Reasoning about Action and Change

A Dynamic Logic Approach

Helmut Prendinger
Int'l Forschungszentrum, Mönchsberg 2a, A-5020 Salzburg, Austria,
Email: Helmut.Prendinger@mh.sbg.ac.at

Gerhard Schurz
Inst. für Philosophie, Abtlg. Logik und Wissenschaftstheorie, Franziskanergasse 1, A-5020 Salzburg, Austria,
Email: Gerhard.Schurz@mh.sbg.ac.at

1. Introduction

Agents situated in real-world environments need to reason about action and change. In order to achieve certain goals agents have to develop plans and keep track of changes, either caused by themselves or by other agents present in the domain. Problems in reasoning about action and change are either aspects of or related to the "frame problem". It was first posed by McCarthy & Hayes (1969) in attempts to formalize consequences of actions *and those facts that remain unaffected by actions.* The problem consists of finding an *efficient* way to formulate axioms (so-called "frame axioms") which state what remains the same after the performance of an action. Conceived as a *combinatorial* problem, a solution to the frame problem has to avoid writing down frame axioms for each fact-action pair (see Reiter, 1980; Georgeff, 1987a). For, even toy domains like a blocks world scenario require a vast number of frame-axioms. For instance, moving a block changes its position but not its color, shape, and so on.

The canonical example to illustrate the frame problem is the "Yale Shooting Problem" ($\mathcal{YSP}$) due to Hanks & McDermott (1987). The shooting scenario is as follows: Mary loads a gun, waits and then shoots at Fred. Here the problem is to derive that Fred is dead as a consequence of Mary's actions. Among other things one has to guarantee that the gun is still loaded after waiting. This is a candidate for a frame axiom: "if the gun is loaded when Mary begins waiting then the gun is still loaded when Mary stops waiting." The frame axiom says that the gun's state of being loaded is unaffected by the action of waiting. Notice that we use "unaffected" here to suggest a formulation of the frame problem for *single-agent* domains. On the other hand, "unchanging" is more appropriate in case of environments with *multiple* agents. It is easy to see that the combinatorial frame problem turns up in both cases.

Unfortunately, most work on the frame problem (see Lifschitz, 1987; Shoham, 1988) is conspicuously vague on the distinction between single-agent and multiagent settings and therefore casual when it comes to constructing models where the ‘right’ things happen (that the gun remains loaded, for instance). As a rule, some kind of nonmonotonic logic is utilized to tackle the problem, usually some sort of *minimization* strategy (see Makinson, 1993). That is, ‘abnormalities’ are minimized such that facts remain unaffected (not changing) in the ‘normal’ course of events (where interference is absent). Global minimization seems inappropriate in reasonably complex environments, especially when other agents may contribute to changes in the domain. To overcome this problem, more selective minimization strategies have been developed, like *scoped* minimization (see Etherington *et al.* ,1991). However, it was just this ‘global nature’ that originally made nonmonotonic approaches so appealing. This is best captured in the so-called *persistence assumption* which states that all facts *usually* persist to hold after the performance of all actions, if not stated otherwise. To the best of our knowledge Georgeff (1987a, p. 73) first noticed that “it is not sensible to try minimizing changes in *all* world properties when told, for example, that a shooting event has occurred.” Of course, it is often not possible to predict what actions other agents will perform. But the formalism should not be overcommitted to the restriction that those actions have not occurred. Accordingly, Georgeff (1987a) identifies the *overcommitment* problem as a second subproblem of the frame problem. A solution to the overcommitment problem requires that most facts are allowed to vary freely, while only those facts which are prerequisites or consequences of the agent’s actions are maintained.

To clarify terminology: like Georgeff (1987a) we use the term “frame problem” in a generalized sense, which consists of several subproblems. One subproblem is the combinatorial frame problem, another one is the overcommitment problem. By now, many other aspects of the frame problem have been identified (see Georgeff, 1987a, 1987b; Shoham, 1988; Ginsberg & Smith, 1988a, 1988b). The *extended prediction* problem concerns the problem of making inferences over extended periods of time. A solution to the extended prediction problem should provide means to state that, for example, under ‘normal’ circumstances the gun remains loaded until Mary fired six times (assume a six-shooter). Note that the combinatorial frame problem may be considered as a subproblem of the extended prediction problem. The *qualification* problem concerns the conditions under which an action is executable. For the shooting action to be executable in the intended way we need to verify—in addition to explicit preconditions (that the gun is loaded)—an immense number of implicit qualifications: that the gun has a firing pin, that there is no magnet present which distracts the bullet, and so on. But if our formalization takes *everything* into account that could block the performance of an action, it will be without practical use. In nonmonotonic logics, both the extended prediction problem and the qualification problem are handled by a strategy which minimizes abnormalities. Observe that a solution to the overcommitment problem is undermined, because thereby independent changes of other facts, for instance, due to activities of other agents are *also* minimized, contrary to observation of a complex world.[1] Again, our criticism only applies to global minimization strategies. On the other hand, we will show that systems which scope reasoning run into similar problems as we do, for we do not accept the persistence assumption.

The point we will make in Section 5 is this: our monotonic approach to the frame problem and (unscoped) nonmonotonic approaches (based on some minimization strategy) are *complementary* in the sense that it is not possible to simultaneously minimize the effort to solve the extended prediction problem (which subsumes the combinatorial problem) *and* give a satisfactory solution to the overcommitment problem.

In this paper, we pursue a monotonic approach to the frame problem and concentrate on the combinatorial problem and the overcommitment problem. We will propose a solution within the framework of *propositional dynamic logic (PDL)*—the modal logic of actions and of computer programs (see Pratt, 1976, 1980; Segerberg 1980; Harel, 1984). It is based on the idea of associating an operator $[\alpha]$ with each action α, the brackets being reminiscent of the box operator $\Box$ of ordinary modal logic (see Hughes & Cresswell, 1984). The reading of a formula $[\alpha]A$ is "after every terminating (halting) execution of α, A is true." *PDL* provides a powerful language for describing *compound* actions such as sequential composition of actions α and β, written $\alpha;\beta$, (non-deterministic) choice between α and β, written $\alpha + \beta$, and (non-deterministic) iteration of α, written α^*. Moreover, test, written $A?$, and 'doing nothing', denoted by λ, are considered as actions.

[1] Remember that early nonmonotonic logics suffered from the overcommitment problem even in single-agent domains. As well known, this gave rise to the 'multiple extensions' problem in temporal projection (see Hanks & McDermott, 1987.

The following readings are standard (see Segerberg, 1980).

$\alpha;\beta$	the action consisting of doing first α and then β
$\alpha+\beta$	the action consisting of doing α or β non-deterministically
α^*	the action consisting of doing α some finite number of times
$A?$	the action consisting of verifying that A holds
λ	'doing nothing' (stationary waiting)

To our knowledge, nobody so far tried to cope with the combinatorial problem and the overcommitment problem in a dynamic logic framework. A notable exception is the paper of Stephan & Biundo (1993) (see also Kautz, 1982; Morreau, 1992). In the language of first-order dynamic logic they propose frame assertions of format $A \supset [\alpha]A$ (where α is allowed to be a compound action). We extend this idea to the *intermediate* states of a plan (typically a sequential composition of actions). Loosely speaking, formulas of type $\mathbf{cpres}(\alpha, A)$ are introduced which are intended to mean that fact A is true at all intermediate states of the execution of (compound) action α. In this way, we may drastically reduce the number of frame axioms. For simplicity, we only consider the propositional case. It will turn out that our resulting logic is a natural alternative to temporal logic which is the major formalism in the planning literature (see Pnueli, 1981; Manna *et al.*, 1993).

The rest of the paper is organized as follows. In Section 2, *PDL* is introduced together with a reminder of the basic steps of its completeness proof. Section 3 shows that *PDL* may be modified in order to deal with multiagent domains in a *uniform* way. The resulting logic, *non-stationary PDL* (*NPDL*), also includes a powerful action concept, the 'any' action. We prove completeness and decidability for *NPDL*. In Section 4, we introduce the concept of *chronological* preservation. By means of this concept we are in a position to reason about the intermediate states of plan execution. Moreover, a border case of chronological preservation, called *terminal* preservation, is proposed. We show that completeness and decidability results carry over to the resulting *extension* of *NPDL*. Section 5 is concerned with *applications* of our framework. First, we demonstrate how the combinatorial problem and the overcommitment problem are solved in our framework. Thereby an example which is more elaborate than the $\mathcal{YSP}$ serves as a testing ground. Here we also address the ramification problem, which is a problem related to the frame problem. Second, several notions important to *planning* are formalized, concerning domain constraints and plan constraints. An example from the manufacturing domain will illustrate these notions. We also show how to encode several notions from *temporal plan theory* (Manna *et al.*, 1993) in our logic. Finally, we compare our solution of the frame problem to other monotonic and (scoped) nonmonotonic solutions.

2. Propositional dynamic logic

2.1. Basic concepts

We assume as given a set $S = \{s, t, ..., s', s'', ...\}$ of possible total states of the world. A *proposition* can be identified with the set of states in which it is true. An *action* α is a binary relation R_α on S, 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. Of course, a final state need not be uniquely determined. Thus the modeling of actions is *semantically non-deterministic*.[2] As mentioned above, expressions of the form $[\alpha]A$ have the informal meaning "whenever the execution of α halts, A is true on termination." Non-terminating 'executions' are said to *fail*. Since we conceive of actions as proceeding from one state to another in a discrete fashion, 'executions' with no final state simply don't count in dynamic logic. In fact, they are no executions at all. By $< \alpha >$*verum* we can express that there exists a terminating execution of α.

Note that choice and iteration in dynamic logic are non-deterministic. This kind of non-determinism may be called "non-determinism with respect to control-flow" (see Harel, 1987). We will call non-determinism of control-flow *procedural* non-determinism. Since dynamic logic provides no machinery to give priority of executing one action over executing the other, choice is non-deterministic. In case of iteration α^*, an action α is performed some non-deterministically chosen finite number $n \geq 0$ of times. We may summarize the distinction between semantical and procedural non-determinism as follows. *Semantical* non-determinism concerns the fact that a (halting) execution of an action α may end up in different states. One reason for this is that there exists always a variety of different ways to perform an action. In Section 3 a second reason is given: other agents may interfere and thus contribute to state-changes even if the agent is passive. It will be seen from the semantic modelling of actions (see below) that a 'multiagent' reading is inappropriate only for two action constructs, test and 'waiting', which

[2] This notion of non-determinism must not be confused with *procedural* non-determinism to be introduced momentarily.

are *stationary* in *PDL*. On the other hand, *procedural* non-determinism between the execution of two actions α and β, for instance, concerns the fact that the modelling of choice in dynamic logic does not fix which action is to be performed.

In the next subsection we proceed with the explication of what is called the *standard* view of propositional dynamic logic (see Goldblatt, 1987).

2.2. Formal syntax and semantics of *PDL*

Language

Let $\mathcal{P}_0 = \{p_1, p_2, ..., p, q, r, ...\}$ and $\mathcal{A}_0 = \{a_1, a_2, ..., a, b, c, ...\}$ be denumerably infinite sets of propositional variables and action variables, respectively. We will use A, A_1, B,... to denote arbitrary formulas and α, α_1, β,... to denote arbitrary terms (denoting actions).

Propositional variables:	$p \in \mathcal{P}_0$
Action variables:	$a \in \mathcal{A}_0$
Formulas:	$A \in \mathcal{L}(PDL)$
Action terms:	$\alpha \in \mathcal{A}$

$$A ::= p \mid \neg A \mid A_1 \vee A_2 \mid [\alpha]A$$

$$\alpha ::= a \mid \alpha_1;\alpha_2 \mid \alpha_1 + \alpha_2 \mid \alpha^* \mid A?$$

We assume the usual definitions of $\wedge$, $\supset$, $\top$ (*verum*), $\bot$ (*falsum*), and so on. For example, $< \alpha > A \stackrel{\text{def}}{=} \neg[\alpha]\neg A$. In particular, we define an action constant λ called *stationary waiting* by $\lambda \stackrel{\text{def}}{=} \top?$.

Semantics

By a *frame* $\mathcal{F}$ we mean a structure $\mathcal{F} = \langle S, \{R_a : a \in \mathcal{A}_0\}\rangle$ such that S is a nonempty set (of world-states), $\{R_a : a \in \mathcal{A}_0\}$ is a set of binary relations, where $R_a \subseteq S \times S$ for each action variable $a \in \mathcal{A}_0$.

A *model* $\mathcal{M}$ based on the frame $\mathcal{F} = \langle S, \{R_a : a \in \mathcal{A}_0\}\rangle$ is a structure $\mathcal{M} = \langle S, \{R_a : a \in \mathcal{A}_0\}, v\rangle$ where v is a function $v : \mathcal{P}_0 \to \text{Pow}(S)$ (Pow(S) is the powerset of S).

A *standard* model $\mathcal{M} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}\}, v\rangle$ is uniquely determined by the model $\langle S, \{R_a : a \in \mathcal{A}_0\}, v\rangle$ through the following conditions that inductively define R_α for compound action terms $\alpha \in \mathcal{A}$: $R_{\alpha;\beta} = R_\alpha \circ R_\beta$ (the relative product of R_α and R_β); $R_{\alpha+\beta} = R_\alpha \cup R_\beta$; $R_{\alpha^*} = (R_\alpha)^\star$ (the reflexive and transitive closure of R_α); and $R_{A?} = \{\langle s, s\rangle : \mathcal{M}, s \models A\}$. From this and the definition of λ it follows that $R_\lambda = \{\langle s, s\rangle : s \in S\}$.

This is exactly the way Goldblatt (1987) proceeds in defining standard models. Fine & Schurz (1995) follow a strategy different from that chosen here. They notice that it is already possible to define 'standard' frames by defining $R_{\alpha;\beta}$, $R_{\alpha+\beta}$, R_{α^*}, and R_λ according to the conditions for standard models. But there is a notable exception: the relations $R_{A?}$ depend on the valuation function and therefore cannot be included in frames.

The concept of *truth of a formula at a state s in a standard model* $\mathcal{M}$ is inductively defined as follows:

- $\mathcal{M}, s \models p$ iff (if and only if) $s \in v(p)$ (for $p \in \mathcal{P}_0$).
- $\mathcal{M}, s \models \neg A$ iff $\mathcal{M}, s \not\models A$.
- $\mathcal{M}, s \models A \vee B$ iff $\mathcal{M}, s \models A$ or $\mathcal{M}, s \models B$.
- $\mathcal{M}, s \models [\alpha]A$ iff $\forall t\,(\langle s, t\rangle \in R_\alpha \Rightarrow \mathcal{M}, t \models A)$.

A formula A is said to be *valid in* a model $\mathcal{M}$ iff A is true at all states s in $\mathcal{M}$; A is *valid on* a frame $\mathcal{F}$ iff A is valid in all models based on $\mathcal{F}$. Moreover, A is *valid with respect to* a class of models $\mathbf{M}$, written $\mathbf{M} \models A$, iff A is valid in all models $\mathcal{M} \in \mathbf{M}$, and A is *valid with respect to* a class of frames $\mathbf{F}$ iff A is valid on all frames $\mathcal{F} \in \mathbf{F}$.

For convenience, we will use the notion of the *truth-set* of a formula A (relative to a model $\mathcal{M}$), $\|A\|^{\mathcal{M}}$, which is defined as $\{s \in S : \mathcal{M}, s \models A\}$.

Logics and axiomatization

An *(action) logic* is any subset $\mathbf{L} \subseteq \mathcal{L}(PDL)$ that contains all instances of the following *axiom* schemes:

Taut	all classical propositional tautologies
K	$[\alpha](A \supset B) \supset ([\alpha]A \supset [\alpha]B)$
Comp	$[\alpha;\beta]A \equiv [\alpha][\beta]A$
Union	$[\alpha+\beta]A \equiv ([\alpha]A \wedge [\beta]A)$
Mix	$[\alpha^*]A \supset (A \wedge [\alpha][\alpha^*]A)$
Ind	$[\alpha^*](A \supset [\alpha]A) \supset (A \supset [\alpha^*]A)$
sTest	$[A?]B \equiv (A \supset B)$

and which is closed under the following *rules of inference*:

MP	from $\vdash A$ and $\vdash A \supset B$ infer $\vdash B$
NEC	from $\vdash A$ infer $\vdash [\alpha]A$
SUBST	from $\vdash A$ infer $\vdash \sigma A$

(where σA is the result of a uniform substitution of arbitrary formulas B for propositional variables p in A).

Propositional dynamic logic is the *smallest* subset $\mathbf{L} \subseteq \mathcal{L}(PDL)$ that contains the above axiom schemes and is closed under the rules of inference **MP** and **NEC**. Note that **SUBST** is not needed since we define *PDL* in terms of axiom schemes rather than instances of schemes.

The deducibility relation $\Gamma \vdash_{\mathbf{L}} A$ between formula sets Γ and formulas A is defined as usual by $\Gamma \vdash_{\mathbf{L}} A$ iff $(\bigwedge \Gamma_f \supset A) \in \mathbf{L}$ for some finite subset $\Gamma_f \subseteq \Gamma$, where $\mathbf{L}$ is a logic and $\bigwedge \Gamma_f$ is the conjunction of Γ_f's elements; with $\bigwedge \emptyset \stackrel{\text{def}}{=} \top$ and $\bigwedge \{A\} \stackrel{\text{def}}{=} A$. As is well known, $\Gamma \vdash_{\mathbf{L}} A$ (so defined) iff A is provable from Γ and $\mathbf{L}$-theorems by using **MP** alone.

Correctness and completeness

A logic $\mathbf{L}$ is *correct* with respect to a class of models $\mathbf{M}$ iff all theorems of $\mathbf{L}$ are valid in all models in $\mathbf{M}$. A logic $\mathbf{L}$ is *(weakly) complete* with respect to a class of models $\mathbf{M}$ iff all formulae which are valid in $\mathbf{M}$ are theorems of $\mathbf{L}$. Finally, a logic $\mathbf{L}$ is *(weakly) determined* (characterised) by a class of models $\mathbf{M}$ iff $\mathbf{L}$ is correct and (weakly) complete with respect to $\mathbf{M}$.

It is well known that *PDL* is (weakly) characterised by the class of all standard models. This is accomplished by employing a canonical model construction and the filtration technique (see Segerberg, 1982; Harel, 1984; Goldblatt, 1987). It is also well known that—by the nature of iteration—*PDL* is *not* compact and therefore not strongly complete with respect to the class of standard models; and moreover, that *PDL* has the *finite model property* which together with the finite axiomatizability of *PDL* implies its decidability.

OBSERVATION 2.1. The notions of correctness and completeness apply similarly to frames. Clearly, *PDL* is also complete with respect to the class of all *frames*. As noted by Fine & Schurz (1995), the star-free fragment of *PDL* is a *definitional extension* of the underlying *stratified multimodal* logic containing only atomic programs terms. (Fine & Schurz speak of atomic program terms rather than action variables). All frame completeness transfer theorems proved there apply to this fragment. This means that if certain atomic programs are characterised by additional axioms (for example **T**, **S4**, **S5**) which are complete for their monomodal frames, the resulting star-free fragment of *PDL* will still be frame-complete.

The next subsection offers the main steps in the standard completeness proof for *PDL*. The proofs are omitted since they are well-known (see Segerberg, 1982; Goldblatt, 1987).

2.3. A determination result for *PDL*

THEOREM 2.1 (Correctness). *PDL is correct with respect to the class* $\mathbf{M}$ *of all standard models, that is, for all formulae* A, $\vdash A$ implies $\mathbf{M} \models A$.

By the *canonical model* $\mathcal{M}^c$ of *PDL* we understand the structure $\mathcal{M}^c = \langle S^c, \{R^c_\alpha : \alpha \in \mathcal{A}\}, v^c\rangle$ where (1) $S^c = \{s \subseteq \mathcal{L}(PDL) : s \text{ is maximally } PDL\text{-consistent}\}$ (in short: *PDL*-maximal), (2) $\langle s, t\rangle \in R^c_\alpha$ iff $\{A : [\alpha]A \in s\} \subseteq t$, and (3) $v^c(p) = \{s \in S^c : p \in s\}$ $(p \in \mathcal{P}_0)$.

One basic step in proving completeness is to establish the 'Truth Lemma' for canonical models which implies the following lemma:

LEMMA 2.1. *A is valid in $\mathcal{M}^c$ iff A is a theorem of PDL.*

Indeed, the preceding claim holds for every normal logic. The important proof which is specific for *PDL* is to verify that $\mathcal{M}^c$ satisfies all the conditions for compound action terms. Unfortunately, this is not completely true.

THEOREM 2.2. *$\mathcal{M}^c$ satisfies all conditions of standard models except* $R^c_{\alpha^*} \subseteq (R^c_\alpha)^\star$.

That is to say, in the canonical model of *PDL*, $(R_\alpha)^\star$ is not the reflexive and transitive closure of R_α since the set $\{[\alpha]^n p : n \geq 0\} \cup \{\neg[\alpha^*]p\}$ can be shown to be *PDL*-consistent. Consequently, *PDL* is *not compact* (and therefore pruned from being strongly complete). Remember that a logic $\mathbf{L}$ is said to be *compact* iff, whenever each finite subset $\Gamma_f \subseteq \Gamma$ is (simultaneously) satisfiable, Γ itself is (simultaneously) satisfiable. Although the canonical model of *PDL* does not satisfy all conditions for standardness, the filtration method produces a new (finite) model that rejects any non-theorem of *PDL* and is in the desired class of standard models. This new model is defined as follows.

The *Fischer–Ladner closure* of a formula set Γ is defined as the smallest set Δ such that the following conditions are satisfied (see Fischer & Ladner, 1979): (1) $\Gamma \subseteq \Delta$; (2) Δ is closed under subformulas; (3) $[\alpha;\beta]B \in \Delta \Rightarrow [\alpha][\beta]B \in \Delta$; (4) $[\alpha+\beta]B \in \Delta \Rightarrow [\alpha]B, [\beta]B \in \Delta$; (5) $[\alpha^*]B \in \Delta \Rightarrow [\alpha][\alpha^*]B \in \Delta$; (6) $[A?]B \in \Delta \;\Rightarrow\; A \in \Delta$.

LEMMA 2.2 (Fischer & Ladner, 1979). *The Fischer–Ladner closure Δ of a finite set Γ is finite.*

Let $\mathcal{M}^c = \langle S^c, \{R^c_\alpha : \alpha \in \mathcal{A}\}, v^c\rangle$ be the canonical model of *PDL*, and let $\Delta \subseteq \mathcal{L}(PDL)$ by any set of formulas closed under the Fischer–Ladner conditions. An *equivalence relation* $\sim_\Delta$ (depending on Δ) is defined on S^c by $s \sim_\Delta t$ iff $s \cap \Delta = t \cap \Delta$. The $\sim_\Delta$-*equivalence class* of s, written $|s|$, is defined as $|s| = \{t \in S^c : s \sim_\Delta t\}$, where the subscript $\sim_\Delta$ in $|s|_{\sim_\Delta}$ is omitted for reasons of readability only. Finally, the *quotient set* S^Δ of S^c modulo $\sim_\Delta$ is defined as the set of $\sim_\Delta$-equivalence classes $|s|$ for all $s \in S^c$, $S^\Delta = \{|s| : s \in S^c\}$. We use notation $\Delta(s)$ for $\{A \in \Delta : A \in s\}$.

Let

$$\mathcal{M}^\Delta = \langle S^\Delta, \{R^\Delta_\alpha : \alpha \in \mathcal{A}^\Delta\}, v^\Delta \rangle$$

be a model of the fragment of the language which appears in Δ and which satisfies the following conditions:

1. S^Δ is defined as above,
2. the set $\{R^\Delta_\alpha : \alpha \in \mathcal{A}^\Delta\}$ satisfies the following conditions:
   a) $\mathcal{A}^\Delta$ is defined as the smallest set of action terms such that
      i) $\{a \in \mathcal{A}_0 : a \,\text{occurs in a member of}\, \Delta\} \subseteq \mathcal{A}^\Delta$.
      ii) $\{A? : A? \,\text{occurs in a member of}\, \Delta\} \subseteq \mathcal{A}^\Delta$.
      iii) $\alpha, \beta \in \mathcal{A}^\Delta \;\Rightarrow\; \alpha;\beta,\; \alpha+\beta,\; \alpha^* \in \mathcal{A}^\Delta$.
   b) for $a \in \mathcal{A}^\Delta_0 = \mathcal{A}^\Delta \cap \mathcal{A}_0$, R^Δ_a is *any* Δ-filtration of R^c_a.
   c) for the remaining action terms,
      i) $R^\Delta_{A?} = \{\langle |s|, |s| \rangle : \mathcal{M}^c, s \models A\}$, if $A? \in \mathcal{A}^\Delta$.
      ii) R^Δ_α for arbitrary $\alpha \in \mathcal{A}^\Delta$ is inductively defined according to the conditions on standard models, that is, $R^\Delta_{\alpha;\beta} = R^\Delta_\alpha \circ R^\Delta_\alpha$; $R^\Delta_{\alpha+\beta} = R^\Delta_\alpha \cup R^\Delta_\alpha$; $R^\Delta_{\alpha^*} = (R^\Delta_\alpha)^\star$.
3. Finally, let $\mathcal{P}^\Delta_0 = \mathcal{P}_0 \cap \Delta$ be the set of propositional variables in Δ. Then, for all $p \in \mathcal{P}^\Delta_0$, $v^\Delta(p) = \{|s| : s \in v^c(p)\}$.

We defined R^Δ_a for $a \in \mathcal{A}^\Delta_0$ to be *any* Δ-filtration of R^c_a. Every Δ-filtration must satisfy the following 'suitability' conditions (see Hughes & Cresswell, 1984; Goldblatt, 1987[3]):

(F1) $\langle s, t \rangle \in R^c_a$ implies $\langle |s|, |t| \rangle \in R^\Delta_a$.

(F2) $\langle |s|, |t| \rangle \in R^\Delta_a$ implies $\{A : [a]A \in \Delta(s)\} \subseteq t$.

THEOREM 2.3. *$\mathcal{M}^\Delta$ is a Δ-filtration of $\mathcal{M}^c$.*

The tedious proof of Theorem 2.3 shows that R^Δ_γ is a Δ-filtration of R^c_γ for all $\gamma \in \mathcal{A}^\Delta$ by induction on the complexity of γ. Note that the respective (Fischer-Ladner) closure conditions are crucial to prove condition (F2). Given Theorem 2.3, the proof of the Filtration Lemma is straightforward by induction on formula complexity.

[3] It is well known that such a filtration always exists, for example, the 'smallest' filtration, confer the definition of $\mathcal{M}^\Delta$ in Subsection 3. 2, item 2. b)*ii)*.

THEOREM 2.4 (Filtration Lemma). *For all $s \in S^c$, all action terms $\alpha \in \mathcal{A}^\Delta$, and all formulas $A \in \Delta$: $\mathcal{M}^c, s \models A$ iff $\mathcal{M}^\Delta, |s| \models A$.*

COROLLARY 2.1. *$\mathcal{M}^\Delta$ is a standard model.*

Corollary 2.1 holds for $\alpha;\beta \in \mathcal{A}^\Delta$, $\alpha + \beta \in \mathcal{A}^\Delta$, and $\alpha^* \in \mathcal{A}^\Delta$ by definition. The only remaining case is $A? \in \mathcal{A}^\Delta$, for which Corollary 2.1 is proved *via* Filtration Lemma (Theorem 2.4) and the definition of $R^\Delta_{A?}$ in $\mathcal{M}^\Delta$; thus $R^\Delta_{A?} = \{\langle |s|, |s| \rangle : \mathcal{M}^\Delta, |s| \models A\}$.

A logic **L** is said to have the *finite model property* iff every **L**-consistent formula A is true at a state in some finite standard model.

THEOREM 2.5 (Finite model property). *PDL has the finite model property.*

Theorem 2.5 follows in a straightforward way: Lemma 2.1 tells us that for every *PDL*-consistent A there is some $s \in S^c$ with $\mathcal{M}^c, s \models A$; Lemma 2.2 guarantees that the Fischer–Ladner closure Δ of the set $\{A\}$ is finite, so $\mathcal{M}^\Delta$ is a finite model which, by Theorem 2.3, is a Δ-filtration of $\mathcal{M}^c$ satisfying the conditions for standardness (see Corollary 2.1), and by the Filtration Lemma, $\mathcal{M}^\Delta, |s| \models A$. In effect, we have a model that is both *finite* and *standard*, and which verifies A.

COROLLARY 2.2 (Weak Completeness). *PDL is (weakly) complete with respect to the class of finite standard models.*

COROLLARY 2.3 (Determination). *PDL is (weakly) determined by the class of finite standard models.*

THEOREM 2.6 (Decidability). *PDL is decidable.*

3. Non-stationary propositional dynamic logic

Almost all interesting domains are populated with other agents. When considering multiple agents, the standard view of dynamic logic runs into problems for essentially two reasons: the standard modelling of ‘doing nothing’ (stationary waiting) and of stationary test are feasible only in single-agent environments. Therefore, *non-stationary* propositional dynamic logic *(NPDL)* is defined below, which overcomes these problems and treats multiagent domains in a *uniform* way. Although actions of other agents are not explicit in the language of the logic, we take them into account semantically.

The restriction of stationary waiting λ to static domains is mirrored semantically by conceiving R_λ as the 'diagonal' of $S \times S$. A state-change during waiting is excluded by this modelling. However, in dynamic environments some other agent may 'cause' a state-transition. Therefore, *non-stationary waiting*, written ω, is introduced as the *dynamic* counterpart to stationary waiting. All we know about R_ω is that $R_\omega \subseteq S \times S$; no further condition can be imposed because the dynamic environment may cause arbitrary state-changes during the waiting of the agent. However, ω makes good sense within preservation constructs to be introduced in the next section.

Analogously to the case of non-stationary waiting we introduce a dynamic counterpart of stationary test, called non-stationary test, written τA. In the modelling of non-stationary test the condition to be verified must hold both at the initial and final state. It is *not* required that the respective states be identical. Generally, non-stationary test seems to be a better match with our intuitions on the verification of a condition. Think, for example, of verifying that a solution is acid by the aid of litmus paper. The crucial point here is that a well known state-change occurs, namely that the paper's color turns into red (the example is Moore's 1985).

However, there is a more subtle point. First, observe that the axiom for stationary test, $[A?]B \equiv (A \supset B)$ (**sTest**), is equivalent to the conjunction of the following three axioms: (1) $B \supset [A?]B$, (2) $\neg A \supset [A?]B$, and (3) $A \wedge B \supset < A? > B$ (note that (1) plus (2) is equivalent to the right-to-left direction of **sTest**, while (3) is equivalent to the left-to-right direction of **sTest**). For our non-stationary test operation we certainly want to keep (2). Observe that (2) is equivalent to $\neg A \supset \neg < A? > \top$ (by modal logic), that is, whenever A is false at the initial state, then there does not exist a terminating execution of $A?$. The contrapositive of this formula is $< A? > \top \supset A$ and its non-stationary analogon will be our axiom **nTest.1** ($< \tau A > \top \supset A$) below. Next, we have to weaken (1): not every proposition remains true after every terminating execution of verifying A, but only A itself—thus we weaken (1) to $[\tau A]A$ (**nTest.2**). The axioms **nTest.1** and **nTest.2** are mirrored semantically by the simple requirement that $R_{\tau A}$ may be any set of states $\langle s, t\rangle$ such that A is both true at s and t (as will be proved soon). Intuitively, the notion τ implies that whenever A is true at a state s and the action "verify A" is applied to s then *if* the action "verify A" terminates, *then* A will also be true at its final state(s). However, the notion τ does not guarantee that the action "verify A" indeed always terminates if applied to a state s where A is true. Hence, the operation τA describes an *unsafe* form of testing: the state-change which occurs during testing A may destroy some of the preconditions of the test operation such that the test has no outcome at all. Imagine that in the above test of acid by means of the litmus paper, the acid's concentration is so high that the litmus paper is destroyed before exhibiting some test result.

For several purposes a stronger notion of non-stationary test is needed, call it τ°: the operation $\tau^\circ A$ always terminates if applied to a state in which A is true. We call this notion *safe non-stationary test*. Formally it is characterized by **snTest.1–2** (as above) plus the additional axiom **snTest.3**: $A \supset < \tau^\circ A > \top$. **snTest.3** is a natural weakening of axiom (3) above for stationary test. Semantically it is mirrored by the additional condition that for each state $s \in S$ where A is true, $R_{\tau^\circ A}$ contains a pair $\langle s, t \rangle$ with initial state s. (This requirement is consistent with the previous requirement that A is true at initial and final state: if A is true at s then at least $\langle s, s \rangle$ is a pair satisfying this requirement.) Observe that **snTest.3** and **snTest.2** together imply $A \supset < \tau^\circ A > A$, and this together with **snTest.1** implies $A \equiv < \tau^\circ A > A$. The operation of safe non-stationary test is stronger than non-stationary test but still much weaker than stationary test. For convenience, we introduce both operations into our extended language.

Finally, a natural construct called *the 'any' action* is included within *NPDL* to enhance the expressiveness of the framework. Following Passy & Tinchev (1991), we will use the letter ν to denote the 'any' action. For reasons that will become clear momentarily, it is called the *universe* program in their paper (see also Gargov *et al.*, 1987; Gargov & Passy, 1990; Goranko & Passy, 1992). The 'any' action will be included in our framework for reasons that are independent of considerations on single-agent or multiagent domains. With help of ν we become able to express that something is true at every state in a generated model, since in generated models R_ν is universal, and hence $R_\alpha \subseteq R_\nu$ for arbitrary action terms (see Passy & Tinchev, 1991). The semantic behavior of R_ν facilitates the formulation of *domain constraints*, that is, formulas which are assumed to be invariant over every performance of every action (see Rosenschein, 1981; Ginsberg & Smith, 1988a; Stephan & Biundo, 1993; in the knowledgebase literature they are also called *integrity constraints*, see Katsuno & Mendelzon, 1991). A more detailed study of this simple yet powerful construct is given in section 5, when we turn to applications of (extended) non-stationary propositional dynamic logic to planning problems.

3.1. Formal syntax and semantics of *NPDL*

Language

The vocabulary of $\mathcal{L}(PDL)$ is enriched by the designated action variable ω, the action constant ν, and the (unary) operators τ and τ°.

$$\alpha ::= a \mid \omega \mid \nu \mid \alpha_1; \alpha_2 \mid \alpha_1 + \alpha_2 \mid \alpha^* \mid A? \mid \tau A \mid \tau^\circ A$$

Here, $A \in \mathcal{L}(NPDL)$, $a \in \mathcal{A}_0$, and $\alpha \in \mathcal{A}$.

Before we turn to semantics, it will be useful to denote the following sets: the set $\mathcal{A}_{at}$ of atomic action terms is defined as $\mathcal{A}_0 \cup \{\nu\}$, and the set $\mathcal{A}_{el}$ of *elementary* action terms is defined as $\mathcal{A}_{at} \cup \{A?, \tau A, \tau^\circ A : A \in \mathcal{L}(NPDL)\}$.

Semantics

An *agent frame* $\mathcal{F}$ is a structure $\mathcal{F} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}_{at}\}, T, T^\bullet \rangle$ where (1) S is a nonempty set of world-states, $\{R_\alpha : \alpha \in \mathcal{A}_{at}\}$ is the set of relations corresponding to atomic action terms, and the following functions T (and $T^\bullet$) mirror (safe) non-stationary test on the level of frames; (2) T is a function $T : \mathrm{Pow}(S) \to \mathrm{Pow}(S \times S)$ such that for all $X \in \mathrm{Pow}(S)$, $TX \subseteq \{\langle s, t \rangle : s \in X, t \in X\}$. (3) $T^\bullet$ is like T except that it meets the further condition that for each $s \in X$ (and $X \in \mathrm{Pow}(S)$), $T^\bullet X$ contains a pair $\langle s, t \rangle$ (for $t \in S$).

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 $\alpha \in \mathcal{A}$.

A *standard agent model* $\mathcal{M}$ based on a standard agent frame $\mathcal{F}$ is a structure $\mathcal{M} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}\}, T, T^\bullet, v \rangle$ where (1) all conditions for standard models are satisfied; (2) $R_{\tau A} = T\|A\|^{\mathcal{M}}$; (3) $R_{\tau^\circ A} = T^\bullet \|A\|^{\mathcal{M}}$ (informally speaking, T ($T^\bullet$) assigns to the proposition denoted by A the action consisting of (safely) non-stationary verifying that A is true); (4) v is a function $v : \mathcal{P}_0 \to \mathrm{Pow}(S)$ as usual.

Observe that the restriction to R_ν-generated frames is harmless in the sense that it produces no new theorems. For, let $\mathcal{M}$ be an agent model which satisfies all the conditions of standardness except that R_ν is just an equivalence relation with $R_\alpha \subseteq R_\nu$ for all $\alpha \in \mathcal{A}$ (hence $\mathcal{F}$ need not be R_ν-generated). For any state $u \in S^{\mathcal{M}}$, let $\mathcal{M}^u$ be the u-R_ν-generated submodel of $\mathcal{M}$. Then it is a well-known fact of modal logic that for all $A \in \mathcal{L}(NPDL)$ and $s \in S^{\mathcal{M}^u}$, $\mathcal{M}^u, s \models A$ iff $\mathcal{M}, s \models A$.

Logics and axiomatization

Non-stationary propositional dynamic logic is defined as the smallest subset $\mathbf{L} \subseteq \mathcal{L}(NPDL)$ which contains all axioms and rules of *PDL*, and in addition contains all instances of the following types of axiom schemes.

– For *(unsafe) non-stationary test*:

nTest.1	$< \tau A > \top \supset A$
nTest.2	$[\tau A]A$

– For *safe non-stationary test*:

snTest.1	$< \tau^\circ A > \top \supset A$
snTest.2	$[\tau^\circ A]A$
snTest.3	$A \supset < \tau^\circ A > \top$

– For the 'any' action:

Any.1	$[\nu]A \supset A$
Any.2	$[\nu]A \supset [\nu][\nu]A$
Any.3	$< \nu > [\nu]A \supset A$
Any.4	$[\nu]A \supset [\alpha]A$

NPDL is (weakly) characterised by the class of all standard agent models, and is decidable. Details are provided in the next subsection.

3.2. A determination result for *NPDL*

We show that the determination result for *PDL* can be modified to account also for non-stationary test and the 'any' action. Thus we build upon previous results and other facts from ordinary modal logic.

THEOREM 3.1 (Correctness). *NPDL is correct with respect to the class of all standard agent models.*

Proof. It is easily verified that **nTest.1–2**, **snTest.1–3** and **Any.1–4** are true in a standard agent model $\mathcal{M}$. For example, in the case of **nTest.1**, assume $\mathcal{M}, s \models < \tau A > \top$. This means that there exists $t \in S^{\mathcal{M}}$ with $sR_{\tau A}t$. By definition of $R_{\tau A}$, $\mathcal{M}, s \models A$. Similarly for **nTest.2**. For **snTest.3**, assume $\mathcal{M}, s \models A$. Since $R_{\tau^\circ A}$ contains at least one pair $\langle s, t \rangle$ for $t \in S$, $\mathcal{M}, s \models < \tau^\circ A > \top$ follows. The proofs of

Any.1–4 are well-known. □

As usual, the *completeness* part of the proof is opened by the definition of canonical structures.

Let $\mathcal{M}^c = \langle S^c, \{R^c_\alpha : \alpha \in \mathcal{A}\}, T^c, T^{\bullet c}, v^c \rangle$ be the *canonical* model of *NPDL*. It is defined like the canonical model of *PDL*. Of course, the definition of $\mathcal{M}^c$ includes the new cases $\alpha = \omega$, $\alpha = \nu$ and $\alpha = \tau B$ or $\tau^\circ B$.

$T^c X$ and $T^{\bullet c} X$ are defined in two steps: (i) For those propositions $X = \|A\|^{\mathcal{M}^c}$ which are denoted by some formula $A \in \mathcal{L}(NPDL)$ we put $T^c \|A\|^{\mathcal{M}^c} = R^c_{\tau A}$ and $T^{\bullet c} \|A\|^{\mathcal{M}^c} = R^c_{\tau^\circ A}$ (ii) In the other case, that is, when there is *no* $A \in \mathcal{L}(NPDL)$ such that $X = \|A\|^{\mathcal{M}^c}$, just let $T^c X$ and $T^{\bullet c} X$ be any subsets of $\{\langle s, t \rangle : s \in X, t \in X\}$, where $T^{\bullet c} X$ satisfies the additional condition that it contains a pair $\langle s, t \rangle$ for each $s \in S^c$.

LEMMA 3.1. *(1)* $R^c_{\tau A} \subseteq \{\langle s, t \rangle : \mathcal{M}^c, s \models A,\ \mathcal{M}^c, t \models A\}$*; (2)* $R^c_{\tau^\circ} A$ *satisfies (1) and for each* $s \in S$ *it contains at least one pair* $\langle s, t \rangle$ *(for* $t \in S^c$*); and (3)* R^c_ν *is an equivalence relation on* S^c *and* $R^c_\alpha \subseteq R^c_\nu$ *for each* $\alpha \in \mathcal{A}$*.*

Proof. (1) *Non-stationary test.* In case (i) of the paragraph preceding Lemma 3.1 we first show that $\langle s, t \rangle \in R^c_{\tau A} \Rightarrow A \in s, t$ (thereby exploiting the power of the Truth Lemma). For s, assume that $\langle s, t \rangle \in R^c_{\tau A}$. We have that $\top \in t$. Therefore $< \tau A > \top \in s$. By **nTest.1** and maximality we have that $(< \tau A > \top \supset A) \in s$. Thus, $A \in s$. For t, suppose that $\langle s, t \rangle \in R^c_{\tau A}$. Then $[\tau A]A \in s$ by maximality, and because $[\tau A]A$ is in the logic (**nTest.2**). Therefore, $A \in t$. In case (ii) of the paragraph preceding Lemma 3.1, the claim holds by definition. In summary, $\mathcal{M}^c$ satisfies $T^c X \subseteq \{\langle s, t \rangle : s \in X, t \in X\}$ in all cases.

(2) *Safe non-stationary test.* It remains to prove the additional semantic condition for the case (i), because in the case (ii) it holds by definition. Assume $s \in S^c$ and $A \in s$. We must show that there exists $t \in S^c$ such that $s R^c_{\tau^\circ A} t$ (for arbitrary formula A) which means that we have to prove that the set $\{B : [\tau^\circ A]B \in s\}$ is *NPDL*-consistent. For reductio, assume the opposite. Hence, by modal logic and maximality of s, s must contain $[\tau^\circ A]\bot$. However, because of **snTest.3** and $A \in s$, s contains also $< \tau^\circ A > \top$, that is, $\neg[\tau^\circ A]\bot$, which contradicts the fact that s is *NPDL*-consistent.

(3) *The 'any' action.* Axioms **Any.1–3** guarantee that R^c_ν is an equivalence relation on S^c (see, for example, Hughes and Cresswell, 1984). The addition of axiom **Any.4** guarantees that $R^c_\alpha \subseteq R^c_\nu$ for all $\alpha \in \mathcal{A}$: assume $s R^c_\alpha t$ and $[\nu]A \in s$. Then, for all $\alpha \in \mathcal{A}$, $[\alpha]A \in s$ by

Any.4 and maximality. Therefore, $A \in t$. □

It is easy to see that *NPDL* (like *PDL* before) is not compact. There is thus no escape from applying the filtration technique again. However, we do not apply filtration directly to $\mathcal{M}^c$, but rather take an R_ν-generated submodel of $\mathcal{M}^c$, call it $\mathcal{M}^{cu}$. The reason is that filtration preserves universality (see Chellas, 1980, p. 103). $\mathcal{M}^{cu}$ is R_ν-universal and equivalent with $\mathcal{M}^c$ for all its states. Hence, by Lemma 3.1 $\mathcal{M}^{cu}$ is a *standard agent* model. Any filtration of $\mathcal{M}^{cu}$ through a given Δ will preserve universality because of the first suitability condition: $sR^{cu}_\nu t \Rightarrow |s|R^\Delta_\nu |t|$.

The definition of the Fischer–Ladner closure Δ of a formula set Γ within $\mathcal{L}(NPDL)$ *has to be extended* vis-à-vis its definition within $\mathcal{L}(PDL)$ by the following conditions: (1) $[\tau A]B \in \Delta \;\Rightarrow\; A \in \Delta$ and (2) $[\tau^\circ A]B \in \Delta \;\Rightarrow\; A \in \Delta$. It is easy to see that the Fischer–Ladner closure Δ of a finite formula set Γ continues to be finite in the new setting (compare Lemma 2.2). The set $\mathcal{A}^\Delta$ is defined as the least set of action terms such that

- $\{\alpha \in \mathcal{A}_{el} : \alpha \,\text{occurs in a member of}\, \Delta\} \subseteq \mathcal{A}^\Delta$.
- $\alpha, \beta \in \mathcal{A}^\Delta \;\Rightarrow\; \alpha;\beta,\, \alpha+\beta,\, \alpha^* \in \mathcal{A}^\Delta$.

In the structure to be introduced momentarily, we write $|A|^\Delta$ for the set $\{|s| \in S^\Delta : \mathcal{M}^{cu}, s \models A\}$, if A is in Δ.

Take $\mathcal{M}^{cu} = \langle S^{cu}, \{R^{cu}_\alpha : \alpha \in \mathcal{A}\}, T^{cu}, T^{\bullet cu}, v^{cu}\rangle$ and let $\Delta \subseteq \mathcal{L}(NPDL)$ be any Fischer–Ladner closed set. Define the model

$$\mathcal{M}^\Delta = \langle S^\Delta, \{R^\Delta_\alpha : \alpha \in \mathcal{A}^\Delta\}, T^\Delta, T^{\bullet\Delta}, v^\Delta\rangle$$

where

1. as before, S^Δ is the quotient set of S^{cu} modulo $\sim_\Delta$,
2. the set $\{R^\Delta_\alpha : \alpha \in \mathcal{A}^\Delta\}$ satisfies the following conditions:
   a) for $a \in \mathcal{A}^\Delta_0$, R^Δ_a is any Δ-filtration of R^{cu}_a,
   b) for action terms $A?, \tau A, \tau^\circ A, \nu \in \mathcal{A}^\Delta$,
      i) $R^\Delta_{A?} = \{\langle |s|, |s|\rangle : \mathcal{M}^{cu}, s \models A\}$ (as before),
      ii) in case of $R^\Delta_{\tau A}$ and $R^\Delta_{\tau^\circ A}$, we take the *smallest* Δ-filtration, that is,

$$R^\Delta_{\tau A} = T^\Delta |A|^\Delta = \{\langle |s|, |t|\rangle : \exists s\prime \in |s| \exists t\prime \in |t| \langle s\prime, t\prime\rangle \in T^c \|A\|^{\mathcal{M}^{cu}}\},$$

      and similar for $R^\Delta_{\tau^\circ A}$ which is defined with help of $T^{\bullet cu}$. Otherwise, for $X \subseteq S^\Delta$ where there is no $A \in \Delta$ such that

$X = |A|^\Delta$, let $T^\Delta X$ and $T^{\bullet\Delta} X$ be subsets of $\{\langle |s|, |t| \rangle : |s| \in X, |t| \in X\}$, where $T^{\bullet\Delta} X$ satisfies the additional condition that for each $|s| \in S^\Delta$ it contains a pair $\langle |s|, |t| \rangle$ (for some $|t| \in S^\Delta$),

iii) R^Δ_ν is any Δ-filtration,

iv) R^Δ_α for compound action terms is inductively defined as before;

3. $v^\Delta(p)$ ($p \in \mathcal{P}_0$) is defined as before.

The next statement extends Theorem 2.3.

THEOREM 3.2. *$\mathcal{M}^\Delta$ is a Δ-filtration of $\mathcal{M}^{cu}$.*

Proof. We show that R^Δ_γ is a Δ-filtration of R^{cu}_γ for all $\gamma \in \mathcal{A}^\Delta$. Thereby, we only consider cases not already covered by Theorem 2.3.

Non-stationary waiting. The case $\gamma = \omega$ is given by the definition of R^Δ_a. Remember that ω is considered as an action variable.

Non-stationary test. We have to show that the filtration conditions (F1) and (F2) are satisfied. (F1): Assume $\tau A \in \mathcal{A}^\Delta$. Suppose that $\langle s, t \rangle \in R^{cu}_{\tau A}$. Then, by definition, $\langle s, t \rangle \in T^{cu} \|A\|^{\mathcal{M}^{cu}}$. The definition of $T^\Delta |A|^\Delta$ yields $\langle |s|, |t| \rangle \in R^\Delta_{\tau A}$. (F2): Assume $|s| R^\Delta_{\tau A} |t|$. Then for some $s\prime \in |s|$ and $t\prime \in |t|$, $\langle s\prime, t\prime \rangle \in T^{cu} \|A\|^{\mathcal{M}^{cu}}$. Hence $\langle s\prime, t\prime \rangle \in R^{cu}_{\tau A}$ by definition. Suppose $[\tau A]A \in \Delta(s)$. We have $s \sim_\Delta s\prime$ and $\mathcal{M}^{cu} \models [\tau A]A$. So A is in $t\prime$. Moreover, because Δ is closed under the Fischer–Ladner conditions, $A \in \Delta$ and so, for $t\prime \sim_\Delta t$, $A \in t$ as desired.

For *safe non-stationary test* the proof is the same.

The 'any' action. Again, this case holds by definition. □

The Filtration Lemma may now be proved as before and reads as follows:

THEOREM 3.3 (Filtration Lemma). *For all $s \in S^{cu}$, and all formulas $A \in \Delta$: $\mathcal{M}^{cu}, s \models A$* iff *$\mathcal{M}^\Delta, |s| \models A$.*

COROLLARY 3.1. *(1) $T^\Delta X \subseteq \{\langle |s|, |t| \rangle : |s| \in X, |t| \in X\}$ (for every $X \subseteq S^\Delta$); (2) $T^{\bullet\Delta} X$ satisfies (1) plus the condition that for all $|s| \in S^\Delta$ it contains a pair $\langle |s|, |t| \rangle$ (for some $|t| \in S^\Delta$); (3) R^Δ_ν is universal on S^Δ; and (4) $\mathcal{M}^\Delta$ is a* standard agent model.

Proof. (1) *Non-stationary test.* Assume that there is some $A \in \Delta$ such that $X = |A|^\Delta$ (in the other case, the claim holds by definition). Assume $\langle |s|, |t| \rangle \in T^\Delta |A|^\Delta$. There exist $s\prime \in |s|$ and $t\prime \in |t|$ such that $\langle s\prime, t\prime \rangle \in T^{cu} \|A\|^{\mathcal{M}^{cu}}$. By Lemma 3.1, $s\prime \in \|A\|^{\mathcal{M}^{cu}}$ and $t\prime \in \|A\|^{\mathcal{M}^{cu}}$, so $|s\prime|, |t\prime| \in |A|^\Delta$ by Theorem 3.3. But $|s| = |s\prime|$, $|t| = |t\prime|$. Therefore, $|s|, |t| \in |A|^\Delta$.

(2) *Safe non-stationary test.* We must show in addition that $T^{\bullet\Delta}|A|^{\Delta}$ contains for each $|s| \in S^{\Delta}$ a pair $\langle |s|, |t| \rangle$ (for some $|t| \in S^{\Delta}$). Given $|s| \in S^{\Delta}$, then $s \in S^{cu}$, and $sR^{cu}_{\tau \circ A}t$ for some $t \in S^{cu}$ by Lemma 3.1, whence by definition (2)(b)(ii), $\langle |s|, |t| \rangle \in T^{\bullet\Delta}|A|^{\Delta}$.

(3) *The 'any' action.* Assume $|s|, |t| \in S^{\Delta}$. Because R^{cu}_{ν} is universal on S^{cu}, $sR^{cu}_{\nu}t$. By suitability condition (F1), $|s|R^{\Delta}_{\nu}|t|$ follows. Hence, R^{Δ}_{ν} is universal on S^{Δ}.

(4) Follows immediately from (1), (2) and Corollary 2.1. This finishes the proof. □

THEOREM 3.4 (Finite model property). *NPDL has the finite model property.*

Proof. Take any *NPDL*-consistent formula A. The relevant steps are as follows: First, by Lemma 2.1 A is true at some state u in the canonical model $\mathcal{M}^c$. Next, let $\mathcal{M}^{cu}$ be the u-R_{ν}-generated submodel of $\mathcal{M}^c$; A is true at u in $\mathcal{M}^{cu}$. Theorems 3.2 and 3.3 guarantee that A is true at $|u|$ in the Δ-filtration $\mathcal{M}^{\Delta}$ of $\mathcal{M}^{cu}$ and Corollary 3.1 (4) tells us that $\mathcal{M}^{\Delta}$ is a standard agent model. Since the Fischer–Ladner closure Δ of $\{A\}$ is finite, $\mathcal{M}^{\Delta}$ is finite. Analogously to Theorem 2.5 we have gained a verifying model for A that exhibits both the features of finiteness and standardness. □

The following statements are consequences of Theorems 3.1 and 3.4.

COROLLARY 3.2 (Determination). *NPDL is (weakly) determined by the given class of finite standard agent models.*

THEOREM 3.5 (Decidability). *NPDL is decidable.*

4. Preservation

Real-world planning often requires that some facts are *protected from* an action, or put differently, *preserved with respect to* an action.[4] Concerning the frame problem, the decisive problem is to offer an *economic* way to reason about what remains true during an action which is composed by sequencing actions. A method that covers that problem efficiently may count as a solution to the *combinatorial* frame problem. In this

[4] On naming: we eventually decided to speak of "preservation with respect to an action" rather than "protection from an action" as only the former seems to support the multiagent reading.

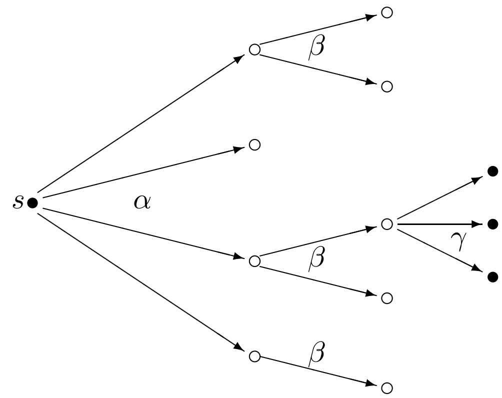

Figure 1. A proposition is terminally *preserved with respect to the sequence of actions α, β and γ. (Filled circles indicate that the proposition is true at the respective states.)*

section, we define the logic resulting from the inclusion of preservation formulas in formal terms. It will be called *extended* non-stationary propositional dynamic logic ($NPDL^+$). Applications to both the frame problem and planning are discussed in the next section.

Depending on whether facts have to hold in a *particular* situation, for example, where they figure as preconditions ('initial conditions') for the execution of an ensuing action, or whether they have to be true at *all* states of an action sequence (as 'boundary conditions'), two concepts of *preservation* are distinguished: *terminal* preservation and *chronological* preservation. These facilities are obtained by introducing formulas of type $\mathbf{tpres}(\alpha, A)$ and $\mathbf{cpres}(\alpha, A)$. In case of terminal preservation the formula $\mathbf{tpres}(\alpha, A)$ is intended to mean that a fact A is terminally preserved with respect to a (possibly compound) action: if A is true when the execution of α begins, A is true upon termination of α. For chronological preservation the formula $\mathbf{cpres}(\alpha, A)$ is intended to mean that a fact A is chronologically preserved with respect to a (possibly compound) action α: whenever A is true initially, A is true not only terminally but at the intermediate states of the execution of α as well.

Fig. 4 illustrates the case of a proposition being terminally preserved with respect to an action consisting of first doing α, then β and finally γ. The filled circles (denoting states) in the figure indicate that the proposition is true at the final states of the relation $R_{\alpha;\beta;\gamma}$ which corresponds to the sequence of actions α, β and γ (in that order). As chronological preservation requires in addition that a proposition be true at intermediate states, application to the sequence $\alpha;\beta;\gamma$ results in the proposition holding at all states, that is, whenever the performance of α, $\alpha;\beta$ or $\alpha;\beta;\gamma$ terminates, the proposition is true. This

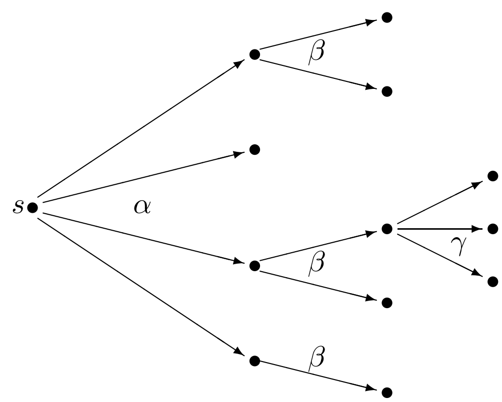

Figure 2. A proposition is chronologically *preserved with respect to the sequence of actions α, β and γ. (Filled circles indicate that the proposition is true at the respective states.)*

case is illustrated in Fig. 4. Note that here all circles are filled. So, 'non-terminating beginnings' of the sequence $\alpha;\beta;\gamma$ count as well.

Several researchers have suggested notions that bear a close similarity to our preservation operators (see Pratt, 1978; Segerberg, 1980). Pratt (1978) introduces formulas **throughout**(α, A) to mean that formula A is true throughout (at every state of) the execution of action α. In Pratt's framework the meaning of an action is realized as a set of *trajectories*, that is, $k+1$-tuples $\langle s_0, ..., s_k\rangle$ of states rather than binary relations. Accordingly, **throughout**(α, A) is true at a state (in a model) iff A holds at every state in every trajectory corresponding to α. Although the conception of the meaning of actions as sets of trajectories seems reasonable in the particular case of a (chronological) preservation operator, it leads to fairly complicated logics in the general case, so-called *process* logics (see Harel, Kozen & Parikh, 1982). In our account, a *definitional extension* of action models (see the function μ_c) does the job. Segerberg (1980) suggests a notion of preservation, written α-**pres**A, which is essentially that of Pratt (1978).

4.1. Formal syntax and semantics *NPDL*$^+$

Language

The operators **tpres** and **cpres** are added to the vocabulary of $\mathcal{L}(NPDL)$.

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

(where $p \in \mathcal{P}_0$, $A \in \mathcal{L}(NPDL^+)$, and $\alpha \in \mathcal{A}$).

For the following, recall the definitions of special term sets already given in the previous section.

Semantics

An *extended standard agent frame* $\mathcal{F}$ is a structure $\mathcal{F} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}_{at}\}, T, T^\bullet, \mu\rangle$ where (1) $\langle S, \{R_\alpha : \alpha \in \mathcal{A}_{at}\}, T, T^\bullet\rangle$ is a standard agent frame; and (2) μ is called the *smoothness* function. It is an *definitional extension* of standard agent frames, a function $\mu : \{R_\alpha : \alpha \in \mathcal{A}_{at}\} \times \mathrm{Pow}(S) \to \mathrm{Pow}(S)$ such that for all $s \in S$, $X \in \mathrm{Pow}(S)$,

$$s \in \mu(R_\alpha, X) \quad \text{iff} \quad s \in X \Rightarrow (\{t : \langle s, t\rangle \in R_\alpha\} \subseteq X).$$

An *extended standard agent model* $\mathcal{M}$ based on $\mathcal{F}$ is a structure

$$\mathcal{M} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}\}, T, T^\bullet, \mu_t, \mu_c, v\rangle$$

where

1. all conditions for standard agent models are satisfied; and

2. the following functions are introduced:

   a) μ_t is the function $\mu_t : \{R_\alpha : \alpha \in \mathcal{A}\} \times \mathrm{Pow}(S) \to \mathrm{Pow}(S)$ 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 $\alpha \in \mathcal{A}_{at}$.

   b) μ_c is the function $\mu_c : \{R_\alpha : \alpha \in \mathcal{A}\} \times \mathrm{Pow}(S) \to \mathrm{Pow}(S)$ such that

   $$\mu_c(R_\alpha, X) = \mu_t(R_\alpha, X)$$

   for all *elementary* action terms $\alpha \in \mathcal{A}_{el}$. For compound action terms, μ_c is inductively defined as follows:

   $$\mu_c(R_{\alpha;\beta}, X) = \mu_c(R_\alpha, X) \cap \{s : \{t : \langle s, t\rangle \in R_\alpha\} \subseteq \mu_c(R_\beta, X)\}$$
   $$\mu_c(R_{\alpha+\beta}, X) = \mu_c(R_\alpha, X) \cap \mu_c(R_\beta, X)$$
   $$\mu_c(R_{\alpha^*}, X) = \bigcap_{n \geq 0} \mu_c(R_{\alpha^n}, X)$$

3. Finally, v is a function $v : \mathcal{P}_0 \to \mathrm{Pow}(S)$.

Terminal preservation (μ_t) takes care only of initial and final states: $s \in \mu_t(R_\alpha, X)$ iff whenever s is an initial state of R_α at which X holds, then X holds at all final states of R_α starting from s. *Chronological preservation* (μ_c) forces X to be preserved also at all intermediate states—it gains its power when applied to compound actions, otherwise it reduces to terminal preservation.[5] Note also that models for $\mathcal{L}(NPDL^+)$ are definitional extensions of models for $\mathcal{L}(NPDL)$ since μ_t and μ_c are defined parameters.

Given any fixed standard agent model $\mathcal{M}$, we define the meanings of formulas and terms by adding the following clauses:

- $\mathcal{M}, s \models \mathbf{tpres}(\alpha, A)$ iff $s \in \mu_t(R_\alpha, \|A\|^{\mathcal{M}})$.
- $\mathcal{M}, s \models \mathbf{cpres}(\alpha, A)$ iff $s \in \mu_c(R_\alpha, \|A\|^{\mathcal{M}})$.

It follows that the following conditions are satisfied by $\mathcal{M}$ (the reference to models is left tacit): For all $s, t \in S$, $\alpha \in \mathcal{A}$:

$$\|\mathbf{cpres}(\alpha;\beta, A)\| = \|\mathbf{cpres}(\alpha, A)\| \cap \{s : \{t : \langle s, t\rangle \in R_\alpha\} \subseteq \|\mathbf{cpres}(\beta, A)\|\}$$
$$\|\mathbf{cpres}(\alpha + \beta, A)\| = \|\mathbf{cpres}(\alpha, A)\| \cap \|\mathbf{cpres}(\beta, A)\|$$
$$\|\mathbf{cpres}(\alpha^*, A)\| = \bigcap_{n \geq 0} \|\mathbf{cpres}(\alpha^n, A)\|$$

Logics and axiomatization

The extended logic is defined as the smallest subset $\mathbf{L} \subseteq \mathcal{L}(NPDL^+)$ which is closed under the rules of *PDL*, contains all axioms of *NPDL* and in addition all instances of the following two new axiom schemes, **tPres** and cPres.

- For *terminal preservation*:

$$\textbf{tPres} \qquad \mathbf{tpres}(\alpha, A) \equiv (A \supset [\alpha]A)$$

The formulation of an appropriate axiom scheme for expressions of the form $\mathbf{cpres}(\alpha, A)$ is more involved. We define I as a function $I : \mathcal{A} \to \mathrm{Pow}(\mathcal{A})$ that assigns to each action term $\alpha \in \mathcal{A}$ the *set* of all action terms that denote *initial strings* of R_α. For *elementary* action terms, $\alpha \in \mathcal{A}_{el}$, we have $I(\alpha) = \{\alpha\}$, that is, an elementary action term has itself as its only initial string. In order to define initial strings for compound action terms, we need the auxiliary definition

$$(\alpha; I(\beta)) \stackrel{\text{def}}{=} \{(\alpha;\gamma) : \gamma \in I(\beta)\}.$$

Then for all $\alpha \in \mathcal{A}$:

[5] A note on naming: *smoothness* is to be understood always with respect to the envisaged facts whose truth is 'preserved'.

$$
\begin{array}{lcl}
I(\alpha;\beta) & = & I(\alpha) \cup (\alpha; I(\beta)) \\
I(\alpha+\beta) & = & I(\alpha) \cup I(\beta) \\
I(\alpha^*) & = & (\alpha^*; I(\alpha))
\end{array}
$$

Hence, the initial strings of $R_{\alpha;\beta}$ are all those of R_α *plus* all those obtained from concatenating initial strings of R_β to R_α; the initial strings of $R_{\alpha+\beta}$ are those of R_α and those of R_β; the initial strings of R_{α^*} are all those obtained by sequencing R_{α^*} with an initial string of R_α.

We now introduce the following axiom scheme

– for *chronological preservation*:

$$
\textbf{cPres} \qquad \textbf{cpres}(\alpha, A) \equiv \left(A \supset \bigwedge_{\beta \in I(\alpha)} [\beta]A\right)
$$

Correctness and completeness

By defining an adequate translation function, we can show that *NPDL*$^+$ is (weakly) characterised by the class of all extended standard agent models, that *NPDL*$^+$ has the finite model property and is decidable.

4.2. A determination result for *NPDL*$^+$

Because of the nature of the proof-procedure pursued in this section it is advisable to distinguish notions of *NPDL* and *NPDL*$^+$. Let $\vdash$ denote the deducibility relation of *NPDL* and $\vdash_+$ the deducibility relation of *NPDL*$^+$. Likewise, models of *NPDL*$^+$ are indexed by "+", and models of *NPDL* have no index. For convenience, we abbreviate $\mathcal{L}(NPDL)$ by $\mathcal{L}$ and $\mathcal{L}(NDPL^+)$ by $\mathcal{L}^+$.

The main steps of the argument are as follows: First, we translate each formula $A \in \mathcal{L}^+$ into a formula $\theta(A) \in \mathcal{L}$ and prove that

$$\vdash_+ A \equiv \theta(A).$$

Then we show that if A is *NPDL*$^+$-consistent then $\theta(A)$ is *NPDL*-consistent. Next, for each model $\mathcal{M}$ for $\mathcal{L}$, $\mathcal{L}$-model for short, we define the corresponding $\mathcal{L}^+$-model $\mathcal{M}^+$ and show that for each $A \in \mathcal{L}^+$, and $s \in S$:

$$\mathcal{M}^+, s \models A \quad \text{iff} \quad \mathcal{M}, s \models \theta(A).$$

For each $A \in \mathcal{L}^+$, the *translation* $\theta : \mathcal{L}^+ \to \mathcal{L}$ is given by the following recursive conditions:

$$\begin{array}{ll}
\theta(p) & = p \qquad (\text{for } p \in \mathcal{P}_0) \\
\theta(\neg A) & = \neg\theta(A) \\
\theta(A \vee B) & = \theta(A) \vee \theta(B) \\
\theta([\alpha]A) & = [\alpha]\theta(A) \\
\theta(\mathbf{tpres}(\alpha, A)) & = \theta(A) \supset [\alpha]\theta(A) \\
\theta(\mathbf{cpres}(\alpha, A)) & = \theta(A) \supset \bigwedge_{\beta \in I(\alpha)} [\beta]\theta(A).
\end{array}$$

We now begin the completeness argument.

LEMMA 4.1. *For all* $A \in \mathcal{L}^+$: $\vdash_+ \theta(A) \equiv A$.

Proof. As well known, the rule of replacement of logically equivalent subformulas holds in all classical modal logics. Now, $\theta(A)$ results from A by a finite number of such replacements of *NPDL*$^+$-equivalent subformulas. □

COROLLARY 4.1. *If* $A \in \mathcal{L}^+$ *is NPDL*$^+$*-consistent, then* $\theta(A)$ *is NPDL-consistent.*

Proof. Suppose otherwise, that $\theta(A)$ is not *NPDL*-consistent. Since all axiom schemes and rules of inference of *NPDL* are also in *NPDL*$^+$, if $\vdash \theta(A) \supset \bot$ then $\vdash_+ \theta(A) \supset \bot$, whence $\vdash_+ A \supset \bot$ follows by Lemma 4.1. Therefore, A is not *NPDL*$^+$-consistent. □

LEMMA 4.2. *For given* $\mathcal{M} = \langle S, \{R_\alpha : \alpha \in \mathcal{A}\}, T, T^\bullet, v\rangle$ *define* $\mathcal{M}^+ = \langle S, \{R_\alpha : \alpha \in \mathcal{A}\}, T, T^\bullet, \mu_t, \mu_c, v\rangle$ *such that* μ_t *and* μ_c *are defined as above (in Section 4.2). Then for all* $A \in \mathcal{L}^+$ *and* $s \in S$: $\mathcal{M}^+, s \models A$ *iff* $\mathcal{M}, s \models \theta(A)$.

Proof. By induction on the formation of A. The case $A = p$ is given because $\mathcal{M}^+$ and $\mathcal{M}$ agree on the valuation function. The cases for $\neg$, $\vee$ and $[\alpha]$ are straightforward. The only critical cases are $A = \mathbf{tpres}(\alpha, B)$ and $A = \mathbf{cpres}(\alpha, B)$.

It is important to note that the operation $+$ is *reversible*. For each $\mathcal{L}^+$-model can be viewed as an extended model $\mathcal{M}^+$ of some $\mathcal{L}$-model, that is, $\mathcal{M}^+$ has a $\mathcal{L}$-reduction.

We start the proof with the case $A = \mathbf{tpres}(\alpha, B)$.

$$\mathcal{M}^+, s \models \mathbf{tpres}(\alpha, B) \iff (s \in \|B\|^{\mathcal{M}^+} \Rightarrow \{t : \langle s, t\rangle \in R_\alpha \subseteq \|B\|^{\mathcal{M}^+}\})$$

by definition of μ_t and the semantic condition for $\mathbf{tpres}(\alpha, B)$-formulas. So

$$\iff (\mathcal{M}^+, s \models B \quad \Rightarrow \quad (\forall t : \langle s, t\rangle \in R_\alpha \Rightarrow \mathcal{M}^+, t \models B)).$$

Applying induction hypothesis we conclude (since $\mathcal{M}$ and $\mathcal{M}^+$ are based on the same frame) that

$$\iff (\mathcal{M}, s \models \theta(B) \quad \Rightarrow \quad (\forall t : \langle s,t\rangle \in R_\alpha \Rightarrow \mathcal{M}, t \models \theta(B)),$$

hence by the truth condition for boxed formulas

$$\iff (\mathcal{M}, s \models \theta(B) \quad \Rightarrow \quad \mathcal{M}, s \models [\alpha]\theta(B))$$

which gives us

$$\iff \mathcal{M}, s \models (\theta(B) \supset [\alpha]\theta(B))$$

and thus, by the recursive clauses for θ,

$$\iff \mathcal{M}, s \models \theta(\mathbf{tpres}(\alpha, B)).$$

The case $A = \mathbf{cpres}(\alpha, B)$ is proved by (nested) induction on the complexity of α.

Elementary action terms. By definition of initial strings for elementary action terms $\alpha \in \mathcal{A}_{el}$, $I(\alpha) = \{\alpha\}$. Hence, the proof is as for $A = \mathbf{tpres}(\alpha, B)$.

Sequential composition. Assume $\alpha = \alpha_1;\alpha_2$. Then $\mathcal{M}^+, s \models \mathbf{cpres}(\alpha_1;\alpha_2, B)$

$$\iff \mathcal{M}^+, s \models \mathbf{cpres}(\alpha_1, B) \;\&\; \forall t: \; sR_{\alpha_1}t \Rightarrow \mathcal{M}^+, s \models \mathbf{cpres}(\alpha_2, B)$$

by the truth condition for $\mathbf{cpres}(\alpha;\beta, A)$-formulas, then

$$\iff \mathcal{M}, s \models \theta(\mathbf{cpres}(\alpha_1, B)) \;\&\; \forall t : sR_{\alpha_1}t \Rightarrow \mathcal{M}, s \models \theta(\mathbf{cpres}(\alpha_2, B))$$

by induction hypothesis, which gives

$$\mathcal{M}, s \models \left(\theta(B) \supset \bigwedge_{\beta \in I(\alpha_1)} [\beta]\theta(B)\right) \quad \& \qquad (1)$$

$$\forall t: \; sR_{\alpha_1}t \quad \Rightarrow \quad \mathcal{M}, t \models \left(\theta(B) \supset \bigwedge_{\gamma \in I(\alpha_2)} [\gamma]\theta(B)\right) \qquad (2)$$

by definition of θ. Since $I(\alpha_1;\alpha_2) = I(\alpha_1) \cup (\alpha_1; I(\alpha_2))$ by definition of initial strings, (1) and (2) yield

$$\iff \mathcal{M}, s \models \left(\theta(B) \supset \bigwedge_{\delta \in I(\alpha_1;\alpha_2)} [\delta]\theta(B)\right),$$

which gives

$$\iff \mathcal{M}, s \models \theta(\mathbf{cpres}(\alpha_1;\alpha_2, B))$$

by definition of θ.

Non-deterministic choice. In a similar way, the proof for the case $\alpha = \alpha_1 + \alpha_2$ uses the truth-clause for $\mathbf{cpres}(\alpha + \beta, B)$-formulas, induction hypothesis, definition of θ and the definition of $I(\alpha_1 + \alpha_2)$.

Non-deterministic iteration. The remaining case is $\alpha = \beta^*$. By the truth condition for $\mathbf{cpres}(\alpha^*, A)$-formulas, $\mathcal{M}^+, s \models \mathbf{cpres}(\beta^*, B)$

$$\iff \mathcal{M}^+, s \models \mathbf{cpres}(\beta^n, B) \quad \text{for all } n \geq 0,$$

and

$$\iff \mathcal{M}, s \models \theta(\mathbf{cpres}(\beta^n, B)) \quad \text{for all } n \geq 0,$$

by induction hypothesis, which yields

$$\iff \mathcal{M}, s \models \left(\theta(B) \supset \bigwedge_{\gamma \in I(\beta^n)} [\gamma]\theta(B)\right) \quad \text{for all } n \geq 0 \tag{3}$$

by definition of θ.

To proceed, we insert the following

LEMMA 4.3 (Initial String Lemma for the star operator). *The following sets are identical:* $\{\langle s,t\rangle \in R_\gamma : \gamma \in I(\beta^n), n \geq 0\} = \{\langle s,t\rangle \in R_\gamma : \gamma \in (\beta^*, I(\beta))\}$.

Proof. For the $\subseteq$-part, assume $\gamma \in I(\beta^n)$. Then γ has the form $(\beta^m; \epsilon)$ for some $\epsilon \in I(\beta)$ and $m \geq 0$. So $R_{\beta^m;\epsilon} \subseteq R_{\beta^*;\epsilon}$. Thus each $\langle s,t\rangle \in R_\gamma$ for such a γ will be in the set at right hand side.

For the $\supseteq$-part, suppose $\langle s,t\rangle \in R_\gamma$ for $\gamma \in (\beta^*, I(\beta))$. Then there will be some $m \geq 0$ such that $\langle s,t\rangle \in R_{\beta^m;\epsilon}$ with $\epsilon \in I(\beta)$. Hence $\langle s,t\rangle \in R_\gamma$ belongs to the set at the left hand side. □

With the help of the preceding lemma we may continue the proof of Lemma 4.2 as follows:

$$(3) \iff \mathcal{M}, s \models \left(\theta(B) \supset \bigwedge_{\gamma \in (\beta^*; I(\beta))} [\gamma]\theta(B)\right)$$

and, since $I(\beta^*) = (\beta^*; I(\beta))$ by definition of initial strings,

$$\iff \mathcal{M}, s \models \left(\theta(B) \supset \bigwedge_{\gamma \in I(\beta^*)} [\gamma]\theta(B)\right),$$

which gives

$$\iff \mathcal{M}, s \models \theta(\mathbf{cpres}(\beta^*, B))$$

by definition of θ. This ends the proof of Lemma 4.2. □

THEOREM 4.1 (Correctness). *NDPL$^+$ is correct with respect to the class of extended standard agent models.*

Proof. We have to show that **tPres** and **cPres** are valid in all *NPDL*$^+$-models. For **tPres**, $\mathcal{M}^+, s \models \mathbf{tpres}(\alpha, A) \iff \mathcal{M}, s \models \theta(\mathbf{tpres}(\alpha, A))$ (by Lemma 4.2) $\iff \mathcal{M}, s \models \theta(A) \supset \theta([\alpha]A)$ (by definition of θ). Then, trivially $\iff \mathcal{M}^+, s \models \theta(A) \supset \theta([\alpha]A)$ (because the valuation function for $\mathcal{L}$-formulas does not depend on the functions μ_t and μ_c), so $\iff \mathcal{M}^+, s \models A \supset [\alpha]A$ (by Lemma 4.1). The argument for **cPres** takes essentially the same steps. □

THEOREM 4.2 (Finite model property). *NPDL$^+$ has the finite model property.*

Proof. Take a *NPDL*$^+$-consistent formula A. Then, by Corollary 4.1, $\theta(A)$ is *NPDL*-consistent. So, $\theta(A)$ is true at a state s in a finite (standard agent) model $\mathcal{M}$ for *NPDL*. Hence, A is true at s in a finite (extended standard agent) model $\mathcal{M}^+$ for *NPDL*$^+$ by Lemma 4.2. □

As immediate consequences we obtain

COROLLARY 4.2 (Determination). *NPDL$^+$ is (weakly) determined by the class of all finite extended standard agent models.*

THEOREM 4.3 (Decidability). *NPDL$^+$ is decidable.*

5. Applications

This section is dedicated to applications of our framework. We do not describe how our approach handles the simple Yale Shooting problem (see Hanks & McDermott, 1987), because our approach yields no significant advantages here as compared, for example, with the handling of this problem in the situation calculus of McCarthy & Hayes (1969). Our approach is profitable as soon as (pre)conditions have to be 'transported' over action sequences of reasonable length. This is the case in *planning*. We give an example which involves more actions and properties than the shooting problem. Then, it is shown that *domain* and *plan constraints* are naturally encoded in *NPDL*$^+$ if the ongoing behavior of a plan is considered. The usefulness of these notions is illustrated by way of an example from the manufacturing domain. Finally, it is shown how the (temporal) properties of *temporal plan theory* (see Manna *et al.*, 1993) are expressed in *NPDL*$^+$. Finally, we compare our solution to some aspects the frame problem to monotonic and nonmonotonic solutions, respectively.

5.1. The frame problem: A TV Example

The following example confronts us with the (nowadays easy) problem of installing a TV set. Consider the following definitions concerning the actions and effects (preconditions) of the installation procedure:

`plug_in`	= (the action consisting of) plugging in the TV set
`connect`	= connecting the signal feed
`press_on`	= pressing the "on" button
`scan`	= scanning the signal for programs (by initiating the automated programming system (APS))
`select`	= selecting a program
`GUIDE`	= a TV guide is at hand
`CURRENT`	= the TV set is provided with a current
`SIGNAL`	= the tuner is provided with a signal
`TV_ON`	= the TV set is on
`PROG`	= the tuner is ready to receive all programs
`FAV_PROG`	= a (favorite) program is selected

The initial situation is described by a single statement,

(INI) $\mathtt{GUIDE}$.

We want to deduce that after every terminating execution of the action sequence consisting of plugging in the TV set, connecting the signal feed, pressing the "on" button, initiating APS, waiting (since waiting became so popular in the AI literature on the frame problem), and finally selecting a program, a (favorite) program is selected.

The assumptions on the preconditions and effects of actions are as follows:

(A1) $[\mathtt{plug\_in}]\mathtt{CURRENT}$

(A2) $[\nu](\mathtt{CURRENT} \supset [\mathtt{connect}]\mathtt{SIGNAL})$

(A3) $[\nu](\mathtt{CURRENT} \wedge \mathtt{SIGNAL} \supset [\mathtt{press\_on}]\mathtt{TV\_ON})$

(A4) $[\nu](\mathtt{CURRENT} \wedge \mathtt{SIGNAL} \wedge \mathtt{TV\_ON} \supset [\mathtt{scan}]\mathtt{PROG})$

(A5) $[\nu](\mathtt{CURRENT} \wedge \mathtt{SIGNAL} \wedge \mathtt{TV\_ON} \wedge \mathtt{PROG} \wedge \mathtt{GUIDE} \supset [\mathtt{select}]\mathtt{FAV\_PROG})$

For instance, the forth assumption states that in all situations we consider as possible, the tuner may receive all programs after initiating the automated programming system, if the TV set is provided with a current, the tuner is provided with a signal and the TV set is on.

Observe that all actions mentioned in the example have effects which figure as preconditions for all ensuing actions. For instance, if current breaks down at some intermediate state, all later actions will not be feasible. Here we have a situation where the **cpres** operator becomes important.

The following conditions ensure that preconditions of subsequent actions are met.

(C1) $[\nu](\mathbf{cpres}(\texttt{connect}; \texttt{press\_on}; \texttt{scan}; \omega, \texttt{CURRENT}))$

(C2) $[\nu](\mathbf{cpres}(\texttt{press\_on}; \texttt{scan}; \omega, \texttt{SIGNAL}))$

(C3) $[\nu](\mathbf{cpres}(\texttt{scan}; \omega, \texttt{TV\_ON}))$

(C4) $[\nu](\mathbf{cpres}(\omega, \texttt{PROG}))$

(C5) $[\nu](\mathbf{tpres}(\texttt{plug\_in}; \texttt{connect}; \texttt{press\_on}; \texttt{scan}; \omega, \texttt{GUIDE}))$

In condition (C4), the **cpres** operator collapses to the **tpres** operator, since ω is an elementary action term. In (C5), we only need **tpres** because the presence of a TV guide is interesting only when the agent selects a program; at the intermediate states of the installation procedure, some other agent may cary the guide to another room, etc.

The problem of installing a TV set may now be solved as follows ($\vdash$ is the deducibility relation of $NPDL^+$):

$$\text{(INI), (A1–5), (C1–5)} \vdash [\texttt{plug\_in}; \texttt{connect}; \texttt{press\_on}; \texttt{scan}; \omega; \texttt{select}]\texttt{FAV\_PROG} \quad (4)$$

The following argument sketch shows the relevant steps.

1. $[\texttt{plug\_in}](\texttt{CURRENT} \supset [\texttt{connect}]\texttt{SIGNAL})$
   (A2); **Any.4**.
2. $[\texttt{plug\_in}][\texttt{connect}]\texttt{SIGNAL}$
   (A1), 1; modal logic.
3. $[\texttt{plug\_in}][\texttt{connect}](\texttt{CURRENT} \wedge \texttt{SIGNAL} \supset [\texttt{press\_on}]\texttt{TV\_ON})$
   (A3); **Any.4**, **Comp**.
4. $[\texttt{plug\_in}][\texttt{connect}]\texttt{CURRENT}$
   (C1), (A1); **Any.4**, **cPres**, and modal logic.
5. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}]\texttt{TV\_ON}$
   2, 3, 4; modal logic.
6. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}](\texttt{CURRENT} \wedge \texttt{SIGNAL} \wedge \texttt{TV\_ON} \supset [\texttt{scan}]\texttt{PROG})$
   (A4); **Any.4**, **Comp**.
7. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}]\texttt{CURRENT}$
   (C1), (A1); **Any.4**, **cPres**, **Comp**, and modal logic.

8. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}]\texttt{SIGNAL}$
(C2), 2; **Any.4**, **cPres**, **Comp**, and modal logic.

9. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}]\texttt{PROG}$
5, 6, 7, 8; modal logic.

10. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}][\omega]\texttt{CURRENT}$
(C1), (A1); **Any.4**, **cPres**, **Comp**, and modal logic.

11. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}][\omega]\texttt{SIGNAL}$
(C2), 2; **Any.4**, **cPres**, **Comp**, and modal logic.

12. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}][\omega]\texttt{TV\_ON}$
(C3), 5; **Any.4**, **cPres**, **Comp**, and modal logic.

13. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}][\omega]\texttt{PROG}$
(C4), 9; **Any.4**, **cPres**, **Comp**, and modal logic.

14. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}][\omega]\texttt{GUIDE}$
(C5), (INI); **Any.1**, **tPres**, **Comp**, and modal logic.

15. $[\texttt{plug\_in}][\texttt{connect}][\texttt{press\_on}][\texttt{scan}][\omega](\texttt{CURRENT} \wedge \texttt{SIGNAL} \wedge \texttt{TV\_ON} \wedge$
$\wedge \texttt{PROG} \wedge \texttt{GUIDE} \supset [\texttt{select}]\texttt{FAV\_PROG})$
(A5); **Any.4**, **Comp**.

16. $[\texttt{plug\_in}; \texttt{connect}; \texttt{press\_on}; \texttt{scan}; \omega; \texttt{select}]\texttt{FAV\_PROG}$
10, 11, 12, 13, 14, 15; **Comp** and modal logic.

In our formalisation, the *combinatorial* problem and the *extended prediction* problem are solved by stating conditions which preserve facts over *sequences* of actions. Thereby we exploit the power of the preservation operators. We will discuss these problems at greater detail momentarily.

Now assume that the initial situation is described by the more elaborate formula

$$\texttt{GUIDE} \wedge \texttt{FRED\_ALIVE} \wedge \texttt{TWEETY\_FLYING}$$

where, apart from the TV guide being at hand, Fred is alive and Tweety is flying. We do *not* force these (possibly unrelated) facts be true after the installation procedure. So other facts may vary freely, dependent on activity of other agents; maybe Fred is shot or Tweety has its wing broken. In this way, the *overcommitment* problem is solved.

Concerning the *qualification* problem, we cannot offer a solution within our framework. We have no means to differentiate the treatment of explicit preconditions from the treatment of qualifications. Consider assumption (A2) of the installation example: the TV set being provided with current figures as an explicit precondition for the action consisting of connecting the signal feed. On the other hand, qualifications which may generally be assumed to hold (the antenna is not broken, there is reception, and so on) have to be made explicit too. If we take into account qualifications, our second premise has the rather unattractive format

$$[\nu](\texttt{CURRENT} \wedge \neg\texttt{BROKEN} \wedge \texttt{RECEPTION} \wedge \ldots \supset [\texttt{connect}]\texttt{SIGNAL}).$$

The problem is that we need to verify all qualifications which is notoriously inefficient. Here, nonmonotonic logics are clearly superior to our approach. They handle qualifications by the (defeasible) assumption that abnormalities (the antenna *is* broken, for instance) do not arise. For a probabilistic justification of this assumption, see Schurz (1994).

A problem related to the frame problem is the *ramification* problem (also called *consistency constraint* problem in Georgeff, 1987a). It is the problem of stating *all* (known) effects of actions. In case of the action of selecting a (favorite) program we should be able to say that the action selects the program without having to specify the consequences of the selection, for example, that watching a (favorite) program makes the person happy. But our formalism seems easily amenable to that problem. By means of the 'any' action we may express the following (self-explaining) formula

$$[\nu](\texttt{FAV\_PROG} \supset \texttt{HAPPY})$$

and thereby avoid to make the fact HAPPY an explicit result of selecting a favorite program. We simply state the fact as a domain constraint (for more discussion on constraints, see Subsection 5.2). A similar route to solve the ramification problem is taken by Ginsberg & Smith (1988a, pp. 170–172).

5.2. Planning

We are now going to demonstrate that $NPDL^+$ is an appropriate tool to formalize concepts essential to *planning*. Planning is a discipline that concentrates on problems adherent to the *specification*, *verification*, and *synthesis* of plans (see Manna & Waldinger, 1980; Rosenschein 1981; Georgeff, 1987b; Stephan & Biundo 1993). Plan *synthesis* concerns the composition ('synthesis') of a plan, usually a compound action term, to achieve some specified goal or goals. The process that checks if the proposed plan meets its specification, is called *verification*. By *specification* one understands certain properties that describe the desired behavior of the plan, in the first place, that the goal is satisfied.

In this subsection, our concern is *specification*. To obtain consistent axiomatizations of planning domains where the 'clean' behavior of a plan is important, *constraints* of essentially two kinds need to be considered. *Domain* constraints make assertions about the whole scenario, while the range of *plan* constraints is restricted to the plan under consideration. Constraints on plans typically express assertions about the *ongoing* behavior of a plan, and consequently regard intermediate states of the execution. It will come as no surprise that domain constraints and plan constraints in $NPDL^+$ are dealt with by the 'any' action and the concept of chronological preservation, respectively.

Domain constraints will have format $[\nu]C$. If C is a term-free formula, the constraint is called *static*, otherwise *dynamic* (see Rosenschein, 1981). For instance, the (dynamic) constraint "if a block is not clear then moving another block to its top fails" is expressed by the formula $[\nu](< \alpha > \top \supset A)$. Plan constraints will be expressed as 'boundary' conditions and defined via the **cpres** operator.

In order to illustrate these notions, we give an example from the manufacturing domain.

A manufacturing example

Consider an agent (robot) working at a car-manufacturing plant. The agent is supplied with a driver and a camera. The agent's task consists in turning screws until they are flush with the car-body.

The manufacturing domain is described as follows:

`screw`	= the action consisting of turning the screw
`FULL`	= the batteries are fully charged
$\mathtt{S}_n$	= the screwhead is n units apart from the car-body

Observe that the screwhead is flush with the car-body, if $n = 0$. Of course, we assume that the $\mathtt{S}_i$'s are mutually exclusive, that is, $\neg(\mathtt{S}_i \wedge \mathtt{S}_j)$ whenever $i \neq j$. We define $\tau^{\circ>} \stackrel{\text{def}}{=} \tau^{\circ}(\neg \mathtt{S}_0)$ and $\tau^{\circ=} \stackrel{\text{def}}{=} \tau^{\circ}(\mathtt{S}_0)$. Hence we assume our camera's operations of testing whether the screwhead is still apart from the car-body or not are *safe*.

The entire plan our robot has to execute is $((\tau^{\circ>}; \mathtt{screw})^*; \tau^{\circ=})$. Clearly this plan, *if* it terminates, will terminate in S_0 (else $\tau^{\circ=}$ will abort). But under which conditions will this plan terminate? The crucial point is to ensure that for both the driver and the camera to work the batteries must be loaded during the *whole* performance of the plan. For convenience, we define $\mathbf{bound}(\alpha, A) \stackrel{\text{def}}{=} A \wedge \mathbf{cpres}(\alpha, A)$, that is, A is a *boundary* condition during the performance of α. Hence our boundary condition is (the plan constraint)

$$\mathbf{bound}((\tau^{\circ>};\texttt{screw})^*;\tau^{\circ=},\texttt{FULL}). \tag{5}$$

We also assume the dynamic domain constraint

$$[\nu]((\texttt{FULL}\wedge\texttt{S}_n)\supset(<\texttt{screw}>\top\wedge[\texttt{screw}]\texttt{S}_{n-1})),\qquad(\text{for all} n>0) \tag{6}$$

that is, the action $\texttt{screw}$ applied to a state where the batteries are full and the screwhead is n units away from the car-body will terminate and after termination the screwhead is $n-1$ units away from the car-body. Since our test operation $\tau^{\circ>}$ is non-stationary, we must explicitly add a further assumption (dynamic domain constraint), namely that the test operation performed by the camera has no effect on the position of the screw, that is,

$$[\nu]\mathbf{tpres}(\tau^{\circ>},\texttt{S}_n).\qquad(\text{for all } n>0) \tag{7}$$

This sounds trivial, but clearly, a defective robot which clashes with the screwdriver whenever it tests whether $n>0$ such that the screw is turned some unit off the car, would never be able to execute the plan.

Within *NPDL*$^+$, we may now prove our desired result, namely

$$(5),(6),(7),\texttt{S}_n\vdash[(\tau^{\circ>};\texttt{screw})^n]\texttt{S}_0\wedge\ <(\tau^{\circ>};\texttt{screw})^n;\tau^{\circ=}>\top, \tag{8}$$

that is, *first*, whenever the action consisting of n times verifying $n>0$ and then turning the screwhead terminates, the screwhead is fixed (flush with the car-body) and *second*, the entire plan of performing this action *some number* of times and then verifying that the screwhead is fixed, will terminate if this number is n.

Here is a sketch of the proof. We first prove by induction on m that

$$<(\tau^{\circ>};\texttt{screw})^m>\top \tag{9}$$
$$[(\tau^{\circ>};\texttt{screw})^m]\texttt{S}_{n-m} \tag{10}$$

follow from the premises (5), (6), (7) and $\texttt{S}_n$ for each $0<m\leq n$.

$\mathbf{m=1:}$ $\texttt{S}_n$ implies (1.a) $<\tau^{\circ>}>\top$ by **snTest.3** and the assumption that the $\texttt{S}_i$'s are mutually exclusive, that is, $\texttt{S}_{n-m}\supset\neg\texttt{S}_0$ for $m<n$. (7) implies (1.b) $[\tau^{\circ>}]\texttt{S}_n$ by **Any.1**, **tPres** and $\texttt{S}_n$. (5) implies (1.c) $[\tau^{\circ>}]\texttt{FULL}$ by the definition of **bound** and the axiom for **cpres**. (1.b) and (1.c) give us (1.d) $[\tau^{\circ>}](\texttt{FULL}\wedge\texttt{S}_n)$ by modal logic. (1.a) and (1.b) yield (1.e) $<\tau^{\circ>}>\texttt{S}_n$ by modal logic. (1.c) and (1.e) give (1.f) $<\tau^{\circ>}>(\texttt{FULL}\wedge\texttt{S}_n)$ by modal logic. (6) implies (1.g) $[\tau^{\circ>}]((\texttt{FULL}\wedge\texttt{S}_n)\supset(<\texttt{screw}>\top\wedge[\texttt{screw}]\texttt{S}_{n-1}))$ by **Any.4**. By standard *PDL*, (1.f) and (1.g) imply $<\tau^{\circ>};\texttt{screw}>\top$, and (1.d) and (1.g) imply $[\tau^{\circ>};\texttt{screw}]\texttt{S}_{n-1}$, which proves the case for $m=1$.

m ⇒ m + 1: We assume (9) and (10) for a *given* m as induction hypothesis and show that (9) and (10) hold for $m+1$. (5) implies (2.) $[(\tau^{\circ>};\mathtt{screw})^m;\tau^{\circ>}]\mathtt{FULL}$ by the definition of **bound**, **cPres** and the *PDL*-theorem $[\alpha^*]A \supset [\alpha^n]A$. (3.) $[(\tau^{\circ>};\mathtt{screw})^m](S_{n-m} \supset <\tau^{\circ>}>\top)$ follows from **snTest.3** by modal logic and the assumption that the $\mathtt{S}_i$'s are mutually exclusive. (9) and (10) imply (4.) $<(\tau^{\circ>};\mathtt{screw})^m>\mathtt{S}_{n-m}$ and (4.) and (3.) imply (5.) $<(\tau^{\circ>};\mathtt{screw})^m;\tau^{\circ>}>\top$ by standard *PDL*. (6.) $[(\tau^{\circ>};\mathtt{screw})^m]\mathbf{tpres}(\tau^{\circ>},\mathtt{S}_n)$ follows from (7) by **Any.4**. (10) and (6.) imply (7.) $[(\tau^{\circ>};\mathtt{screw})^m;\tau^{\circ>}]\mathtt{S}_{n-m}$ by the axiom for **tpres**. Now, (5.) and (7.) imply (8.) $<(\tau^{\circ>};\mathtt{screw})^m;\tau^{\circ>}>\mathtt{S}_{n-m}$ by modal logic. From (2.), (8.) and modal logic it follows that (9.) $<(\tau^{\circ>};\mathtt{screw})^m;\tau^{\circ>}>(\mathtt{FULL}\wedge\mathtt{S}_{n-m})$. (6) implies (10.) $[(\tau^{\circ>};\mathtt{screw})^m;\tau^{\circ>}]((\mathtt{FULL}\wedge \mathtt{S}_{n-m}) \supset (<\mathtt{screw}>\top \wedge [\mathtt{screw}]\mathtt{S}_{n-m-1}))$ by **Any.4**. Finally, by standard *PDL*, (9.) and (10.) imply $<(\tau^{\circ>};\mathtt{screw})^{m+1}>\top$ and (7.) and (10.) imply $[(\tau^{\circ>};\mathtt{screw})^{m+1}]\mathtt{S}_{n-m-1}$, which was to prove.

Putting **m = n**, (10) gives us $[(\tau^{\circ>};\mathtt{screw})^n]\mathtt{S}_0$, the *first* conjunct we wanted to prove; and this together with (9) gives (11.) $<(\tau^{\circ>};\mathtt{screw})^n>\mathtt{S}_0$. **snTest.3** implies (12.) $[(\tau^{\circ>};\mathtt{screw})^n](\mathtt{S}_0 \supset <\tau^{\circ=}>\top)$ by modal logic. Finally, (11.) and (12.) yield $<(\tau^{\circ>};\mathtt{screw})^n;\tau^{\circ=}>\top$, the *second* conjunct we wanted to prove. □

Temporal plan theory

Recently, Manna *et al.* (1993) stressed the importance of 'safety' properties which are required to be true a *intermediate* states of plan execution. They call them *temporal properties*, and the proposed theory *temporal plan theory*. In our framework, *plan constraints* figure as dynamic logic counterparts to temporal properties (see also Rosenschein, 1981). As opposed to domain constraints, the range of plan constraints is restricted to the plan under consideration. Manna *et al.* (1993) organize temporal properties into a hierarchy of several classes, whereby each class is associated with a characteristic formula scheme. Within *NPDL*$^+$, all canonical schemes mentioned in Manna *et al.* (1993) are easily adapted. Recall our definition of a boundary condition, $\mathbf{bound}(\alpha, A) \stackrel{\mathrm{def}}{=} A \wedge \mathbf{cpres}(\alpha, A)$.

Let $\pi = \alpha_1;...;\alpha_j;\alpha_k;...;\alpha_n$ be a fixed plan such that for all $1 \leq i \leq n$, α_i is an elementary action term, β an arbitrary action term, and A, B term-free formulas.

Safety	$\mathbf{bound}(\pi, A)$
Guarantee	$A \vee \bigvee_{\gamma \in I(\pi)[\gamma]A}$
Obligation	the disjunction of *safety* and *guarantee*
Response	$\mathbf{bound}(\pi, A \supset <\beta> B)$
Persistence	$[\alpha_1; ...; \alpha_j]\mathbf{bound}(\alpha_k; ...; \alpha_n, A)$
Reactivity	the disjunction of *response* and *persistence*

The *safety* property has already been used, in the TV installation example and in the manufacturing example. *Guarantee* asserts that A is initially true or holds after some initial string of R_π, and *obligation* is the disjunction of the *safety* and the *guarantee* property. The *response* property states that always during the performance of the plan, where A holds, some state satisfying B is attainable via action β. *Persistence* is of importance when the agent aims to achieve several goals simultaneously. It may be the case that a certain subgoal A results from (successfully) executing the subplan $\alpha_1; ...; \alpha_j$, and after the final state of this subplan (also called a *protection point* by Waldinger, 1977), A is preserved throughout the rest of the plan, say, $\alpha_k; ...; \alpha_n$. The disjunction of the two previous properties is called *reactivity* (see also Pnueli, 1981).

5.3. Comparision to monotonic solutions

As mentioned in the Introduction, a solution to the combinatorial frame problem has to avoid writing down frame axioms of format $p \supset [\alpha]p$ for *each* fact-action pair p-α ($p \in \mathcal{P}_0$ a propositional variable, $\alpha \in \mathcal{A}_{el}$ an elementary action term). In the original formulation of McCarthy & Hayes (1969), for a domain with m propositional variables and n (elementary) action terms, $m \times n$ frame axioms are needed. [6] Recently, more efficient monotonic solutions have been proposed in the situation calculus. Schubert (1990) needs at most $2 \times m$ (explanation closure) axioms and Reiter's (1991) axiomatization of action and frame axioms requires $m + n$ axioms in total (m the number of fluents and n the number of actions).

If we make no assumptions on the domain, the number of frame axioms needed in our approach is equally large as in the original formulation of McCarthy & Hayes (1969). In the worst case, we have to

[6] McCarthy & Hayes (1969) formulate the frame problem in the language of situation calculus. In the terminology of situation calculus, propositional variables are called propositional *fluents*.

state $m \times n$ assertions of the form

$$[\nu]\mathbf{tpres}(\alpha, p)$$

where $p \in \mathcal{P}_0$, $\alpha \in \mathcal{A}_{el}$. With help of nondeterministic choice we may reduce this number to m axioms of the form

$$[\nu]\mathbf{tpres}\left(\sum_{\alpha \in \mathcal{A}_{el}} \alpha, p\right).$$

The above axioms will guarantee that every state of affairs specified in terms of $\mathcal{P}_0$ is preserved over every sequence of actions from $\mathcal{A}_{el}$.

However, this is rather unrealistic. In *planning* contexts, the termination (or success) of actions always depends on the outcomes of previous actions. Therefore, the ordering of elementary actions is essential for the success of the plan, while most permutations of these actions will lead into nonsensical or at least non-terminating action sequences. For instance, the TV installation problem is of that sort. As another example, take the problem of attaching a new wheel to a car. First, the car is lifted by means of a (lifting) jack. Second, the wheel is removed from the car. Third, a new wheel is attached to the car. Finally, the jack is removed. It is important to perform the individual actions in exactly that order. For instance, if the jack is removed after removing the wheel, the whole plan aborts.

In the context of planning, therefore, preservation axioms are not required for arbitrary action sequences, including all permutations, but only for that sequence of actions which constitutes the plan (or in the worst case, for all subsequences of that sequence). As seen from the proof in the TV installation example, we formulate our preservation conditions by the combination of the 'any' and the **cpres** operator. Assume $\alpha_1; ...; \alpha_n$ is the sequence defining the entire plan.[7] In the best case, we have just one precondition to be preserved over the entire plan. This holds in the manufacturing example (condition 5, batteries are fully charged). In the worst case, we have a distinct precondition p_i necessary for the termination of each action α_i and all of its successor actions.[8]

Then our preservation requirements are expressed by n axioms of the form

$$[\nu]\mathbf{cpres}(\alpha_i; ...; \alpha_n, p_i) \qquad (1 \leq i \leq n).$$

[7] Note that in our framework the α_i's need not be elementary but may themselves be composed.

[8] This is the simplest case; instead of p_i we might have also have a conjunction of propositional variables $p_1^i \wedge ... \wedge p_{m_i}^i$.

Typically, p_i will be the outcome of action α_{i-1}; but conditions p_i may also be externally given (like `GUIDE` in our TV example).

5.4. Comparision to nonmonotonic solutions

Nonmonotonic logics circumvent the computational problem of stating a large number of frame axioms by introducing a 'blanket' frame axiom which covers all (atomic) facts and (elementary) actions. If one abstracts from the different syntactical appearance of the blanket frame axiom in nonmonotonic logics, an informal rendering of the frame assertion might read (see Ginsberg, 1991):

(FA) if a fact p is true in a situation s and the action α is not abnormal with respect to p when performed in s, then p is still true after termination of α.

An action α is called *abnormal with respect to p in s* if α reverses the truth-value of p. According to the policy of *causal* minimization (of abnormalities), a fact changes its truth-value if *and only if* a terminating action (performed by the agent) causes it to do so (see Lifschitz, 1987). The circumscriptive approach (a nonmonotonic logic based on some sort of minimization) became very popular, since it solves the frame problem correctly and is robust as regards various aspects of the frame problem (see the 'standard' solution of Baker, 1991). After all, the circumscriptive policy usually minimizes the extension of the 'abnormality' predicate which seems unnatural to us when multiagent domains are considered.

Our approach does not depend on normality assumptions. On the contrary, we only force certain facts to persist, generally those which are preconditions to ensuing actions or effects of actions while other facts may *vary* due to activity of other agents.

Etherington *et al.* (1991) observe that nonmonotonic reasoning mechanisms fix the truth-value of too many facts. Hence, they introduce a methodology of *scoped* nonmonotonic reasoning which *restricts* the scope of reasoning to some pre-defined extension of a predicate. If applied to the blanket frame axiom (FA), their approach does not suffer from the overcommitment problem: instead of minimizing 'globally', that is, minimize changes of *all* facts if not forced otherwise, scoped circumscription only minimizes abnormalities concerning *certain* properties $p \in \mathcal{P}_{scope}$. Of course, a scope adequate for a specific problem has to be determined in advance. In the context of nonmonotonic solutions to the frame problem, Miller & Shanahan (1994) began to make up criteria how to *actually* determine scope. For instance, the criterion of *causal independence* states that actions can only affect facts (fluents in their terminology) within a specific region. In our approach, a process similar to the identification of a scope consists in finding a set of facts which are to be preserved over sequences of actions.

We may now substantiate the claim made in the Introduction: in our monotonic approach and (unscoped) nonmonotonic logic (possibly based on the minimization strategy) there is an opposite trade-off between a satisfactory solution to the overcommitment problem and the problems of extended prediction and qualification. To solve the extended prediction problem, nonmonotonic logic employs a blanket frame axiom which by default forces the domain to be *maximally* stable, that is, all facts remain unchanged over actions. Deviations from stability must be made explicit by stating *abnormality* assertions. In this way, instability is introduced to the domain. In scoped nonmonotonic logic, in addition to instability, *ignorance* of the truth-value of certain facts can be introduced by scope restrictions. We approach the extended prediction problem from the opposite side. If we are to maintain facts over actions, we have to state frame axioms explicitly. As a consequence of the monotonic character of our approach, we remain ignorant on facts not (syntactically) appearing in the frame axioms. We have the choice of formulating weaker or stronger frame axioms. Weak frame axioms allow for independent activity of other agents, that is, a *minimum* amount of facts is preserved. On the other hand, stronger frame axioms impose increasing stability on the domain.

Observe that (chronological) preservation formulas are very compact representations of frame axioms. Frame axioms which 'transport', say an atomic fact over an elementary action are *monotonic* consequences of $[\nu](\mathbf{cpres}(\alpha, A))$ format formulas (α a compound action term, A an arbitrary formula). Thereby, our approach avoids compuational difficulties resulting from repeatedly applying a blanket frame axiom. For a different strategy to avoid the computational problem with complex facts and sequences of actions, see Ginsberg (1991).

Acknowledgements

The authors would like to thank the anonymous referees for their helpful comments. Helmut Prendinger was supported by a grant from the *Int'l Forschungszentrum Salzburg*.

References

Baker, Andrew B., 1991, "Nonmonotonic reasoning in the framework of situation calculus," *Artificial Intelligence*, **49**, 5–23.

Chellas, Brian F., 1980, *Modal Logic: an introduction*, Cambridge: Cambridge University Press.

Etherington, David W. and Kraus, Sarit and Perlis, Donald, 1991, "Nonmonotonicity and the scope of reasoning," *Artificial Intelligence*, **52**, 221–261.

Fine, Kit and Schurz, Gerhard, 1995, "Transfer theorems for stratified multimodal logics," In J. Copeland, editor, *Logic and Reality. Essays in Pure and Applied Logic. In Memory of Arthur Prior*, Oxford University Press.

Fischer, Michael J. and Ladner, Richard E., 1979, "Propositional dynamic logic of regular programs," *Journal of Computer and System Sciences*, **18**, 194–211.

Gargov, George and Passy, Solomon, 1990, "A note on boolean modal logic," In Petio Petrov Petkov, editor, *Proceedings of the Summer School and Conference on Mathematical Logic*, Bulgaria, 299–309.

Gargov, George, Passy, Solomon and Tinchev, Tinko, 1987, "Modal environment for boolean speculations," In Dimiter G. Skordev, editor, *Proceedings of an Advanced International Summer School and Conference on Mathematical Logic and Its Applications, 1986*, Bulgaria, 253–263.

Georgeff, Michael P., 1987a, "Many agents are better than one," In F. M. Brown, editor, *Proceedings of the 1987 Workshop on the Frame Problem in Artificial Intelligence*, Lawrence, Kansas, 59–75.

Georgeff, Michael P., 1987b, "Planning," *Annual Review of Computer Science*, **2**, 359–400.

Ginsberg, Matthew L., 1991, "Computational considerations in reasoning about action," In James Allen, Richard Fikes, and Erik Sandevall, editors, *Proceedings of the Second International Conference On Principles of Knowledge Representation and Reasoning*, Cambridge, Massachusetts, 250–261.

Ginsberg, Matthew L. and Smith, David E., 1988a, "Reasoning about actions I: a possible worlds approach," *Artificial Intelligence*, **35**, 165–195.

Ginsberg, Matthew L. and Smith, David E., 1988b, "Reasoning about actions II: the qualification problem," *Artificial Intelligence*, **35**, 311–342.

Goldblatt, Robert, 1987, *Logics of Time and Computation*, Center for the Study of Language and Information Lecture Notes 7, Leland Stanford Junior University, Stanford.

Goranko, V. and Passy, S., 1992, "Using the universal modality: Gains and questions," *Journal of Logic and Compution*, **2**(1), 5–30.

Hanks, St. and McDermott, D., 1987, "Nonmonotonic logic and temporal projection," *Artificial Intelligence*, **33**, 379–412.

Harel, David, 1984, "Dynamic logic," In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic. II: Extensions of Classical Logic*, Reidel, Boston, MA, 497–604.

Harel, David, 1987, *Algorithmics. The Spirit of Computing*, Addison-Wesley Publishing Company, Inc., Wokingham–England, Reading–Massachusetts.

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

Hughes, G.E. and Cresswell, M.J., 1984, *A Companion to Modal Logic*, Methuen & Methuen, London, New York.

Katsuno, Hirofumi and Mendelzon, Alberto O., 1991, "Propositional knowledge base revision and minimal change," *Artificial Intelligence*, **52**, 263–294.

Kautz, Henry A., 1982, "A first-order dynamic logic for planning," Technical report, Computer Systems Research Group (CSRG), University of Toronto.

Lifschitz, Vladimir, 1987, "Formal theories of action," In Matthew L. Ginsberg, editor, *Readings in Nonmonotonic Reasoning*, Morgan Kaufmann Publishers, Inc., Los Altos, California, 410–432.

Makinson, David, 1993, "Five faces of minimality," *Studia Logica*, **52**, 339–379.

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

Manna, Zohar and Waldinger, Richard, 1980, "A deductive approach to program synthesis," In B. L. Webber and N.J. Nilsson, editors, *Readings in Artificial Intelligence*, Morgan Kaufmann Publishers, Inc., Los Altos, California, 1981, 141–172. Originally published in 1980.

McCarthy, John and Hayes, Patrick J., 1969, "Some philosophical problems from the standpoint of artificial intelligence," In B. L. Webber and N.J. Nilsson, editors, *Readings in Artificial Intelligence*, Morgan Kaufmann Publishers, Inc., Los Altos, California, 1981, 431–450. Originally published in 1969.

Miller, Bob and Shanahan, Murray, 1994, "Narratives in the situation calculus," *Journal of Logic and Computation*, **4**(5), 513–530.

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

Morreau, Michael, 1992, "Planning from first principles," In Peter Gärdenfors, editor, *Belief Revision*, Cambridge University Press, Cambridge, 204–219.

Passy, Solomon and Tinchev, Tinko, 1991, "An essay in combinatory dynamic logic," *Information and Computation*, **93**, 263–332.

Pnueli, Amir, 1981, "The temporal semantics of concurrent programs," *Theoretical Computer Science*, **13**, 45–60.

Pratt, Vaughan R., 1976, "Semantical considerations on Floyd-Hoare logic," In *Proceedings of the 17th Annual IEEE Symposium on Foundations of Computer Science, 1976*, 109–121.

Pratt, Vaughan R., 1978, "A practical decision method for propositional dynamic logic," In *Proceedings 10th Annual Association for Computing Machinery (ACM) Symposium on Theory of Computing*, 326–337.

Pratt, Vaughan R., 1980, "Applications of modal logic to programming," *Studia Logica*, **39**, 257–274.

Reiter, Raymond, 1980, "A logic for default reasoning," *Artificial Intelligence*, **13**, 81–132.

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

Rosenschein, Stanley J., 1981, "Plan synthesis: a logical perspective," In *Proceedings of the Seventh International Joint Conference of Artificial Intelligence*, Vancouver, 331-337.

Schubert, Lenhart, 1990, "Monotonic solution of the frame problem in the situation calculus: an efficient method for worlds with fully specified actions," In H. E. Kyburg, Jr., R. P. Loui, and G. N. Carlson, editors, *Knowledge Representation and Defeasible Reasoning*, Kluwer Academic Publishers, Dordrecht, Boston, London, 1990, 23–67.

Schurz, Gerhard, 1994, "Probabilistic justification of default reasoning," In B. Nebel and L. Dreschler-Fischer, editors, *KI-94, 18th German Annual Conference on Artificial Intelligence: Advances in Artificial Intelligence*, Springer, 248–259.

Segerberg, Krister, 1980, "Applying modal logic," *Studia Logica*, **39**, 275–295.

Segerberg, Krister, 1982, "A completeness theorem in the modal logic of programs," In Tadeusz Traczyk, editor, *Universal Algebra and Applications*, volume 9 of *Banach Center Publications*, Warsaw, 31–46.

Shoham, Yoav, 1988, *Reasoning about Change. Time and Causation from the Standpoint of Artificial Intelligence.* The MIT Press, Cambridge, Massachusetts, London, England.

Stephan, Werner and Biundo, Susanne, 1993, "A new logical framework for deductive planning," In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence ICJAI-93*, Chambëry, France, 32–38.

Waldinger, Richard, 1977, "Achieving several goals simultaneously," In B. L. Webber and N. J. Nilsson, editors, *Readings in Artificial Intelligence*, Morgan Kaufmann Publishers, Inc., Los Altos, California, 1981, 250–271. Originally published in 1977.