Dynamic Action Logic with Information States

(Extended Abstract)

Helmut Prendinger*
University of California, Irvine
Irvine, CA 92717–4555
Email: hprendin@benfranklin.hnet.uci.edu
Fax: (714) 653–2305

Abstract

We present a modal logic approach to reasoning about actions and agents where propositional dynamic logic is used to describe plans and reasoning about the agent's beliefs is performed in a modal system as well. In order to reason about changes of the agent's belief state as a result of executing actions, a logic of belief is *embedded* within dynamic logic. The proposed extension of (standard) dynamic logic introduces various action constructs that make the system applicable to problems related to the frame problem and the update problem. *Sense* actions are introduced as a special kind of 'belief-producing' actions, following the work of Moore [11] and Scherl and Levesque [17, 18]. Unlike [18], we allow for destructive sense actions and interfering actions of other agents. Although the work described in this paper considers only a single agent, the generalization of the framework to the specification and verification of multi-agent systems seems feasible due to the flexibility provided by modal logic. Formal properties of the resulting system of belief and action can be proved by a methodology introduced by Finger and Gabbay [3].

Keywords: action and change, knowledge and perception, belief update.

1 Introduction

In this paper[1], we propose a modal logic approach to describe actions and agents; we present a theory of actions within *dynamic* logic [5, 4, 13] and identify the mental state of the agent with a set of *beliefs*. More elaborate frameworks of mental categories can be found in [1, 14, 19] and [21].

Propositional dynamic logic (*PDL*) is the modal logic of actions [5] where formulas $[\alpha]A$ have the intended meaning "after every terminating execution of action α, formula A is true." Standard *PDL* can express compound actions such as sequential composition of actions, $\alpha;\beta$, (non-deterministic) choice between actions, $\alpha+\beta$, iteration, α^*, and test, $A?$. Our extension of *PDL*, called dynamic *action* logic (*PDAL*) [12, 13], provides additional action constructs tailored to deal with problems of reasoning about action and change.

Although there exists extensive literature dealing with logics of action and logics of belief (and knowledge) seperately, approaches combining both notions are almost lacking. A notable exception is a recent paper of Scherl and Levesque [18] (an earlier version appeared as [17]) that extends Moore's [11] work on knowledge and action. Moore [11] observes that there is a mutual dependence between knowledge and action: the agent needs certain knowledge in

---

[*] On leave from University of Salzburg, Austria.

[1] Revised and updated version of [12].

order to take reasonable actions, on the other hand, the execution of actions may modify the agent's knowledge. Scherl and Levesque [18] call actions that affect the agent's mental state *knowledge-producing* actions. For instance, performing an acidity test by means of litmus paper [11] 'produces knowledge' in the sense that the agent may thereby determine whether a solution is acid or not. In [18] knowledge is treated as a relation (fluent) that satisfies appropriate first-order conditions. We take the weaker notion of belief to describe agents, whereby a logic of belief is *embedded* within dynamic logic, following a methodology developed by Finger and Gabbay [3].

The paper is organized as follows. The next section introduces dynamic action logic. In Section 3, we show how dynamic logic (as the *global* logic) can be 'externally' applied to a modal logic of belief (as the *local* logic). In Section 4, our system is compared to the approach of Scherl and Levesque [18]. Finally, in Section 5, we discuss related work and hint at topics for further research.

2 Dynamic Action Logic

Dynamic action logic is a general framework to reason about action and plans that covers certain aspects of temporal logic [5], process logic [7], and even deontic logic [2]. In dynamic logic, an action α is a binary relation R_α on a set $S = \{s, t, ...\}$ of possible total states of the world; that is, a set of ordered pairs $\langle s, t \rangle$ of states where s is the initial state of some execution of α and t is the final state. Since a final state in general not uniquely determined the modeling of actions is (semantically) non-deterministic: first, an action can always be performed in different ways, and second, multiple agents may interfere and thus contribute to state-changes as well.

In order to deal with problems related to action and change, *PDAL* adds various action constructs to (standard) *PDL*.

Test actions. The test action $A?$ (stationary test, verification of A) is complemented with two non-stationary versions, τA and $\tau^\circ A$, that allow for state-changes while a test is performed. The latter one is called *safe* non-stationary test since the operation $\tau^\circ A$ always terminates if A holds initially.[2]

Update actions. Updates are 'change-recording' actions that are used when we want to bring the database up to date because the world described by it has changed [8]. In *PDAL* update actions $\diamond A$ ("updating with A") cannot be confused with revisions—another form of database adaption—since only the former adaption involves a state-transition. The relation corresponding to $\diamond A$ does not restrict $R_{\diamond A}$ such that only minimal transitions $\langle s, t \rangle \in R_{\diamond A}$ are considered where t differs from s 'as little as possible'. Although our concept of update is rather simple, we obtain appropriate solutions for the 'book-and-magazine problem' [8] and for the problem of disjunctive updates [20] (see [12] for details).

The 'any' action. The 'any' action ν is used to make assertions about facts that are intended to be invariant over every performance of every action, for instance, in the formulation of *domain constraints.*

Preservation formulas. Formulas **cpres**(α, A) (chronological preservation) say that if a fact A holds when the execution of a (typically compound) action α is initiated then A is true at all intermediate states of the performance of α. For convenience, formulas **tpres**(α, A) (terminal preservation) are introduced as a special case. By means of chronological preservation formulas we may encode *plan constraints*, that is, constraints that ensure the 'clean' behavior of a plan by imposing restrictions on accessible states. In *temporal plan theory* [9] this is achieved

---

[2] Linder *et al.* [22] prefer the notion *confirmations* for dynamic logic tests. Below, we will show how *PDAL* tests can be not only 'truthful' but also 'informative'; yet, tests here do not meet the requirement of 'minimal change' (see [22] for an explanation of these criteria).

by stating temporal (liveness and safety) properties. For instance, the property response states that always during the performance of a plan $\alpha_1, ..., \alpha_n$ where A holds, there exists a state which is attainable via β and satisfies B. Within *PDAL*, response is expressed as $\mathbf{bound}(\alpha_1, ..., \alpha_n, A \rightarrow < \beta > B)$ where $\mathbf{bound}(\alpha, A) \stackrel{\text{def}}{=} A \wedge \mathbf{cpres}(\alpha, A)$. For a dynamic rendering of other temporal properties, see [13].

We start with the description of the *language* of *PDAL* [12]. Let $\mathcal{P}_0 = \{p, q, r, ...\}$ be a denumerably infinite set of propositional variables and $\mathcal{A}_0 = \{a, b, c, ...\}$ be a denumerably infinite set of action variables. A, B,... and α, β,... are used to denote arbitrary formulas and arbitrary terms, respectively. The vocabulary of $\mathcal{L}(PDL)$ is enriched by the action constant ν and the (unary) operators τ, τ° and $\diamond$; moreover the operators **tpres** and **cpres** are added to the vocabulary of $\mathcal{L}(PDL)$.

$$A ::= p \mid \neg A \mid A_1 \vee A_2 \mid [\alpha]A \mid \mathbf{tpres}(\alpha, A) \mid \mathbf{cpres}(\alpha, A)$$
$$\alpha ::= a \mid \nu \mid \alpha_1; \alpha_2 \mid \alpha_1 + \alpha_2 \mid \alpha^* \mid A? \mid \tau A \mid \tau^\circ A \mid \diamond A$$

(where $p \in \mathcal{P}_0$, $A \in \mathcal{L}(PDAL)$, and $\alpha \in \mathcal{A}$. We assume the usual definitions of $\wedge$, $\rightarrow$, $\leftrightarrow$, $\top$ (*verum*), and $\bot$ (*falsum*); $<\alpha> A \stackrel{\text{def}}{=} \neg[\alpha]\neg A$.

We turn to the *semantics* of *PDAL*. An *agent frame* $\mathcal{F}$ is a structure $\mathcal{F} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}_{at}\}, T, T^\bullet, U, \mu\rangle$ such that S is a nonempty set of states, $\{R_\alpha : \alpha \in \mathcal{A}_{at}\}$ is the set of relations corresponding to *atomic* action terms $\mathcal{A}_{at} \stackrel{\text{def}}{=} \mathcal{A}_0 \cup \{\nu\}$. The functions T ($T^\bullet$) and U mirror (safe) non-stationary test and update on the level of frames, respectively; $T : \text{Pow}(S) \rightarrow \text{Pow}(S \times S)$ such that for all $X \in \text{Pow}(S)$, $TX \subseteq \{\langle s,t\rangle : s \in X, t \in X\}$; $T^\bullet$ is like T except that it meets the further condition that for each $s \in X$, $T^\bullet X$ contains a pair $\langle s,t\rangle$ (for $t \in S$); $U : \text{Pow}(S) \rightarrow \text{Pow}(S \times S)$ such that for all $X, Y \in \text{Pow}(S)$, $UX \subseteq \{\langle s,t\rangle : t \in X\}$ and $U(X \cup Y) = UX \cup UY$. A *standard* agent frame is an agent frame $\mathcal{F}$ where R_ν is universal. Hence $\mathcal{F}$ is R_ν-generated and $R_\alpha \subseteq R_\nu$ for all $\alpha \in \mathcal{A}$. μ is called the *smoothness* function. It is an definitional extension of standard agent frames, a function $\mu : \{R_\alpha : \alpha \in \mathcal{A}_{at}\} \times \text{Pow}(S) \rightarrow \text{Pow}(S)$ such that for all $s \in S$, $X \in \text{Pow}(S)$: $s \in \mu(R_\alpha, X)$ iff $s \in X \Rightarrow (\{t : \langle s,t\rangle \in R_\alpha\} \subseteq X)$. A standard agent *model* $\mathcal{M}$ based on $\mathcal{F}$ is a structure $\mathcal{M} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}\}, T, T^\bullet, U, \mu_t, \mu_c, v\rangle$ where R_α is inductively defined for compound action terms $\alpha \in \mathcal{A}$: $R_{\alpha;\beta} = R_\alpha \circ R_\beta$; $R_{\alpha+\beta} = R_\alpha \cup R_\beta$; $R_{\alpha^*} = (R_\alpha)^\star$; and $R_{A?} = \{\langle s,s\rangle : \mathcal{M}, s \models A\}$. In addition to the conditions for *PDL* we have $R_{\tau A} = T\|A\|^{\mathcal{M}}$, $R_{\tau^\circ A} = T^\bullet\|A\|^{\mathcal{M}}$, and $R_{\diamond A} = U\|A\|^{\mathcal{M}}$; $\|A\|^{\mathcal{M}}$ is defined as $\{s \in S : \mathcal{M}, s \models A\}$. Moreover, the following functions are introduced:

1. μ_t is the function $\mu_t : \{R_\alpha : \alpha \in \mathcal{A}\} \times \text{Pow}(S) \rightarrow \text{Pow}(S)$ such that
$$s \in \mu_t(R_\alpha, X) \quad \text{iff} \quad s \in X \Rightarrow (\{t : \langle s,t\rangle \in R_\alpha\} \subseteq X).$$
Note that $\mu_t(R_\alpha, X) = \mu(R_\alpha, X)$ for *atomic* action terms $\alpha \in \mathcal{A}_{at}$.

2. μ_c is the function $\mu_c : \{R_\alpha : \alpha \in \mathcal{A}\} \times \text{Pow}(S) \rightarrow \text{Pow}(S)$ such that
$$\mu_c(R_\alpha, X) = \mu_t(R_\alpha, X)$$
for all *elementary* action terms $\alpha \in \mathcal{A}_{el}$ where $\mathcal{A}_{el} \stackrel{\text{def}}{=} \mathcal{A}_{at} \cup \{A?, \tau A, \tau^\circ A, \diamond A : A \in \mathcal{L}(PDAL)\}$. For compound action terms, μ_c is inductively defined appropriately, for instance, $\mu_c(R_{\alpha;\beta}, X) = \mu_c(R_\alpha, X) \cap \{s : \{t : \langle s,t\rangle \in R_\alpha\} \subseteq \mu_c(R_\beta, X)\}$.

Finally, v is a function $v : \mathcal{P}_0 \rightarrow \text{Pow}(S)$.

Truth of a formula at a state s in a model $\mathcal{M}$ is inductively defined as follows: $\mathcal{M}, s \models p$ iff $s \in v(p)$ (for $p \in \mathcal{P}_0$); the propositional connectives are treated as usual; $\mathcal{M}, s \models [\alpha]A$ iff

$\forall t\,(\langle s,t\rangle \in R_\alpha \Rightarrow \mathcal{M},t \models A)$; $\mathcal{M},s \models \mathbf{tpres}(\alpha, A)$ iff $s \in \mu_t(R_\alpha, \|A\|^{\mathcal{M}})$; $\mathcal{M},s \models \mathbf{cpres}(\alpha, A)$ iff $s \in \mu_c(R_\alpha, \|A\|^{\mathcal{M}})$.

As an *axiomatic system PDAL* contains all instances of the following axiom schemes of *PDL*

Taut	all classical propositional tautologies
$\mathbf{K}_{[\alpha]}$	$[\alpha](A \rightarrow B) \rightarrow ([\alpha]A \rightarrow [\alpha]B)$
Comp	$[\alpha;\beta]A \leftrightarrow [\alpha][\beta]A$
Union	$[\alpha+\beta]A \leftrightarrow ([\alpha]A \wedge [\beta]A)$
Mix	$[\alpha^*]A \rightarrow (A \wedge [\alpha][\alpha^*]A)$
Ind	$[\alpha^*](A \rightarrow [\alpha]A) \rightarrow (A \rightarrow [\alpha^*]A)$
sTest	$[A?]B \leftrightarrow (A \rightarrow B)$

and additional schemes for (safe) non-stationary test, update, the 'any' action, and the notions of terminal and chronological preservation.

nTest.1	$< \tau A > \top \rightarrow A$
nTest.2	$[\tau A]A$
snTest.1	$< \tau^\circ A > \top \leftrightarrow A$
snTest.2	$[\tau^\circ A]A$
Update.1	$[\diamond A]A$
Update.2	$[\diamond(A \vee B)]C \leftrightarrow ([\diamond A]C \wedge [\diamond B]C)$
Any.1	$[\nu]A \rightarrow A$
Any.2	$[\nu]A \rightarrow [\nu][\nu]A$
Any.3	$< \nu > [\nu]A \rightarrow A$
Any.4	$[\nu]A \rightarrow [\alpha]A$
tPres	$\mathbf{tpres}(\alpha, A) \leftrightarrow (A \rightarrow [\alpha]A)$
cPres	$\mathbf{cpres}(\alpha, A) \leftrightarrow \left(A \rightarrow \bigwedge_{\beta \in I(\alpha)} [\beta]A\right)$

and which is closed under the following rules of inference:

MP	from $\vdash A$ and $\vdash A \rightarrow B$ infer $\vdash B$
NEC	from $\vdash A$ infer $\vdash [\alpha]A$

The last mentioned axiom scheme makes use of the concept of *initial strings* of R_α where I is a function $I : \mathcal{A} \rightarrow \mathrm{Pow}(\mathcal{A})$ that assigns to each action term $\alpha \in \mathcal{A}$ the set of all action terms that denote initial strings of R_α.

The following theorem can be proved by a canonical model construction and the filtration technique [5] (for details, see [12, 13]).

THEOREM 2.1
(1) PDAL has the finite model property; (2) PDAL is correct and (weakly) complete with respect to the class of finite standard agent models; and (3) PDAL is decidable.

3 Embedding a Logic of Belief in *PDAL*

When the epistemic part of the logic is embedded inside dynamic logic we get the following picture as illustrated in Fig. 1: the logic of belief is the *local* logic while dynamic logic is the *global* logic that describes the evolution of the agent's beliefs as a result of the execution of actions. Note that states $s \in S$ which are primitives in the modeling of actions in dynamic logic are now better termed *information states* (or structured states).

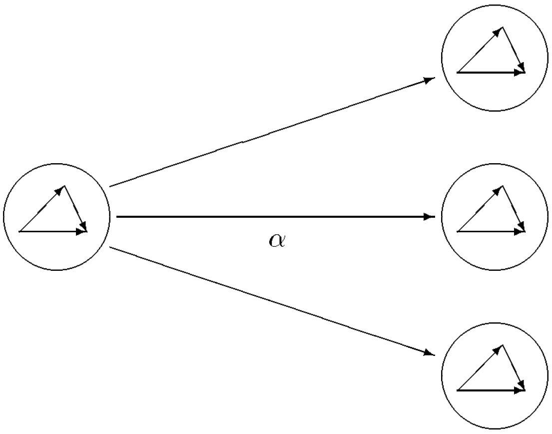

Figure 1: *Relation* R_α *on* information *states corresponding to action* α*.*

Concerning the logic of belief B, we decided for the popular modal system **KD45** (or *weak* **S5**) that is considered an appropriate logic to capture the propositional attitude of *belief* [6]. **KD45** contains both the 'positive introspection' axiom **4** and the 'negative introspection' axiom **5**; moreover, by axiom **D**, the agent is assumed to have no inconsistent beliefs, although the agent is not required to only believe true facts (a distinguished feature of belief as opposed to knowledge, see [6]). Semantically, we assume as given a set $W = \{w, w', ...\}$ of (possible) *worlds*. The relation $R \subseteq W \times W$ picks out those worlds $\{w' : \langle w, w' \rangle \in R\}$ which are possible (consistent) wrt. what the agent believes in world w. An agent believes A, written $\mathbf{B}A$, if A is true in all worlds the agent considers possible relative to the world it inhabits.

We proceed with the more formal development which closely resembles the approach in [3] but adds a *dynamic* dimension to the logic of belief B. Let $\mathcal{P}_0$ be as above.

$$A ::= p \mid \neg A \mid A_1 \vee A_2 \mid \mathbf{B}A$$

where $p \in \mathcal{P}_0$ and $A \in \mathcal{L}(B)$. Let $ML(B)$ denote the set of *monolithic formulas* of $\mathcal{L}(B)$ [3], that is, formulas A that are not built up by any of the Boolean connectives.

$$A ::= D \mid \neg A \mid A_1 \vee A_2 \mid [\alpha]A \mid \mathbf{tpres}(\alpha, A) \mid \mathbf{cpres}(\alpha, A)$$

such that $D \in ML(B)$, $A \in \mathcal{L}(PDAL[B])$ and $\alpha \in \mathcal{A}$ (defined analogously as above). Note that, for instance, $\mathbf{B}[\alpha]p$ is not in the language.

A *belief model* $\mathcal{M}_B$ is a structure $\mathcal{M}_B = \langle W, R, l, w_c \rangle$ such that W is a set of worlds and R is a binary relation on W. We take R to be a serial, transitive and Euclidean relation. l is a function $l : \mathcal{P}_0 \to \mathrm{Pow}(W)$. Finally, $w_c \in W$ is the current world from which observations are made. The truth condition for belief formulas is $\mathcal{M}_B, w \models \mathbf{B}A$ iff $\forall t\ (\langle w, w' \rangle \in R \Rightarrow \mathcal{M}, w' \models A)$. The additional parameter w_c guarantees that either $\mathcal{M}_B \models A$ or $\mathcal{M}_B \models \neg A$ for all $A \in \mathcal{L}(B)$. Now consider a frame $\mathcal{F}_{PDAL} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}_{at}\} \rangle$ of *PDAL* and a function

$g : S \to \mathbf{M}_B$ where $\mathbf{M}_B$ is a class of belief models. By a model of $PDAL[B]$ we mean a structure $\mathcal{M}_{PDAL[B]} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}_{at}\}, g\rangle$. The truth conditions are inductively defined as usual; for $A \in ML(B)$ we have

$$\mathcal{M}_{PDL[B]}, s \models A \text{ iff } g(s) = \mathcal{M}_B \text{ and } \mathcal{M}_B \models A.$$

B contains all instances of the following axiom schemes and is closed under the rules of inference **MP** and $\mathbf{NEC}_B$ (from $\vdash_B A$ infer $\vdash_B \mathbf{B}A$).

Taut	all classical propositional tautologies
$\mathbf{K}_B$	$\mathbf{B}(A \to B) \to (\mathbf{B}A \to \mathbf{B}B)$
D	$\neg\mathbf{B}\bot$
4	$\mathbf{B}A \to \mathbf{B}\mathbf{B}A$
5	$\neg\mathbf{B}A \to \mathbf{B}\neg\mathbf{B}A$

$PDAL[B]$ contains all instances of axiom schemes and rules of $PDAL$ and the additional rule of inference [3].

PRESERVE from $\vdash_B A$ infer $\vdash_{PDL[B]} A$ (for all $A \in \mathcal{L}(B)$).

It is well known that B is determined by the class of belief models and is decidable [6]. The following statement can be proved by means of the techniques developed in [3].

Theorem 3.1
(1) $PDAL[B]$ is correct and (weakly) complete with respect to the class of finite standard agent models and (2) $PDAL[B]$ is decidable.

4 Comparision with Scherl and Levesque

Scherl and Levesque [18] utilize the situation calculus to reason about action and change. They combine the notions of knowledge and action by conceiving of knowledge as a (relational) fluent that can be affected by a special kind of actions, so-called *knowledge-producing* actions such as sense actions. A solution to the frame problem is obtained by extending Reiter's [15] solution for ordinary (non-knowledge-producing) actions. We introduce *sense* actions σP ($P \in \mathcal{L}(PL)$) to $PDAL$ where $R_{\sigma P} \subseteq \{\langle s, t\rangle : s \in \|P\| \Rightarrow t \in \|\mathbf{B}P\|\}$ and add

Sense $P \to [\sigma P]\mathbf{B}P \qquad P \in \mathcal{L}(PL)$

as an axiom scheme. Although the approach of Scherl and Levesque and our approach are similar in spirit, important differences exist. This will be shown by means of the 'acidity test' example (see also [11]).

Example 4.1 (acidity test)
Consider an agent that is supposed to find out about the acidity of a solution by means of litmus paper. The point is that the agent is able to *deduce* the unobservable precondition, whether the solution is acid or alkaline, from the observable result, whether the paper turned red or blue. We write `ACID` for "the solution is acid" and `RED` for "the litmus paper is red". Assume the following three statements as premises.

(1) $\texttt{ACID} \to [\tau\texttt{ACID}]\texttt{RED}$

(2) $[\tau\texttt{ACID};\sigma\texttt{RED}]\mathbf{B}(\texttt{RED} \to \texttt{ACID})$

(3) $\texttt{ACID}$

For instance, the second premise says that after verifying that the solution is acid and then sensing that the paper turned red, the agent believes that there is a connection between the color the paper and the solution being acid. With the proof theory of *PDAL[B]* we get

$$(1),(2),(3) \vdash [\tau\mathtt{ACID};\sigma\mathtt{RED}]\mathbf{B}\mathtt{ACID}$$

Note that we cannot deduce $[\tau\mathtt{ACID};\sigma\mathtt{RED}]\mathtt{ACID}$, so destructive sense actions are covered as well. For non-destructive sensing we might add the condition $[\nu]\mathbf{bound}(\sigma\mathtt{RED},\mathtt{ACID})$. In this case, the agent will even *know* that the solution is acid, by $\mathbf{K}A \overset{\text{def}}{=} \mathbf{B}A \wedge A$, where $\mathbf{K}$ is the knowledge operator of [18]. On the other hand, Scherl and Levesque prove (by means of successor state axioms for each fluent) that knowledge-producing actions never change the (objective) state of the world. As a consequence of this theorem ([18, Th. 1]) belief-producing actions cannot be easily incorporated into their framework, since for those actions it is possible to have, for instance, $\mathbf{B}\mathtt{ACID} \wedge \neg\mathtt{ACID}$ as a resulting state. Although a significant reduction of frame axioms can be achieved by utilizing the **cpres** operator, our approach lacks the frame axiom reducing power of the successor state axioms. Whenever some fact is to be preserved over a sequence of actions, we have to state that explicitly. Take, for instance, the notion of *memory* [18]: if a fact P is known at some state then it will be known at all successor states, given that no action changes the truth-value of P. Let $\alpha_1;...;\alpha_n$ be an arbitrary plan such that $<\alpha_i> \top$ $(1 \leq i \leq n)$. Then memory in *PDAL* can be expressed as follows.

$$[\nu]\mathbf{cpres}(\alpha_1;...;\alpha_n,[\tau^\circ\neg P]\bot \to \mathbf{B}P) \qquad (\mathit{memory})$$

Reiter [16] claims that our approach does not count as a solution to the frame problem (for some arguments in favor of our 'solution', see [13]). In fact, De Giacomo and Lenzerini [4] adopted Reiter's solution to the frame problem for dynamic logic.

5 Conclusion

Admittedly, our theory of agents is still in a very preliminary state. Within our setting, the mental state of an agent consists of a single component, belief. In addition to belief and time, Shoham [19] takes into account *obligation* and *capability* as modalities, and Cohen and Levesque [1] *goals* and *intention*. Eventually, we want to relax the restriction that mental modalities are embedded within dynamic logic. This is not even true for Scherl and Levesque's [18] theory of knowledge and action that has the equivalents of both $[\alpha]\mathbf{K}A$ and $\mathbf{K}[\alpha]A$ in the logic. In the case of Example 4.1 (acidity test) we could provide a solution without admitting formulas of format $\mathbf{B}[\alpha]A$.

The most urgent research objective is to extend *PDAL* to the *multi-agent* case [23] and to integrate inter-agent *communication* [19, 10]. For instance, we want to consider 'liveness' properties such as

$$[\nu](\mathbf{request}(\alpha) \wedge \mathbf{B}\mathtt{INTEGRITY} \to <\nu^*> \mathbf{perform}(\alpha)),$$

when the agent is requested to perform the action α and believes that the sender is integer, then the agent will eventually perform α; or

$$[\nu](\mathbf{receive}(P) \wedge \mathbf{B}\mathtt{INTEGRITY} \to <\nu^*> \mathbf{B}P),$$

the agent will eventually believe a proposition P, given it is informed about P and believes in the sender's integrity. Due to its generality and flexibility, we expect our modal logic approach to be especially apt to deal with the specification and verification of multi-agent systems. An initial treatment of communicative actions is given by the introduction of sense actions.

References

[1] Philip R. Cohen and Hector J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[2] F. P. M. Dignum and J.-J. Ch. Meyer. Negations of transactions and their use in the specification of dynamic and deontic integrity constraints. In M. Z. Kwiatkowska, M. W. Shields, and R. M. Thomas, editors, *Proceedings of the International BCS-FACS Workshop on Semantics for Concurrency*, pages 61–80, University of Leicester, UK, 1990.

[3] Marcelo Finger and Dov M. Gabbay. Adding a temporal dimension to a logic system. *Journal of Logic, Language and Information*, 1:203–233, 1992.

[4] Giuseppe De Giacomo and Maurizio Lenzerini. PDL-based framework for reasoning about actions. In *Proceedings 4th Congress of the Italian Association for Artificial Intelligence (AI*LA-95)*, LNAI 992, pages 103–114, 1995.

[5] Robert Goldblatt. *Logics of Time and Computation.* Center for the Study of Language and Information Lecture Notes No. 7. Leland Stanford University, Stanford, second, revised and expanded edition, 1992.

[6] Joseph Y. Halpern and Yoram Moses. A guide to the modal logics of knowledge and belief: preliminary draft. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 480–490, 1985.

[7] David Harel, Dexter Kozen, and Rohit Parikh. Process logic: expressiveness, decidability, completeness. *Journal of Computer and System Sciences*, 25:144–170, 1982.

[8] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In James Allen, Richard Fikes, and Erik Sandevall, editors, *Proceedings of the Second International Conference On Principles of Knowledge Representation and Reasoning*, pages 387–394, Cambridge, Massachusetts, 1991.

[9] Zohar Manna, Massimo Paltrinieri, and Richard Waldinger. Temporal planning. In *Workshop on Reasoning about Action and Change at IJCAI-93*, pages 37–41, Chambéry, France, 1993.

[10] Daniel Marcu, Yves Lespérance, Hector J. Levesque, Fangzhen Lin, Raymond Reiter, and Richard B. Scherl. Distributed software agents and communication in the situation calculus. In *Proceedings Intelligent Computer Communication Conference (ICC-95)*, 1995.

[11] Robert C. Moore. A formal theory of knowledge and action. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal Theories of the Commonsense World*, pages 317–358. Ablex Publishing Corporation, Norwood, New Jersey, 1985.

[12] Helmut Prendinger. Modal logic for modelling actions and agents. In *Proceedings International Conference on Formal and Applied Practical Reasoning (FAPR-96)*, LNAI 1085, pages 525–539, 1996.

[13] Helmut Prendinger and Gerhard Schurz. Reasoning about action and change. a dynamic logic approach. *Journal of Logic, Language, and Information*, 5:209–245, 1996.

[14] Anand S. Rao and Michael P. Georgeff. Modeling rational agents within a BDI-architecture. In James Allen, Richard Fikes, and Erik Sandevall, editors, *Proceedings of the Second International Conference On Principles of Knowledge Representation and Reasoning*, pages 473–484, Cambridge, Massachusetts, 1991.

[15] Raymond Reiter. The frame problem in the situation calculus: a simple solution (sometimes) and a completeness result for goal regression. In Vladimir Lifschitz, editor, *Artificial Intelligence and Mathematical Theory of Computation. Papers in Honor of John McCarthy*, pages 359–380. Academic Press, Inc., Boston, New York, 1991.

[16] Raymond Reiter. Personal communication, 1996.

[17] Richard B. Scherl and Hector J. Levesque. The frame