

# Exploiting Landmarks for Hybrid Planning

Mohamed Elkawkagy and Pascal Bercher  
and Bernd Schattenberg and Susanne Biundo<sup>1</sup>

**Abstract.** Very recently, the well-known concept of landmarks has been adapted from the classical planning setting to hierarchical planning. It was shown how a pre-processing step that extracts local landmarks from a planning domain and problem description can be used in order to prune the search space that is to be explored before the actual search is performed. This pruning technique eliminates all branches of the task decomposition tree, for which can be proven that they will never lead to a solution. In this paper, we investigate this technique in more detail and extend it by introducing search strategies which use these local landmarks in order to guide the planning process more effectively towards a solution. Our empirical evaluation shows that the pre-processing step dramatically improves performance because dead ends can be detected much earlier than without pruning and that our search strategies using the local landmarks outperform many other possible search strategies.

## 1 Introduction

In recent years, the exploitation of knowledge gained by pre-processing a planning domain and/or problem description has proven to be an effective means to reduce planning effort. Various pre-processing procedures, like effect relaxation [3, 13], abstractions [11], and landmarks [19] have been proposed for classical planning, where they serve to compute strong search heuristics. However, pre-processing techniques can also be used to perform some pruning of the search space before the actual search is performed. Very recently, different techniques have been introduced which restrict the domain and problem description of an Hierarchical Task Network (HTN) problem to a smaller subset, since some parts of the domain description might be irrelevant for the given problem at hand [5, 10]. In this paper, we investigate our previously introduced landmark technique [5] in more detail, which uses local landmarks to prune the search space that is to be explored before the actual search is performed. We further investigate, how search strategies can take advantage of these extracted local landmarks.

While the use of landmark tasks is a novelty in hierarchical planning, it is a familiar concept in classical state-based planning. There, landmarks are *facts* that have to hold in some intermediate state of every plan that solves the problem. The concept was introduced by Porteous et al. [19] and further developed by Hoffmann et al. [14] and Zhu and Givan [26], where landmarks and orderings between them are extracted

from a planning graph of the relaxed planning problem. Other strands of research arranged landmarks into groups of intermediate goals to be achieved [24] and extended the landmark concept to so-called disjunctive landmarks [9, 18]. A disjunctive landmark is a set of literals any of which has to be satisfied in the course of a valid plan. A generalization of landmarks resulted in the notion of so-called action landmarks [16, 25]. An action is an action landmark if it occurs in every solution plan. Most of the recent landmark approaches use landmark information to compute heuristic functions for a forward searching planner [16, 20] and investigate their relations to critical-path-, relaxation-, and abstraction-heuristics [12]. In summary, it turned out that the use of landmark information can significantly improve the performance of classical state-based planners.

In hierarchical planning, *landmarks* are mandatory abstract or primitive tasks, i.e. tasks that have to be performed by any solution plan. *Local landmarks* are abstract or primitive tasks that are mandatory, given their parent task is mandatory (where a parent task is the abstract task that introduced the local landmark by decomposition). That is, a local landmark is also a landmark if its parent is one, too. For an initial task network that states a current planning problem, a pre-processing procedure computes the corresponding local landmarks. It does so by systematically inspecting the methods that are eligible to decompose the relevant abstract tasks. Beginning with the (landmark) tasks of the initial network, the procedure follows the way down the decomposition hierarchy until no further abstract tasks qualify as local landmarks. Using the precondition and effects of primitive tasks, one can perform a relaxed reachability test [8]. A failure indicates that the method which introduced the primitive task is no longer eligible. If the tested primitive task was a local landmark, we can even get further: its parent abstract task can never be decomposed into a solution because one of its local landmarks cannot be achieved. Hence, this abstract (parent) task can also be pruned without the need of inspecting the primitive tasks in the other methods for this abstract task. Being able to prune useless regions of the search space this way, a hierarchical planner performs significantly better than it does without exploiting the local landmark information.

Before introducing the local landmark extraction procedure for hierarchical planning in Section 3, we will briefly review HTN planning in general and our underlying framework and planning procedure in particular (Section 2). Afterwards, Section 4 shows how the information about local landmarks can be used during planning. It presents experimental results from a set of benchmark problems of the *UM-Translog* [1] and

---

<sup>1</sup> Institute of Artificial Intelligence, Ulm University, D-89069 Ulm, Germany, email: *forename.surname@uni-ulm.de*

*Satellite* domains, which give evidence for the considerable performance increase gained by pre-processing the planning problem to prune unnecessary parts and by the use of the novel search strategies using the local landmarks. The paper ends with possible extensions to our approach (Section 5) and with some concluding remarks in Section 6.

## 2 Formal Framework

Hierarchical Task Network (HTN) planning is based on the concepts of tasks and methods [6]. Abstract tasks represent compound activities like making a business trip or transporting certain goods to a specific location. Primitive tasks correspond to classical planning operators. Hierarchical domain models hold a number of methods for each abstract task. Each method provides a task network, also called partial plan, which specifies a pre-defined (abstract) solution of the corresponding abstract task. A planning problem consists of finding a decomposition of the initial task network, using the tasks and methods provided by the domain model. Thus, the planning problem is solved by incrementally decomposing the abstract tasks in the initial task network until it contains only primitive tasks and is consistent w.r.t. their ordering and causal structure. The decomposition of an abstract task by an appropriate method replaces the abstract task by the partial plan specified by the respective method.

Our approach [2] relies on a *hybrid* planning framework [7, 15], which combines HTN planning with concepts of partial-order-causal-link (POCL) planning. The resulting systems integrate task decomposition with explicit causal reasoning. Therefore, they are able to use predefined standard solutions like in pure HTN planning and can thus benefit from the landmark technique we will introduce below; they can also develop (parts of) a plan from scratch or modify a default solution (i.e., a method’s task network) in cases where the initial state deviates from the presumed standard. It is this flexibility that makes hybrid planning particularly well suited for real-world applications [4, 7].

In our framework, a task network or partial plan  $P = \langle S, \prec, V, C \rangle$  consists of a set of plan steps  $S$ , i.e., (partially) instantiated task schemata, augmented with a unique label to differentiate between multiple occurrences of the same task. We denote by  $Tasks(S)$  the set of (partially) instantiated task schemata in  $S$ , i.e.,  $S$  without labels. It also contains a set of ordering constraints  $\prec$  that impose a partial order on the plan steps, a set of variable constraints  $V$ , and a set  $C$  of causal links. Variable constraints are (in-) equations between variables or between variables and constants. A causal link  $s_i \rightarrow_{\varphi} s_j$  indicates that the precondition  $\varphi$  of plan step  $s_j$  is an effect of plan step  $s_i$  and is *supported* this way. A domain model  $D = \langle T, M \rangle$  includes a set of task schemata and a set of decomposition methods. A task schema  $t(\bar{\tau}) = \langle prec(t(\bar{\tau})), add(t(\bar{\tau})), del(t(\bar{\tau})) \rangle$  specifies the preconditions as well as the positive and negative effects of a task. Preconditions and effects are sets of literals and  $\bar{\tau} = \tau_1 \dots \tau_n$  are the task parameters. In the hybrid setting, both primitive and abstract tasks show preconditions and effects. This enables the use of POCL planning operations even on abstract levels. However, in this paper we restrict our language to pure HTN; preconditions and effects are thus omitted for abstract tasks. A method  $m = \langle t, P \rangle$  relates an abstract task  $t$  to a

partial plan  $P$ , which represents an (abstract) solution or “implementation” of the task. In general, a number of different methods are provided for each abstract task. Please note that no application conditions are associated with the methods, as opposed to other representatives of HTN-style planning. A planning problem  $\Pi = \langle D, s_{init}, P_{init} \rangle$  includes a domain model  $D$ , an initial state  $s_{init}$ , and  $P_{init}$ , which represents an initial partial plan. Please note, that in our *hybrid* planning framework, one can also specify a goal state. However, since we restrict ourselves in this paper to pure HTN planning, the goal state is omitted.

Based on these strictly declarative specifications of planning domains and problems, hybrid and HTN planning is performed by refining an initial partial plan  $P_{init}$  of  $\Pi$  stepwise until a partial plan  $P = \langle S, \prec, V, C \rangle$  is obtained that satisfies the following solution criteria:

1.  $P$  is a refinement of  $P_{init}$ , i.e., it is a successor of the initial plan in the induced search space (cf. Definition 1),
2. each precondition of a plan step in  $P$  is supported by a causal link in  $C$ ,
3. the ordering and variable constraints are consistent, i.e., the ordering does not induce cycles on the plan steps and the (in-) equations of variable constraints are free of contradiction,
4. none of the causal links in  $C$  is threatened, i.e., for each causal link  $s_i \rightarrow_{\varphi} s_j$  the ordering constraints ensure that no plan step  $s_k$  with effect  $\neg\varphi$  can be ordered between plan steps  $s_i$  and  $s_j$ , and
5. all plan steps in  $S$  are primitive tasks.

Before we present our planning algorithm in more detail, we define the search space induced by the HTN planning problem  $\Pi$ . Refinement steps include the decomposition of abstract tasks by appropriate methods, the insertion of causal links to support open preconditions of plan steps as well as the insertion of ordering and variable constraints. We call such a refinement step a *plan modification*.

**Definition 1** (Induced Search Space). *Let  $\mathcal{P}_{\Pi} = \langle \mathcal{V}, \mathcal{E} \rangle$  be the directed acyclic graph which represents the (possibly infinite) search space induced by a planning problem  $\Pi = \langle D, s_{init}, P_{init} \rangle$ . Then, the set of vertexes  $\mathcal{V}$  is the set of plans in the induced search space and the set of edges  $\mathcal{E}$  corresponds to the set of plan used modifications. By abuse of notation, we write  $P \in \mathcal{P}_{\Pi}$  to state  $P \in \mathcal{V}$ . The root of  $\mathcal{P}_{\Pi}$  is the initial plan of  $\Pi$ ; thus,  $P_{init} \in \mathcal{P}_{\Pi}$ . The direct successors of a plan  $P \in \mathcal{P}_{\Pi}$  are all Plans  $P'$ , such that  $P'$  resulted from  $P$  by applying a plan modification  $m$  to  $P$ . Then,  $m \in \mathcal{E}$ .*

Now, we present our planning algorithm (Algorithm 1) which takes the initial plan of the planning problem  $\Pi$  as an input and refines it stepwise until a solution is found. Our algorithm performs an informed search, guided by so-called search strategies, in the search space induced by the HTN planning problem  $\Pi$  (cf. Definition 1).

The fringe of the algorithm is a plan sequence  $\langle P_1 \dots P_n \rangle$  ordered by the used search strategy. It contains all non-visited plans that are direct successors of visited non-solution plans. According to the used search strategy, a plan  $P_i$  leads more quickly to a solution than plans  $P_j$  for  $j > i$ . The current plan under consideration is always the first plan of the fringe.

---

**Algorithm 1:** Planning Algorithm

---

**Input** : The sequence  $\text{Fringe} = \langle P_{init} \rangle$ .**Output:** A solution or *fail*.

```

1 while  $\text{Fringe} = \langle P_1 \dots P_n \rangle \neq \varepsilon$  do
2    $F \leftarrow f^{\text{FlawDet}}(P_1)$ 
3   if  $F = \emptyset$  then return  $P_1$ 
4    $\langle m_1 \dots m_{n'} \rangle \leftarrow f^{\text{ModOrd}}(\bigcup_{f \in F} f^{\text{ModGen}}(f))$ 
5    $\text{succ} \leftarrow \langle \text{app}(m_1, P_1) \dots \text{app}(m_{n'}, P_1) \rangle$ 
6    $\text{Fringe} \leftarrow f^{\text{PlanOrd}}(\text{succ} \circ P_2 \dots P_n)$ 
7 return fail

```

---

The planning algorithm loops as long as no solution is found and there are still plans to refine (line 1). Hence, in line 2, the flaw detection function  $f^{\text{FlawDet}}$  calculates all flaws of the current plan. A flaw is a plan component that violates a solution criterion. For instance, in the HTN planning setting, (the occurrence of) an abstract task is a flaw. If no flaws can be found, the plan is a solution and is returned (line 3). In line 4, all plan modifications are calculated by the modification generating function  $f^{\text{ModGen}}$ , which address all found flaws. Afterwards, the modification ordering function  $f^{\text{ModOrd}}$  orders all these modifications according to a given strategy. Finally, this fringe is updated in two steps: First, the plans resulting from applying the modifications are calculated (line 5) to be inserted at the head of the fringe in line 6. Afterwards, the plan ordering function  $f^{\text{PlanOrd}}$  takes the updated fringe and orders it according to its strategy. This step can also be used in order to discard some plans (i.e., to delete some plans permanently from the fringe). This is useful for plans which contain unresolvable flaws like an inconsistent ordering of tasks. If the fringe becomes empty, no solution exists and *fail* gets returned.

In contrast to other systems, which *implicitly* define their search strategy by their search procedure, our approach – implemented in the planning environment PANDA [21] (Planning and Acting in a Network Decomposition Architecture) – *explicitly* defines the search strategy: It is the result of the combination of the used modification and plan ordering functions. Let us take a look at a simple example strategy for clarification: To perform a depth first strategy, the plan ordering strategy has to be the identity (i.e.,  $f^{\text{PlanOrd}}(\bar{P}) = \bar{P}$  for any plan sequence  $\bar{P}$ ), whereas the modification ordering strategy  $f^{\text{ModOrd}}$  can be arbitrary (but decides, which branches to visit first). Thus, the plan ordering strategy is used to prioritize the plans; several strategies can be concatenated for tie-breaking. The plan ordering strategy uses also its input sequence for tie breaking: If two plans are still invariant after application of the plan ordering function, the order given in the input is used.

Many different plan ordering strategies<sup>2</sup> have been described and evaluated in our previous work [21, 22, 23]. In this work, we will only sketch those used in the experiments (cf. Subsection 4.1).

### 3 Local Landmark Extraction

For a given hierarchical planning problem  $\Pi = \langle D, s_{init}, P_{init} \rangle$ , *global landmarks*<sup>3</sup> are the tasks that occur in every sequence of decompositions leading from the initial task network  $P_{init}$  to a solution plan. However, we do not calculate (global) landmarks, but – what we call – *local landmarks*. Local landmarks are landmarks with respect to a given abstract task. We will define them more formally in this section. The local landmark extraction is performed using a so-called *task decomposition tree* (TDT) of  $\Pi$ . Figure 1 depicts such a tree schematically. The TDT of  $\Pi$  is an AND/OR tree that represents all possible ways to decompose the abstract tasks of  $P_{init}$  by methods in  $D$  until a primitive level is reached or a task is encountered that is already included somewhere in the TDT. Each level of a TDT consists of two parts: a task and a method level. Method nodes are AND nodes, because their children are the tasks that occur in the partial plan of the respective method, all of which have to be performed in order to apply the corresponding method. Task nodes, on the other side, are OR nodes, because their children are the methods that can be used to decompose the respective task. To avoid loops, each abstract task is decomposed only once in the TDT; hence, all but one identical and fully grounded abstract tasks become leaf nodes in the TDT. Other leaf nodes are the primitive tasks. A TDT is built by forward chaining from the (grounded) abstract tasks in the initial task network until all nodes of the fringe are leaf nodes. The root node on level 0 is an artificial method node that represents the initial partial plan  $P_{init}$ .




Figure 1. A schematic task decomposition tree.

Before we can formally define landmarks and local landmarks, we first need to define plan and solution sequences, respectively.

**Definition 2** (Plan and Solution Sequences). *For a planning problem  $\Pi$  and a given plan  $P \in \mathcal{P}_\Pi$ , let  $\text{Seq}_\Pi(P)$  be the set of all plan sequences in  $\mathcal{P}_\Pi$  rooted in that plan, i.e.,  $\text{Seq}_\Pi(P) = \{\langle P_1 \dots P_n \rangle \mid P = P_1, P_1 \in \mathcal{P}_\Pi \text{ and } P_{i+1} \text{ is a direct successor of } P_i \in \mathcal{P}_\Pi \text{ for all } 1 \leq i < n\}$ .*

*The set of all solution sequences rooted in  $P$  is then  $\text{SolSeq}_\Pi(P) = \{\langle P_1 \dots P_n \rangle \in \text{Seq}_\Pi(P) \mid P_n \text{ solution of } \Pi, n \geq 1\} \subseteq \text{Seq}_\Pi(P)$ .*

**Definition 3** (Landmark). *A landmark is a grounded (i.e., fully instantiated) task that occurs in every sequence of decompositions leading from the initial task network to a solution. That is, the task  $t$  is called a landmark of  $\Pi = \langle D, s_{init}, P_{init} \rangle$ ,*

<sup>2</sup> In previous work, we called the plan and modification ordering functions plan and modification *selection* functions, respectively.

<sup>3</sup> In the following, we will only call them *landmarks*.

if for every sequence  $\langle P_1 \dots P_n \rangle \in \text{SolSeq}_{\Pi}(P_{init})$  there is an  $1 \leq i \leq n$ , such that  $t \in \text{Tasks}(P_i)$ .

Whereas a landmark has to occur in every decomposition sequence of a solution (which is rooted in the initial plan), a *local* landmark only has to occur in each solution sequence rooted in a plan containing a specific task  $t$ .

**Definition 4** (Local Landmark of an Abstract Task). *For a given grounded abstract task  $t$ , let  $\mathcal{P}_{\Pi}(t)$  be the set of all plans in  $\mathcal{P}_{\Pi}$  containing  $t$ , i.e.,  $\mathcal{P}_{\Pi}(t) = \{P \in \mathcal{P}_{\Pi} | t \in \text{Tasks}(P)\}$ .*

*We call the grounded task  $t'$  a local landmark of  $t$ , if for all  $P \in \mathcal{P}_{\Pi}(t)$  holds, that for all sequences  $\langle P_1 \dots P_n \rangle \in \text{SolSeq}_{\Pi}(P)$  there is a  $P_j$  with  $j > 1$  such that  $t' \in \text{Tasks}(P_j)$ .*

We use the next definition to calculate all tasks that occur in all available methods for the same abstract task.

**Definition 5** (Common Task Set Operator  $\widehat{\cap}$ ). *Let  $t$  be an abstract task in the TDT and  $m_i = \langle t, \langle S_i, \prec_i, V_i, C_i \rangle \rangle$  and  $m_j = \langle t, \langle S_j, \prec_j, V_j, C_j \rangle \rangle$  two of its methods in the TDT. That is, both  $t$  and its methods are fully grounded. Then, the Common Task Set Operator  $\widehat{\cap}$  of  $m_i$  and  $m_j$  is defined as*

$$m_i \widehat{\cap} m_j = \text{Tasks}(S_i) \cap \text{Tasks}(S_j)$$

Using this definition, we can calculate the intersection of an abstract task  $t$ ,  $I(t)$ , by intersecting all available methods. Obviously, the tasks contained in  $I(t)$  are local landmarks, because these tasks are contained in all solution sequences that are rooted in a plan containing  $t$ . It is also notable, that all tasks in  $I(t)$  are local landmarks of  $t$  if  $t$  is not contained in any solution sequence<sup>4</sup> (i.e., if for all  $\langle P_1 \dots P_n \rangle \in \text{SolSeq}_{\Pi}(P_{init})$  holds, that there is no  $P_i, 1 \leq i \leq n$  such that  $t \in \text{Tasks}(P_j)$ ).

However, not all local landmarks of an abstract task can be detected that way because not all local landmarks have to be in such an intersection.

We would also like to emphasize, that local landmarks are in general no landmarks. This is obvious, because one can calculate the local landmarks of an abstract task which is not contained in all valid decompositions (or even in *any* valid decomposition) of the initial plan.

Based on the definition of the common task set operator, we will now define the remaining task set operator which calculates the set of tasks in which two (grounded) methods differ.

**Definition 6** (Remaining Task Set Operator  $\widehat{\cup}$ ). *Let  $t$  be an abstract task in the TDT and  $m_i = \langle t, \langle S_i, \prec_i, V_i, C_i \rangle \rangle$  and  $m_j = \langle t, \langle S_j, \prec_j, V_j, C_j \rangle \rangle$  two of its methods in the TDT. Then, the Remaining Task Set Operator  $\widehat{\cup}$  of  $m_i$  and  $m_j$  is defined as*

$$m_i \widehat{\cup} m_j = \{\text{Tasks}(S_i) \setminus (m_i \widehat{\cap} m_j), \text{Tasks}(S_j) \setminus (m_i \widehat{\cap} m_j)\}$$

Analogously to the intersection  $I(t)$  of an abstract task  $t$ , we can define its remaining tasks  $R(t)$ , by applying the remaining task set operator to all methods of  $t$  in the TDT.  $I(t)$  and  $R(t)$  can be regarded as a partition of the methods of  $t$  in the TDT, i.e., it holds:  $\{I(t) \cup r | r \in R(t), \text{if } R(t) \neq \emptyset \text{ or } r = \emptyset, \text{ else}\} = \{\text{Tasks}(P) | \text{there is a method } m = \langle t, P \rangle \text{ in the TDT}\}$ .

The landmark extraction algorithm (Algorithm 2) calculates for each abstract task occurring in the TDT these two

<sup>4</sup> In fact, all grounded tasks  $t'$  are local landmarks of  $t$  if  $t$  is not contained in any solution sequence.

**Table 1.** A schematic landmark table, showing in each line an ground instance of an abstract task, the intersection of its decompositions and the remaining task sets.

abstr. Tasks	Intersection	Remaining
$t_1$	$I(t_1)$	$R(t_1)$
$t_2$	$I(t_2)$	$R(t_2)$
$\vdots$	$\vdots$	$\vdots$

sets and stores it into a so-called landmark table. Table 1 shows such a landmark table schematically. The algorithm takes a TDT<sup>5</sup>, which is computed before the algorithm is called, as input and returns a landmark table after its termination.

Intuitively, the algorithm simply tests all primitive tasks for relaxed reachability, starting with the initial plan (the root of the TDT) and proceeding level by level of the TDT. If a task can be proven unreachable, the method introducing this task is pruned from the TDT and all its sub-nodes (and so forth). After all infeasible methods of an abstract task  $t$  have been pruned from the TDT, this task, its intersection, and the remaining tasks are stored into the landmark table.

Now, we will take a look, how this is achieved by our algorithm: First, the landmark table and a set for backward propagation get initialized (line 1). Afterwards, each abstract task, which is not yet stored into the landmark table is considered level by level of the TDT (line 2 to 4). For the current abstract task at hand, line 6 to 8 calculate the intersection and the remaining tasks in the yet unpruned TDT according to Definition 5 and 6. In line 8, we subtract the empty set from  $R(t)$ , because we are only interested in the tasks, that are actually remaining; if there are no remaining tasks,  $R(t)$  should be empty, instead of containing an empty set. After the tasks introduced by decomposition of  $t$  have been partitioned into  $I(t)$  and  $R(t)$ , these sets are analyzed for infeasibility. This test is performed by a relaxed reachability analysis. First, we study the primitive tasks of  $I(t)$  (line 9). If such a task can be proven to be infeasible, all methods of  $t$  become obsolete and can hence be pruned from the TDT<sup>6</sup> (line 10 and 12). After this test, each remaining task set is tested for reachability. If an infeasible task can be found, only this specific method gets pruned from the TDT (line 13 to 17). If something was pruned, the loop (line 5 to 18) enters another cycle, because the set  $I(t)$  might have grown. If no more pruning is possible, the intersection and remaining task sets for  $t$  are stored into the landmark table in line 19. When storing an entry in line 21, it is checked whether the stored abstract task is feasible or not (an abstract task is infeasible if it does not have any methods left, i.e., if  $I(t)$  and  $R(t)$  are empty). If some abstract task could actually be proven infeasible, it is stored for backward propagation, because again all methods containing this abstract task can be pruned from the TDT and from the landmark table. Finally, if all abstract tasks are checked, the backward propagation procedure is called with the current landmark table and TDT in line 22.

<sup>5</sup> We use the indefinite article, because only the task decomposition graph is unique, whereas the resulting task decomposition tree depends on the chosen order in which tasks get decomposed.

<sup>6</sup> In the presented algorithm, the remaining task sets would still be tested, which is obviously not necessary. However, for the sake of readability, we did not handle this case in the algorithm.



**Procedure propagate** takes as input the already filled landmark table, the possibly pruned TDT and a set **infeasible** of abstract tasks which have been proved infeasible due to no remaining methods in the TDT. It works tail-recursively and returns the final landmark table as soon as no propagation is possible (line 1). To this end, it first takes and removes some arbitrary task  $t'$  from the set **infeasible**. Because this abstract task was proven infeasible, all methods containing it have to be removed from the TDT. As a consequence of this pruning, the intersection and remaining task sets have to be updated; additionally, further propagation can now be possible. To calculate the methods that can possibly be pruned, all parent tasks of  $t'$  are identified (line 3). Then, for all these parents (line 4), the respective methods are removed in line 5. Because methods were removed, the intersection and the remaining task sets could have changed again. Hence, they are recalculated in line 6 to 8. Next, the old landmark table entry of the current parent  $t$  is removed and replaced by the new one (line 9). In line 11, it is tested again, whether the new landmark table entry corresponds to an infeasible abstract task. If so, it is put into the set **infeasible** for later testing. The procedure is then called with the modified parameters in line 1.

Without a formal proof, we want to mention that **Algorithm 2** (i.e., the initial landmark table calculation as well as the backward propagation) always terminates. For the first part of the algorithm, this is easy to see because both loop conditions (line 2 and 3) cannot be modified within the loops. For the second part, i.e., the propagate procedure, we have to show that the set **infeasible** becomes empty eventually. This is the case because each task gets inserted at most once and will be removed at some point.

After the algorithm terminated, the TDT does not have to be considered anymore. All necessary information is encoded in the landmark table.

As we have already pointed out, we only calculate *local* landmarks. That is, given a landmark table entry  $(t, I(t), R(t))$ ,  $I(t)$  contains some of the local landmarks of  $t$ , which, in general, don't have to be actual landmarks because  $t$  was not proven to be a landmark. However, all local landmarks of the abstract tasks in the task level 1 of the TDT are also actual landmarks, because all tasks in the task level 1 are those contained in the initial task and hence landmarks. Thus, if we restrict our local landmark extraction procedure to calculate only the local landmarks of tasks which are (local) landmarks by themselves, all tasks in the  $I(t)$  sets in the local landmark table returned are actual landmarks, too. These landmarks are, however, of limited use because every decomposition contains them anyway. Thus, a "guiding" towards these landmarks as done in classical planning does not bring any benefit.

## Example

In order to illustrate our landmark extraction technique, let us consider a simple example from the UM-Translog domain [1]. Assume a package  $P_1$  is at location  $L_1$  in the initial state and we would like to transport it to a customer location  $L_3$  in the same city. Figure 2 shows a part of the task decomposition tree of this example.

The local landmark extraction algorithm detects that the

---

### Algorithm 2: Local landmark Extraction Algorithm

---

**Input** : A task decomposition tree TDT.

**Output**: The filled landmark table LT.

---

```

1 LT ← ∅, infeasible ← ∅
2 for i ← 1 to TDT.maxDepth() do
3   foreach abstract task t in level i of TDT do
4     if LT contains an entry for t then continue
5     repeat
6       Let M be the methods of t in the TDT.
7       I(t) ← ⋂m∈M m
8       R(t) ← (⋃m∈M m) \ {∅}
9       foreach primitive task t' ∈ I(t) do
10        if t' can be proven infeasible then
11          remove all m ∈ M from the TDT,
12          including all sub-nodes.
13          break
14        foreach remaining task set r ∈ R(t) do
15          foreach primitive task t' ∈ r do
16            if t' can be proven infeasible then
17              remove the method m = ⟨t, P⟩, with
18              Tasks(P) = I(t) ∪ r from the TDT,
19              including all sub-nodes.
20              continue
21      until no method was removed from TDT
22      LT ← LT ∪ {(t, I(t), R(t))}
23      if I(t) = R(t) = ∅ then
24        infeasible ← infeasible ∪ {t}
25  return propagate(LT, TDT, infeasible)

```

---

first level in the TDT contains only one abstract task  $t = transport(P_1, L_1, L_3)$  and that there is only one method,  $Pi.ca.de$ , that can decompose the task into a partial plan, which contains the subtasks  $pickup(P_1)$ ,  $carry(P_1, L_1, L_3)$ , and  $deliver(P_1)$ .

$I(t)$  becomes  $\{pickup(P_1), carry(P_1, L_1, L_3), deliver(P_1)\}$  and  $R(t) = \emptyset$ . The current abstract task and the sets  $I(t)$  and  $R(t)$  are entered as the first row of the landmark table as shown in Table 2.

The landmark extraction algorithm then takes the (unchanged) TDT to investigate the next tree level. The abstract tasks to be inspected on this level are  $pickup(P_1)$  and  $carry(P_1, L_1, L_3)$ . The primitive task  $deliver(P_1)$  is tested and considered executable. Suppose, the task  $t = pickup(P_1)$  is chosen first in line 3 of algorithm 2. As shown in Figure 2, the TDT accounts for three methods to decompose this task:  $Pickup.hazardous$ ,  $Pickup.normal$ , and  $Pickup.valuable$ . By computing the common task set and the remaining task sets we get  $I(t) = \{collect.fees(P_1)\}$ , and  $R(t) = \{\{have.permit(P_1)\}, \{collect.insurance(P_1)\}\}$ . At this point, the relaxed reachability analysis is performed. First,  $collect.fees(P_1)$  is being tested, because it is contained in the intersection  $I(t)$ . Suppose, this task can not be proven to be infeasible. Then, each primitive task in each set  $r \in R(t)$  has to be checked. Assume the primitive task  $have.permit(P_1)$  is feasible, whereas  $collect.insurance(P_1)$  is not. The method  $Pickup.valuable$  is therefore deleted from the TDT. After an

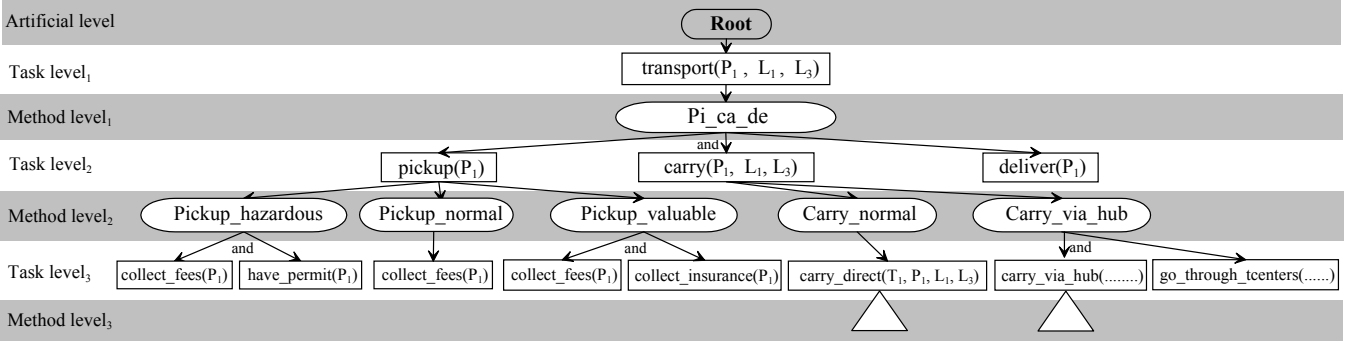


Figure 2. Part of the TDT for the transportation task

Table 2. Example landmark table containing the first three entries for the transportation task illustrated in Figure 2. The sets in the right most column are indexed by the method’s name that contains its tasks.

abstr. Task	Intersection	Remaining
$\text{transport}(P_1, L_1, L_3)$	$\{\text{pickup}(P_1), \text{carry}(P_1, L_1, L_3), \text{deliver}(P_1)\}$	$\emptyset$
$\text{pickup}(P_1)$	$\{\text{collect\_fees}(P_1)\}$	$\{\{\text{have\_permit}(P_1)\}_{\text{Pickup\_hazardous}}\}$
$\text{carry}(P_1, L_1, L_3)$	$\{\{\text{carry\_direct}(T_1, P_1, L_1, L_3)\}\}$	$\emptyset$

### Procedure propagate(LT, TDT, infeasible)

**Input** : A landmark table LT, a task decomposition tree TDT, possibly pruned, and a set of abstract tasks **infeasible**, which have been proved infeasible.

**Output**: the updated landmark table LT, in which methods are pruned that contain infeasible abstract tasks.

```

1 if infeasible = ∅ then return LT
2 infeasible ← infeasible \ {t'}, where t' ∈ infeasible.
3 parents ← {t|(t, I(t), R(t)) ∈ LT, t' ∈ I(t) ∪ ⋃_{r ∈ R(t)} r}
4 foreach t ∈ parents do
5   Remove all methods from the TDT, that contain t' in
   its plan, i.e., all m = ⟨t, P⟩ with t' ∈ Tasks(P).
6   Let M be the methods of t in the TDT.
7   I(t) ← ⋂_{m ∈ M} m
8   R(t) ← (⋂_{m ∈ M} m) \ {∅}
9   LT ← (LT \ {(t, I'(t), R'(t)) ∈ LT}) ∪ {(t, I(t), R(t))}
10  if I(t) = R(t) = ∅ then
11    infeasible ← infeasible ∪ {t}
12 return propagate(LT, TDT, infeasible)

```

additional iteration in which  $I(t)$  and  $R(t)$  get recalculated, the current abstract task  $t = \text{pickup}(P_1)$ , the set  $I(t)$ , and the modified set  $R(t)$  are added to the landmark table as depicted in the second line of Table 2. From the fact that  $R(t)$  contains only one set  $r$ , we can conclude that there is another method with no remaining tasks (if there were no such method, the tasks of  $r \in R(t)$  would be contained in  $I(t)$ ).

In the second iteration (line 3) the abstract task  $t = \text{carry}(P_1, L_1, L_3)$  is considered. The methods  $\text{Carry\_normal}$  and  $\text{Carry\_via\_hub}$  are available to decompose this task. We obtain  $I(t) = \emptyset$  and  $R(t) = \{\{\text{carry\_direct}(T_1, P_1, L_1, L_3)\}\}$ ,

$\{\text{carry\_via\_hub}(\dots), \text{go\_through\_tcenters}(\dots)\}$ . Suppose the primitive task  $\text{go\_through\_tcenters}(\dots)$  is infeasible. The subtree with root  $\text{carry\_via\_hub}(\dots)$  has then to be removed from the TDT. Because the TDT was changed, the iteration (line 5 to 18) enters another cycle. Because there is now only one method left,  $I(t)$  now contains all tasks of this remaining method. Hence, the current abstract task  $t = \text{carry}(P_1, L_1, L_3)$  together with the modified  $I(t)$  and  $R(t)$  are added to the landmark table as depicted in the last line of Table 2.

## 4 Landmark Exploitation

The information about landmarks can be exploited in two ways: The first is to deduce heuristic guidance from the knowledge about which tasks have to be decomposed on refinement paths that lead towards a solution. But before we investigate into this matter, we will present a second way of landmark exploitation, namely the reduction of domain models or, more precisely, the transformation of a universal domain model into one that includes problem-specific pruning information.

### 4.1 Domain Model Reduction

During the construction of the landmark table, the feasibility check and the consecutive propagation of its result into the abstract task level lead to a pruning of the task decomposition tree. The result of this analysis implies that if a method is removed from the TDT during the operation of our landmark extraction algorithm, it can be safely ignored as a refinement option during plan generation.

We consequently supply our refinement generating module with the landmark table for the current planning problem and verify for every incoming abstract task flaw, which of the methods specified in the domain model are reasonably applicable. However, the landmark table is built from grounded tasks, while the plan generation procedure operates on lifted

instances for which the final grounding is yet to be computed. We therefore calculate all groundings of the abstract task at hand that are consistent with the current variable constraints and match these grounded tasks  $t$  with the entries in the landmark table. The union of the (lifted) method schemata that constitute the (grounded) instances in the Remaining Task Sets  $R(t)$  is the set of method schemata that we consider for application to the currently flawed abstract task. Obviously, the earlier a task is addressed in the planning process, the less variable constraints are typically introduced in the partial plan, the more task groundings are implied by the lifted instance, and consequently the less probable is one of its methods pruned by this technique.

In order to quantify the effect of this landmark exploitation technique, we have performed several benchmark tests on the UM-Translog domain with various different search strategies. Table 3 shows the domain model sizes after our pruning process. According to this table, the pruning technique achieves a reduction of the number of abstract task instances that ranges between 33% and 43%, while the reduction of the number of inapplicable methods per instance varied between 27% and 41%.

**Table 3.** This table shows the remaining sizes of the domain model after our reduction for typical problems from the *UM-Translog* domain. On all problems that are grouped together the same reduction was achieved.

Problem	abstr. Tasks (of 21)	Methods (of 51)
<i>Regular Truck Problems</i>		
Hopper Truck, Auto Truck, Regular Truck.3 Locations Regular Truck.2 Region, Regular Truck.1, Regular Truck.2,	12 (57%)	30 (59%)
<i>Various Truck Type Problems</i>		
Flatbed Truck, Armored-R-Truck	12 (57%)	32 (63%)
<i>Traincar Problems</i>		
Auto Traincar, Mail Traincar, Auto Traincar bis, Refrigerated Regular Traincar	14 (67%)	32 (63%)
<i>Airplane Problem</i>		
Airplane	14 (67%)	37 (73%)

In theory, it is quite intuitive that a reduced domain model leads to an improved performance of the planning system. It is however hard to predict the actual effect the pruning information on the grounded instance level has on the lifted computations, in particular taking into account that the landmark table typically contains a number of “distracting” local landmarks that are located on non-solution paths. In order to quantify the practical performance gained by the hierarchical landmark technique, we therefore conducted a series of experiments in the PANDA planning environment [21]. The planning strategies we used are representatives from the rich portfolio provided by PANDA, which has been discussed in previous work [21, 22, 23]. We briefly review the ones on which we based our experiments.

As was already mentioned in Section 2, the search strategy

is encoded by the combination of the modification and plan ordering functions. We distinguish ordering principles that are based on a prioritization of certain flaw or modification classes and strategies that opportunistically choose from the presented set. We call the latter ones *flexible strategies*.

Representatives for *inflexible* strategies are the classical HTN strategy patterns that try to balance task expansion with respect to other plan refinements. The *SHOP strategy*, like the system it is named after [17], prefers task expansion for the abstract tasks in the order in which they are to be executed. The expand-then-make-sound (*ems*) modification ordering strategy alternates task expansion modifications with other classes, resulting in a “level-wise” concretion of all plan steps. The third type of classical HTN strategies – the preference of expansion as it has been realized in the UMCP system [6] – has been omitted in this survey because it trivially benefits from the reduced method set.

As for the *flexible* modification orderings, we included the well-established Least Committing First (*lcf*) paradigm, a generalization of POCL strategies that prioritizes those modifications higher that address flaws for which the smallest number of alternative solutions has been proposed. From previous work on planning strategy development we deployed two HotSpot-based strategies: HotSpots denote those components in a plan that are referred to by multiple flaws, thereby quantifying to which extent solving one deficiency may interfere with the solution options for coupled components. The Direct Uniform HotSpot (*du*) strategy consequently avoids those modifications which address flaws that refer to HotSpot plan components. While the *du* heuristic takes all flaws uniformly into account when calculating their interference potential, the Direct Adaptive HotSpot (*da*) strategy puts problem-specific weights on the binary combinations of flaw types that occur in the plan. The strategy adapts to a repeated occurrence of type combinations by increasing their weights: If abstract task flaws happen to coincide with causal threats, their combined occurrence becomes more important for this plan generation episode. As a generalization of singular HotSpots to commonly affected areas of plan components, the HotZone (*hz*) modification ordering takes into account connections between HotSpots and tries to give modifications that deal with these clusters a low priority.

Plan ordering functions control the traversal through the refinement space that is provided by the modification ordering functions. The strategies in our experimental evaluation were based on the following five components: The least commitment principle on the plan ordering level is represented in two different ways, namely the Fewer Modifications First (*fmf*) strategy, which prefers plans for which a smaller number of refinement options has been announced, and the Less Constrained Plan (*lcp*) strategy, which is based on the ratio of plan steps to the number of variable constraints and causal links in the plan.

The HotSpot concept can be lifted on the plan ordering level: The Fewer HotZone (*fhz*) strategy prefers plans with fewer HotZone clusters. The rationale for this search principle is to focus on plans in which the deficiencies are more closely related and that are hence candidates for an early decision concerning the compatibility of the refinement options. The fourth strategy operates on the HotSpot principle implemented on plan modifications: the Fewer Modification-

based HotSpots (*f<sub>mh</sub>*) function summarizes for all refinement-operators that are proposed for a plan the HotSpot values of the corresponding flaws. It then prefers those plans for which the ratio of plan modifications to accumulated HotSpot values is less. By doing so, this search schema focuses on plans that are expected to have less interfering refinement options.

Finally, since our framework’s representation of the SHOP strategy solely relies on modification ordering, a depth first plan selection is used for constructing a simple hierarchical ordered planner (that is, the plan ordering function is the identity function).

It is furthermore important to mention that our strategy functions can be combined into selection cascades in which succeeding components decide on those cases for which the result of the preceding ones is a tie: With  $s_1 \circ s_2$  we denote, that the strategy  $s_1$  is applied first and afterwards strategy  $s_2$  for tie-breaking.

We have built five combinations from the components above, which can be regarded as representatives for completely different approaches to plan development. Please note that the resulting strategies are general domain-independent planning strategies, which means that they are not tailored to the application of domain model reduction by pre-processing in any way.

We ran our experiments on two distinguished planning domains. The *Satellite* domain is an established benchmark in the field of non-hierarchical planning. It is inspired by the problem of managing scientific stellar observations by earth-orbiting instrument platforms. Our encoding of this domain regards the original primitive operators as implementations of abstract observation tasks, which results in a domain model with 3 abstract and 5 primitive tasks, related by 8 methods. The second domain is known as *UM-Translog*, a transportation and logistics model originally written for HTN planning systems. We adopted its type and decomposition structure to our representation which yielded a deep expansion hierarchy in 51 methods for decomposing 21 abstract tasks into 48 different primitive ones.

We have chosen the above domain models because of the problem characteristics they induce: *Satellite* problems typically become difficult when modeling a repetition of observations, which means that a small number of methods is used multiple times in different contexts of a plan. The evaluated scenarios are thus defined as observations on one or two satellites. *UM-Translog* problems, on the other hand, typically differ in terms of the decomposition structure, because specific transportation goods are treated differently, e.g., toxic liquids in trains require completely different methods than transporting regular packages in trucks. We consequently conducted our experiments on qualitatively different problems by specifying various transportation means and goods.

Table 4 and 5 show the runtime behavior of our system in terms of the size of the average search space and CPU time consumption for the problems in the *UM-Translog* and *Satellite* domains, respectively. The size of the search space is measured in the number of plans visited for obtaining the first solution. Reviewing the overall result, it is quite obvious that the landmark pre-processing pays off in all strategy configurations and problems. It does so in terms of search space size as well as in terms of runtime. The only exception to this is the problem concerning air freight, on which using the

pruned domain model has a measurable negative effect (decrease of performance of 18%). In two configurations in the easiest *Satellite* problem the search space cannot be reduced but a negligible overhead is introduced by pre-processing.

The average performance improvement over all strategies and over all problems in the *UM-Translog* domain is about 40% as is documented in Table 4. The biggest gain is achieved in the transportation tasks that involve special goods and transportation means, e.g., the transport of auto-mobiles, frozen goods, and mail via train saves between 53% and 71%. In general, the flexible strategies profit from the landmark technique, which gives further evidence to the previously obtained results that opportunistic planning strategies are very powerful general-purpose procedures and in addition offer potential to be improved by pre-processing methods. The SHOP-style strategy cannot take that much advantage of the reduced domain model, because it does not adapt its focus on the reduced method alternatives. It continues to address the abstract tasks in the order of their intended execution, regardless of the opportunities that the changes in the method structure may offer. We believe, however, that there may be other possibilities for a SHOP strategy to take into account the reduced domain models.

The *Satellite* domain does not benefit significantly from the landmark technique due to its shallow decomposition hierarchy. We are, however, able to solve problems for which the participating strategies do not find solutions within the given resource bounds otherwise.

## 4.2 Landmark-Aware Strategies

Our landmark-aware strategies are based on the idea that the refinement options, which are basically stored in the remaining task set column of the landmark table, estimate an upper bound for the number of expansion refinements that an abstract task requires before a solution is found. In the previous example (see Table 2), the implementation options for the abstract task *pickup* can be completely explored via the *Pickup\_hazardous* and *Pickup\_normal* methods. This heuristic is only a rough estimation for the “expansion effort” because the table may contain tasks that turn out to be un-achievable and it does not take into account the refinement effort it takes to make an implementation operationable on the primitive level. For our first strategies, we assume that all methods deviate more or less to the same amount in terms of both factors. We will see that this simplification already yields a heuristic with good performance.

Our first modification ordering function  $lm_1$  is defined as follows:

**Definition 7** (Landmark-Aware Ordering  $lm_1$ ). *Let  $f_i$  and  $f_j$  be two abstract task flaws in a plan  $P$  and let  $t_i$  and  $t_j$  be ground instances of the abstract tasks that are compatible with the (in-) equations in the variable constraints of  $P$  and that are referenced by  $f_i$  and  $f_j$ , respectively. Furthermore, let the landmark table contain corresponding entries  $(t_i, I(t_i), R(t_i))$  and  $(t_j, I(t_j), R(t_j))$ .*

*The modification ordering function  $lm_1$  then orders a plan modification  $m_i$  before  $m_j$  if and only if  $m_i$  addresses  $f_i$ ,  $m_j$  addresses  $f_j$ , and  $|R(t_i)| < |R(t_j)|$  holds.*

This strategy implements a rationale that is similar to the



**Table 4.** Results for the *UM-Translog* domain. The column *pruned* refers to the reduced domain models, whereas *unpruned* refers to the original ones (cf. Table 3). The tests were run with the planning environment PANDA [21] on a machine with a 3 GHz CPU and 256 MB Heap memory for the Java VM. *Space* refers to the number of created plans and *Time* refers to the used time in seconds including pre-processing. Values are the averages of three runs. Dashes indicate that no solution was found within a limitation of 5,000 created nodes and a time limit of 150 minutes.

Problem	Modification ordering function $f^{\text{ModOrd}}$	Plan ordering function $f^{\text{PlanOrd}}$	unpruned		pruned	
			Space	Time	Space	Time
Hopper Truck	lcf ◦ hz	fmh ◦ fmf	72	147	41	95
	lcf ◦ ems	fmh ◦ fmf	101	211	72	174
	lcf ◦ du	fhz ◦ fmf	75	155	46	99
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	71	143	54	115
	SHOP Strategy		160	323	89	212
Flatebed Truck	lcf ◦ hz	fmh ◦ fmf	81	182	58	40
	lcf ◦ ems	fmh ◦ fmf	120	269	90	216
	lcf ◦ du	fhz ◦ fmf	96	216	54	129
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	130	299	69	162
	SHOP Strategy		243	595	98	257
Auto Truck	lcf ◦ hz	fmh ◦ fmf	119	301	85	236
	lcf ◦ ems	fmh ◦ fmf	191	443	114	298
	lcf ◦ du	fhz ◦ fmf	129	314	92	251
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	183	469	157	413
	SHOP Strategy		226	558	164	433
Regular Truck 3_Location	lcf ◦ hz	fmh ◦ fmf	149	377	73	203
	lcf ◦ ems	fmh ◦ fmf	234	613	105	206
	lcf ◦ du	fhz ◦ fmf	241	483	131	370
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	190	458	115	307
	SHOP Strategy		163	479	146	406
Regular Truck 2_Region	lcf ◦ hz	fmh ◦ fmf	70	142	42	98
	lcf ◦ ems	fmh ◦ fmf	106	216	81	182
	lcf ◦ du	fhz ◦ fmf	83	160	46	105
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	75	152	54	122
	SHOP Strategy		146	283	106	241
Regular Truck.1	lcf ◦ hz	fmh ◦ fmf	72	149	41	92
	lcf ◦ ems	fmh ◦ fmf	109	225	78	179
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	74	153	54	120
	lcf ◦ du	fhz ◦ fmf	84	173	46	104
	SHOP Strategy		409	911	80	177
Regular Truck.2	lcf ◦ hz	fmh ◦ fmf	–	–	275	1237
	lcf ◦ ems	fmh ◦ fmf	–	–	293	1144
	lcf ◦ du	fhz ◦ fmf	753	2755	295	1262
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	–	–	787	3544
	SHOP Strategy		–	–	926	4005
Mail Traincar	lcf ◦ hz	fmh ◦ fmf	380	1241	89	221
	lcf ◦ ems	fmh ◦ fmf	590	1805	138	313
	lcf ◦ du	fhz ◦ fmf	559	1450	64	160
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	93	213	70	171
	SHOP Strategy		832	1911	121	274
Refrigerated Regular Traincar	lcf ◦ hz	fmh ◦ fmf	384	1240	89	215
	lcf ◦ ems	fmh ◦ fmf	634	1861	138	315
	lcf ◦ du	fhz ◦ fmf	446	1074	64	159
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	92	198	70	172
	SHOP Strategy		777	1735	173	353
Auto Traincar bis	lcf ◦ hz	fmh ◦ fmf	342	1137	144	421
	lcf ◦ ems	fmh ◦ fmf	460	1425	177	477
	lcf ◦ du	fhz ◦ fmf	365	1044	107	328
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	357	958	278	770
	SHOP Strategy		541	1282	247	963
Airplane	lcf ◦ hz	fmh ◦ fmf	164	507	141	435
	lcf ◦ ems	fmh ◦ fmf	142	413	167	471
	lcf ◦ du	fhz ◦ fmf	257	749	200	621
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	280	777	240	700
	SHOP Strategy		335	821	150	450

**Table 5.** Results for the *Satellite* domain. The column *pruned* refers to the reduced domain models, whereas *unpruned* refers to the original ones. The tests were run with the planning environment PANDA [21] on a machine with a 3 GHz CPU and 256 MB Heap memory for the Java VM. *Space* refers to the number of created plans and *Time* refers to the used time in seconds including pre-processing. Values are the averages of three runs. Dashes indicate that no solution was found within a limitation of 5,000 created nodes and a time limit of 150 minutes.

Problem	Modification ordering function $f^{\text{ModOrd}}$	Plan ordering function $f^{\text{PlanOrd}}$	unpruned		pruned	
			Space	Time	Space	Time
1obs- 1sat 1 mode	lcf ◦ hz	fmh ◦ fmf	38	41	37	42
	lcf ◦ ems	fmh ◦ fmf	46	51	46	53
	lcf ◦ du	fhz ◦ fmf	67	72	67	72
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	58	62	53	60
	SHOP Strategy		61	67	57	61
2obs- 1sat 1 mode	lcf ◦ hz	fmh ◦ fmf	602	788	539	708
	lcf ◦ ems	fmh ◦ fmf	964	1631	903	1428
	lcf ◦ du	fhz ◦ fmf	1135	1319	901	1030
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	1468	1699	1216	1474
	SHOP Strategy		251	270	237	264
2obs- 2sat 1 mode	lcf ◦ hz	fmh ◦ fmf	–	–	–	–
	lcf ◦ ems	fmh ◦ fmf	–	–	–	–
	lcf ◦ du	fhz ◦ fmf	–	–	2821	3353
	hz ◦ lcf	fhz ◦ lcp ◦ fmf	–	–	–	–
	SHOP Strategy		–	–	1406	1780

least commitment principle of the *lcf*, because it favors those plan refinements that impose less successor plans, that means, it reduces the effective branching factor of the search space. We note, that the proper choice of the grounded task instances  $t_i$  and  $t_j$  in the above definition is crucial for the actual performance, because the plan modifications typically operate on the lifted abstract tasks and method definitions. For our first experiments, we implemented a random choice on the compatible grounded landmark table entries, future work will however focus on a better informed candidate selection.

While the above heuristic focuses on the very next level of refinement, the following definition also takes into account estimates for subsequent refinement levels thus minimizing the number of refinement choices until no more decompositions are necessary.

**Definition 8** (Indirect Landmark-Aware Ordering  $lm_2$ ). *Let  $f_i$  and  $f_j$  be two abstract task flaws in a plan  $P$  and let  $t_i$  and  $t_j$  be ground instances of the abstract tasks that are compatible with the (in-) equations in the variable constraints of  $P$  and that are referenced by  $f_i$  and  $f_j$ , respectively.*

*Furthermore, let  $R^*(t)$  be the transitive closure of a recursive traversal of the landmark table that begins in  $t$ . More formally:  $R^*(t) = \{r \mid r \in R(t) \text{ for } (t, I(t), R(t)) \in \text{LT} \text{ or } r \in R(t') \text{ for } (t', I(t'), R(t')) \in \text{LT}, t' \in r', \text{ and } r' \in R^*(t)\}$ .*

*The modification ordering function  $lm_2$  then orders a plan modification  $m_i$  before  $m_j$  if and only if  $m_i$  addresses  $f_i$ ,  $m_j$  addresses  $f_j$ , and  $|R^*(t_i)| < |R^*(t_j)|$  holds.*

We would like to point out that  $R^*$  is always finite due to the finiteness of the landmark table, even for cyclic method definitions.

The results of our first experimental evaluation in the UM-Translog domain are given in Table 6: The two landmark-aware heuristic functions  $lm_1$  and  $lm_2$  do outperform the other strategies on practically all problems in terms of both, size of the explored search space and computational time. We believe

that it is because of the relatively unreliable random choice of grounded candidates for the lifted task instances that the supposedly better informed  $lm_2$  does not consistently perform better than  $lm_1$ . We will address this crucial issue in future work by focusing the computational methods for lifting the landmark information: We will investigate into, for example, calculating the average remaining task set sizes of the compatible ground task instances, use the minimal set sizes for consistently underestimating the effort (analog to admissible heuristics), and the like.

## 5 Outlook

We have empirically shown that pruning the planning problem can significantly reduce the explored search space. This pruning relies on a relaxed reachability analysis of fully grounded primitive tasks. So far, a very basic reachability test has been used, which only tests for unsatisfied rigid predicates<sup>7</sup>. In future work, we will use a more elaborated technique as, for instance, explained by Fox and Long [8].

Our empirical evaluation has also shown the success of the two introduced search strategies which use the calculated local landmarks in order to guide the search process. As has been shown earlier [21, 22, 23], many different search strategies can be developed. Future work will introduce additional landmark strategies and discuss the results in more detail.

The techniques discussed in this paper directly apply to hierarchical planning. However, there are various extensions possible that apply to *hybrid* planning. One of the main differences between those two approaches is that in hybrid planning, not only primitive tasks show preconditions and effects, but also abstract tasks. This allows to test even the abstract

<sup>7</sup> Rigid predicates are predicates that are interpreted state-independently and hence their truth value cannot be changed by actions. Their usage is therefore restricted to preconditions and the initial state specification.

**Table 6.** Results for the *UM-Translog* domain. The table shows the impact of the used modification ordering functions on the planning process. In all experiments, except in the SHOP case, the plan ordering function *fmf* was used. In the case of SHOP, the plan and modification ordering functions were used that simulate SHOP’s search process. The tests were run with the planning environment PANDA [21] on a machine with a 3 GHz CPU and 256 MB Heap memory for the Java VM. *Space* refers to the number of created plans and *Time* refers to the used time in seconds including pre-processing. Values are the averages of three runs. Dashes indicate that no solution was found within a limitation of 5,000 created nodes and a time limit of 150 minutes. The best result for a given problem is emphasized bold, the second best italic.

Mod. ordering function $f^{\text{ModOrd}}$	Hopper Truck Space	Hopper Truck Time	Flatbed Truck Space	Flatbed Truck Time	Auto Truck Space	Auto Truck Time	Reg. Truck 3 Location Space	Reg. Truck 3 Location Time
lcf	55	118	<i>62</i>	179	155	470	162	463
hz	55	121	159	399	197	527	191	473
lm <sub>1</sub>	<i>52</i>	<i>111</i>	63	<i>155</i>	<b>133</b>	<b>329</b>	<b>145</b>	<b>374</b>
lm <sub>2</sub>	<b>51</b>	<b>109</b>	<b>61</b>	<b>144</b>	<i>135</i>	<i>462</i>	154	430
ems	147	295	1571	3797	405	976	211	507
da	144	352	99	237	644	2077	239	562
du	101	224	1047	2601	459	1304	1508	4097
SHOP	89	212	98	257	164	433	<i>146</i>	<i>406</i>

Mod. ordering function $f^{\text{ModOrd}}$	Reg. Truck 2.Region Space	Reg. Truck 2.Region Time	Regular Truck.1 Space	Regular Truck.1 Time	Regular Truck.2 Space	Regular Truck.2 Time	Mail Traincar Space	Mail Traincar Time
lcf	78	173	127	222	327	1278	79	209
hz	<i>55</i>	<i>117</i>	<i>55</i>	<i>137</i>	–	–	81	224
lm <sub>1</sub>	62	135	<b>53</b>	<b>122</b>	<i>291</i>	<i>1172</i>	<b>75</b>	<b>184</b>
lm <sub>2</sub>	<b>52</b>	<b>112</b>	65	142	<b>266</b>	<b>1162</b>	<i>78</i>	<i>205</i>
ems	127	262	114	235	–	–	879	1806
da	114	257	148	352	723	2560	641	2031
du	160	460	117	258	–	–	424	1090
SHOP	106	241	83	190	926	4005	121	274

Mod. ordering function $f^{\text{ModOrd}}$	Refrigerated Regular Traincar Space	Refrigerated Regular Traincar Time	Auto Traincar bis Space	Auto Traincar bis Time	Airplane Space	Airplane Time
lcf	90	225	227	926	247	798
hz	<i>76</i>	<i>196</i>	701	1616	345	1323
lm <sub>1</sub>	<b>72</b>	<b>180</b>	<i>183</i>	<i>608</i>	<b>142</b>	<b>441</b>
lm <sub>2</sub>	89	212	<b>158</b>	<b>543</b>	189	676
ems	500	1048	2558	6447	784	2517
da	588	1958	184	705	172	620
du	307	775	1390	4018	643	2134
SHOP	173	353	247	963	<i>150</i>	<i>450</i>

tasks for reachability. Another difference is that hybrid planning problems also specify a goal state that has to be accomplished. Using this goal state, one can use techniques from classical planning in order to generate classical (action) landmarks which can then be used in the hybrid setting.

## 6 Conclusion

We have presented an effective landmark technique for hierarchical planning. It analyzes the planning problem by pre-processing the underlying domain and prunes those regions of the search space where a solution cannot be found. Our experiments on a number of representative hierarchical planning domains and problems give reliable evidence for the practical relevance of our approach. The best performance gain could be achieved for problems with a deep hierarchy of tasks. Our technique is domain- and strategy-independent and can help any hierarchical planner to improve its performance. We have also introduced two search strategies which use the local landmarks in order to guide the search process more efficiently towards a solution. In our empirical evaluation, both strategies outperform the other strategies we chose for comparison in most cases.

## ACKNOWLEDGEMENTS

This work is done within the Transregional Collaborative Research Centre SFB/TRR 62 “Companion-Technology for Cognitive Technical Systems” funded by the German Research Foundation (DFG).

## REFERENCES

- [1] Scott Andrews, Brian Kettler, Kutluhan Erol, and James A. Hendler, ‘UM Translog: A planning domain for the development and benchmarking of planning systems’, Technical Report CS-TR-3487, University of Maryland, (1995).
- [2] Susanne Biundo and Bernd Schatttenberg, ‘From abstract crisis to concrete relief (a preliminary report on combining state abstraction and HTN planning)’, in *Proc. of the 6th European Conference on Planning (ECP 2001)*, pp. 157–168. Springer Verlag, (2001).
- [3] Blai Bonet and Héctor Geffner, ‘Planning as heuristic search’, *Artificial Intelligence*, **129**, 5–33, (2001).
- [4] Luis Castillo, Juan Fdez-Olivares, and Antonio González, ‘On the adequacy of hierarchical planning characteristics for real-world problem solving’, in *Proc. of the 6th European Conference on Planning (ECP 2001)*, (2001).
- [5] Mohamed Elkawkagy, Bernd Schatttenberg, and Susanne Biundo, ‘Landmarks in hierarchical planning’, in *Proc. of the 20th European Conference on Artificial Intelligence (ECAI 2010)*, (2010).
- [6] Kutluhan Erol, James Hendler, and Dana S. Nau, ‘UMCP: A sound and complete procedure for hierarchical task-network planning’, in *Proc. of the 2nd International Conference on Artificial Intelligence Planning Systems (AIPS 1994)*, pp. 249–254, (1994).
- [7] Tara A. Estlin, Steve A. Chien, and Xuemei Wang, ‘An argument for a hybrid HTN/operator-based approach to planning’, in *Proc. of the 4th European Conference on Planning: Recent Advances in AI Planning*, pp. 182–194, (1997).
- [8] Maria Fox and Derek Long, ‘The automatic inference of state invariants in TIM’, *Journal of Artificial Intelligence Research (JAIR)*, **9**, 367–421, (1998).
- [9] Peter Gregory, Stephen Cresswell, Derek Long, and Julie Porteous, ‘On the extraction of disjunctive landmarks from planning problems via symmetry reduction’, in *Proc. of the 4th International Workshop on Symmetry and Constraint Satisfaction Problems*, eds., W. Harvey and Z. Kiziltan, pp. 34–41, (2004).
- [10] Ronny Hartanto and Joachim Hertzberg, ‘On the benefit of fusing DL-reasoning with HTN-planning’, in *Advances in Artificial Intelligence, Proc. of the 32nd German Conference on Artificial Intelligence (KI 2009)*, eds., Bärbel Mertsching, Marcus Hund, and Zaheer Aziz, pp. 41–48, (2009).
- [11] Patrik Haslum, Blai Bonet, and Héctor Geffner, ‘New admissible heuristics for domain-independent planning’, in *Proc. of the Twentieth National Conference on Artificial Intelligence*, volume 3, pp. 1163–1168, (2005).
- [12] Malte Helmert and Carmen Domshlak, ‘Landmarks, critical paths and abstractions: What’s the difference anyway?’, in *Proc. of the 19th International Conference on Automated Planning and Scheduling (ICAPS 2009)*, pp. 162–169, (2009).
- [13] Jörg Hoffmann and Berhard Nebel, ‘The FF planning system: Fast plan generation through heuristic search’, *Journal of Artificial Intelligence Research*, **14**, 253–302, (2001).
- [14] Jörg Hoffmann, Julie Porteous, and Laura Sebastia, ‘Ordered landmarks in planning’, *Journal of Artificial Intelligence Research*, **22**, 215–278, (2004).
- [15] Subbarao Kambhampati, Amol Mali, and Biplav Srivastava, ‘Hybrid planning for partially hierarchical domains’, in *Proc. of the 15th National Conference on Artificial Intelligence*, pp. 882–888. American Association for Artificial Intelligence (AAAI Press), (1998).
- [16] Erez Karpas and Carmel Domshlak, ‘Cost-optimal planning with landmarks’, in *Proc. of the 21st International Joint Conference on Artificial Intelligence (IJCAI 2009)*, pp. 1728–1733, (2009).
- [17] Dana S. Nau, Yue Cao, Ammon Lotem, and Héctor Muñoz-Avila, ‘SHOP: Simple hierarchical ordered planner’, in *Proc. of the 16th International Joint Conference on Artificial Intelligence (IJCAI 1999)*, pp. 968–975, (1999).
- [18] Julie Porteous and Stephen Cresswell, ‘Extending landmarks analysis to reason about resources and repetition’, in *In Proc. of the 21st Workshop of the UK Planning and Scheduling Special Interest Group (PLANSIG 2002)*, pp. 45–54, (2002).
- [19] Julie Porteous, Laura Sebastia, and Jörg Hoffmann, ‘On the extraction, ordering, and usage of landmarks in planning’, in *Proc. of the 6th European Conference on Planning (ECP 2001)*, eds., A. Cesta and D. Borrajo, pp. 37–48, (2001).
- [20] Silvia Richter, Malte Helmert, and Matthias Westphal, ‘Landmarks revisited’, in *Proc. of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI 2008)*. AAAI Press, (2008).
- [21] Bernd Schatttenberg, *Hybrid Planning & Scheduling*, Ph.D. dissertation, Ulm University, Germany, 2009.
- [22] Bernd Schatttenberg, Julien Bidot, and Susanne Biundo, ‘On the construction and evaluation of flexible plan-refinement strategies’, in *Advances in Artificial Intelligence, Proc. of the 30th German Conference on Artificial Intelligence (KI 2007)*, eds., Joachim Hertzberg, Michael Beetz, and Roman Englert, volume 4667 of *Lecture Notes in Artificial Intelligence*, pp. 367–381. Springer, (2007).
- [23] Bernd Schatttenberg, Andreas Weigl, and Susanne Biundo, ‘Hybrid planning using flexible strategies’, in *Advances in Artificial Intelligence, Proc. of the 28th German Conference on Artificial Intelligence (KI 2005)*, pp. 249–263. Springer-Verlag Berlin Heidelberg, (2005).
- [24] Laura Sebastia, Eva Onaindia, and Eliseo Marzal, ‘Decomposition of planning problems’, *AI Communications*, **19**(1), 49–81, (2006).
- [25] Vincent Vidal and Héctor Geffner, ‘Branching and pruning: An optimal temporal POCL planner based on constraint programming’, *Artificial Intelligence*, **170**, 298–335, (2006).
- [26] Lin Zhu and Robert Givan, ‘Landmark extraction via planning graph propagation’, in *Proc. of the ICAPS 2003 Doctoral Consortium*, pp. 156–160, (2003).