions} \text{ when } L_1, \dots, L_k$$

An example of such a reaction would be

$$\text{assert event } runaway @ \text{reactions} \text{ when } danger$$

where *runaway* is the name of an action. We can also reactively and temporarily block actions by using *LUPS* commands of the form:

$$\text{assert event } not \omega @ \text{reactions} \text{ when } L_1, \dots, L_k$$

Planner. In [2], it was shown how planning can be achieved by means of abduction in *LUPS* specified scenarios. Such an abduction based planner uses the *Object Knowledge Base* together with the *Intentions* (actions already set to be executed in the future), *Capabilities*, and the *Common Behaviour Rules*, to achieve a plan for a given goal, consistent with the current set of intentions. A *LUPS* command for the *Planner*, with such an abductive planner represented by the procedure call *AbductPlan*(*Goal*, *Plan*, *Cond*), would be:

$$\begin{aligned} &\text{assert } plan(Goal, Plan, Cond) @plans \\ &\text{when } goal(Goal, T, -, 1) @goals, AbductPlan(Goal, Plan, Cond) \end{aligned}$$

Other planners, possibly more efficient and/or complex can be used, namely partial planners, with the appropriate *LUPS* commands to interface them.

Goal manager. Goals can be asserted by other sub-agents, some of them originating in other agents. These goals may conflict with one another. The task of the *Goal Manager* is to deal with such issues. It manipulates the *Goals* structure, being able to merge goals, set priorities, delete goals, etc. An example of a command for this sub-agent, specifying that if two goals are incompatible, the one with lower priority should be removed, is:

$$\begin{aligned} &\text{retract } goal(G_1, T_1, A_1, P_1) @goals \\ &\text{when } goal(G_1, T_1, A_1, P_1) @goals, goal(G_2, T_2, A_2, P_2) @goals, \\ &\quad incomp(G_1, G_2), P_1 < P_2 \end{aligned}$$

Scheduler. The *Scheduler* determines the *Intentions* of the agent, based on the current state of the agent. The *Scheduler* acts whenever there are pending reactions or goals and plans. More than one specialized scheduling procedure may exist. For example, if there were three procedures, depending on the need to combine reactions and plans or just to schedule one of them, the following *LUPS* commands would select their execution:

$$\begin{aligned} &\text{assert } II @ \text{intentions} \text{ when } goal(G, -, -, -) @goals, plan(G, -, -) @plans, \\ &\quad not X @ \text{Reactions}, SchedulePlans(II). \\ &\text{assert } II @ \text{intentions} \text{ when } not goal(-, -, -, -) @goals, X @ \text{reactions}, \\ &\quad ScheduleReactions(II). \\ &\text{assert } II @ \text{intentions} \text{ when } goal(G, -, -, -) @goals, plan(G, -, -) @plans, \\ &\quad X @ \text{reactions}, CombineSchedule(II). \end{aligned}$$

Other sub-agents. Other sub-agents can concurrently exist in the *MINERVA* architecture. Fig. 1 mentions a *Contradiction Remover* and a *Learner* sub-agent, but many others can be specified. Either fully specified by means of *LUPS* commands or with *LUPS* serving simply as an interface to procedure calls or legacy code, a great variety of sub-agents can be fully integrated into this architecture. If tight control is needed, *LUPS* also encompasses the specification of control sub-agents. For examples of the application of *LUPS* in several domains the reader is referred to [4].

6 Related Work

The characteristic of *MINERVA* that makes it most distinct from other existing agent architectures is that its knowledge is represented by non-monotonic theories, expressed by logic programs, with default negation in both the heads and bodies of clauses, organized according to acyclic digraphs. This permits the use of the richer knowledge representation capabilities afforded by such logic programs, for example to represent the relations amongst the agents themselves in a MAS, as well as for a direct and simple mechanism to represent and reason about dynamic environments where the governing rules change with time. It has been shown [1] that updating non-monotonic theories is quite different from updating theories represented in classical propositional logic, governed by the update postulates in [8], and far different from the mere assertion and deletion of formulas. To the best of our knowledge, *MINERVA* is the only agent architecture that proposes to incorporate such update capabilities. Most other agent architectures incorporate some update facilities, but all in the line of the update postulates of [8], or by simple assertion and retraction of formulas. We should also make clear at this point that the assertions and retractions specified by the *LUPS* commands differ substantially from those approaches employing simple assertion and retraction of formulas. The former are dealt with by the semantics of *MDLP*, which further rejects other existing rules on the grounds of the newly asserted ones and may as well reinstate other previously rejected rules, while the latter only deal with “physical” assertions and retractions of formulas, the semantics being determined according to the new set of formulas.

Clearly, not all situations require such update features as our’s and, in those situations, other existing agent proposals may perform better. Nevertheless, when the governing rules change dynamically, new rules possibly invalidating existing rules, as for example in legal reasoning scenarios, a semantics such as the one provided for *MDLP* will be required. *MDLP* additionally caters for a clear framework, with a precise semantics, to represent societies of hierarchically related agents.

In [7], Hindriks et al. propose the agent programming language *3APL*, together with a transition system based operational semantics. Knowledge is represented as a set of first order formulas (although the authors claim it could be any logical language) in a database and goals are sequences of imperative and logic constructs. On top of these there are the so-called practical reasoning rules to specify changes to the agent state. *3APL*, as claimed by the authors, intends to be a language to serve both as a sound basis for theoretical work on agents and as a vehicle for practical applications. The approach to the design of *3APL* is based on the premise that an agent programming language

amounts to providing a way to specify the dynamics of the mental state of an agent, be it its knowledge and beliefs or its goals, intentions, and commitments. This intention and premise are shared by *MINERVA* in the sense that the first is the ground for the choice of non-monotonic logic programming in general as the basic framework, and the second amounts to the motivation for using multi-dimensional dynamic logic programming and *LUPS* as knowledge representation and evolution specification language. In some sense, we can compare the *3APL* database to the object knowledge base of a *MINERVA* agent and the *3APL* practical reasoning rules to the *LUPS* commands in *MINERVA*. The differences being that, in *MINERVA*, updates to the environment (and to the agent mental state) that can bring about contradictory theories (if formulas are simply regarded as all having the same strength) are automatically dealt with by *LUPS* and *MDLP*. As noted above, updates are not simple assertions and retractions of formulas. In particular, a newly asserted rule is not equal in “power” to an old rule (in that case it would be a revision and not an update) and does not simply replace it. Katsumo and Mendelzon [8] study the updates for theories in classical propositional logic. Updating a belief base depends on the representation language. In particular, updates of non-monotonic knowledge bases, for example represented by logic programs under the stable models semantics, require a substantially different mechanism from those of classical propositional logic based representations, as discussed in [1]. It seems that, unlike the authors claim, the belief base used by *3APL* agents cannot be represented by an arbitrary logic language. In particular, their belief base cannot use the representational capabilities of default negation. *MINERVA* has been specifically designed to use logic programs with default negation both in the heads and in the bodies of clauses as the basic representation mechanism. Programs written in such a language are arranged according to several acyclic digraphs whose semantics is precisely characterized by *MDLP*. The use of *LUPS*, which has an imperative reading, provides for the mental state transitions by means of the necessary update operations. In *3APL* there is no notion of sub-agents concurrently changing the mental state of the agent with possibly contradictory information. In *MINERVA*, these contradictions are solved by the semantics of *MDLP* according to a specified hierarchy between those sub-agents. Like in *3APL*, action specifications in *LUPS* are static inasmuch as *LUPS* commands cannot be updated. We are currently extending *LUPS* programs with self-update capabilities to overcome this drawback, and to enrich the behaviour specification power of *LUPS*.

MetateM (and *Concurrent MetateM*), introduced by Fisher [5], is a programming language based on the notion of direct execution of temporal formulae, primarily used to specify reactive behaviours of agents. It shares similarities with the *LUPS* language inasmuch as both employ rules to represent a relation between the past and the future i.e., each rule in *MetateM* and in *LUPS* consists of conditions about the past (or present) and a conclusion about the future. While the use of temporal logics in *MetateM* allows for the specification of rather elaborate temporal conditions, something for which *LUPS* was not designed for, the underlying *MDLP* semantics of *LUPS* allows the specification of agents capable of being deployed in dynamic environments where the governing laws change over time. If we move to the more general class of logic programs with non-monotonic default negation both in the premises and conclusions of clauses, *LUPS* and *MDLP* directly provide an update semantics needed to resolve contradictions naturally

arising from conflicting rules acquired at different time instants, something apparently not possible in *MetateM*. This partially amounts to the distinction between updating theories represented in classical logic and those represented by non-monotonic logic programs (cf. [1]).

The *IMPACT* agent architecture, introduced by Subrahmanian et al. [15], provides a framework to build agents on top of heterogeneous sources of knowledge. To “agentize” such sources, the authors introduce the notion of agent program written over a language of so-called code-calls, which can be seen as encapsulations of whatever the source represents, and the actions the agent can execute. Such agent programs and their semantics resemble logic programs extended with deontic modalities. In *IMPACT*, at every state transition, the agent determines a set of actions to be executed, obeying some notion of deontic consistency, which simultaneously corresponds to a basic change in the agent’s mental state. *IMPACT* agents do not incorporate the rule based update mechanisms of *MLN \mathcal{E} RVA* although, in principle, it appears to be possible to specify a set of code-calls for the object knowledge base of a *MLN \mathcal{E} RVA* agent, accessing an *MDLP* meta-interpreter, and defining the actions (corresponding to the *LUPS* commands) that change the MDLP. In other words, one could agentize a *MLN \mathcal{E} RVA* agent to become an *IMPACT* agent.

7 Conclusions

We began this article by extolling the virtue and the promise of Computational Logic, most adeptly in the guise of Logic Programming, for adumbrating the issues of, and the solutions for, an integrated multi-agent architecture with a precisely defined declarative semantics. As we near the end of our protracted exposition, we are confident to have brought the point home. To wit, our architectural scaffolding rests on a sound and wide basis, whose core are the evolving updatable agents and their attending, and multi-dimensionally configured, knowledge bases.

The basic architecture affords us too, we purport to have shown, with the elasticity and resilience to further support a spate of crucial ancillary functionalities, in the form of additional specialized agents, via compositional, communication, and procedural mechanisms. The circum-exploration of the territories under purview has hardly started, but the means of locomotion appear to be already with us.

Logic Programming is, without a doubt, the privileged melting pot for the articulate integration of functionalities and techniques, pertaining to the design and mechanization of complex systems, addressing ever more demanding and sophisticated computational abilities. For instance, consider again those reasoning abilities mentioned in previous sections. Forthcoming rational agents, to be realistic, will require an admixture of any number of them to carrying out their tasks. No other computational paradigm affords us with the wherewithal for their coherent conceptual integration. And, all the while, the very vehicle that enables testing its specification, when not outright its very implementation.

Acknowledgements

We would like to thank the POCTI project FLUX. João Leite is partially supported by PRAXIS XXI scholarship no. BD/13514/97. Thanks are due to the anonymous referees for their remarks which helped us in improving the paper.

References

1. J. J. Alferes, J. A. Leite, L. M. Pereira, H. Przymusinska, and T. Przymusinski. Dynamic updates of non-monotonic knowledge bases. *Journal of Logic Programming*, 45(1-3):43–70, 2000. A short version titled *Dynamic Logic Programming* appeared in A. Cohn and L. Schubert (eds.), *KR'98*, Morgan Kaufmann.
2. J. J. Alferes, J. A. Leite, L. M. Pereira, and P. Quaresma. Planning as abductive updating. In D. Kitchin, editor, *Proceedings of the AISB'00 Symposium on AI Planning and Intelligent Agents*, pages 1–8. AISB, 2000.
3. J. J. Alferes, L. M. Pereira, H. Przymusinska, and T. Przymusinski. LUPS : A language for updating logic programs. *Artificial Intelligence*, 2001. To appear. A short version appeared in M. Gelfond, N. Leone and G. Pfeifer (eds.), *LPNMR-99*, LNAI 1730, Springer.
4. J. J. Alferes, L. M. Pereira, T. Przymusinski, H. Przymusinska, and P. Quaresma. An exercise with dynamic logic programming. In L. Garcia and M. Chiara Meo, editors, *Proceedings of the 2000 Joint Conference on Declarative Programming (AGP-00)*, 2000.
5. M. Fisher. A survey of concurrent METATEM: The language and its applications. In Dov M. Gabbay and Hans Jürgen Ohlbach, editors, *Proceedings of the 1st International Conference on Temporal Logic*, volume 827 of *LNAI*, pages 480–505, Berlin, 1994. Springer.
6. M. Gelfond and V. Lifschitz. Action languages. *Linkoping Electronic Articles in Computer and Information Science*, 3(16), 1998.
7. K. V. Hindriks, F. S. de Boer, W. van der Hoek, and J. C. Meyer. Formal semantics for an abstract agent programming language. In Munindar P. Singh, Anand Rao, and Michael J. Wooldridge, editors, *Proceedings of the 4th International Workshop on Agent Theories, Architectures, and Languages (ATAL-97)*, volume 1365 of *LNAI*, pages 215–230, Berlin, 1998. Springer.
8. Hirojumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In James Allen, Richard Fikes, and Erik Sandewall, editors, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, pages 387–394, San Mateo, CA, USA, 1991. Morgan Kaufmann Publishers.
9. J. A. Leite, J. J. Alferes, and L. M. Pereira. Multi-dimensional dynamic knowledge representation. In Thomas Eiter, Wolfgang Faber, and Mirosaw Truszczyński, editors, *Proceedings of the 6th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR-01)*, volume 2173 of *LNAI*, pages 365–378. Springer, 2001. A preliminary version appeared in F. Sadri and K. Satoh (Eds.) *Procs. of CLIMA'00*.
10. J. A. Leite, J. J. Alferes, and L. M. Pereira. On the use of multi-dimensional dynamic logic programming to represent societal agents' viewpoints. In P. Brazdil and A. Jorge, editors, *Progress in Artificial Intelligence, 10th Portuguese International Conference on Artificial Intelligence (EPIA-01)*, volume 2258 of *LNAI*. Springer, 2001.
11. I. Niemelä and P. Simons. Smodels: An implementation of the stable model and well-founded semantics for normal LP. In J. Dix, U. Furbach, and A. Nerode, editors, *LPNMR-97*, volume 1265 of *LNAI*, pages 420–429, Berlin, 1997. Springer.
12. 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>.

13. F. Sadri and F. Toni. Computational logic and multiagent systems: A roadmap, 1999. Available from <http://www.compulog.org>.
14. D. De Schreye, M. Hermenegildo, and L. M. Pereira. Paving the roadmaps: Enabling and integration technologies, 2000. Available from <http://www.compulog.org/net/Forum/Supportdocs.html>.
15. V. S. Subrahmanian, Piero Bonatti, Jürgen Dix, Thomas Eiter, Sarit Kraus, Fatma Ozcan, and Robert Ross. *Heterogeneous Agent Systems*. MIT Press/AAAI Press, Cambridge, MA, USA, 2000.
16. XSB-Prolog. The XSB logic programming system, version 2.0, 1999. Available at <http://www.cs.sunysb.edu/sbprolog>.