

Planning from Second Principles

Jana Koehler
German Research Center for AI (DFKI)
Stuhlsatzenhausweg 3
D-66123 Saarbrücken
Germany
e-mail: koehler@dfki.uni-sb.de

Abstract

Planning from second principles by reusing and modifying plans is one way of improving the efficiency of planning systems. In this paper, we study it in the general framework of deductive planning and develop a logical formalization of planning from second principles, which relies on a systematic decomposition of the planning process.

Deductive inference processes with clearly defined semantics formalize each of the subtasks a second principles planner has to address. Plan modification, which comprises matching and adaptation tasks, is based on a deductive approach yielding provably correct modified plans. Description logics are introduced as query languages to plan libraries, which leads to a novel and efficient solution to the indexing problem in case-based reasoning.

Apart from sequential plans, this approach enables a planner to reuse and modify complex plans containing control structures like conditionals and loops.

1 Introduction

Planning from first principles generates plans from “scratch”. The planner inspects its set of available actions and tries to construct a plan that achieves the desired goal with respect to specified preconditions. A serious limitation of first principles planners is the invariable nature of the planning process over time: If a planner receives exactly the same planning problem, it will repeat exactly the same planning operations. In other words, it is unable to benefit from experience that can be drawn from previous planning processes.

Approaches to *planning from second principles* try to overcome this limitation of planning from scratch by reusing and modifying previously generated plans. From a complexity-theory point of view, we cannot hope to prove efficiency gains in the worst case [Nebel and Koehler, 1995], since the reuse of plans comprises subtasks that are not computationally easier than plan generation. But in many practical applications it seems to be more reasonable to reuse existing plans than generating a new one. The current state of the art comprises a variety of approaches that tackle the problems from a cognitive point of view (cf. [Kolodner, 1993] for a summary of approaches) or in the framework of STRIPS-based planning, cf. [Kambhampati and Hendler, 1992; Hanks and Weld, 1992a; Veloso, 1994].

In using a deductive framework, we present a formal approach to planning from second principles, which makes no commitments to particular planning formalisms and application domains. We formalize the whole reuse process in a unique logical framework including the *retrieval* of candidate plans from a library as well as the problem of plan *modification*.

Plan modification is based on deductive inference processes that yield provably correct modified plans. It comprises two subtasks: First, the task of *matching* an old and a new planning problem in order to determine their similarities and differences. Secondly, the task of *refitting* the reused plan to accommodate new requirements. As a new issue in plan modification, we discuss the reuse and refitting of control structures occurring in plans, like case analyses and loops, which introduce qualitatively new problems.

As for the plan library, we propose a *hybrid* knowledge representation formalism linking the planning logic with a *description logic*. In this approach, description logics are introduced as query languages to large knowledge bases or case libraries. Their use leads to well-defined abstraction, retrieval, and update procedures that possess theoretical and practical properties of interest. In particular, description logics enable us to develop efficient and complete approximation algorithms for the matching problem [Nebel and Koehler, 1995], which are guaranteed to retrieve all plans from the library that solve a planning problem.

Finally, the formal framework allows us to prove important properties like the correctness and completeness of the underlying inference procedures. Besides this, the approach provides the foundation for the implemented plan reuse system MRL¹, which has been developed as an integrated part of the PHI planner [Bauer *et al.*, 1993].

The paper is organized as follows: We begin in Section 2 with a summary of the logical formalisms that are used by MRL. The section also contains a short introduction into deductive planning as well as a short overview of the PHI planner. Section 3 introduces a four-phase model as the foundation of second principles planning. The model supports a temporal view as well as a task-specific view of the second principles planning process. A logical formalization of each phase provides the theoretical basis for the system MRL that is described in the subsequent sections. Sections 4 and 6 are devoted to the inference procedures working on the plan library, while in Section 5 the deductive approach to plan modification is presented. Finally, in Section 7 we review related work and propose a systematic categorization of the various principles and design decisions underlying second principles planners. We summarize the main properties of MRL in the light of this categorization.

2 Formal Preliminaries

Deductive planning is a longstanding variant of artificial intelligence planning, whose origins go back to QA3 [Green, 1969]. To generate plans deductively, constructive proofs of formal plan specifications are performed, i.e., “to construct a plan that will meet a specified condition, one proves the existence of a state in which the condition is true”, cf. [Manna and Waldinger, 1987], page 14. Usually, this requires us to constructively prove plan specifications of the form

$$\forall s_0 \forall a \exists z Q[s_0, a, z]$$

where s_0 denotes the initial state, a is an argument or input parameter, and z is a planvariable representing the plan term that has to be constructed [Manna and Waldinger, 1987].

Two properties of deductive planners are particularly interesting when studying planning from second principles: First, plans are provably correct. Once provided with a correct axiomatization of a particular application domain and the actions which can be performed in this domain, a deductive planner generates plans that are guaranteed to work. Preserving this property during plan modification is a real challenge—how can we ensure that removing, reordering or adding actions leads to a sound plan, that solves the planning problem at hand?

Secondly, deductive planning has been closely related to program synthesis right from its origins. Plans are viewed as programs and consequently, they contain control structures like *if-then-else* and *while*-loops. While it is very rare that classical planners generate such complex plans, many deductive planning systems are able to generate control structures in plans. The retrieval and modification of plans containing conditionals and loops is therefore another challenge we are going to address.

A deductive planning system—like any other classical planning system—is faced with a search space of enormous complexity and controlling search is a difficult problem. For a deductive planner, this means controlling the inferences in the underlying logic in such a way that plans can be constructed automatically. This involves enabling a system to correctly “guess” appropriate instantiations of planvariables and to provide it with mechanisms that decide which inference rules apply to certain logical formulae. The difficulty of this task led to a temporary abandonment of deductive techniques. Many (but not all) classical planners, which have been developed in the meanwhile, sacrifice soundness by ignoring frame or ramification problems in order to reduce search complexity.

¹ MRL stands for *Modification and Reuse in Logic*.

Recently, deductive planning has seen a renaissance. On one hand, various planning logics that allow for an efficient solution to the frame problem have been carefully devised, e.g., [Biundo *et al.*, 1992; Reiter, 1993]. On the other hand, new ways of controlling a deductive planner by using *tactics* have been developed, e.g., [Traverso *et al.*, 1992; Biundo *et al.*, 1992]. The use of tactics in deductive planning is inspired by *tactical theorem proving* [Constable, 1986; Heisel *et al.*, 1990; Paulson, 1990]. Tactics support the declarative representation of control knowledge and guide the inference so that plans can be automatically constructed when proving the plan specification. In practice, this implies giving up completeness in order to purchase tractability. As we will demonstrate, tactics also play an important role in implementing plan modification.

The examples we are going to use to illustrate planning from second principles are taken from the PHI planner. PHI has been designed to perform plan generation and plan recognition tasks in command language environments, e.g., software systems. It provides a logic-based kernel that can be used to develop intelligent help systems supporting users of software. A prototype application of PHI is the UNIX *mail domain* where objects like *messages* and *mailboxes* are manipulated by actions like *read*, *delete*, and *save*.

The system uses the so-called Logical Language for Planning LLP [Biundo *et al.*, 1992] as the underlying planning formalism. The logic LLP reflects the specific requirements of command language environments. For example, the basic actions which occur in plans are the elementary statements of the application system language. Furthermore, control structures like loops and conditionals are available as defined operators, which allows to describe complex user actions on the level of the logical formalism.

2.1 The Logical Language for Planning LLP

LLP is a modal temporal logic with interval-based semantics, which combines features of *choppy logic* [Rosner and Pnueli, 1986] with a *temporal logic for programs* [Kröger, 1987]. Interval-based temporal logics have been proposed as appropriate formalisms to describe the behavior of programs or plans, e.g., [Moszkowski and Manna, 1983; Gabbay, 1989]. Plans can be decomposed into successively smaller periods or intervals of, e.g., subplans or actions. The intervals provide a convenient framework for introducing quantitative timing details. State transitions can be characterized by properties relating the initial and final values of variables over intervals of time.

The basis of LLP is a many-sorted first-order language with equality. It distinguishes *local variables*, the value of which may vary from state to state and *global variables* which are the usual logical variables. Local variables are borrowed from *programming logics* where they correspond to program variables. LLP provides the modal operators \circ (*next*), \diamond (*sometimes*), \square (*always*), and the binary modal operator $;$ (*chop*). In the following, we shortly review the main properties of LLP as introduced in [Biundo *et al.*, 1992].

2.1.1 Syntax and Semantics

A state σ_i is a pair $\sigma_i = (\sigma_i^1, \sigma_i^2)$ where σ_i^1 is a valuation assigning domain elements to local variables, while σ_i^2 indicates the command to be executed in state σ_i . Note that only the values of local variables may change from state to state. Function and predicate symbols are *rigid*, i.e., their interpretation does not vary over time. An interval w is a nonempty finite or infinite sequence of states $\langle \sigma_0 \sigma_1 \dots \rangle$. W denotes the set of all intervals. The length of an interval w is defined as

$$|w| = \begin{cases} \omega, & \text{if } w \text{ is infinite} \\ n, & \text{if } w = \langle \sigma_0 \sigma_1 \dots \sigma_n \rangle \end{cases}$$

Observe that $|w| = 0$ iff $w = \langle \sigma_0 \rangle$ is a *singleton* containing only one state. Intuitively, the length of an interval does not represent the number of states this interval contains, but the number of possible state transitions. The *immediate accessibility* on intervals is defined as the subinterval relationship R with

$$w R w' \text{ iff } w = \langle \sigma_0 \sigma_1 \sigma_2 \dots \rangle \text{ and } w' = \langle \sigma_1 \sigma_2 \dots \rangle.$$

The relation R is not serial, i.e., $\forall w \exists w' w R w'$ does not hold since an interval of length zero has no successor. R^* denotes the transitive and reflexive closure of R . The *composition* is defined as a partial function over the set of intervals W :

$$w \circ w' = \begin{cases} w, & \text{if } w \text{ is infinite} \\ \langle \sigma_0 \dots \sigma_{n-1} \sigma_n \sigma_{n+1} \dots \rangle, & \text{if } w = \langle \sigma_0 \dots \sigma_{n-1} \sigma_n \rangle \text{ and} \\ & w' = \langle \sigma_n \sigma_{n+1} \dots \rangle \end{cases}$$

Global variables are interpreted by mapping them to domain elements using a valuation function. The value of a local variable in an interval w for a particular interpretation is given by its value in the initial state of the interval. The satisfiability relation \models for modal-free formulae is defined as in classical first-order logic. F and T denote the propositional constants *false* and *true*, respectively. The special-purpose predicate ϵx , which takes a command term as the argument, denotes the action to be executed, i.e., $w \models_{\mathcal{I}} \epsilon x(t)$ iff $\mathcal{I}(t) = \sigma_0^2$. For the modal operators we define:

- $w \models_{\mathcal{I}} \bigcirc \phi$ iff $w' \models_{\mathcal{I}} \phi$ for all $w' \in W$ with $w R w'$
- $w \models_{\mathcal{I}} \diamond \phi$ iff $w' \models_{\mathcal{I}} \phi$ for some $w' \in W$ with $w R^* w'$
- $w \models_{\mathcal{I}} \square \phi$ iff $w' \models_{\mathcal{I}} \phi$ for all $w' \in W$ with $w R^* w'$
- $w \models_{\mathcal{I}} \phi ; \psi$ iff there are $w', w'' \in W$, with $w = w' \circ w''$, w' finite and $w' \models_{\mathcal{I}} \phi$ and $w'' \models_{\mathcal{I}} \psi$

For example, $\bigcirc F$ holds in an interval w iff w has length 0, i.e., it is a singleton. More generally, $\bigcirc^n F$ holds in w iff w has at most n states, that is iff w has at most length $n-1$. A formula $\phi \wedge \neg \bigcirc F \wedge \bigcirc \bigcirc F ; \bigcirc \square \psi$ holds in an interval $\langle \sigma_0 \sigma_1 \sigma_2 \sigma_3 \dots \rangle$ if

- ϕ holds in the subinterval $\langle \sigma_0 \sigma_1 \rangle$ and
- $\bigcirc \square \psi$ holds in the subinterval $\langle \sigma_1 \sigma_2 \sigma_3 \dots \rangle$, i.e., ψ holds in all subintervals $\langle \sigma_n \dots \rangle$ with $n \geq 2$.

2.1.2 Plans and Plan Specifications in LLP

Certain types of formulae in LLP are viewed as plans and plan specifications.

Definition 1 Let t be a command term, ϵ an atomic formula, ϕ and ψ consistent plan formulae. Plans are all formulae of the form²:

- $\epsilon x(t)$,
- $\phi ; \psi$,
- **if** ϵ **then** ϕ **else** ψ ,
- **while** ϵ **do** ϕ **od** ; ψ .

The conditional **if** ϵ **then** ϕ **else** ψ stands for the formula $[\epsilon \rightarrow \phi] \wedge [\neg \epsilon \rightarrow \psi]$. The **while** operator is defined by the following axiom:³

$$\mathbf{while} \ \epsilon \ \mathbf{do} \ \phi \ \mathbf{od} \ ; \ \psi \ \leftrightarrow \ \mathbf{if} \ \epsilon \ \mathbf{then} \ \phi \ ; \ [\mathbf{while} \ \epsilon \ \mathbf{do} \ \phi \ \mathbf{od} \ ; \ \psi] \ \mathbf{else} \ \psi$$

²Currently, PHI uses an extended class of plan formulae comprising nonlinear plans that contain temporal abstractions (“*sometimes* execute an action”). For the purpose of this paper, we restrict the class of plan formulae to those plans that are used in the examples.

³To complete the recursive definition of **while**, an extra semantic condition is necessary, which amounts to a smallest fixpoint construction as in [Stephan and Biundo, 1995] for example. Since in our application the formula ϕ is restricted to be a valid plan formula not containing any **while** structure, the operational view of the **while** definition is sufficient.

The atomic actions available to the planner are the elementary commands of the UNIX mail system. They are axiomatized like assignment statements in programming logics. Changes of state caused by executing an action are reflected in a change of the values of local variables, which represent the mailboxes in the domain under consideration. For example, the axiomatization of the *delete*-command which deletes a message x in a mailbox $mbox$ ⁴ reads

$$\forall x [\text{open_flag}(mbox) = T \wedge \text{delete_flag}(msg(x, mbox)) = F \wedge \\ \text{ex}(\text{delete}(x, mbox)) \\ \rightarrow \text{O delete_flag}(msg(x, mbox)) = T]$$

The state of a mailbox is represented with the help of *flags*. As a precondition, the *delete*-command requires that the mailbox $mbox$ is open, i.e., its *open_flag* yields the value *true* (T) and that the message x has not yet been deleted, i.e., its *delete_flag* yields the value *false* (F).⁵ As an effect, the action sets the *delete_flag* of message x in mailbox $mbox$ to the value *true* in the next state. In general, PHI uses more general second-order axiom schemata that are instantiated during the proof process. The axiom schemata describe effects as well as frame conditions, which leads to a representational solution of the frame problem, since only one axiom schemata is necessary for each action [Biundo *et al.*, 1992].

Definition 2 *Plan specifications are universally quantified LLP formulae of the form*

$$\text{Plan} \wedge \text{preconditions} \rightarrow \diamond \text{goal},$$

i.e., if the Plan is carried out in the initial state where the preconditions hold then a state will be achieved satisfying goal. Plan is a metalogical variable standing for the plan formula that has to be generated by constructively proving the plan specification, i.e., Plan meets the specification iff Plan \wedge preconditions $\rightarrow \diamond$ goal is true.

Plan specification formulae have to obey various syntactic restrictions. Preconditions are described by a modal-operator free first-order formula containing negation, conjunction, disjunction, and a limited form of implication and universal quantification. Goals may be described by formulae containing nested *sometimes* operators, like $\diamond[\phi \wedge \diamond\psi]$, conjunction, and also a limited form of implication and universal quantification.

Let us consider three specifications and example plans that will be used throughout this paper. The first specification S_{P1} specifies the planning problem: “read and delete a message m in the mailbox *mybox*”. As preconditions, we assume that the mailbox *mybox* has already been opened and that the message m has not yet been deleted.

$$\mathbf{SP1} : \quad \begin{array}{l} \text{Plan}_{P1} \wedge \\ \text{open_flag}(mybox) = T \wedge \text{delete_flag}(msg(m, mybox)) = F \\ \rightarrow \diamond [\text{read_flag}(msg(m, mybox)) = T \wedge \\ \diamond [\text{delete_flag}(msg(m, mybox)) = T]] \end{array}$$

SP1 specifies temporary goals with the help of nested *sometimes* operators, i.e., goals that have to be achieved at some point and not necessarily in the end. It requires the message to be read first and then deleted. The plan **P1** solving this specification is a simple sequence containing the actions *type* and *delete*.

$$\mathbf{P1}: \quad \text{ex}(\text{type}(m, mybox)); \text{ex}(\text{delete}(m, mybox))$$

In the second specification S_{P2} , we have the same goals as specified in S_{P1} , but formulated here as a conjunctive goal.

$$\mathbf{SP2} : \quad \begin{array}{l} \text{Plan}_{P2} \wedge \text{delete_flag}(msg(x, mbox)) = F \\ \rightarrow \diamond [\text{read_flag}(msg(x, mbox)) = T \wedge \\ \text{delete_flag}(msg(x, mbox)) = T] \end{array}$$

⁴Constants begin with capital letters, while variables are written in lower case.

⁵The reader may note the difference between the *boolean* constants T and F , which are assigned as values to local variables and the *propositional* constants \top and \perp , which are formulae as in OOF for example.

$$\frac{\Gamma \Rightarrow \text{O}\phi \wedge \neg\text{O}F, \Delta}{\Gamma \Rightarrow \diamond\phi, \Delta} \text{ sometimes_to_next}$$

Definition 5 An axiom is a sequent of form $\Gamma, A \Rightarrow A, \Delta$, i.e., antecedent and succedent contain at least one common formula.

If a sequent derivation of a formula does not terminate in a formulae set comprising only axioms, then the proof of this formula has failed. But if there is a subtree of the derivation tree which contains only axioms as leaves, then we have a partial proof stating the validity of a subformula.

2.2 Description Logics

A main problem during planning from second principles is the identification of an appropriate reusable plan. This can be achieved by *matching* a given plan specification formula against a set of plan specification formulae stored in a plan library. As we have seen, plan specifications are very complex logical descriptions. Therefore, we cannot hope to find an *efficient* matching algorithm comparing two plan specifications by a *partial* or *best match*. The main idea is therefore, to shift from the *source* formalism LLP to a *target* formalism, which can be used to represent *abstracted* plan specifications. As we will show, a *partial* match in the source formalism can be easily reduced to an *exact* match in the target formalism.

This approach uses *concept description logics* (also called terminological logics) as the ideal target formalism for the representation of abstract knowledge. Furthermore, we introduce description logics as *query languages* to plan libraries by formalizing the retrieval problem as a *classification* task. With that, description logics lead to a novel and efficient solution to the *indexing problem* in case-based reasoning.

2.2.1 Syntax and Semantics

Description logics comprise a whole family of logical languages with different degrees of expressivity and inference services of different computational complexity.⁷ In the following, we define a language of restricted expressivity which is sufficient for the representation of abstracted plan specifications. It can be considered as a subset of various well-known description logics such as *ALC* [Schmidt-Schauß and Smolka, 1991], *CLASSIC* [Borgida and Patel-Schneider, 1994], and *KRIS* [Baader and Hollunder, 1991]. The basic “building blocks” are *concepts* and *roles*, which denote subsets of the objects of the application domain \mathcal{D} and binary relationships over \mathcal{D} connecting objects, respectively. The concept \top (*top*) denotes the whole domain, while \perp (*bottom*) denotes the empty set. Concepts are defined intensionally in terms of descriptions that specify the properties an object must satisfy to belong to the concept.

Definition 6 Let C and D be syntactical variables for concept descriptions and R, R_1, R_2 for role names. The following concept descriptions can be formed:

$$\begin{array}{ll} C \sqcap D & (\text{conjunction}) \\ C \sqcup D & (\text{disjunction}) \\ \neg C & (\text{negation}) \\ \exists R.C & (\text{existential role restriction}) \\ R_1 \circ R_2 & (\text{role composition/role chain}) \end{array}$$

Besides the construction of complex concepts, new concept descriptions can be defined with the help of terminological axioms.

Definition 7 Let A be a concept name and D a concept description, then $A \doteq D$ is a terminological axiom.

Primitive concepts are all concepts which are not contained on left sides of terminological axioms. They remain undefined in the terminology and can be considered as the basic terms on which other concept descriptions can be built.

⁷For a good introduction into description logics see for example [Nebel, 1990a].

The formal model theoretic semantics of description logics uses the set of objects, the domain \mathcal{D} , for the interpretation of concept descriptions. Concepts denote a (sub)set of objects, while each role denotes a set of object pairs. These sets are called *extensions* of concepts and roles.

Definition 8 An interpretation \mathcal{I} consists of a domain \mathcal{D} and an extension function \mathcal{E} mapping each concept name A to a subset $\mathcal{E}[A]$ from \mathcal{D} and each role name R to a binary relation $\mathcal{E}[R]$ over \mathcal{D} .

The terminological relationships between concept descriptions can be used to determine the necessary relationships between their extensions.

Definition 9 Let C be the set of concept symbols and R be the set of role symbols. \mathcal{D} is an arbitrary set and \mathcal{E} a function:

$$\mathcal{E} = \begin{cases} C & \rightarrow 2^{\mathcal{D}} \\ R & \rightarrow 2^{\mathcal{D} \times \mathcal{D}} \end{cases}$$

\mathcal{E} is an extension function iff the following equations hold:

$$\begin{aligned} \mathcal{E}[\top] &= \mathcal{D} \\ \mathcal{E}[\perp] &= \emptyset \\ \mathcal{E}[C \sqcap D] &= \mathcal{E}[C] \cap \mathcal{E}[D] \\ \mathcal{E}[C \sqcup D] &= \mathcal{E}[C] \cup \mathcal{E}[D] \\ \mathcal{E}[\neg C] &= \mathcal{D} \setminus \mathcal{E}[C] \\ \mathcal{E}[\exists R.C] &= \{x \in \mathcal{D} \mid \text{exists } y \in \mathcal{D} \langle x, y \rangle \in \mathcal{E}[R] \text{ and } y \in \mathcal{E}[C]\} \\ \mathcal{E}[R_1 \circ R_2] &= \{x, y \in \mathcal{D} \mid \text{exists } z \in \mathcal{D} \langle x, z \rangle \in \mathcal{E}[R_1] \text{ and } \langle z, y \rangle \in \mathcal{E}[R_2]\} \end{aligned}$$

The interpretation of a terminological axiom is the equation $\mathcal{E}[A] = \mathcal{E}[D]$.

Definition 10 A concept C is satisfiable (also denoted as consistent) iff it has a nonempty extension $\mathcal{E}[C] \neq \emptyset$.

2.2.2 Subsumption and Classification

The main inference procedure provided in terminological systems is *subsumption*, which determines whether one concept description is more general than another.

Definition 11 Let \mathcal{T} be a terminology and C, D be concepts. D subsumes C in \mathcal{T} ($C \sqsubseteq_{\mathcal{T}} D$) iff $\mathcal{E}[C] \subseteq \mathcal{E}[D]$ holds in all models of \mathcal{T} .

The computation of all subsumption relationships between a set of concept descriptions is called *classification*. Since we intend to formalize retrieval as a classification task, we are interested in grounding it on a complete and efficient subsumption algorithm, which runs in polynomial time depending on the number of concept forming operators in a concept description. Due to the inherent intractability of subsumption [Nebel, 1990b], we can either give up completeness, or confine concept descriptions to a subset of admissible expressions, in order to obtain an algorithm with the desired properties. Giving up completeness is problematic in this application, because inability to detect existing subsumption relations may lead to incorrect behavior of the retrieval algorithm:

- Existing solutions may not be found in the library. This can lead to an undesirable computational overhead in second principles planning, because the system does not reuse the best available plan, which may result in needless modification effort.
- Uncontrolled growth of the plan library may occur. Identical abstracted specifications are added to the library, because the incomplete subsumption algorithm is unable to recognize their equivalence.

Consequently, concept descriptions are restricted to so-called *admissible concepts* in conjunctive normal form, for which a sound, complete, and polynomial-time subsumption algorithm exists.

Definition 12 Let R be a composition of primitive roles (role chain), and C be a primitive concept atom. Concept terms of form $\exists R.C$ and $\exists R.\neg C$ are called primitive components.

Definition 13 A consistent concept description D is admissible iff D is a conjunction of concept descriptions d_i , where each d_i is restricted to be a disjunction of primitive components.

Restricting the language to admissible concept descriptions makes the usual expansion and normalization steps unnecessary. In the general case of more expressive descriptions, the relevant part of a terminology has to be transformed into a normal form before the computation of subsumption can start. Defined concepts have to be expanded by their definition using terminological axioms. This step is not necessary for admissible concepts, because they are already given in a normal form, and all concept terms are restricted to be primitive. In fact, admissible concepts define a subset of propositional logic. Primitive components can be treated as atomic units during the computation of subsumption. The following rule set taken from [Givan and McAllester, 1992] defines a sound and complete subsumption algorithm for admissible concepts that runs in polynomial time.

$$\begin{array}{rcl}
E \sqsubseteq_{\mathcal{T}} C, E \sqsubseteq_{\mathcal{T}} D & \rightarrow & E \sqsubseteq_{\mathcal{T}} C \sqcap D \\
C \sqsubseteq_{\mathcal{T}} E & \rightarrow & C \sqcap D \sqsubseteq_{\mathcal{T}} E \\
C \sqsubseteq_{\mathcal{T}} E, D \sqsubseteq_{\mathcal{T}} E & \rightarrow & C \sqcup D \sqsubseteq_{\mathcal{T}} E \\
E \sqsubseteq_{\mathcal{T}} C & \rightarrow & E \sqsubseteq_{\mathcal{T}} C \sqcup D \\
C & \sqsubseteq_{\mathcal{T}} & C
\end{array}$$

3 A Four-Phase Model of Second Principles Planning

Reasoning from second principles proceeds in a basic cycle of *problem input—activation of previous solutions—adaptation*, cf. [Riesbeck and Schank, 1989]. Besides these three phases, we consider a fourth and final phase in the cycle, which updates the memory of the second principles planner.

(I) Plan Determination: The current plan specification is the input to the plan-determination phase. Taking this specification, the retrieval process starts by mapping the LLP formula into a concept description, which represents the search key to the plan library.

The plan library contains a collection of *plan entries* that are extracted from previously solved planning problems. A plan entry provides comprehensive information about a planning problem and its solution, e.g., the specification of the problem describing initial and goal states, the plan which was generated as a solution for it, and information that is extracted from the plan generation process. Each plan entry possesses an index, which is represented as a concept description. The position of a plan entry is determined by the position of its index in the subsumption hierarchy.

Retrieval classifies the current search key in the subsumption hierarchy of indices, and returns a set of reuse candidates. Ranking heuristics are applied in order to determine the best candidate.

(II) Plan Interpretation: Plan interpretation attempts to prove that the current plan specification formula S_{new} is a logical consequence of the reused plan specification formula S_{old} . If the attempt succeeds, the reused plan specification is sufficient for the current one, which means solving the old planning problem will solve the current planning problem. This means that the reused plan solves S_{new} directly, and no modification of it is necessary. If the attempt fails, a plan skeleton is constructed. A plan skeleton provides an entry point into the search space of possible plans. It represents an incomplete solution to the current planning problem, because it may contain “placeholders” for subplans achieving open subgoals, which the reused plan is unable to achieve. The plan skeleton keeps any actions of the reused plan that were determined as reusable during plan interpretation, and in which variables are appropriately instantiated with object parameters taken from the current plan specification.

(III) Plan Refitting: The third phase completes the plan skeleton to a *correct plan* with the help of an interleaved process of plan verification and generative planning. Similar to a generative planner, plan refitting selects a subgoal to work on and successively reuses subplans from the plan skeleton to construct the final plan. During this process, actions and control structures in the plan skeleton may be deleted, added or reordered.

(IV) Plan-Library Update: Planning from second principles terminates with a *plan-library update*, during which a new plan entry is constructed from three sources of information: the current plan specification, the plan which was generated by modifying an existing plan, and information that is extracted from the proof tree which was constructed as a result of plan refitting. The plan entry is related to the current search key, which serves as its index in the plan library. The modified plan is now available to subsequent planning processes.

The four-phase model describes a temporal view on the reuse process. The phases I to III are necessary to generate a plan by reusing an existing one. They are also distinguished by other authors who sometimes denote them as *retrieval - matching - adaptation* phases, cf. [Hanks and Weld, 1995]. The fourth phase comprises the maintenance of the plan library. The formalization groups those phases together, which perform similar tasks. Operations on the plan library provide the basis for phases I and IV, while plan interpretation and refitting (phases II and III) work on plan specifications and are summarized as *plan modification*.

Figure 1 shows the architecture of the MRL system. The system comprises four modules, each of which implements a phase of second principles planning.

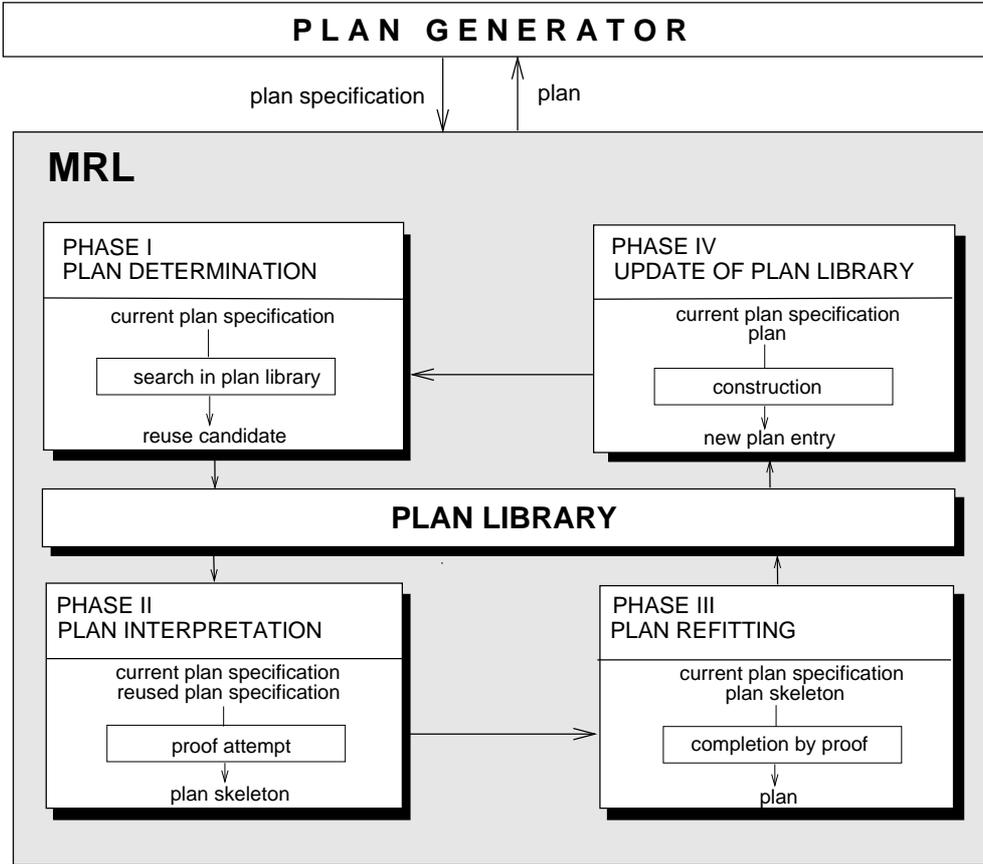


Figure 1: Architecture of the MRL system

3.1 Formalizing Plan Modification

The input to plan modification is the current plan specification formula S_{new} of form

$$Plan_{new} \wedge pre_{new} \rightarrow \diamond goal_{new}$$

and the plan specification formula S_{old} from the reuse candidate of form

$$\text{Plan}_{\text{old}} \wedge \text{pre}_{\text{old}} \rightarrow \diamond \text{goal}_{\text{old}}$$

The planvariable Plan_{new} stands for the plan (formula) we want to determine by reusing the existing plan (formula) \mathbf{P}_{old} that is available as a correct instantiation of the planvariable Plan_{old} . Plan modification has now to answer the question of whether \mathbf{P}_{old} meets S_{new} , i.e., if

$$\mathbf{P}_{\text{old}} \wedge \text{pre}_{\text{new}} \rightarrow \diamond \text{goal}_{\text{new}}$$

is true. In this case, \mathbf{P}_{old} can be reused immediately, otherwise it has to be refitted to meet new requirements. Instead of directly proving the validity of this formula, we show sufficient conditions between both plan specifications according to Theorem 1.

Theorem 1 \mathbf{P}_{old} meets S_{new} if $Ax \models S_{\text{old}} \rightarrow S_{\text{new}}$.

This means, if the new plan specification formula is a logical consequence of the old one under the given domain theory Ax , solving the old planning problem with plan \mathbf{P}_{old} is sufficient for solving S_{new} . Ax contains the set of action axiom schemata as well as the domain constraints.⁸

Since plan specifications contain formal descriptions of initial and goal states, we can show $Ax \models S_{\text{old}} \rightarrow S_{\text{new}}$ by proving sufficient relations between preconditions and goals according to Theorem 2:

Theorem 2 $Ax \models S_{\text{old}} \rightarrow S_{\text{new}}$ if

$$Ax \models \text{pre}_{\text{new}} \rightarrow \text{pre}_{\text{old}} \text{ and } Ax \models \diamond \text{goal}_{\text{old}} \rightarrow \diamond \text{goal}_{\text{new}}.$$

This means, we have to prove that the preconditions required by the old plan are satisfied in the current initial state and that the goals achieved by the old plan are sufficient for the currently required goals. If these relationships between initial and goal state specifications hold, we know that

- \mathbf{P}_{old} is applicable in pre_{new} and
- \mathbf{P}_{old} achieves at least all of the goals required in goal_{new} .

The validity of both theorems can easily be demonstrated by a sequent proof of $Ax \models S_{\text{old}} \rightarrow S_{\text{new}}$. We start with the initial sequent

$$Ax, \mathbf{P}_{\text{old}} \wedge \text{pre}_{\text{old}} \rightarrow \diamond \text{goal}_{\text{old}} \Rightarrow \mathbf{P}_{\text{old}} \wedge \text{pre}_{\text{new}} \rightarrow \diamond \text{goal}_{\text{new}}$$

in which both planvariables are already instantiated with the reused plan formula. Applying sequent rules for the logical connectives \wedge and \rightarrow to both sides of the sequent leads to the following three proof tasks⁹

$$(1a) \ Ax, \mathbf{P}_{\text{old}}, \text{pre}_{\text{new}} \Rightarrow \mathbf{P}_{\text{old}}, \diamond \text{goal}_{\text{new}}$$

$$(1b) \ Ax, \mathbf{P}_{\text{old}}, \text{pre}_{\text{new}} \Rightarrow \text{pre}_{\text{old}}, \diamond \text{goal}_{\text{new}}$$

$$(1c) \ Ax, \mathbf{P}_{\text{old}}, \text{pre}_{\text{new}}, \diamond \text{goal}_{\text{old}} \Rightarrow \diamond \text{goal}_{\text{new}}$$

Sequent (1a) is a logical axiom since antecedent and succedent contain the plan \mathbf{P}_{old} as a common formula. Sequents (1b) and (1c) lead to the two subproof tasks as required in Theorem 2. Of course, sequents (1b) and (1c) would also be valid if $Ax, \text{pre}_{\text{new}} \Rightarrow \text{goal}_{\text{new}}$ can be proved, i.e., if the current plan specification is a tautology where the goal is already satisfied in the initial state. But this will seldom hold and therefore, there is not much sense in trying to prove this.

The reader may now wonder why it makes more sense to prove relations between preconditions and goals instead of directly proving the validity of $\mathbf{P}_{\text{old}} \wedge \text{pre}_{\text{new}} \rightarrow \diamond \text{goal}_{\text{new}}$. Both approaches are possible in principle. We decided for the former since the subproofs are easier. For preconditions in particular, we in many cases have to perform a first-order proof involving formulae of a rather simple syntactic structure. This allows us to use complete proof tactics that run in polynomial time relative to the length of the formulae. Directly proving the instantiated plan specification formula means to split the plan

⁸An example of a domain constraint in the mail domain is the formula $\forall x [\text{open_flag}(\text{mailbox}) = F \rightarrow \text{delete_flag}(\text{msg}(x, \text{mailbox})) = F]$, i.e., no closed mailbox can contain deleted messages.

⁹The sequent proof can be found in detail in the Appendix.

formula into subformulae in such a way, that each of the subgoals from $goal_{new}$ can be shown to follow from a (sub)plan of P_{old} and the current preconditions pre_{new} . This requires the application of correct sequent rules for the splitting of plan formulae and goals, which often requires additional complicated proof tasks involving frame axioms.

The reader may note that an analogue to Theorem 2 can be obtained for syntactically different plan specifications, where plans are represented as terms as is usual in deductive planning. In this case, planvariables and plan terms occur as additional arguments in the goal-state specifications.

3.2 Formalizing Library Retrieval and Update

In principle, the inference procedures working on the plan library can be formalized in the same way as plan modification. A plan solving the current planning problem can be determined by finding a candidate that is applicable in the current initial state and achieves all of the current goals, i.e., by proving sufficient conditions between preconditions and goals. But apart from the complexity problems we would encounter, this is too restrictive because such a search process can only retrieve solutions. A failed proof would not tell us which of the plans is the best candidate for plan modification.

Research in case-based reasoning proposes to solve this problem by computing so-called *partial matches*.¹⁰ Partial matches are computed between the indices of a case, which are “labels . . . that designate under what conditions each of the case can be used to make useful inferences” [Kolodner, 1993], page 20. A pre-indexing technique identifies a privileged set of features to organize cases in the case library. Usually, indices are restricted to be vectors of propositional atoms in order to make the retrieval and matching problem tractable.¹¹

There are several disadvantages of pre-indexing and partial matching. As [Anderson *et al.*, 1994] have observed, indexing based on a fixed vocabulary hinders flexibility in retrieval and flattens the representation of cases into simple feature vectors. Besides this, these indexing schemes restrict the case library to have a *tree*-like structure usually represented as a *discrimination network* [Feigenbaum, 1963]. Finally, partial matches require a “relaxation” on the inference relation for which clear semantics can rarely be given and thus soundness of the retrieval algorithm is impossible to prove.

To overcome these serious limitations, we developed a novel solution to the indexing problem based on description logics. As a main advantage, description logics offer a more expressive representation language beyond sets of atoms. Simple feature vectors can be replaced by logical formulae, which may involve relational and functional descriptions. As shown in [Borgida and Patel-Schneider, 1994], such concept descriptions can be interpreted as labeled, directed multigraphs. Subsumption as the basic inference procedure in description logics formalizes the matching of these graph structures when indices have to be compared. The formal properties of the “matcher” can easily be investigated by proving formal properties of the subsumption algorithm. *Soundness* guarantees that the retrieved candidate meets the search criterion. *Completeness* ensures that all matching candidates are in the retrieval set. The *complexity* of the subsumption algorithm decides whether an efficient retrieval algorithm is available. As a further advantage, case libraries are indexed on a more general *lattice structure* provided by the subsumption hierarchy instead of a *tree* structure [Koehler, 1994a].

Finally, the pre-indexing problem can be resolved, because no set of features has to be identified in advance. Instead, the indexing vocabulary is computed from the logical representation of cases (here plan specification formulae) with the help of an *encoding scheme* ω . The encoding scheme is defined as a *mapping* from LLP formulae to concept descriptions. The formal basis for ω is the model-theoretic semantics of both logics. In a first step, an LLP plan specification formula is *equivalently* translated into a first-order formula using the method described in [Koehler and Treinen, 1995]. The resulting first-order formula is then replaced by an *abstracted* first-order formula.

Definition 14 *Given two first-order formulae ϕ and $abs(\phi)$, the formula $abs(\phi)$ is an abstraction of ϕ iff*

$$\mathcal{M} \models \phi \rightarrow abs(\phi)$$

¹⁰See [Kolodner, 1993] for a summary of the state of the art in case-based reasoning.

¹¹See [Nebel and Koehler, 1995] for a complexity-theoretic analysis of the matching problem.

holds in all models \mathcal{M} .

For example, the proposition A is an abstraction of $A \wedge B$. The result of abstraction is a formula in a sublanguage of first-order logic which can be translated into a concept description by preserving equivalence again. Here, we exploit the fact that description logics can be seen as sublanguages of first-order logic [Brachmann and Levesque, 1984]. A very desirable property of encoding schemes is the so-called *monotonicity property*:

Definition 15 Given two plan specification formulae S_{old} , S_{new} and their respective encoding concepts $\omega(S_{old})$, $\omega(S_{new})$, an encoding scheme ω satisfies the monotonicity property iff

$$Ax \models S_{old} \rightarrow S_{new} \text{ then } \omega(S_{old}) \sqsubseteq_{\mathcal{T}} \omega(S_{new})$$

This means, if a plan specification S_{new} is a logical consequence of a plan specification S_{old} in a domain axiomatization Ax , then $\omega(S_{new})$ must subsume $\omega(S_{old})$. In other words, the encoding scheme preserves an existing subset relationship between the set of models of the plan specification formulae as a subset relationship between the extensions of the concept descriptions. The monotonicity property ensures that existing solutions are found in the plan library, when a *complete* subsumption algorithm is used.

The *abstraction* of formulae as part of the encoding guarantees that an *exact* match as computed by the subsumption algorithm corresponds to a *partial* match between the original formulae. To compute an abstraction, we define a set of abstraction rules of form $\phi \rightarrow abs(\phi)$, which replace a formula ϕ by a *weaker* formula $abs(\phi)$. Examples of abstraction rules are $A \wedge B \rightarrow A$ or $\forall x p(x) \rightarrow \exists x p(x)$. An analysis of the abstraction rules allows us to draw conclusions about the degree of partial matches. Furthermore, when using different sets of abstraction rules, different degrees of partial matches can be performed by a retrieval algorithm.

Finally, an encoding scheme provides a formalization of *reasoning by approximation*. The retrieval algorithm approximates the relationship of logical consequence between plan specifications when computing subsumption between their encoding concept descriptions.

Figure 2 summarizes the formal framework. It bases planning from second principles on deductive inference processes. Plan modification is formalized by proving sufficient conditions between preconditions and goals in the underlying planning logic. A formalization of the inference procedures working on the plan library is obtained by computing their approximation in a description logic.

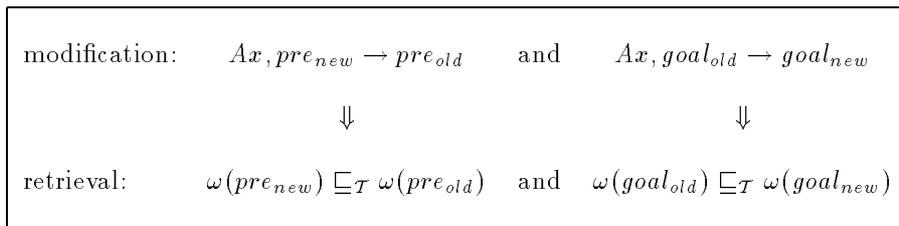


Figure 2: The logical framework for planning from second principles

The idea of exploiting relationships between preconditions and goals can be found in other approaches as well. Hanks and Weld [Hanks and Weld, 1992a] write that “retrieval takes the problem’s *initial* and *goal conditions* and finds in the plan library a plan that has worked under circumstances *similar* to those posed by the current problem”. The basic approach described by Hammond [Hammond, 1990] is “to find a past plan in memory that satisfies as many of the *most important goals* as possible”. Plan modification as formalized by Kambhampati and Hendler [Kambhampati and Hendler, 1992] relies on marking “the *differences* between the *initial* and *goal state specifications*”.

In the remaining part of the paper we show how this formal framework serves as a theoretical foundation for the implementation of well-defined inference procedures in the second principles planner MRL. We obtain a system of predictable behavior, since theoretical properties like soundness, completeness, and efficiency of the inference procedures are provable.

$$p(A) \rightarrow \exists x p(x)$$

For example, applying this rule to a formula $sender(msg(x, mbox)) = Joe$ leads to $\exists s sender(msg(x, mbox)) = s$ which abstracts from the specific sender of a mail. Abstraction from *temporal information* is slightly more complicated. To give the reader an impression of the temporal abstraction process we consider the formula

$$\begin{aligned} &\diamond [read_flag(msg(m, mybox)) = T \wedge \\ &\quad \diamond [delete_flag(msg(m, mybox)) = T]] \end{aligned}$$

The goal of temporal abstraction is to abstract from the ordering of subgoal states. In the example formula, the nested *sometimes* operators specify such an ordering. A simple rule, which accomplishes the desired abstraction is

$$\diamond [p \wedge \diamond q] \rightarrow \diamond p \wedge \diamond q$$

Such a modal formula can be equivalently translated into a first-order formula following the relational translation method by [Ohlbach, 1991]. The translation adds an additional predicate representing the accessibility relation on intervals into the object language. Local variables, which are the only fluents in LLP, are translated into unary functions that are equipped with an interval argument w_i [Koehler and Treinen, 1995]. Another abstraction rule is applied to eliminate the additional predicates introduced by the translation. In the example, we would obtain (w_1, w_2 are interval variables)

$$\begin{aligned} read_flag(msg(m, mybox(w_1))) &= T \wedge \\ delete_flag(msg(m, mybox(w_2))) &= T \end{aligned}$$

The result of the abstraction process is a first-order formula containing negation, conjunction, and disjunction. Each literal is then mapped to a primitive component, i.e., an existential role restriction of the form $\exists R.C$ or $\exists R.\neg C$, while \wedge and \vee are mapped to \sqcap and \sqcup , respectively.

The structure of a term like $read_flag(msg(m, mybox(w_1)))$ is reflected in the composition of roles. The unary function $mybox$ is of type $interval \rightarrow mailbox$ and is abstracted by a binary relation $interval \times mailbox$. The binary function msg is of type $mailbox \times integer \rightarrow message$, i.e., it takes a mailbox and an integer as arguments and returns the message that can be found in the mailbox at the position indicated by the integer. Thus, this function is abstracted by the composition of binary relations $mailbox \times integer \circ integer \times message$. The unary function $read_flag$ is of type $message \rightarrow boolean$, i.e., we abstract it by a binary relation of type $message \times boolean$. Consequently, for the formula $read_flag(msg(m, mybox(w_1))) = T$ we obtain the concept description

$$\exists mailbox \circ position \circ message \circ read_flag.TRUE$$

After the encoding process has been completed, the conjunctive normal form of the concept description is computed. Of course, the computational effort for this operation grows exponentially with the formula length. But remember that the subsumption algorithm is only complete for concepts in conjunctive normal form. Nevertheless, for pragmatic reasons it is less costly to compute the normal form only once during the encoding process instead of computing it several times during the classification of an index. In many cases, the normalization will not be necessary because many planning tasks involve only *sets* of atomic subgoals.

The encoding scheme ω used in MRL leads to the following encodings of the specifications SP_1 to SP_3 .

$$\begin{aligned} \omega(pre_{SP_1}): & \quad \exists mailbox \circ open_flag.TRUE \sqcap \\ & \quad \exists mailbox \circ position \circ message \circ delete_flag.FALSE \\ \omega(goal_{SP_1}): & \quad \exists mailbox \circ position \circ message \circ read_flag.TRUE \sqcap \\ & \quad \exists mailbox \circ position \circ message \circ delete_flag.TRUE \\ \omega(pre_{SP_2}): & \quad \exists mailbox \circ position \circ message \circ delete_flag.FALSE \\ \omega(goal_{SP_2}): & \quad \exists mailbox \circ position \circ message \circ read_flag.TRUE \sqcap \\ & \quad \exists mailbox \circ position \circ message \circ delete_flag.TRUE \end{aligned}$$

$$\begin{aligned}
\omega(pre_{SP_3}): & \exists \text{ mailbox} \circ \text{open_flag}.\text{TRUE} \sqcap \\
& [\exists \text{ mailbox} \circ \text{position} \circ \text{message} \circ \text{sender}.\neg\text{SENDER} \sqcup \\
& \exists \text{ mailbox} \circ \text{position} \circ \text{message} \circ \text{delete_flag}.\text{FALSE}] \\
\omega(goal_{SP_3}): & [\exists \text{ mailbox} \circ \text{position} \circ \text{message} \circ \text{sender}.\neg\text{SENDER} \sqcup \\
& \exists \text{ mailbox} \circ \text{position} \circ \text{message} \circ \text{delete_flag}.\text{TRUE}] \sqcap \\
& [\exists \text{ mailbox} \circ \text{position} \circ \text{message} \circ \text{sender}.\neg\text{SENDER} \sqcup \\
& \exists \text{ mailbox} \circ \text{position} \circ \text{message} \circ \text{read_flag}.\text{TRUE}]
\end{aligned}$$

The expressiveness of admissible concepts is sufficient to represent the mail domain adequately. The reader may note that this property may not generalize to other application domains for which different encoding schemes must be defined. In some cases, this can require to use more expressive concept languages for which subsumption becomes intractable. In this situation, we have either to give up completeness, or perform further abstractions leading to simplified formulae remaining within a tractable sublanguage.

The general idea of using description logics as query languages to case libraries seems to be widely applicable. Given a logical description of a case, i.e., a logical formula, it is possible to map it to some weaker logical formula, which can be interpreted as a concept description. Nevertheless, the development of encoding schemes mapping logical formulae to concept descriptions is a creative process. Its mechanization is an interesting subject for further research.

The encoding scheme used in MRL satisfies the monotonicity property as formulated by Definition 15. Thus, the retrieval algorithm is guaranteed to find existing solutions. Note that the inverse of the monotonicity property does not hold in general. A plan retrieved from the library will not, with certainty, provide a solution to the new planning problem. This reflects *reasoning by approximation*. The retrieval algorithm approximates the relationship between the plan specifications when comparing their abstractions. Thereby, it extends the computed set of candidates.

4.2 Weakening Retrieval

The retrieval algorithm takes the encoding of preconditions and goals of the current plan specification in order to determine its position in the subsumption hierarchy. By testing

$$\omega(pre_{new}) \sqsubseteq_{\mathcal{T}} \omega(pre_{old}) \quad \text{and} \quad \omega(goal_{old}) \sqsubseteq_{\mathcal{T}} \omega(goal_{new})$$

a set of subsumed plan specifications is determined. The plans which meet these specifications are possible reuse candidates. If no plan is returned, i.e., if the encoded current specification only subsumes the bottom concept, no directly reusable plan is contained in the plan library. In this situation, a second principles planner is faced with the decision of either:

- to give up further reuse attempts and plan from scratch, or
- to weaken the retrieval criterion and accept that any reuse candidate has to be modified.

The principal problem is to anticipate the modification effort. Planning from second principles makes sense when the effort for retrieval and modification is lower than the effort for planning from scratch. Unfortunately, it is impossible to decide this *before* we start retrieval. However, practical experiments demonstrated that modifying a plan is often less costly than generating one from scratch [Hanks and Weld, 1995; Koehler, 1994b]. Weakening retrieval and searching for plans that have to be modified is therefore a practical decision even under efficiency considerations.

Retrieval based on classification offers two principle ways to weaken the search criterion, which are to classify according to goals *or* preconditions *only*:

$$\omega(pre_{new}) \sqsubseteq \omega(pre_{old}) \quad \text{or} \quad \omega(goal_{old}) \sqsubseteq \omega(goal_{new})$$

In the first case, retrieval searches for a plan achieving the current goals by accepting that its preconditions may not be satisfied in the current initial state. In the latter case, retrieval searches for a plan which is applicable in the current initial state, but will not achieve all of the currently required goals.¹²

¹²In a working system it seems to be a good restriction to implement only one of the possible approaches to weaken retrieval in order to improve retrieval efficiency. Here, we discuss both possibilities in order to demonstrate how retrieval

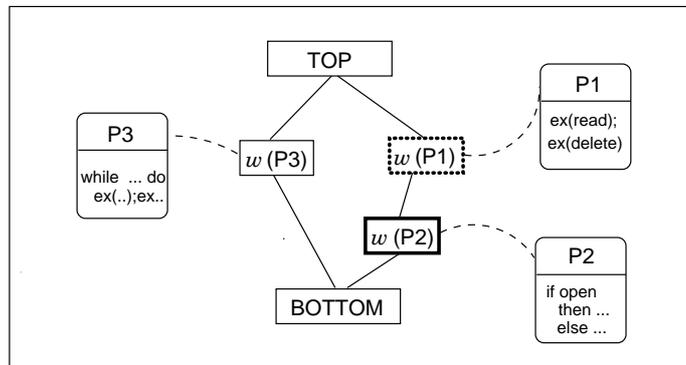


Figure 3: The example library

Figure 3 shows the small example library obtained for the three plan specifications under consideration. Obviously, $\omega(S_{P_2})$ subsumes only the bottom concept, i.e., no directly reusable plan can be retrieved. Therefore, a weaker retrieval algorithm is activated searching for an $\omega(pre_{old_i})$ that subsumes $\omega(pre_{S_{P_2}})$ or for an $\omega(goal_{old_i})$ that is subsumed by $\omega(goal_{S_{P_2}})$. Classification of the encoded preconditions fails in retrieving a candidate as well, while classification of the encoded goals is successful for $\omega(goal_{S_{P_1}})$ since $\omega(goal_{S_{P_1}}) \sqsubseteq_{\mathcal{T}} \omega(goal_{S_{P_2}})$ holds. Therefore, the plan **P1** attached to $\omega(S_{P_1})$ is activated as a reuse candidate. Since strong retrieval failed, we know that **P1** cannot represent a solution to the current planning problem S_{P_2} . We expect it to achieve all of the current goals, but we know that its preconditions are not satisfied in the current initial state. Thus, plan refitting has to start as will be described in Section 5.

4.3 Ranking of Plans

Only one candidate plan has been retrieved from the plan library in the example under consideration. But in general, retrieval will determine several appropriate reuse candidates. Consequently, a *ranking* is needed for the candidates in order to determine the best one.

The subsumption hierarchy, which is a directed acyclic graph, provides an easy way to rank candidates by computing their distance to the current concept description. The best candidate has the shortest path to $\omega(S_{new})$ in the graph. If several candidates have a shortest path, ranking heuristics are used to approximate the optimization and modification effort, respectively.

Strong retrieval returns plans that are supposed to be applicable in the initial state and to achieve *at least* all of the current goals. This implies that the candidate set may contain plans which achieve superfluous goals, i.e., goals that are currently unnecessary. Actions achieving these goals can be eliminated from the reused plan by making attempts at optimizing it. Thus, the ranking of candidates is based on an estimation of the *optimization effort* for each candidate, i.e., the number of superfluous actions that have to be eliminated from the candidate plan. The heuristic estimates the number of atomic subgoals that are achieved by a candidate plan but that are not required in the current plan specification. It assumes that this number reflects the minimal number of primitive actions that have to be eliminated from the candidate plan. Therefore, the plan with the smallest number is selected as the best reuse candidate and sent to plan modification. If several candidates receive the same ranking value, one of them is selected arbitrarily.

Definition 16 Let $\omega(S_{old_1}), \dots, \omega(S_{old_n})$ be the indices of candidates retrieved by strong retrieval for $\omega(S_{new})$. The heuristic compares the encoding of goals of the candidates $\omega(goal_{S_{old_1}}), \dots, \omega(goal_{S_{old_n}})$ with the encoding of the current goal $\omega(goal_{S_{new}})$. The set of primitive components that occur in a concept C is denoted by $P[C]$. The cardinality of the set $P[C]$ is as usually denoted by $|P[C]|$.

based on classification works. MRL applies weak retrieval only to preconditions, i.e., it requires plans to be applicable in the current initial state as a heuristic to reduce the refitting effort for control structures during plan modification.

The optimization effort for each candidate is defined as

$$OPT_{\omega(goal_{S_{old_1}})} = |P[\omega(goal_{S_{old_1}})] \setminus P[\omega(goal_{S_{new}})]|$$

The ranking heuristic selects a plan needing minimal optimization effort.

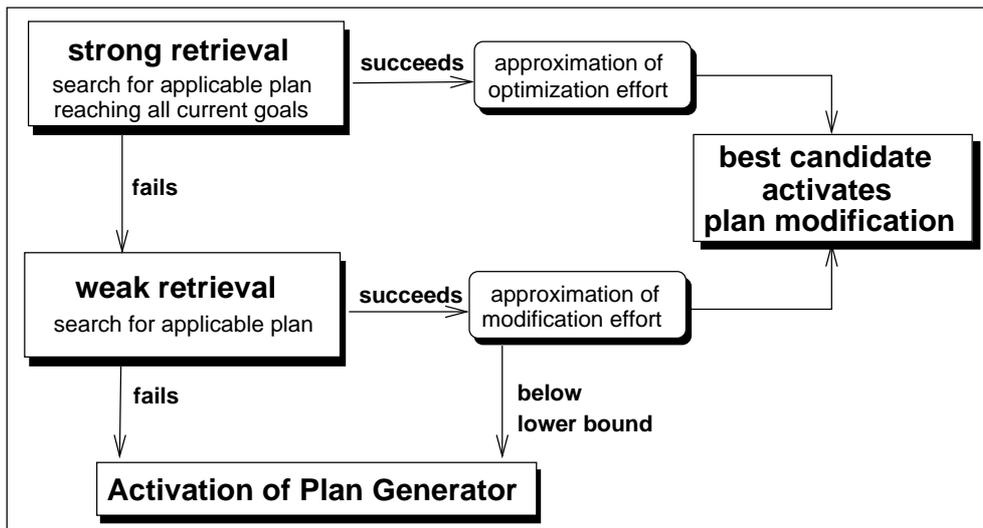


Figure 4: Interaction between first and second principles planning

Weak retrieval returns candidate plans that are either supposed to be applicable in the initial state or to achieve the desired goals, i.e., we have to expect that every candidate has to be modified. Consequently, the heuristic estimates the effort to *modify* each candidate in the retrieval set by computing the intersection of $\omega(goal_{S_{new}})$ with $\omega(goal_{S_{old_1}}), \dots, \omega(goal_{S_{old_n}})$, which approximates the number of *current* atomic subgoals that are achieved by each candidate. cf. [Koehler, 1994a].

Furthermore, the ranking heuristic verifies whether the ranking value of the best candidate exceeds a lower bound: it requires that at least half of the primitive components in $\omega(goal_{S_{new}})$ must be contained in $\omega(goal_{S_{old_1}})$. If this condition is satisfied, the ranking heuristic assumes that the best candidate achieves at least half of the current atomic subgoals. If no candidate receives a ranking value which exceeds the lower bound, all candidates are rejected because their modification effort is estimated as too expensive. In this situation, plan determination reports a failure and planning from scratch with the PHI planner is activated.

The ranking heuristics guide the interaction between planning from first and second principles, see Figure 4.

5 Correct Modification of Complex Plans

Plan modification is based on deductive inference processes which lead to modified plans that are provably correct. As introduced in Section 3, it proceeds in two phases. First, *plan interpretation* computes a plan skeleton and secondly, *plan refitting* completes the plan skeleton to a correct plan that meets the current specification.

In the following, we apply the formal approach to plan modification as defined in Section 3 to the example under consideration and discuss deductive plan modification in MRL.

5.1 Plan Interpretation

Plan interpretation receives two sources of input:

1. The current plan specification for which a plan has to be generated.
2. The best reusable plan, which the determination phase could identify in the plan library, together with its specification.

It takes the two plan specifications and tries to prove the required relations between preconditions and goals:

$$Ax \models pre_{new} \rightarrow pre_{old} \text{ and } Ax \models \diamond goal_{old} \rightarrow \diamond goal_{new}$$

In this example, proving the applicability of the reused plan **P1** in the current initial stated as specified in SP_2 requires the proof of Sequent 1:

$$\begin{aligned} delete_flag(msg(x, mbox)) = F \\ \Rightarrow open_flag(mybox) = T \wedge delete_flag(msg(m, mybox)) = F \end{aligned} \quad (1)$$

Starting point for the *goal proof* is Sequent 2:

$$\begin{aligned} \diamond [read_flag(msg(m, mybox)) = T \wedge \\ \diamond [delete_flag(msg(m, mybox)) = T]] \\ \Rightarrow \diamond [read_flag(msg(x, mbox)) = T \wedge \\ delete_flag(msg(x, mbox)) = T] \end{aligned} \quad (2)$$

In the following, we discuss both sequent proofs as they are performed by the proof tactics used during plan interpretation [Koehler, 1994c]. The tactics run in polynomial time on the length of the input formula. On one hand, this enables plan interpretation to *efficiently* compute an *entry point* into the search space of plans. On the other hand, this implies that the tactic is *incomplete* in the sense that it cannot compute a *maximal plan skeleton* which has been shown to be a PSPACE-hard problem [Nebel and Koehler, 1995].

The precondition proof for the example sequent is very simple because of the simple syntactic structure of the formulae. The first rule that is successfully applied to Sequent 1 is rule $r\wedge$.¹³

$$\bullet \frac{\Gamma, A, B \Rightarrow \Delta}{\Gamma, A \wedge B \Rightarrow \Delta} \quad l\wedge \qquad \bullet \frac{\Gamma \Rightarrow A, \Delta \quad \Gamma \Rightarrow B, \Delta}{\Gamma \Rightarrow A \wedge B, \Delta} \quad r\wedge$$

The rule splits Sequent 1 into Sequents 3 and 4:

$$delete_flag(msg(x, mbox)) = F \Rightarrow open_flag(mybox) = T \quad (3)$$

$$delete_flag(msg(x, mbox)) = F \Rightarrow delete_flag(msg(m, mybox)) = F \quad (4)$$

While Sequent 3 cannot be reduced to an axiom, Sequent 4 can be closed, i.e., it leads to an axiom under the substitution $\{x/m, mbox/mybox\}$. In order to obtain an appropriate instantiation of the reused plan, variables in the reused specification S_{old} are substituted by terms which occur in the current specification S_{new} . Furthermore, different variables must be mapped to different terms, i.e., the substitutions must be *injective*. Injectivity may not always be required, but it is a safe condition ensuring that a proper instantiation of the reuse candidate is computed during the proof. The reader may note that an instantiation of variables in sequents during a sequent proof is only possible when quantifier rules are applied. Plan specification formulae are implicitly universally quantified, i.e., when proving $S_{old} \Rightarrow S_{new}$ in the sequent calculus we remove the universal quantifiers using the rules

¹³A survey of all sequent rules that are used in this paper can be found in the Appendix.

$$\bullet \frac{\Gamma, A[c/x] \Rightarrow \Delta}{\Gamma, \forall x A \Rightarrow \Delta} l\forall \qquad \bullet \frac{\Gamma \Rightarrow \Delta, A[a/x]}{\Gamma \Rightarrow \Delta, \forall x A} r\forall^{14}$$

and have to “guess” the appropriate instantiation. Of course, this is unacceptable in an implemented prover due to the resulting computational overhead. Therefore, the instantiation is delayed until we know which instantiation is appropriate, i.e., which one will lead to a proof of the sequent. The restrictions we pose on the instantiations of the leaf sequents ensure that only those instantiations are computed that can be introduced with the help of quantifier rules.

Sequents 3 and 4 are the leaves of the derivation tree, because no further rules are applicable. Since only one of the leaves is an axiom, the tactic did not find a valid proof of the reused plan’s preconditions.

When trying to prove that the reused plan achieves all of the current goals the prover has to cope with *sometimes* operators and therefore additionally uses the following sequent rules:

$$\bullet \frac{\Gamma^*, A \Rightarrow \Delta^*}{\Gamma, \diamond A \Rightarrow \Delta} l\diamond \qquad \bullet \frac{\Gamma \Rightarrow A, \Delta}{\Gamma \Rightarrow \diamond A, \Delta} r\diamond$$

With Γ^* and Δ^* : $\Gamma^* \stackrel{\text{df}}{=} \{\Box B \mid \Box B \in \Gamma\}$ and $\Delta^* \stackrel{\text{df}}{=} \{\diamond B \mid \diamond B \in \Delta\}$.

The proof proceeds recursively over the *sometimes* operators in both goal specifications in order to compare every temporary subgoal state specified in $goal_{old}$ with each of the temporary subgoal states from $goal_{new}$. First, the tactic applies rule $l\diamond$ to Sequent 2 followed by rule $r\diamond$ and obtains Sequent 5:

$$\begin{aligned} & read_flag(msg(m, mybox)) = T \wedge \\ & \quad \diamond [delete_flag(msg(m, mybox)) = T] \\ \Rightarrow & \\ & read_flag(msg(x, mbox)) = T \wedge delete_flag(msg(x, mbox)) = T \end{aligned} \tag{5}$$

Now, rule $l\wedge$ is applied to Sequent 5 followed by rule $r\wedge$, which leads to Sequents 6 and 7

$$\begin{aligned} & read_flag(msg(m, mybox)) = T, \\ & \quad \diamond [delete_flag(msg(m, mybox)) = T] \\ \Rightarrow & read_flag(msg(x, mbox)) = T \end{aligned} \tag{6}$$

$$\begin{aligned} & read_flag(msg(m, mybox)) = T, \\ & \quad \diamond [delete_flag(msg(m, mybox)) = T] \\ \Rightarrow & delete_flag(msg(x, mbox)) = T \end{aligned} \tag{7}$$

Sequent 6 can also be closed under the substitution $\{x/m, mbox/mybox\}$, i.e., the current subgoal $read_flag(msg(x, mbox)) = T$ has been successfully proved. The system proceeds on Sequent 7 and removes the remaining \diamond operator with the help of rule $l\diamond$ which leads to Sequent 8

$$delete_flag(msg(m, mybox)) = T \Rightarrow \emptyset \tag{8}$$

The formula $delete_flag(msg(x, mbox)) = T$ from the succedent of Sequent 7 disappears in Sequent 8 because it does not occur in the scope of a \diamond operator. Thus, the tactic fails in proving the remaining subgoal. The reason for this failure is obvious: the current goal specification requires the two subgoals to be achieved in the same state, while the reused goal specification only requires the two subgoals to be achieved one after the other. Of course, deleting a message preserves the effect that the message has been read, i.e., the reused plan that first reads the message and then deletes it also leads to a final state where the message has been read and deleted. But we have no way to derive this fact from the original plan specification formula. This is a motivation for a completion process of plan specification formulae that is described in Section 6.

¹⁴Eigenvariable condition: a must not occur in the conclusion of $r\forall$.

5.2 Plan Refitting

Proof tactics are always designed to terminate. In addition, they are considered as decision procedures: If a tactic does not result in a proof tree, it is assumed that no proof is possible and that a falsifying valuation for some of the leaves has been obtained. Two situations are possible after the termination of a proof tactic in the sequent calculus:

1. A proof tree has been constructed, i.e., the leaves of the tree describe a set of logical axioms from which the original formula follows. In this case the original formula was proved to be valid.
2. No proof tree has been found and the assumption is made that no proof is possible and that a counter-example tree has been constructed.

This assumption is a safe condition ensuring the soundness of plan modification. Remember that the tactics are incomplete, i.e., when a tactic terminates with a failure it might either be the case that the formula is invalid or that the formula is valid, but the tactic failed to find a proof.

Assuming that the formula is invalid ensures that the correctness of a plan is verified during plan refitting. Thereby, it prevents the reuse of plans that are not provably correct with respect to the current plan specification.

The proof tactics guarantee that the leaves of a counter-example tree contain only atomic formulae. The falsifying valuation makes:

- All *old* atomic goals (in the example from S_{P_1}) true, however some of the atomic formulae which describe *current goals* (in the example from S_{P_2}) are valued as false. These falsified goals are interpreted as those current goals that are not achieved by the reused plan (in the example by **P1**).
- All atomic formulae describing *current* preconditions (in the example from S_{P_2}) true, but some of the *old* preconditions (in the example from S_{P_1}) false. These falsified preconditions are interpreted as those preconditions of the reused plan (in the example of **P1**) that do not hold in the current initial state.

Plan **P1** must be modified by constructing a plan skeleton from it, because it was neither possible to prove that its preconditions are satisfied nor that it achieves all of the currently required goals. First, the reused plan is instantiated with the substitutions $\{x/m, mbox/mybox\}$ computed during plan interpretation leading to

P1': $ex(type(x, mbox)); ex(delete(x, mbox))$

Plan refitting concludes from the failed proofs that the (instantiated) precondition $open_flag(mbox) = T$ required by **P1** does not hold in the initial state and that the current goal $delete_flag(msg(x, mbox)) = T$ is not achieved by it. Furthermore, **P1** achieved a subgoal

$$delete_flag(msg(m, mybox)) = T$$

that is not contained in the axiom that was obtained from Sequent 6, which was constructed during the goal proof. Thus, plan refitting concludes that the action $ex(delete(m, mybox))$ achieving this subgoal is (at least at the current position where it occurs) superfluous and can be removed from the plan skeleton. This analysis of the result of plan interpretation leads to the following modification operations that have to be performed on the instantiated plan **P1'**:

1. A planvariable **Plan₁** has to be introduced in front of the reused plan. It represents a subplan achieving the missing *precondition*

$$open_flag(mbox) = T$$

2. The *superfluous* action $ex(delete(x, mbox))$ is removed from the plan skeleton.
3. A planvariable representing the subplan for the open *subgoal*

$$delete_flag(msg(x, mbox)) = T$$

must be introduced into the plan skeleton. In order to determine the position in the skeleton where this planvariable has to be added, the current goal-state specification must be analyzed with the help of the PHI planner.

The planvariable in the current plan specification **SP2** is instantiated with the preliminary plan skeleton **P1''**. It serves as a starting point for plan refitting:

$$\mathbf{P1''}: \text{Plan}_1 ; ex(type(x, mbox))$$

$$\mathbf{SP2'}: \begin{array}{l} \text{Plan}_1 ; ex(type(x, mbox)) \wedge delete_flag(msg(x, mbox)) = F \\ \rightarrow \diamond [read_flag(msg(x, mbox)) = T \wedge \\ delete_flag(msg(x, mbox)) = T] \end{array}$$

In a first step, a subplan to replace **Plan1** has to be generated. Plan refitting applies the rule *effect_intro* [Biundo and Dengler, 1996] to $S_{P2'}$ and introduces the missing precondition $open_flag(msg(mbox)) = T$ as the new subgoal $goal_{new}$:

$$\frac{pre, \text{Plan}_1 \Rightarrow \diamond [goal_{new} \wedge OF \wedge pre'] \quad pre', \text{Plan}_2 \Rightarrow \diamond goal}{pre, \text{Plan}_1 ; \text{Plan}_2 \Rightarrow \diamond goal} \text{effect_intro}$$

It obtains two subplan specifications 9 and 10 where **Plan2** is instantiated with the action $ex(type(x, mbox))$ taken from the plan skeleton. The precondition pre' can be instantiated after a plan for **Plan1** has been generated and frame conditions have been computed.

$$\begin{array}{l} delete_flag(msg(x, mbox)) = F, \text{Plan}_1 \\ \Rightarrow \diamond [open_flag(msg(mbox)) = T \wedge OF \wedge pre'] \end{array} \quad (9)$$

$$\begin{array}{l} pre', ex(type(x, mbox)) \\ \Rightarrow \diamond [read_flag(msg(x, mbox)) = T \wedge \\ delete_flag(msg(x, mbox)) = T] \end{array} \quad (10)$$

The proof of subplan specification 9 leads to a conditional plan because there is no atomic action available in the domain axiomatization that achieves the required goal under the given precondition. Plan refitting applies the rule *if_intro* [Biundo and Dengler, 1996] to instantiate the planvariable **Plan1** with a case analysis:

$$\frac{pre, \mathbf{if}(cond, \text{Plan}_A, \text{Plan}_B) \Rightarrow \diamond goal}{pre, \text{Plan} \Rightarrow \diamond goal} \text{if_intro}$$

In the example, the conditional *cond* is instantiated with the missing precondition $open_flag(mbox) = T$ that plan interpretation failed to prove:

$$\text{Plan}_1 := \mathbf{if} \ open_flag(mbox) = T \mathbf{then} \ \text{Plan}_3 \\ \mathbf{else} \ \text{Plan}_4$$

Applying the rule *if_splitting* [Biundo and Dengler, 1996]

$$\frac{pre, cond, \text{Plan}_A \Rightarrow \diamond goal \quad pre, \neg cond, \text{Plan}_B \Rightarrow \diamond goal}{pre, \mathbf{if}(cond, \text{Plan}_A, \text{Plan}_B) \Rightarrow \diamond goal} \text{if_splitting}$$

to Sequent 9, plan refitting obtains the following subplan specifications:

$$\begin{array}{l} delete_flag(msg(x, mbox)) = F, open_flag(mbox) = T, \text{Plan}_3 \\ \Rightarrow \diamond [open_flag(mbox) = T \wedge OF \wedge pre'] \end{array} \quad (11)$$

$$\begin{aligned}
& delete_flag(msg(x, mbox)) = F, \neg open_flag(mbox) = T, Plan_4 \\
& \Rightarrow \diamond [open_flag(mbox) = T \wedge \text{OF} \wedge pre']
\end{aligned} \tag{12}$$

Plan₃ can be instantiated with the empty action $ex(empty_action)$, because the desired subgoal holds already in the initial state.

The formula $\neg open_flag(mbox) = T$ in Sequent 12 is equivalent to the formula $open_flag(mailbox) = F$ and thus, Plan₄ is instantiated with the action instance $ex(open(mbox))$ which opens the mailbox and starts a mail session:

$$\begin{aligned}
& \forall x [open_flag(mailbox) = F \wedge ex(open(mailbox)) \\
& \quad \rightarrow \text{O} open_flag(mailbox) = T]
\end{aligned}$$

With that, the following conditional plan is obtained as an instantiation of Plan₁:

$$Plan_1 := \mathbf{if} \ open_flag(mbox) = T \ \mathbf{then} \ ex(empty_action) \\
\quad \quad \quad \mathbf{else} \ ex(open(mbox))$$

Now, the precondition pre' in Sequent 10 can be instantiated with the formula

$$delete_flag(msg(x, mbox)) = F \wedge open_flag(mbox) = T$$

The proof of subplan specification 10 proceeds as an interleaved process of plan generation and plan verification. A tactic for the ordering of conjunctive goals is activated [Biundo and Dengler, 1996] which decides to achieve the subgoal $read_flag(msg(x, mbox)) = T$ before the subgoal $delete_flag(msg(x, mbox)) = T$, since deleting a message destroys the possibility of reading it subsequently. The first subgoal is isolated with the help of the *effect_split* rule [Biundo and Dengler, 1996]:

$$\frac{pre, Plan_A \Rightarrow \diamond [goal_1 \wedge \text{OF} \wedge pre'] \quad pre', Plan_B \Rightarrow \diamond [goal_2]}{pre, Plan_A ; Plan_B \Rightarrow \diamond [goal_1 \wedge goal_2]} \ \mathit{effect_split}$$

This rule requires a sequential composition of two planvariables that can be split such that the first planvariable represents a subplan achieving the first subgoal, while the second planvariable represents a *subplan* achieving the remaining subgoals. But the planvariable Plan₂ introduced by the *effect_intro* rule has been instantiated with the *single* atomic action $ex(type(x, mbox))$ in specification 10. Thus, this instantiation must be withdrawn and plan refitting sets Plan₂ to Plan₅; Plan₆.

We obtain two subplan specifications 13 and 14

$$\begin{aligned}
& open_flag(msg(mbox)) = T \wedge delete_flag(msg(x, mbox)) = F, \\
& Plan_5 \\
& \Rightarrow \diamond [read_flag(msg(x, mbox)) = T \wedge \text{OF} \wedge pre'']
\end{aligned} \tag{13}$$

$$\begin{aligned}
& pre'', Plan_6 \\
& \Rightarrow \diamond [delete_flag(msg(x, mbox)) = T]
\end{aligned} \tag{14}$$

The first subgoal $read_flag(msg(x, mbox)) = T$ (Sequent 13) has successfully been proven during plan interpretation. Consequently, the action from the plan skeleton $ex(type(x, mbox))$ achieving this subgoal is reused as an instantiation of the planvariable Plan₅:

$$\begin{aligned}
& open_flag(msg(mbox)) = T \wedge delete_flag(msg(x, mbox)) = F, \\
& ex(type(x, mbox)) \\
& \Rightarrow \diamond [read_flag(msg(x, mbox)) = T \wedge \text{OF} \wedge pre'']
\end{aligned} \tag{15}$$

The instantiation can be successfully verified by plan refitting. The precondition pre'' in Sequent 14 is instantiated with

2. Are control structures introduced by the modification strategy if this is required by the refitting process?

versus

Are control structures only introduced if the current planning problem requires a plan containing control structures?

The treatment of control structures in a second principles planner requires to make these decisions carefully and to take into consideration specific requirements from the application domain. The MRL system provides the reuse component of the PHI planner which is working in a help-system application. Here, plans are generated to provide active help to users of software environments [Bauer *et al.*, 1993]. This means that plans are required to meet exactly the user’s goals and to be as simple as possible. Therefore, control structures are only reused in a restricted way in the implemented system MRL. They are introduced into the modified plan or preserved in the plan skeleton only if the current planning problem requires the generation of a plan containing control structures.

An unrestricted reuse of control structures can lead to the following problems:

- Reused control structures are not guaranteed to correspond to the requirements of the current planning situation. This can result in *over-complicated plans*. For example, a case analysis makes the execution of a plan more complicated because a test on the conditional has to be performed during execution time. Thus, a case analysis should only be introduced into a plan skeleton when the current planning problem requires us to generate a conditional plan.
- Plans can achieve *unintended side-effects*. Plan refitting makes some attempts at optimizing a reused plan by removing superfluous actions from it, but it is not able to generate *optimal* plans because this is usually harder than generating an arbitrary plan. Superfluous control structures render the problem worse. For example, an iterative plan which achieves a particular goal for all objects satisfying a precondition could in principle be reused to satisfy the goal for only one of the objects. As an example, the reader may think of reusing a plan achieving the goal: “delete all my files in directory *x*”. This also achieves the goal: “delete file *x.ps* in directory *x*”. Without any attempts at optimizing the reused plan by removing the superfluous iterative control structure a drastic and harmful side effect is achieved.

Restricting the reuse of control structures as in MRL is one way of coping with these problems. Further research is necessary in order to identify other solutions.

6 Updating the Plan Library

The plan library is updated dynamically in MRL. Figure 5 summarizes the hybrid representation of the library based on a description logic and a planning logic. The description logic provides the library with an indexing structure based on the subsumption hierarchy of encoded plan specifications, which are represented as concept descriptions and serve as *indices* to *plan entries*. A *plan entry* contains information about a successfully solved planning problem: the plan, the plan specification, and information extracted from the planning process that has led to this plan. This information is represented in the planning logic. A plan entry can be retrieved from and inserted into the library over the *index* to which it is linked. Each index represents an abstract class of planning problems in the application domain. Several plan entries can be linked to the same index when their specifications are encoded by equivalent concept descriptions. In the current implementation, only one plan entry is stored on a “first come—first serve” basis as an instance of the abstract class of planning problems represented by the index. This avoids redundant entries in the plan library. Planning problems belonging to the same abstract class can be solved by modifying the plan stored in the plan entry.

The system starts with the initial plan library containing only the indices *top* and *bottom*. A new plan entry is added to the library under the following conditions:

- No reusable plan has been found and the planner has to generate a plan from first principles.
- The reused plan had to be modified.

The plan library is not updated when:

- a library plan directly solves a current planning problem,
- the index of the current planning problem is already contained in the library.

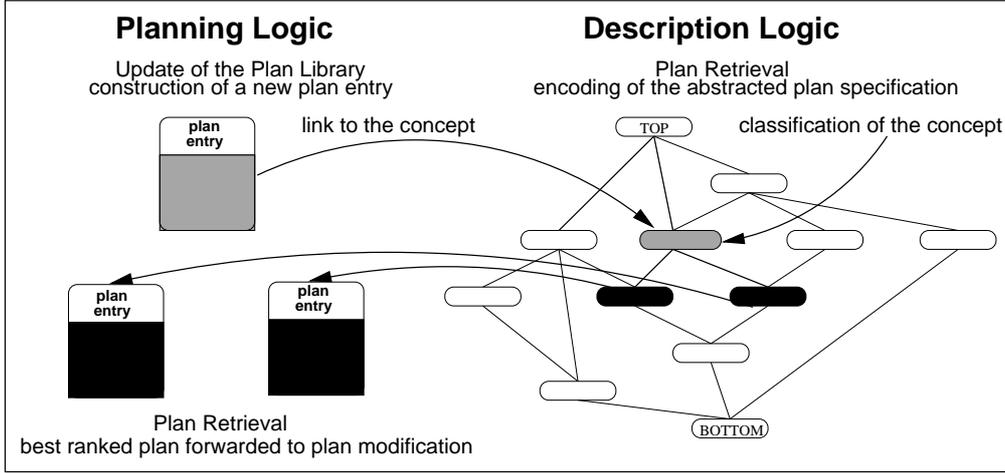


Figure 5: Plan entries in the plan library

Let us continue the example from Section 5. According to the above mentioned conditions, the plan library is updated because the reused plan has been modified. Three sources of information are available for the construction of the plan entry:

1. the current plan specification,
2. the modified plan that meets the specification,
3. the proof tree that has been constructed during plan refitting.

The index of the plan entry has already been computed during plan determination. The current plan specification is *completed* before it is added to the plan entry. By completion we understand the computation of the *weakest preconditions* and *strongest goals* of a plan. This means, those preconditions are eliminated from the plan specification which are not necessary for the plan and those goals are added to the plan specification that a plan can achieve as side-effects.

To determine the weakest preconditions and strongest goals of a plan, the planning process that has led to the plan is analyzed. In particular, action axiom schemata that have been applied during the proof are investigated. They specify the *necessary preconditions* of an action and the *effects* it achieves, see Section 2. If the completion process leads to a changed plan specification formula, the encoding of the plan specification is repeated, because a different concept description may result from it and thus, the position of the plan in the library can change caused by altered subsumption relationships.

In the example under consideration, the completion of plan specification S_{P_2} leads to a disjunctive precondition reflecting the complete case analysis that has been introduced into the plan with the help of the *if_intro* rule:

$$\begin{aligned}
 & \text{Plan}_{P_2} \wedge \\
 & [\text{delete_flag}(\text{msg}(x, \text{mbox})) = F \wedge \text{open_flag}(\text{mbox}) = T] \vee \\
 & [\text{delete_flag}(\text{msg}(x, \text{mbox})) = F \wedge \text{open_flag}(\text{mbox}) = F] \\
 & \rightarrow \diamond [\text{read_flag}(\text{msg}(x, \text{mbox})) = T \wedge \\
 & \quad \text{delete_flag}(\text{msg}(x, \text{mbox})) = T]
 \end{aligned}$$

7 Related Work

The implementation of a second principles planner based on the formal framework as introduced in Section 3 requires design decisions that specify how planning from second principles proceeds in detail. In this section, we discuss the most important of these decisions underlying MRL and relate the system to other approaches.

Meta Level versus Object Level

Planning from second principles can proceed on a *meta level* or on an *object level*. On the *object level*, previously generated plans are directly reused to solve the current planning problem. This means that the plans as the *objects* of the planning process provide the basis for planning from second principles. Reuse on the *meta level* means to “recycle” knowledge extracted from previous planning processes that represents planning experience in the form of planning tactics, heuristics or strategies.

The commitment of a particular planner to one of these levels is a fundamental design decision. A commitment to the object level leads to case-based planners and reuse systems, e.g., PRIAR [Kambhampati and Hendler, 1989], CHEF [Hammond, 1990] and SPA [Hanks and Weld, 1992a]. A commitment to the meta level leads to adaptive and reactive systems based on learning techniques, e.g., PRODIGY [Minton, 1988] and GRASSHOPPER [Leckie and Zuckerman, 1993].

MRL proceeds mainly on the object level because it relies on the reuse of stored plans. Meta-level knowledge is reused, e.g., when plan refitting is supplied with information about preconditions on which case analyses have to be performed, see the example in Section 5.

Skeletal Plan Refinement versus Flexible Modification

When planning from second principles proceeds on the object level, plans are modified in order to construct the desired plan from them. Plan modification can be implemented as *skeletal plan refinement* [Friedland and Iwasaki, 1985] or as *flexible modification* [Kambhampati, 1990; Hanks and Weld, 1992a]. Skeletal plan refinement computes an appropriate ground-level instantiation for each operator occurring in the abstract skeleton. The admissible modification operations are restricted to *instantiation*, but they can proceed in several hierarchical steps and backtracking may occur. The modified plan is obtained as an instance of the skeleton.

Flexible modification as implemented in MRL admits a variety of operations on plans, like the *deletion* and *addition* of operators and control structures.

Skeletal refinement occurs in MRL when the current plan specification has successfully been proved to be a logical consequence of the reused plan specification. In this situation, an instantiation of the library plan will solve the current planning problem and plan modification can be restricted to easy-to-compute substitutions.

Transformation-based versus Generation-based

The modification of a plan can be done with the help of transformations [McDermott, 1978; Hammond, 1990; Beetz and McDermott, 1994] or by extending a first principles planner with the ability to modify plans [Kambhampati, 1989; Hanks and Weld, 1992a; Veloso, 1994].

Transformation-based approaches execute a plan in a simulated environment. Failures are classified in a failure hierarchy and resolved by activating transformations on the plan. This approach requires a prediction of all possible failures, i.e., a proof of the *completeness* of the failure hierarchy and the available transformation rules, which is hard to achieve.¹⁷ Further problems are related to the *soundness* and *termination* of the transformations. Transformations resolving a failure may introduce other failures, which makes it difficult to ensure that the transformation process does not loop and that the transformed plan is sound, i.e., that it solves the current planning problem.

To overcome these problems, a generation-based approach has been introduced in the PRIAR system [Kambhampati, 1989]. The proof of the completeness of plan modification with respect to the planner

¹⁷As an example see the incompleteness proof of CHEF in [Hanks and Weld, 1992b].

is trivial since plan modification can rely on plan generation as a “fall-back” possibility. Soundness and termination are also easy to ensure if the underlying first principles planner possesses these properties. The modification of a plan in MRL proceeds *generation-based*. MRL computes a *plan skeleton* and sends it to plan refitting for completion, which interacts with the generative PHI planner. The plan skeleton preserves those control structures and actions that are assumed to be reusable. The extension of a skeleton to a correct plan requires *flexible* modification operations, which *add*, *delete* or *reorder* operators and control structures. The correctness of deductive plan refitting, which completes the skeleton, ensures that the modified plan is *sound*. Planning knowledge represented by the plan skeleton guides plan refitting and dynamically constrains the search space.

MRL is “complete” with respect to the planner because plan refitting can “fall back” on plan generation. The system is incomplete in the sense that it will not always find a plan if there is one because the use of tactics makes the underlying LLP theorem prover incomplete.

Conservative versus Non-Conservative

A desirable property of plan modification is *conservatism*, which means to “produce a plan ... by minimally modifying [the original plan]” [Kambhampati and Hendler, 1992]. Minimal modification of a plan implies to preserve the maximal number of applicable operators in a plan skeleton. A critical analysis of conservatism in [Nebel and Koehler, 1993] shows that the computation of such maximal plan skeletons is PSPACE-hard. Therefore, implemented systems including MRL are non-conservative. In order to ensure efficiency of the plan modification process, they rely on polynomial approximations, for example proof tactics for plan interpretation that run in polynomial time, which compute an “entry point” into the search space of possible plans as made explicit by Hanks and Weld [Hanks and Weld, 1992a]. This entry point cannot be guaranteed to be the best, but it is the best the approximation algorithm can compute. It is an open problem whether the maximal applicable subplan is efficiently approximable within a constant ratio. Recent results for similar problems [Bäckström, 1994; Selman, 1994] seem to hint at a negative result.

The Plan Library

Recently, the representation of plans based on terminological knowledge-representation systems has led to several approaches, which extend description logics with new application-oriented representational primitives for the representation of actions and plans.

One such extension is the system RAT [Heinsohn *et al.*, 1991] which is based on *KRIS* [Baader *et al.*, 1992]. RAT implements reasoning about plans by inferences in the underlying description logic. The system simulates the execution of plans, verifies their applicability in particular situations, and solves tasks of temporal projection.

An application of description logics to tasks of plan recognition is developed in T-REX [Weida and Litman, 1994]. Plans in T-REX may contain conditions and iterations as well as non-determinism in the form of disjunctive actions.

The plan library can be *static* as well as *dynamic* in MRL. A static library comprises user-predefined typical plans. The system retrieves these plans for reuse, but does not add new plans to the library. A dynamic plan library grows during the lifetime of the system, i.e., MRL starts with an empty library and incrementally adds new plan entries to it.

The main advantage in using a description logic as a query language to the plan library as in MRL lies in the novel solution to the indexing problem and in the theoretically well-founded properties of the retrieval algorithm. For the first time, retrieval guarantees that *solutions* are found in a library in *polynomial* time. This contrasts to approaches that are restricted to retrieve “reasonable similar past cases ... within limited bounded resources” (cf. [Veloso, 1994], p. 103). Furthermore, an indexing of plan libraries based on the lattice structure provided by the *subsumption* hierarchy overcomes problems occurring in indexing schemes based on discrimination networks. On one hand, discrimination networks fail in indexing complex plan specifications because they are restricted to cope with conjunctions of literals. On the other hand, retrieval algorithms working on discrimination networks are often faced with an exponentially growing input set. For example, given a goal state containing n atomic subgoals, the retrieval algorithm developed in [Veloso, 1994] first searches a plan covering these n subgoals. If

this fails, it computes all subsets of subgoals of cardinality $n - 1$, then $n - 2$ and so on until it takes the atomic subgoals as input. This means, the retrieval algorithm takes the power set of n except the empty set as input in the worst case, which is $2^n - 1$. Retrieval based on concept descriptions avoids such problems because existing relationships between sets of subgoals are reflected in the subsumption hierarchy.

8 Conclusion

We have presented a logic-based approach to planning from second principles, which relies on a systematic decomposition of the planning process with the help of a four-phase model. Deductive inference processes with clearly defined semantics formalize each phase. The formal model is independent of a particular planning formalism and makes no commitments to application domains, implementational details, the nature of plans or planning situations.

Plan modification yields provably correct modified plans and enables a second principles planner to reuse plans containing control structures like conditionals and iterations.

Reusable plans are retrieved from a dynamically updated plan library using a description logic as query language to the library. The plan library can be indexed basing on a lattice structure and retrieval is formalized using a KL-ONE like classifier which is guaranteed to find existing solutions.

The formal framework has led to an implemented system with predictable behavior. Furthermore, in contrast to heuristic approaches, theoretical properties like the correctness, completeness and efficiency of the inference procedures can be proved.

Acknowledgement

This work was supported by the German Ministry for Research and Technology (BMFT) under contract ITW 9000 8 as part of the PHI project. I want to thank the members of the PHI group, Mathias Bauer, Susanne Biundo, Dietmar Dengler, and Gabriele Paul for their interest in my work and for fruitful discussions. Furthermore, I am indebted to Wolfgang Wahlster for his advice and support, and to Bernhard Nebel and Hans-Jürgen Ohlbach for commenting on the draft version. The anonymous reviewers made very detailed comments helping me during the revision of the paper.

Appendix

Proofs of Theorems

Theorem 1 \mathbf{P}_{old} meets S_{new} if $Ax \models S_{old} \rightarrow S_{new}$.

Theorem 2 $Ax \models S_{old} \rightarrow S_{new}$ if

$$Ax \models pre_{new} \rightarrow pre_{old} \text{ and } Ax \models \diamond goal_{old} \rightarrow \diamond goal_{new}.$$

Proof:

We have to prove the following sequent

$$Ax, \mathbf{P}_{old} \wedge pre_{old} \rightarrow \diamond goal_{old} \Rightarrow \mathbf{P}_{old} \wedge pre_{new} \rightarrow \diamond goal_{new} \quad (17)$$

in which both planvariables are already instantiated with the reused plan formula \mathbf{P}_{old} . We start by applying the sequent rule $r \rightarrow$ which leads to Sequent 18. To this sequent, the rule $l \rightarrow$ is applied leading to Sequents 19 and 20.

$$\bullet \frac{\Gamma \Rightarrow A, \Delta \quad \Gamma, B \Rightarrow \Delta}{\Gamma, A \rightarrow B \Rightarrow \Delta} l \rightarrow \quad \bullet \frac{\Gamma, A \Rightarrow B, \Delta}{\Gamma \Rightarrow A \rightarrow B, \Delta} r \rightarrow$$

$$Ax, \mathbf{P}_{\text{old}} \wedge pre_{old} \rightarrow \diamond goal_{old}, \mathbf{P}_{\text{old}} \wedge pre_{new} \Rightarrow \diamond goal_{new} \quad (18)$$

$$Ax, \mathbf{P}_{\text{old}} \wedge pre_{new} \Rightarrow \mathbf{P}_{\text{old}} \wedge pre_{old}, \diamond goal_{new} \quad (19)$$

$$Ax, \mathbf{P}_{\text{old}} \wedge pre_{new}, \diamond goal_{old} \Rightarrow \diamond goal_{new} \quad (20)$$

Sequents 19 and 20 have to be expanded further. We start with Sequent 19 and apply the rule $l\wedge$, followed by $r\wedge$ and obtain Sequents 21 and 22.

$$Ax, \mathbf{P}_{\text{old}}, pre_{new} \Rightarrow \mathbf{P}_{\text{old}}, \diamond goal_{new} \quad (21)$$

$$Ax, \mathbf{P}_{\text{old}}, pre_{new} \Rightarrow pre_{old}, \diamond goal_{new} \quad (22)$$

To Sequent 20, the rule $l\wedge$ is applied leading to Sequent 23.

$$Ax, \mathbf{P}_{\text{old}}, pre_{new}, \diamond goal_{old} \Rightarrow \diamond goal_{new} \quad (23)$$

The original sequent is proved iff the three sequents are axioms, i.e., antecedent and succedent contain a common formula. This holds for Sequents 21 immediately, while Sequents 22 and 23 lead to the two subproofs

$$Ax, pre_{new} \Rightarrow pre_{old} \quad \text{and} \quad Ax, \diamond goal_{old} \Rightarrow \diamond goal_{new}$$

■

Summary of Sequent Rules

$$\Gamma, A \Rightarrow A, \Delta$$

$$\frac{\Gamma \Rightarrow A, \Delta \quad \Gamma, B \Rightarrow \Delta}{\Gamma, A \rightarrow B \Rightarrow \Delta} \quad l\rightarrow$$

$$\frac{\Gamma, A \Rightarrow B, \Delta}{\Gamma \Rightarrow A \rightarrow B, \Delta} \quad r\rightarrow$$

$$\frac{\Gamma, A, B \Rightarrow \Delta}{\Gamma, A \wedge B \Rightarrow \Delta} \quad l\wedge$$

$$\frac{\Gamma \Rightarrow A, \Delta \quad \Gamma \Rightarrow B, \Delta}{\Gamma \Rightarrow A \wedge B, \Delta} \quad r\wedge$$

$$\frac{\Gamma, A[c/x] \Rightarrow \Delta}{\Gamma, \forall x A \Rightarrow \Delta} \quad l\forall$$

$$\frac{\Gamma \Rightarrow \Delta, A[a/x]}{\Gamma \Rightarrow \Delta, \forall x A} \quad r\forall$$

$$\frac{\Gamma^*, A \Rightarrow \Delta^*}{\Gamma, \diamond A \Rightarrow \Delta} \quad l\diamond$$

$$\frac{\Gamma \Rightarrow A, \Delta}{\Gamma \Rightarrow \diamond A, \Delta} \quad r\diamond$$

$$\Gamma^* \stackrel{\text{df}}{=} \{\Box B \mid \Box B \in \Gamma\}$$

$$\Delta^* \stackrel{\text{df}}{=} \{\Diamond B \mid \Diamond B \in \Delta\}$$

- $\frac{pre, \text{Skeleton} \Rightarrow \diamond goal}{pre, \text{Plan} \Rightarrow \diamond goal} \quad seq_inst$

- $\frac{pre, Plan_1; Plan_2 \Rightarrow \Diamond goal}{pre, Plan \Rightarrow \Diamond goal} \textit{sequence_assumption}$
- $\frac{pre, Plan_1 \Rightarrow \Diamond[goal_1 \wedge \text{OF} \wedge pre'] \quad pre', Plan_2 \Rightarrow \Diamond goal_2}{pre, Plan_1; Plan_2 \Rightarrow \Diamond[goal_1 \wedge goal_2]} \textit{effect_split}$
- $\frac{pre, Plan_1 \Rightarrow \Diamond[goal_{new} \wedge \text{OF} \wedge pre'] \quad pre', Plan_2 \Rightarrow \Diamond goal}{pre, Plan_1; Plan_2 \Rightarrow \Diamond goal} \textit{effect_intro}$
- $\frac{pre, \mathbf{if}(cond, Plan_A, Plan_B) \Rightarrow \Diamond goal}{pre, Plan \Rightarrow \Diamond goal} \textit{if_intro}$
- $\frac{pre, cond, Plan_A \Rightarrow goal \quad pre, \neg cond, Plan_B \Rightarrow \Diamond goal}{pre, \mathbf{if}(cond, Plan_A, Plan_B) \Rightarrow \Diamond goal} \textit{if_splitting}$

References

- [Anderson *et al.*, 1994] W. Anderson, J. Hendler, M. Evett, and B. Kettler. Massively parallel matching of knowledge structures. In H. Kitano and J. Hendler, editors, *Massively Parallel Artificial Intelligence*, pages 53–71. MIT Press, Cambridge, MA, 1994.
- [Baader and Hollunder, 1991] F. Baader and B. Hollunder. KRIS: Knowledge representation and inference system. *SIGART Bulletin*, 2(2):8–15, 1991.
- [Baader *et al.*, 1992] F. Baader, B. Hollunder, B. Nebel, H.-J. Profitlich, and E. Franconi. An empirical analysis of optimization techniques for terminological representation systems, or making KRIS get a move on. In Nebel *et al.* [1992], pages 270–281.
- [Bäckström and Sandewall, 1994] C. Bäckström and E. Sandewall, editors. *Current Trends in AI Planning*. IOS Press, Amsterdam, Washington, Tokyo, 1994.
- [Bäckström, 1994] C. Bäckström. Finding least constrained plans and optimal parallel executions is harder than we thought. In Bäckström and Sandewall [1994].
- [Bajjci, 1993] R. Bajjci, editor. *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, 1993.
- [Bauer *et al.*, 1993] M. Bauer, S. Biundo, D. Dengler, J. Koehler, and G. Paul. PHI - a logic-based tool for intelligent help systems. In Bajjci [1993], pages 460–466.
- [Beetz and McDermott, 1994] M. Beetz and D. McDermott. Improving robot plans during their execution. In Hammond [1994], pages 7–12.
- [Biundo and Dengler, 1996] S. Biundo and D. Dengler. The logical language for planning LLP. Research Report, German Research Center for Artificial Intelligence, 1996.
- [Biundo *et al.*, 1992] S. Biundo, D. Dengler, and J. Koehler. Deductive planning and plan reuse in a command language environment. In Neumann [1992], pages 628–632.
- [Borgida and Patel-Schneider, 1994] A. Borgida and P. Patel-Schneider. A semantics and complete algorithm for subsumption in the CLASSIC description logic. *Journal of Artificial Intelligence Research*, 1:277–308, 1994.
- [Brachmann and Levesque, 1984] R. Brachmann and H. Levesque. The tractability of subsumption in frame based description languages. In *Proceedings of the 4th National Conference of the American Association for Artificial Intelligence*, pages 34–37. MIT Press, 1984.
- [Cohn, 1994] A. Cohn, editor. *Proceedings of the 11th European Conference on Artificial Intelligence*. John Wiley & Sons, Chichester, New York, 1994.

- [Constable, 1986] R. Constable. *Implementing Mathematics with the Nuprl Proof Development System*. Prentice Hall, 1986.
- [Dengler, 1994] D. Dengler. An adaptive deductive planning system. In Cohn [1994], pages 610–614.
- [Doyle *et al.*, 1994] J. Doyle, E. Sandewall, and P. Torasso, editors. *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kaufmann, 1994.
- [Feigenbaum, 1963] E. Feigenbaum. The simulation of natural learning behavior. In E. Feigenbaum and J. Feldman, editors, *Computers and Thought*. Mc Graw-Hill, New York, 1963.
- [Friedland and Iwasaki, 1985] P. Friedland and Y. Iwasaki. The concept and implementation of skeletal plans. *Journal of Automated Reasoning*, 1:161–208, 1985.
- [Gabbay, 1989] D. Gabbay. Declarative past and imperative future: Executable temporal logic for imperative systems. In H. Barringer and A. Pnueli, editors, *Proceedings of the Colloquium on Temporal Logic in Specification*, number 398 in Lecture Notes in Computer Science, pages 402–450. Springer, 1989.
- [Gallier, 1987] J. Gallier. *Logic for Computer Science: Foundations of Automatic Theorem Proving*. John Wiley and Sons, New York, 1987.
- [Givan and McAllester, 1992] R. Givan and D. McAllester. New results on local inference relations. In Nebel *et al.* [1992], pages 403–412.
- [Green, 1969] C. Green. Application of theorem proving to problem solving. In D. Walker and L. Norton, editors, *Proceedings of the 1st International Joint Conference on Artificial Intelligence*, pages 219–239, 1969.
- [Hammond, 1990] K. Hammond. Explaining and repairing plans that fail. *Artificial Intelligence*, 45:173–228, 1990.
- [Hammond, 1994] K. Hammond, editor. *Proceedings of the 2nd International Conference on Artificial Intelligence Planning Systems*. AAAI Press, Menlo Park, 1994.
- [Hanks and Weld, 1992a] S. Hanks and D. Weld. Systematic adaptation for case-based planning. In J. Hendler, editor, *Proceedings of the 1st International Conference on Artificial Intelligence Planning Systems*, pages 96–105. Morgan Kaufmann, 1992.
- [Hanks and Weld, 1992b] S. Hanks and D. Weld. The systematic plan adaptor: A formal foundation of case-based planning. Technical Report 92-09-04, University of Washington, Department of Computer Science and Engineering, 1992.
- [Hanks and Weld, 1995] S. Hanks and D. Weld. A domain-independent algorithm for plan adaptation. *Journal of Artificial Intelligence Research*, 2:319–360, 1995.
- [Heinsohn *et al.*, 1991] J. Heinsohn, D. Kudenko, B. Nebel, and H.-J. Profitlich. Integration of action representation in terminological logics. In C. Peltason, K. Luck, and C. Kindermann, editors, *Proceedings of the Terminological Logic Users Workshop*. KIT-Report 95, TU Berlin, Germany, 1991.
- [Heisel *et al.*, 1990] M. Heisel, W. Reif, and W. Stephan. Tactical theorem proving in program verification. In *Proceedings of the 10th International Conference on Automated Deduction*, number 449 in Lecture Notes in Artificial Intelligence, pages 117–131. Springer, Berlin, 1990.
- [Kambhampati and Hendler, 1989] S. Kambhampati and J. Hendler. Control of refitting during plan reuse. In N. Sridharan, editor, *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 943–949. Morgan Kaufmann, 1989.
- [Kambhampati and Hendler, 1992] S. Kambhampati and J. Hendler. A validation-structure-based theory of plan modification and reuse. *Artificial Intelligence*, 55:193 – 258, 1992.

- [Kambhampati, 1989] S. Kambhampati. Flexible reuse and modification in hierarchical planning: A validation-structure-based approach. PhD Thesis MD 207 42-3411, University of Maryland, 1989.
- [Kambhampati, 1990] S. Kambhampati. A theory of plan modification. In *Proceedings of the 8th National Conference of the American Association for Artificial Intelligence*, pages 176–182. MIT Press, 1990.
- [Koehler and Treinen, 1995] J. Koehler and R. Treinen. Constraint deduction in an interval-based temporal logic. In M. Fisher and R. Owen, editors, *Executable Modal and Temporal Logic*, volume 897 of *Lecture Notes in Artificial Intelligence*, pages 103–117. Springer, 1995.
- [Koehler, 1994a] J. Koehler. An application of terminological logics to case-based reasoning. In Doyle et al. [1994], pages 351–362.
- [Koehler, 1994b] J. Koehler. Avoiding pitfalls in case-based planning. In Hammond [1994], pages 104–109.
- [Koehler, 1994c] J. Koehler. Correct modification of complex plans. In Cohn [1994], pages 605–609.
- [Koehler, 1994d] J. Koehler. Flexible plan reuse in a formal framework. In Bäckström and Sandewall [1994], pages 171–184.
- [Kolodner, 1993] J. Kolodner. *Case-Based Reasoning*. Morgan Kaufman, 1993.
- [Kröger, 1987] F. Kröger. *Temporal Logic of Programs*. Springer, Heidelberg, 1987.
- [Leckie and Zuckerman, 1993] C. Leckie and I. Zuckerman. An inductive approach to learning search control rules for planning. In Baijci [1993], pages 1100–1105.
- [Manna and Waldinger, 1987] Z. Manna and R. Waldinger. A theory of plans. In M. Georgeff and A. Lansky, editors, *Reasoning about Actions and Plans: Proceedings of the 1986 Workshop*, pages 11–45. Morgan Kaufmann, 1987.
- [McDermott, 1978] D. McDermott. Planning and acting. *Cognitive Science*, 2:71–109, 1978.
- [Minton, 1988] S. Minton. Quantitative results concerning the utility of explanation-based learning. In *Proceedings of the 7th National Conference of the American Association for Artificial Intelligence*, pages 564–569. MIT Press, 1988.
- [Moszkowski and Manna, 1983] B. Moszkowski and Z. Manna. Reasoning in interval temporal logic. In E. Clarke and D. Kozen, editors, *Proceedings of the Conference on Logics of Programs*, Lecture Notes in Computer Science. Springer, 1983.
- [Nebel and Koehler, 1993] B. Nebel and J. Koehler. Plan modification versus plan generation: A complexity-theoretic perspective. In Baijci [1993], pages 1436–1441.
- [Nebel and Koehler, 1995] B. Nebel and J. Koehler. Plan reuse versus plan generation: A theoretical and empirical analysis. *Artificial Intelligence - Special Issue on Planning and Scheduling*, 76:427–454, 1995.
- [Nebel et al., 1992] B. Nebel, W. Swartout, and C. Rich, editors. *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*. Morgan Kaufmann, 1992.
- [Nebel, 1990a] B. Nebel. *Reasoning and Revision in Hybrid Representation Systems*. Number 422 in Lecture Notes in Artificial Intelligence. Springer, 1990.
- [Nebel, 1990b] B. Nebel. Terminological reasoning is inherently intractable. *Artificial Intelligence*, 43(2):235–249, 1990.
- [Neumann, 1992] B. Neumann, editor. *Proceedings of the 10th European Conference on Artificial Intelligence*. John Wiley & Sons, Chichester, New York, 1992.

- [Ohlbach, 1991] H.-J. Ohlbach. Semantics-based translation methods for modal logics. *Journal of Logic and Computation*, 1(5):691–775, 1991.
- [Paulson, 1990] L. Paulson. Isabelle: The next 700 theorem provers. In P. Odifredi, editor, *Logic and Computer Science*. Academic Press, 1990.
- [Reiter, 1993] R. Reiter. Proving properties of states in the situation calculus. *Artificial Intelligence*, 64:337–351, 1993.
- [Riesbeck and Schank, 1989] C. Riesbeck and R. Schank. *Inside Case-based Reasoning*. Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1989.
- [Rosner and Pnueli, 1986] R. Rosner and A. Pnueli. A choppy logic. In *Symposium on Logic in Computer Science*, Cambridge, Massachusetts, 1986.
- [Schmidt-Schauß and Smolka, 1991] M. Schmidt-Schauß and G. Smolka. Attributive concept descriptions with complements. *Artificial Intelligence*, 48:1–26, 1991.
- [Selman, 1994] B. Selman. Near-optimal plans, tractability, and reactivity. In Doyle et al. [1994], pages 521–529.
- [Stephan and Biundo, 1995] W. Stephan and S. Biundo. Multilevel refinement planning in an interval-based temporal logic. In *Proceedings of the 7th Portuguese Conference on Artificial Intelligence*, volume 990 of *Lecture Notes in Artificial Intelligence*, pages 483–487. Springer, 1995.
- [Traverso *et al.*, 1992] P. Traverso, A. Cimatti, and L. Spalazzi. Beyond the single planning paradigm: introspective planning. In Neumann [1992], pages 643–647.
- [Veloso, 1994] M. Veloso. *Planning and Learning by Analogical Reasoning*. Number 886 in *Lecture Notes in Artificial Intelligence*. Springer, 1994.
- [Wallen, 1989] L. Wallen. *Automated Deduction in Nonclassical Logics*. MIT Press, Cambridge, London, 1989.
- [Weida and Litman, 1994] R. Weida and D. Litman. Subsumption and recognition of heterogeneous constraint networks. In *Proceedings of the Tenth IEEE Conference on Artificial Intelligence for Applications*. IEEE Press, 1994.