On Reasonable and Forced Goal Orderings and their Use in an Incremental Planning Algorithm

Jana Koehler
Schindler Lifts, S. A.
R & D Technology Management
6031 Ebikon, Switzerland
jana_koehler@ch.schindler.com

Jörg Hoffmann
Institute for Computer Science
Albert Ludwigs University
Am Flughafen 17
79110 Freiburg, Germany
hoffmann@informatik.uni-freiburg.de

Abstract

The paper addresses the problem of computing goal orderings, which is one of the longstanding issues in AI planning. It makes two new contributions. First, it formally defines and discusses two different goal orderings, which are called the *reasonable* and the *forced* ordering. Both orderings are defined for simple STRIPS operators as well as for more complex ADL operators supporting negation and conditional effects. The complexity of these orderings is investigated and their practical relevance is discussed. Secondly, two different methods to compute reasonable goal orderings are developed. One of them is based on planning graphs, while the other investigates the set of actions directly. Finally, it is shown how the ordering relations, which have been derived for a given set of goals $\mathcal{G}$, can be used to compute a so-called *goal agenda* that divides $\mathcal{G}$ into an ordered set of subgoals. Any planner can then, in principle, use the goal agenda to incrementally plan for increasing sets of subgoals. This can lead to an exponential complexity reduction, as the solution to a complex planning problem is found by solving easier subproblems. Since only a polynomial overhead is caused by the goal agenda computation, a potential exists to dramatically speed up planning algorithms as we demonstrate in the empirical evaluation, where we use this method in the IPP planner.

Technical Report No. 132

January 2000

Contents

1 Introduction **1**

2 Ordering Relations between Atomic Goals **1**

3 The Computation of Goal Orderings **6**
3.1 Reasonable Goal Orderings based on GRAPHPLAN 6
3.2 Reasonable Goal Orderings derived by a Fast Heuristic Method 9
3.3 On the Non-Existence of Forced Goal Orderings 16
3.4 An Extension of Goal Orderings to ADL Actions 20

4 The Use of Goal Orderings During Planning **22**
4.1 The Goal Agenda 23
4.2 Extension of Goal Orderings to Goal Sets 25
4.3 An Incremental Planning Algorithm 26

5 Empirical Results **28**
5.1 Comparison of $\leq_h$ and $\leq_e$ 29
5.2 Influence of Goal Orderings on the Performance of IPP and Interaction with RIFO 31
5.2.1 The Blocksworld 31
5.2.2 The Tyreworld 34
5.2.3 The Tower of Hanoi 35

6 Summary and Comparison to Related Work **37**

7 Outlook **40**

1 Introduction

How to effectively plan for interdependent subgoals has been in the focus of AI planning research for a very long time. Starting with the early work in ABSTRIPS [Sac74] or in [Cha87], quite a number of approaches have been presented and the complexity of the problems has been studied. But until today, planners have made only some progress in solving bigger planning instances and scalability of classical planning systems is still a problem.

In this paper, we focus on the following problem: Given a set of conjunctive goals, can we define and detect an ordering relation over subsets from the original goal set? To arrive at such an ordering relation over subsets, we first focus on the atomic facts contained in the goal set. We formally define two closely related ordering relations over such atomic goals, which we call *reasonable* and *forced* ordering, and study their complexity. It turns out that both are very hard to decide.

Consequently, we introduce two efficient methods that can both be used to approximate reasonable goal orderings. The definitions are first given for simple STRIPS domains, where the desired theoretical properties can be easily proven. Afterwards, we extend our definitions to ADL operators handling conditional effects and negative preconditions, and discuss why we do not invest any effort in trying to find forced orderings.

We show how a set of ordering relations between atomic goals can be used to divide the goal set into disjunct subsets, and how these subsets can be ordered with respect to each other. The resulting sequence of subsets comprises the so-called *goal agenda*, which can then be used to control an incremental planning algorithm.

The method, called **G**oal **A**genda **M**anager, is implemented in the context of the IPP planning system, where we show its potential of exponentially reducing computation times on certain planning domains.

The paper is organized as follows: Section 2 introduces and motivates reasonable and forced goal orderings. Starting with simple STRIPS operators, they are formally defined, and their complexity is investigated.

In Section 3, we present two methods, which compute an approximation of the reasonable ordering and discuss both orderings from a more practical point of view. The section concludes with an extension of our definitions to ADL operators having conditional effects.

Section 4 shows how a planning system can benefit from ordering information by computing a goal agenda that guides the planner. We define how subsets of goals can be ordered with respect to each other and discuss how a goal agenda can affect the theoretical properties, in particular the completeness of a planning algorithm.

Section 5 contains the empirical evaluation of our work, showing results that we obtained using the goal agenda in IPP, and in Section 6 we summarize our approach in the light of related work. The paper concludes with an outlook on possible future research directions in Section 7.

2 Ordering Relations between Atomic Goals

For a start, we only investigate simple STRIPS domains just allowing sets of atoms to describe states, the preconditions, and the add and delete lists of operators.

Definition 1 (State) *The set of all ground atoms is denoted with P. A* state $s \in 2^P$ *is a subset of ground atoms.*

Note that all states are assumed to be *complete*, i.e., we always know for an atom p whether $p \in s$ or $p \notin s$ holds. We also assume that all operator schemata are grounded, i.e., we only talk about *actions*.

Definition 2 (Strips Action) *A* STRIPS action o *has the usual form*

$$pre(o) \;\Rightarrow\; ADD\;\; add(o)\;\; DEL\;\; del(o)$$

where $pre(o)$ are the preconditions of o, $add(o)$ is the Add list of o and $del(o)$ is the Delete list of the action, each being a set of ground atoms. We also assume that $del(o) \cap add(o) = \emptyset$. The result of applying a STRIPS action to a state is defined as usual:

$$Result(s, o) := \begin{cases} (s \cup add(o)) \setminus del(o) & pre(o) \subseteq s \\ s & otherwise \end{cases}$$

In the first case, where $pre(o) \subseteq s$, the action is said to be applicable *in s. The result of applying a sequence of more than one action to a state is recursively defined as*

$$Result(s, \langle o_1, \ldots, o_n \rangle) := Result(Result(s, \langle o_1, \ldots, o_{n-1} \rangle), o_n).$$

Definition 3 (Planning Problem) *A* planning problem $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ *is a triple where $\mathcal{O}$ is the set of actions, and $\mathcal{I}$ (the initial state) and $\mathcal{G}$ (the goals) are sets of ground atoms. A* plan $\mathcal{P}$ *is an ordered sequence of actions. If all actions in a plan are taken out of a certain action set $\mathcal{O}$, we denote this by writing $\mathcal{P}^{\mathcal{O}}$.*

Note that we define a plan to be a sequence of *actions*, not a sequence of *parallel steps*, as it is done in [BF97], for example. This makes the subsequent theoretical investigation more readable. The results directly carry over to parallel plans.

Given two atomic goals A and B, various ways to define an ordering relation over them can be imagined. First, one can distinguish between *domain-specific* and *domain-independent* goal ordering relations. But although domain-specific orderings can be very effective, they need to be redeveloped for each single domain. Therefore, one is in particular interested in domain-independent ordering relations having a broader range of applicability. Secondly, following [HMAW99], one can distinguish the *goal selection* and the *goal achievement* order. The first ordering determines in which order a planner works on the various atomic goals, while the second one determines the order, in which the solution plan achieves the goals. In this paper, we compute an ordering of the latter type. In the incremental planning approach that we propose later in the paper, both orderings coincide, though. The goals that are achieved first in the plan are those that the planner works on first.

The following scenario motivates how an achievement order for goals can be possibly defined. Given two atomic goals A and B, for which a solution plan exists, let us assume the planner has just achieved the goal A, i.e., it has arrived at a state $s_{(A,\neg B)}$, in which A holds, but B does not yet hold. Now, if there exists a plan that is executable in $s_{(A,\neg B)}$ and achieves B without ever deleting A, a solution has been found. If no such plan can be found, then two possible reasons exist:

1. The problem is unsolvable—achieving A first leads the planner into a deadlock situation. Thus, the planner is *forced* to achieve B before or simultaneously with A.

2. The only existing solution plans have to destroy A temporarily in order to achieve B. But then, A should not be achieved first. Instead, it seems to be *reasonable* to achieve B before or simultaneously with A for the sake of shorter solution plans.

In the first situation, the ordering "B before or simultaneously with A" is *forced* by inherent properties of the planning domain. In the second situation, the ordering "B before or simultaneously with A" appears to be *reasonable* in order to avoid non-optimal plans. Consequently, we will define two goal orderings, called the *forced* and the *reasonable* ordering. For the sake of clarity, we first give some more basic definitions.

Definition 4 (Reachable State) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem and let P be the set of ground atoms that occur in the problem. We say that a state $s \subseteq P$ is* reachable*, iff there exists a sequence $\langle o_1, \ldots, o_n \rangle$ out of actions in $\mathcal{O}$ for which $s = Result(\mathcal{I}, \langle o_1, \ldots, o_n \rangle)$ holds.*

Definition 5 (Generic State $s_{(A,\neg B)}$) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem. By $s_{(A,\neg B)}$, we denote all reachable states in which A has just been achieved, but B is still* FALSE*, i.e., $B \notin s_{(A,\neg B)}$ and $s_{(A,\neg B)} = Result(\mathcal{I}, \langle o_1, \ldots, o_n \rangle)$ with $A \in add(o_n)$.*

Definition 6 (Reduced Action Set $\mathcal{O}_A$) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A \in \mathcal{G}$ be an atomic goal. By $\mathcal{O}_A$ we denote the set of all actions that do not delete A, i.e., $\mathcal{O}_A = \{o \in \mathcal{O} \mid A \notin del(o)\}$.*

We are now prepared to define what we exactly mean by *forced* and *reasonable* goal orderings.

Definition 7 (Forced Ordering $\leq_f$) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A, B \in \mathcal{G}$ be two atomic goals. We say that there is a* forced ordering *between B and A, written $B \leq_f A$, if and only if*

$$\forall\, s_{(A,\neg B)} : \quad \neg\exists\, \mathcal{P}^{\mathcal{O}} : \; B \in Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}})$$

If Definition 7 is satisfied, then each plan achieving A and B must achieve B before or simultaneously with A, because otherwise it will encounter a deadlock, rendering the problem unsolvable.

Definition 8 (Reasonable Ordering $\leq_r$) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A, B \in \mathcal{G}$ be two atomic goals. We say that there is a* reasonable ordering *between B and A, written $B \leq_r A$, if and only if*

$$\forall\, s_{(A,\neg B)} : \quad \neg\exists\, \mathcal{P}^{\mathcal{O}_A} : \; B \in Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}_A})$$

Definition 8 gives $B \leq_r A$ the meaning that if, after the goal A has been achieved, there is no plan anymore that achieves B without—at least temporarily—destroying A, then B is a goal prior to A.

We remark that obviously $B \leq_f A$ implies $B \leq_r A$, but not vice versa. We also make a slightly less obvious observation at this point: The formulae in Definitions 7 and 8 use a universal quantification over states $s_{(A,\neg B)}$. If in a planning problem there is no such state at all, the formulae are satisfied and the goals A and B get ordered, i.e., $B \leq_f A$ and $B \leq_r A$ follow, resp. In this case, however, there is not much information gained by a goal ordering between A and B, because *any* sequence of actions will achieve B prior or simultaneously with A—A cannot be achieved with B still being FALSE. Thus in this case, the ordering relations $B \leq_f A$ and $B \leq_r A$ are *trivial* in the sense that no planner would invest much effort in considering the goals A and B ordered the other way round anyway.

Definition 9 (Trivial Ordering Relation) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A, B \in \mathcal{G}$ be two atomic goals. An ordering relation $B \leq_f A$ or $B \leq_r A$ is called* trivial *iff there is no state $s_{(A,\neg B)}$.*

In this paper, we will usually consider forced and reasonable goal orderings as non-trivial orderings and make the distinction explicit only if we have to do so.

Both Definition 7 and 8 seem to deliver promising candidates for an achievement order. Unfortunately, both are very hard to test: it turns out that their corresponding decision problems are PSPACE hard.

Theorem 1 *Let* F_ORDER *denote the following problem:*

Given two atomic facts A and B, as well as an action set $\mathcal{O}$ and an initial state $\mathcal{I}$, does $B \leq_f A$ hold ?

Deciding F_ORDER *is* PSPACE*-hard.*

Proof: The proof proceeds by polynomially reducing PLANSAT [Byl94]—the decision problem of whether there exists a solution plan for a given arbitrary STRIPS planning instance—to the problem of deciding F_ORDER.

Let $\mathcal{I}$, $\mathcal{G}$, and $\mathcal{O}$ denote the initial state, the goal state, and the action set in an arbitrary STRIPS instance. Let A, B, and C be new atomic facts not contained in the instance so far. We build a new action set and initial state for our F_ORDER instance by setting

$$\mathcal{O}' := \mathcal{O} \cup \left\{ \begin{array}{lll} o_{I_1} & = \{C\} & \Rightarrow \text{ ADD } \{A\} \text{ DEL } \{C\}, \\ o_{I_2} & = \{A\} & \Rightarrow \text{ ADD } \mathcal{I} \text{ DEL } \{A\}, \\ o_G & = \mathcal{G} & \Rightarrow \text{ ADD } \{B\} \text{ DEL } \emptyset \end{array} \right\}$$

and

$$\mathcal{I}' := \{C\}$$

With these definitions, the only way of achieving A is by applying $o_{\mathcal{I}_1}$ to $\mathcal{I}'$. According to Definition 5, the only state $s_{(A,\neg B)}$ is $\{A\}$. Thus, the argumentation goes as follows:

$$\begin{array}{lll} B \leq_f A & \Leftrightarrow \forall\, s_{(A,\neg B)} :\ \neg\exists\, \mathcal{P}^{\mathcal{O}'} :\ B \in Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}'}) & \text{Definition 7} \\ & \Leftrightarrow \neg\exists\, \mathcal{P}^{\mathcal{O}'} :\ B \in Result(\{A\}, \mathcal{P}^{\mathcal{O}'}) & \{A\} \text{ is the only reachable state } s_{(A,\neg B)} \\ & \Leftrightarrow \neg\exists\, \mathcal{P}^{\mathcal{O}} :\ \mathcal{G} \subseteq Result(\mathcal{I}, \mathcal{P}^{\mathcal{O}}) & \text{from the definition of } \mathcal{O}' \\ & \Leftrightarrow \text{no solution plan exists for } \mathcal{I}, \mathcal{G} \text{ and given } \mathcal{O} & \end{array}$$

Thus, the complement of PLANSAT can be polynomially reduced to F_ORDER. As PSPACE = co-PSPACE, we are done. ∎

Theorem 2 *Let* R_ORDER *denote the following problem:*

Given two atomic facts A and B, as well as an action set $\mathcal{O}$ and an initial state $\mathcal{I}$, does $B \leq_r A$ hold ?

Deciding R_ORDER *is* PSPACE*-hard.*

Proof: The proof proceeds by polynomially reducing PLANSAT to R_ORDER.

Let $\mathcal{I}$, $\mathcal{G}$, and $\mathcal{O}$ be the initial state, the goal state, and the action set in an arbitrary STRIPS planning instance. Let A, B, C, and D be new atomic facts not contained in the instance so far. We define the new action set $\mathcal{O}'$ by setting

$$\mathcal{O}' := \mathcal{O} \cup \left\{ \begin{array}{lll} o_{I_1} & = \{C\} & \Rightarrow \text{ ADD } \{A, D\} \text{ DEL } \{C\}, \\ o_{I_2} & = \{A, D\} & \Rightarrow \text{ ADD } \mathcal{I} \text{ DEL } \{D\}, \\ o_G & = \mathcal{G} & \Rightarrow \text{ ADD } \{B\} \text{ DEL } \emptyset \end{array} \right\}$$

and the new initial state by

$$\mathcal{I}' := \{C\}$$

With these definitions, the only way to achieve A is by applying o_{I_1} to $\mathcal{I}'$, so per Definition 5 the only state $s_{(A,\neg B)}$ is $\{A, D\}$. Thus, we can argue

$B \leq_r A$

$\Leftrightarrow \forall\ s_{(A,\neg B)}\ \neg\exists\ \mathcal{P}^{\mathcal{O}'_A} :\ B \in Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}'_A})$ Definition 8

$\Leftrightarrow \neg\exists\ \mathcal{P}^{\mathcal{O}'_A} :\ B \in Result(\{A, D\}, \mathcal{P}^{\mathcal{O}'_A})$ $\{A, D\}$ is the only reachable state $s_{(A,\neg B)}$

$\Leftrightarrow \neg\exists\ \mathcal{P}^{\mathcal{O}'} :\ B \in Result(\{A, D\}, \mathcal{P}^{\mathcal{O}'})$ no action in $\mathcal{O}'$ deletes A

$\Leftrightarrow \neg\exists\ \mathcal{P}^{\mathcal{O}}$ such that $\mathcal{G} \subseteq Result(\mathcal{I}, \mathcal{P}^{\mathcal{O}})$ from the definition of $\mathcal{O}'$

$\Leftrightarrow$ no solution plan exists for $\mathcal{I}, \mathcal{G}, \mathcal{O}$

Thus, the complement of PLANSAT can be polynomially reduced to R_ORDER. With PSPACE = co-PSPACE, we are done. ∎

Consequently, finding necessary and forced ordering relations between atomic goals is already as hard as the original planning problem, and it appears unlikely that a planner will gain any advantage from doing that. A possible way out of the dilemma is to define new ordering relations, which can be decided in polynomial time and which are, ideally, sufficient for the existence of reasonable or forced goal orderings. In the following, we introduce two such orderings.

3 The Computation of Goal Orderings

In this section, we will

1. define a goal ordering $\leq_e$, which can be computed using GRAPHPLAN's exclusivity information about facts. We prove that this ordering is sufficient for $\leq_r$ and that it can be decided in polynomial time (the subscript "e" stands for "efficient").
2. define a goal ordering $\leq_h$, which is computed based on a heuristic method that is much faster than the computation based on GRAPHPLAN, but also delivers powerful goal ordering information (the subscript "h" stands for "heuristic").
3. discuss that almost all of the "more natural" planning domains do not contain forced orderings, i.e., $\leq_f$ will fail in providing a reasonable problem decomposition.
4. show how the orderings we defined can be extended to handle more expressive ADL operators.

3.1 Reasonable Goal Orderings based on GRAPHPLAN

A goal ordering is always computed for a specific planning problem involving an initial state $\mathcal{I}$, a goal set $\mathcal{G} \supseteq \{A, B\}$, and the set $\mathcal{O}$ of all ground actions. In order to develop an efficient computational method, we proceed in two steps now:

1. We compute more knowledge about the generic state $s_{(A,\neg B)}$.
2. We define the relation $\leq_e$ and investigate its theoretical properties. In particular, we prove that $\leq_e$ implies $\leq_r$.

The state $s_{(A,\neg B)}$ represents states that are reachable from $\mathcal{I}$, and in which A has been achieved, but B does not hold. Given this information about $s_{(A,\neg B)}$, one can derive additional knowledge about it. In particular, it is possible to determine a subset of atoms F, of which one definitely knows that $\mathsf{F} \cap s_{(A,\neg B)} = \emptyset$ must hold.

This set can be easily computed using GRAPHPLAN [BF95, BF97]. Starting from $\mathcal{I}$ with $\mathcal{O}$, the planning graph is built until the graph has *leveled* off at some time step. The proposition level at this time step represents a set of states, which is a superset of all states that are reachable from $\mathcal{I}$ when applying actions from $\mathcal{O}$. All atoms, which are marked as *mutually exclusive* [BF97] of A in this level can never hold in a state satisfying A. Thus, they cannot hold in $s_{(A,\neg B)}$. We denote this set with F^A_{GP}—the False set with respect to A returned by GRAPHPLAN.[1]

$$\mathsf{F}^A_{GP} := \{p \mid p \text{ is exclusive of } A \text{ after the graph has leveled off}\} \tag{1}$$

Note that the planning graph is only grown once for a given $\mathcal{I}$ and $\mathcal{O}$, but can be used to determine the F^A_{GP} sets for all atomic goals $A \in \mathcal{G}$.

Lemma 1 *$\mathsf{F}^A_{GP} \cap s_A = \emptyset$ holds for all states s_A satisfying $A \in s_A$ that are reachable from $\mathcal{I}$ using actions from $\mathcal{O}$.*

[1] We assume the reader to be familiar with GRAPHPLAN, because this planning system is very well known in the planning research community. Otherwise, [BF95, BF97] provide the necessary background.

Proof: Follows immediately from the definitions of "level-off" and "two propositions being mutual exclusive" given in [BF97]. ∎

We now provide a simple test which is sufficient for the existence of a reasonable ordering $B \leq_r A$ between two atomic goals A and B.

Definition 10 (Efficient Ordering $\leq_e$) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G} \supseteq \{A, B\})$ be a planning problem. Let F^A_{GP} be the* False *set for A. The ordering $B \leq_e A$ holds if and only if*

$$\forall\, o \in \mathcal{O}_A : B \in add(o) \Rightarrow pre(o) \cap \mathsf{F}^A_{GP} \neq \emptyset$$

This means, B is ordered before A if the reduced action set only contains actions, which either do not have B in their add lists or if they do, then they require a precondition which is contained in the False set. Such preconditions can never hold in a state satisfying A and thus, these actions will never be applicable.

Theorem 3

$$B \leq_e A \Rightarrow B \leq_r A$$

Proof: Assume that $B \not\leq_r A$, i.e., $B \in Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}_A})$ for a reachable state $s_{(A,\neg B)}$ with $A \in s_{(A,\neg B)}$, $B \notin s_{(A,\neg B)}$ and a Plan $\mathcal{P}^{\mathcal{O}_A} = \langle o_1, \ldots, o_n \rangle$ where $o_i \in \mathcal{O}_A$ for $1 \leq i \leq n$.

As $A \notin del(o_i)$ for all i (Definition 6), we have

$$A \in Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_i \rangle) \text{ for } 0 \leq i \leq n$$

and, with Lemma 1,

$$\mathsf{F}^A_{GP} \cap Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_i \rangle) = \emptyset \text{ for } 0 \leq i \leq n \tag{2}$$

Furthermore, as $B \notin s_{(A,\neg B)}$, but $B \in Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_n \rangle)$, there must be a step which makes B true, i.e.,

$$\exists 1 \leq k \leq n : B \notin Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_{k-1} \rangle) \wedge B \in Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_k \rangle)$$

For this step, we obviously have $B \in add(o_k)$ and consequently, with the definition of $B \leq_e A$, $pre(o_k) \cap \mathsf{F}^A_{GP} \neq \emptyset$. Now, as o_k must be applicable in the state where it is executed (otherwise it would not add anything to this state), the preconditions of o_k must hold, i.e., $pre(o_k) \subseteq Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_{k-1} \rangle)$. This immediately leads to $\mathsf{F}^A_{GP} \cap Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_{k-1} \rangle) \neq \emptyset$, which is a contradiction to Equation (2). ∎

Obviously, the ordering $\leq_e$ can be decided in polynomial time.

Theorem 4 *Let* E_ORDER *denote the following problem:*

Given two atomic facts A and B, as well as an action set $\mathcal{O}$, does $B \leq_e A$ hold ?

Then, E_ORDER *can be decided in polynomial time:* E_ORDER $\in$ P.

Proof: The computation of F^A_{GP} takes only polynomial time, which follows directly from the results in [BF97]. Testing $B \leq_e A$ then only involves looking at all actions in $\mathcal{O}$, and rejecting them if they either

- delete A or
- have a precondition which is an element of F^A_{GP}.

So we have an additional running time for the test which is dominated by $|\mathcal{O}| * |P|^2$, where P denotes the set of all ground atoms. ■

Let us consider the following example, which illustrates the computation of $\leq_e$ using a common representational variant of the *blocksworld* with actions to **stack**, **unstack**, **pickup**, and **putdown** blocks:

pickup(?ob)
clear(?ob) on-table(?ob) arm-empty() ⇒ *ADD holding(?ob)*
DEL clear(?ob) on-table(?ob) arm-empty().

putdown(?ob)
holding(?ob) ⇒ *ADD clear(?ob) arm-empty() on-table(?ob)*
DEL holding(?ob).

stack(?ob,?underob)
clear(?underob) holding(?ob) ⇒ *ADD arm-empty() clear(?ob) on(?ob,?underob)*
DEL clear(?underob) holding(?ob).

unstack(?ob,?underob)
on(?ob,?underob) clear(?ob) arm-empty() ⇒ *ADD holding(?ob) clear(?underob)*
DEL on(?ob,?underob) clear(?ob) arm-empty().

Given the simple task of stacking three blocks:

initial state: on-table(a) on-table(b) on-table(c)
goal state: on(a,b) on(b,c)

is there a reasonable ordering between the two atomic goals? Intuitively, the *blocksworld* domain possesses a very natural goal ordering, namely that the planner should start building each tower from the bottom to the top and not the other way round.[2]

Let us first investigate whether the relation $on(a,b) \leq_e on(b,c)$ holds. Vividly speaking, it asks whether it is still possible to stack block a on b after $on(b,c)$ has been achieved. As a first step, we run GRAPHPLAN to find out which atoms are exclusive of $on(b,c)$ when the planning graph, wich corresponds to this problem, has leveled off. The result is

$$\mathsf{F}^{on(b,c)}_{GP} = \{\textit{clear(c), on-table(b), holding(c), holding(b), on(a,c), on(c,b), on(b,a)}\}$$

One observes immediately that these atoms can never be true in a state that satisfies $on(b,c)$.

[2]Note that the goals do not specify where the block c has to go, but leave this to the planner.

Secondly, we remove all ground actions which delete $on(b,c)$ (in this case, only the action **unstack(b,c)** satisfies this condition) and obtain the reduced action set $\mathcal{O}_{on(b,c)}$.

Now we are ready to test if $on(a,b) \leq_e on(b,c)$ holds. The only action, which can add $on(a,b)$ is **stack(a,b)**. It has the preconditions $holding(a)$ and $clear(b)$, neither of which is a member of $\mathsf{F}_{GP}^{on(b,c)}$. The test fails and we get $on(a,b) \not\leq_e on(b,c)$.

Now we test whether $on(b,c) \leq_e on(a,b)$ holds. GRAPHPLAN returns the following False set:

$$\mathsf{F}_{GP}^{on(a,b)} = \{clear(b),\ on\text{-}table(a),\ holding(b),\ holding(a),\ on(a,c),\ on(c,b),\ on(b,a)\}$$

The action **unstack(a,b)** is not contained in $\mathcal{O}_{on(a,b)}$ because it deletes $on(a,b)$. The only action which adds $on(b,c)$ is **stack(b,c)**. It needs the preconditions $clear(c)$ and $holding(b)$. The second precondition $holding(b)$ is contained in the set of false facts, $holding(b) \in \mathsf{F}_{GP}^{on(a,b)}$ and thus, we derive $on(b,c) \leq_e on(a,b)$. Altogether, we have $on(a,b) \not\leq_e on(b,c)$ and $on(b,c) \leq_e on(a,b)$, which correctly reflects the intuition that b needs to be stacked onto c *before* a can be stacked onto b.

Although $\leq_e$ appears to impose very strict conditions on a domain in order to derive a reasonable goal ordering, it succeeds in finding reasonable goal orderings in all available test domains in which such orderings exists. For example, in the *tyreworld*, in *bulldozer* problems, in the *shopping problem* [RN95], the *fridgeworld*, the *glass* domain, the *tower of hanoi* domain, the *link-world*, and the *woodshop*. Its only disadvantage are the computational resources it requires, since building planning graphs, while being theoretically polynomial, is a quite time- and memory-consuming thing to do.[3]

Therefore, the next section presents a fast heuristic computation of goal orderings, which analyzes the domain actions directly and does not need to build planning graphs anymore.

3.2 Reasonable Goal Orderings derived by a Fast Heuristic Method

One can analyze the available actions directly using a method we will call *Direct Analysis* (DA). It determines an initial value for F by computing the intersection of all delete lists of all actions which contain A in their add list, as defined in the following equation.

$$\mathsf{F}_{DA}^{A} := \bigcap_{o \in \mathcal{O},\ A \in add(o)} del(o) \tag{3}$$

The atoms in this set are all FALSE in a state where A has just been achieved: they are deleted from the state description independently of the action that is used to add A. As a short example, let us consider the two actions

$$\Rightarrow \text{ ADD } \{A\} \text{ DEL } \{C, D\}$$

[3] More recent implementations of planning graphs, which are for example developed for STAN [FL99] and IPP 4.0 do not build the graphs explicitly anymore and are magnitudes of order faster than the original GRAPHPLAN implementation, but still the computation of the planning graph takes almost all the time that is needed to determine the $\leq_e$ relations.

$$\Rightarrow \text{ ADD } \{A, C\} \text{ DEL } \{D\}$$

Only the atom D is deleted by both actions, and thus D is the only element initially contained in F^A_{DA}.

However, Equation (3) only says that when A is added then the atoms from F^A_{DA} will be deleted. It does not say anything about whether it might be possible to reestablish atoms in F^A_{DA}. One can easily imagine that actions exist, which leave A true, and at the same time add such atoms. If this is the case, there are reachable states in which A and atoms from F^A_{DA} hold.

Now, our goal is to derive an ordering relation that can be easily computed, and that ideally, like the $\leq_e$ relation, is sufficient for the $\leq_r$ relation. Therefore, we want to make sure that the atoms in F^A_{DA} are really FALSE in any state after A has been achieved. We arrive at an approximation of atoms that remain FALSE by performing a fixpoint reduction on the F^A_{DA} set, removing those atoms that are *achievable* in the following sense.

Definition 11 (Achievable Atoms) *An atom p is* achievable *from a state s given an action set $\mathcal{O}$ (written $\mathcal{A}(s, p, \mathcal{O})$) if and only if*

$$p \in s \ \vee \ \exists\, o \in \mathcal{O} : p \in add(o) \ \wedge \ \forall\, p' \in pre(o) : \ \mathcal{A}(s, p', \mathcal{O})$$

The definition says that an atom p is achievable from a state s if it holds in s, or if there exists an action in the domain, which adds p and whose preconditions are all achievable from s. This is a necessary condition for the existence of a plan $\mathcal{P}^{\mathcal{O}}$ from s to a state where p holds.

Lemma 2 $\exists\, \mathcal{P}^{\mathcal{O}} : \ p \in Result(s, \mathcal{P}^{\mathcal{O}}) \Rightarrow \mathcal{A}(s, p, \mathcal{O})$

Proof: The atom p must either already be contained in the state s, or it has to be added by a step o out of $\mathcal{P}^{\mathcal{O}}$. In the second case, all preconditions of o need to be established by $\mathcal{P}^{\mathcal{O}}$ in the same way. So p and all preconditions of the step, which adds it, are achievable in the sense of Definition 11. ■

There are two obvious difficulties with Definition 11: First, $p \in s$ must be tested. With complete knowledge about the state s, this should not cause any problems. In our case, however, we only have the *generic* state $s_{(A,\neg B)}$ and cannot decide whether an arbitrary atom is contained in it or not. Secondly, we observe an infinite regression over preconditions, which must be tested for achievability.

As for the first problem, it turns out that it is a good heuristic to simply assume $p \notin s$, i.e., no test is performed at all. As for the second problem, in order to avoid infinite looping of the "achievable"-test, one needs to terminate the regression over preconditions at a particular level. The point in question is how far to regress? A quick approximation simply decides "achievable" after the first recursive call.

Definition 12 (Possibly Achievable Atoms) *An atom p is* possibly achievable *given an action set $\mathcal{O}$ (written $p\mathcal{A}(p, \mathcal{O})$) if and only if*

$$\exists\, o \in \mathcal{O} : p \in add(o) \ \wedge \ \forall\, p' \in pre(o) : \ \exists\, o' \in \mathcal{O} : p' \in add(o')$$

holds, i.e., there is an action that adds p and all of its preconditions are add effects of other actions in $\mathcal{O}$.

If the assumption is justified that none of the atoms p is contained in the state s, then being possibly achievable is a necessary condition for being achievable.

Lemma 3 *Let s be a state for which $p \notin s$ and also $\forall o \in \mathcal{O} : p \in add(o) \Rightarrow pre(o) \cap s = \emptyset$ holds. Then we have*

$$\mathcal{A}(s, p, \mathcal{O}) \Rightarrow p\mathcal{A}(p, \mathcal{O})$$

Proof: From $\mathcal{A}(s, p, \mathcal{O})$ and $p \notin s$, we know that there is a step $o \in \mathcal{O}$, $p \in add(o)$, with $\forall\, p' \in pre(o)\ \mathcal{A}(s, p', \mathcal{O})$. We also know that $pre(o) \cap s = \emptyset$, so for each $p' \in pre(o)$ there must be an achiever $o' \in \mathcal{O} : p' \in add(o')$. ∎

The condition that all of the facts p must not be contained in the state s seems to be rather rigid. Nevertheless, the condition of being possibly achievable delivers good results on all of the benchmark domains and is easy to decide. We can now use this test to both

- perform a fixpoint reduction on the set F^A_{DA} and
- decide whether an atomic goal B should be ordered before A.

The fixpoint reduction, as depicted in Figure 1, uses the approximative test $p\mathcal{A}(f, \mathcal{O}^*)$ to remove facts from F^A_{DA} that can be achieved. It finds *all* these facts under certain restrictions, see below. As a side effect of the fixpoint algorithm, we obtain the set $\mathcal{O}^*$ of actions that our method assumes to be applicable after a state $s_{(A,\neg B)}$. We then order B before A iff it cannot possibly be achieved using these actions.

```
F* := F^A_DA
O* := O_A \ {o | F* ∩ pre(o) ≠ ∅}
fixpoint_reached := FALSE
while ¬fixpoint_reached
        fixpoint_reached := TRUE
        for f ∈ F*
            if pA(f, O*) then
               F* := F* \ {f}
               O* := O_A \ {o | F* ∩ pre(o) ≠ ∅}
               fixpoint_reached := FALSE
            endif
        endfor
endwhile
return F*, O*
```

Figure 1: Quick, heuristic fixpoint reduction of the set F^A_{DA}.

The computation checks whether atoms of F^*, which is initially set to F^A_{DA}, are possibly achievable using only those actions, which do not delete A and do not require atoms from F^* as a precondition. Achievable atoms are removed from F^*, and $\mathcal{O}^*$ gets updated accordingly. If in one iteration, F^* does not change, the fixpoint is reached, i.e., F^* will

not further decrease and $\mathcal{O}^*$ will not further increase—the final sets F^* of false facts and $\mathcal{O}^*$ of applicable actions are returned.

As already mentioned, the intention behind the fixpoint procedure is the following: Starting from a state $s_{(A,\neg B)}$, we want to know which facts can become TRUE without destroying A, and consequently, which actions can become applicable. In the first step, only actions that do not use any of the facts in F^A_{DA} are applicable, as all those facts are deleted from the state description when A is added. However, such actions may make facts in F^A_{DA} TRUE, so we want to remove those facts from F^A_{DA}. If we manage to find *all* the facts that can be made TRUE without destroying A, then the final set F^* will contain only facts that are never fulfilled in a state reachable from $s_{(A,\neg B)}$ without destroying A. In this case, the final action set $\mathcal{O}^*$ will contain *all* the actions that can be applied after $s_{(A,\neg B)}$, and we can safely use this action set to determine whether another goal B can still be achieved or not.

However, as we only use the approximative test $p\mathcal{A}(f, \mathcal{O}^*)$ with $f \in \mathsf{F}^*$ to find out if a fact in the current F^* is achievable, there may be facts which are achievable without destroying A, but which remain in the set F^*. This could exclude actions from the set $\mathcal{O}^*$ which can be safely applied after $s_{(A,\neg B)}$. Under certain restrictions, however, we can prove that this will not happen. In order to do so, we need to impose a restriction on the particular state $s_{(A,\neg B)}$, in which we achieved the goal A: If none of the preconditions of actions, which add facts contained in F^A_{DA}, occur in the state $s_{(A,\neg B)}$, then the fixpoint procedure will remove all facts from F^A_{DA} that are achievable without destroying A. We will use this property of the fixpoint procedure later to show that our heuristic ordering relation approximates reasonable orderings.

Lemma 4 *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A \in \mathcal{G}$ be an atomic goal. Let $s_{(A,\neg B)}$ be a reachable state where A has just been achieved. Let $\mathcal{P}^{\mathcal{O}_A} = \langle o_1, \ldots, o_n \rangle$ be a sequence of actions that do not destroy A. Let F^* be the set of facts that is returned by the fixpoint computation depicted in Figure 1. If we have*

$$\forall f \in \mathsf{F}^A_{DA} : \forall o \in \mathcal{O}_A : f \in add(o) \Rightarrow pre(o) \cap s_{(A,\neg B)} = \emptyset \qquad (*)$$

then no fact in F^* holds in the state that is reached by applying $\mathcal{P}^{\mathcal{O}_A}$, i.e.,

$$Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}_A}) \cap \mathsf{F}^* = \emptyset$$

Proof:

Let F^*_j and $\mathcal{O}^*_j$ denote the state of the fact and action sets, respectively, after j iterations of the algorithm depicted in Figure 1. As F^* only decreases during the computation, we have $\mathsf{F}^* \subseteq \mathsf{F}^*_j$ for all j. Let $s_0, \ldots, s_n$ denote the sequence of states that are encountered when executing $\mathcal{P}^{\mathcal{O}_A} = \langle o_1, \ldots, o_n \rangle$ in $s_{(A,\neg B)}$, i.e., $s_0 = s_{(A,\neg B)}$ and $s_i = Result(s_{i-1}, \langle o_i \rangle)$ for $0 \leq i \leq n$. We can assume that each action o_i is applicable in state s_{i-1}, i.e., $pre(o_i) \subseteq s_{i-1}$. Otherwise, o_i does not cause any state transition, and we can skip it from $\mathcal{P}^{\mathcal{O}_A}$. Obviously, we have $s_n = Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}_A})$, so we need to show that $s_n \cap \mathsf{F}^* = \emptyset$. The proof proceeds per induction over the length n of $\mathcal{P}^{\mathcal{O}_A}$.

$\underline{n = 0}$: $\mathcal{P}^{\mathcal{O}_A} = \langle\rangle$ and $s_n = s_0 = s_{(A,\neg B)}$. All facts in F^A_{DA} are deleted from the state description when A is added, so we have $s_n \cap \mathsf{F}^A_{DA} = \emptyset$. As $\mathsf{F}^A_{DA} = \mathsf{F}^*_0$ and $\mathsf{F}^* \subseteq \mathsf{F}^*_0$, the proposition follows immediately.

$\underline{n \to n+1}$: $\mathcal{P}^{\mathcal{O}_A} = \langle o_1, \ldots, o_n, o_{n+1} \rangle$. From the induction hypothesis, we know that $s_i \cap \mathsf{F}^* = \emptyset$ for $0 \leq i \leq n$. What we need to show is $s_{n+1} \cap \mathsf{F}^* = \emptyset$.

Let j be the number of the fixpoint iteration where the intersection $\mathsf{F}^*_j \cap \bigcup_{i=0,\ldots,n} s_i$ becomes empty, i.e., j denotes the iteration in which the intersections of all the states $s_i, i \leq n$ with F^*_j are empty the first time. Such an iteration exists, because all the intersections $s_i \cap \mathsf{F}^*$ with $i \leq n$ are empty.

Now each action $o_i, 1 \leq i \leq n+1$ is applicable in state s_{i-1}, i.e., $pre(o_i) \subseteq s_{i-1}$, and thus $pre(o_i) \cap \mathsf{F}^*_j = \emptyset$ for all the actions o_i in $\mathcal{P}^{\mathcal{O}_A}$. Therefore, all these actions are contained in $\mathcal{O}^*_j$, as this set contains all the actions out of $\mathcal{O}_A$ whose intersection with F^*_j is empty. Let us focus on the facts in the state s_{n+1}. All these facts are achieved by executing $\mathcal{P}^{\mathcal{O}_A}$ in $s_{(A,\neg B)}$. In other words, there is a plan from $s_{(A,\neg B)}$ to each of these facts. As we have just seen, this plan consists out of actions in $\mathcal{O}^*_j$. Applying Lemma 2 to all the facts $p \in s_{n+1}$ using $s_{(A,\neg B)}$ and $\mathcal{P}^{\mathcal{O}_A}$ $(= \mathcal{P}^{\mathcal{O}^*_j})$, we know that all facts p are achievable using actions from $\mathcal{P}^{\mathcal{O}^*_j}$.

$$\forall p \in s_n : \mathcal{A}(s_{(A,\neg B)}, p, \mathcal{O}^*_j)$$

We will now show that those facts $f \in s_n$ we are interested in, namely the F facts that are added by o_{n+1} and that are still contained in F_j, are also *possibly* achievable using actions from $\mathcal{O}^*_j$. Let f be a fact $f \in s_{n+1}$, $f \in \mathsf{F}^*_j$. We apply Lemma 3 using $s_{(A,\neg B)}$, f, and $\mathcal{O}^*_j$. We *can* apply Lemma 3 as obviously $f \notin s_{(A,\neg B)}$, and as $\forall o \in \mathcal{O}^*_j : f \in add(o) \Rightarrow pre(o) \cap s_{(A,\neg B)} = \emptyset$ by prerequisite $(*)$. With $\mathcal{A}(s_{(A,\neg B)}, p, \mathcal{O}^*_j)$, we arrive at

$$\forall f \in s_n \cap \mathsf{F}^*_j : p\mathcal{A}(f, \mathcal{O}^*_j)$$

What remains to be proven is that all these facts f will be removed from F^* during the fixpoint computation. With the argumentation above, it is sufficient to show that all the facts $f \in s_n \cap \mathsf{F}^*_j$ will get tested for $p\mathcal{A}(f, \mathcal{O}^*_j)$ in iteration $j+1$ of the fixpoint computation. These tests will succeed and lead to $s_n \cap \mathsf{F}^*_{j+1} = \emptyset$, yielding, as desired, $s_n \cap \mathsf{F}^* = \emptyset$. Remember that $\mathsf{F}^*_{j+1} \supseteq \mathsf{F}^*$. There are two cases, which we need to consider:

1. $j = 0$: all intersections $s_i \cap \mathsf{F}^*_0$ are initially empty, i.e., $s_i \cap \mathsf{F}^A_{DA} = \emptyset$ for $0 \leq i \leq n$. In this case, all facts $f \in s_n \cap \mathsf{F}^A_{DA}$ are tested for $p\mathcal{A}(f, \mathcal{O}^*_0)$ in iteration $j+1 = 1$ of the fixpoint computation.

2. $j > 0$: in this case, at least one of the intersections $s_i \cap \mathsf{F}^*_j$ became empty in iteration j by definition of j, i.e., at least one fact was removed from F^* in this iteration. Therefore, the fixpoint has not been reached yet, and the computation performs at least one more iteration, namely iteration $j+1$. All facts in F^*_j will be tested in this iteration, in particular all facts $f \in s_n \cap \mathsf{F}^*_j$.

With these observations, the induction is complete and the proposition is proven. ■

Let us illustrate the fixpoint computation with a short example:

initial state: $\emptyset$
goals: $\{ A, B \}$

actions: **op1**: $\Rightarrow$ *ADD* { *A* } *DEL* { *C, D* }
op2: $\Rightarrow$ *ADD* {*A, C* } *DEL* { *D* }
op3: { *C* } $\Rightarrow$ *ADD* { *D* }
op4: { *D* } $\Rightarrow$ *ADD* { *B* }

When assuming that A has been achieved, we obtain $\mathsf{F}^* = \mathsf{F}^A_{DA} = \{D\}$ as the initial value of the False set, since D is the only atom that both actions delete when adding A. Figure 2 illustrates a hypothetical planning process. Starting in the empty initial state and trying to achieve A first, we get two different states $s_{(A,\neg B)}$ in which A holds. The atom D does not hold in any of them and thus in both states, no action is applicable that requires D as a precondition. This excludes **op4** from $\mathcal{O}_A$, yielding the initial action set $\mathcal{O}^* = \{\mathbf{op1}, \mathbf{op2}, \mathbf{op3}\}$. Now, **op4** is the only action that can add B. Therefore, if we used this action set to see if B can still be achieved, we would find that this is *not* the case. Consequently, without performing the fixpoint computation, we would order B before A. But as can be seen in Figure 2, this would not be a reasonable ordering: there is the plan $\langle$**op3** ,**op4**$\rangle$ that achieves B from the state $s_{(A,\neg B)} = \mathit{Result}(\mathcal{I}$,**op2**$)$ without destroying A.

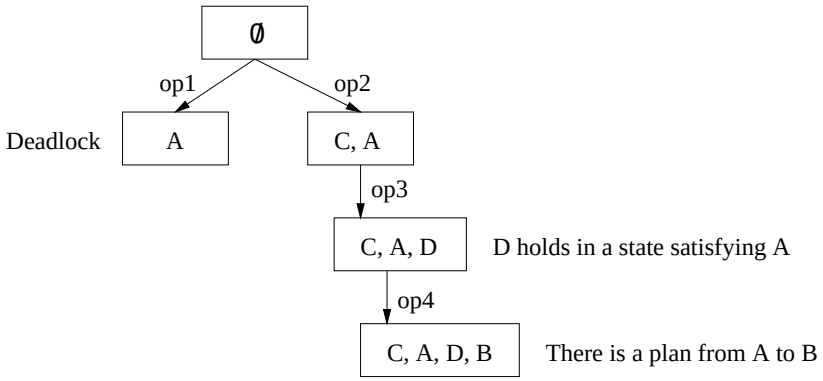

Figure 2: An example illustrating why we need the fixpoint computation.

The fixpoint computation works us around this problem as follows: There is the action **op3**, which can add the precondition D of **op4** without deleting A. When checking $p\mathcal{A}(D, \mathcal{O}^*)$ in the first iteration, the fixpoint procedure finds this action. It then checks whether the preconditions of **op3** are achievable in the sense that they are added by another action. This is the case since the only precondition C is added by **op2**. Thus, D is removed from F^*, which becomes empty now. The action **op4** is put back into the set $\mathcal{O}^*$, which now becomes identical with the action set $\mathcal{O}_A$. This set, in turn, is identical with the original action set $\mathcal{O}$ as no action deletes A. The fixpoint process terminates and B will not be ordered before A as it can be achieved by the action **op4**. This correctly reflects the fact that there exists a plan from the state $s_{(A,\neg B)} = \mathit{Result}(\mathcal{I}, \langle\mathbf{op2}\rangle) = \{C, A\}$ to a state that satisfies B without destroying A.

As has already been said, we now simply order B before A, if it is not possibly achievable using the action set that resulted from the fixpoint computation. The ordering relation $\leq_h$ (where h stands for "heuristic") obtained in this way approximates the reasonable goal ordering $\leq_r$.

Definition 13 (Heuristic Ordering $\leq_h$) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G} \supseteq \{A, B\})$ be a planning problem. Let $\mathcal{O}^*$ be the set of actions that is obtained from $\mathcal{O}$ by performing the fixpoint computation shown in Figure 1.*

The ordering $B \leq_h A$ holds if and only if

$$\neg p\mathcal{A}(B, \mathcal{O}^*)$$

If A has been reached in a particular state $s_{(A,\neg B)}$ where the assumptions made by the fixpoint computation and by the test for $p\mathcal{A}(B, \mathcal{O}^*)$ are justified, then being not possibly achievable is a sufficient condition for the non-existence of a plan to B that does not temporarily destroy A.

Theorem 5 *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A, B \in \mathcal{G}$ be two atomic goals. Let $s_{(A,\neg B)}$ be a reachable state where A has just been achieved, but B is still false, i.e., $B \notin s_{(A,\neg B)}$. Let F^* and $\mathcal{O}^*$ be the sets of facts and actions, respectively, that are derived by the fixpoint computation shown in Figure 1. If we have*

$$\forall f \in \mathsf{F}^A_{DA} \cup \{B\} : \forall o \in \mathcal{O}_A : f \in add(o) \Rightarrow pre(o) \cap s_{(A,\neg B)} = \emptyset \qquad (**)$$

then we have

$$\neg p\mathcal{A}(B, \mathcal{O}^*) \Rightarrow \neg \exists \mathcal{P}^{\mathcal{O}_A} : \; B \in Result(s_{(A,\neg B)}, \mathcal{P}^{\mathcal{O}_A})$$

Proof: Assume that there is a plan $\mathcal{P}^{\mathcal{O}_A} = \langle o_1, \ldots, o_n \rangle$ that does not destroy A, but achieves B, i.e., $B \in Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_n \rangle)$. With the restriction of $(**)$ to the facts in F^A_{DA}, Lemma 4 can be applied to each action sequence $\langle o_1, \ldots, o_{i-1} \rangle$ yielding $Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_{i-1} \rangle) \cap \mathsf{F}^* = \emptyset$. Consequently, each o_i is either

- not applicable in $Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_{i-1} \rangle)$,
- or its preconditions are contained in $Result(s_{(A,\neg B)}, \langle o_1, \ldots, o_{i-1} \rangle)$, yielding $pre(o_i) \cap \mathsf{F}^* = \emptyset$.

In the first case, we simply skip o_i as it does not have any effects. In the second case, $o_i \in O^*$ follows. Thus, we have a plan constructed out of actions in O^* that achieves B from $s_{(A,\neg B)}$. Applying Lemma 2 leads us to $\mathcal{A}(s_{(A,\neg B)}, B, O^*)$. We have $B \notin s_{(A,\neg B)}$. We also know, from $(**)$ with respect to B, as $\mathcal{O}^* \subseteq \mathcal{O}_A$, that $\forall o \in \mathcal{O}^* : B \in add(o) \Rightarrow pre(o) \cap s_{(A,\neg B)} = \emptyset$ holds. Therefore, we can now apply Lemma 3 and arrive at $p\mathcal{A}(B, \mathcal{O}^*)$, which is a contradiction. ∎

We return to the *blocksworld* example and show how the computation of $\leq_h$ proceeds. Let us first investigate whether $on(a, b) \leq_h on(b, c)$. The initial value for $\mathsf{F}^{on(b,c)}_{DA}$ is obtained from the delete list of the **stack(b,c)** action, which is the only one that adds this goal.

$$\mathsf{F}^{on(b,c)}_{DA} = \{clear(c), holding(b)\}$$

Intuitively, it is immediately clear that neither of these facts can ever hold in a state where $on(b, c)$ is true: if b is on c, then c is not clear and the gripper cannot hold b. It turns out that the fixpoint computation respects this intuition and leaves the set $\mathsf{F}^{on(b,c)}_{DA}$ unchanged, yielding $\mathsf{F}^* = \{clear(c), holding(b)\}$. We do not repeat the fixpoint process in detail here, because it can be reconstructed from Figure 1 and the details are not

necessary for understanding how the correct ordering relations are derived. In short, for both facts there are achievers in the reduced action set, but all of them need preconditions for which no achiever is available. For example, $holding(b)$ can be achieved by either an **unstack** or a **pickup** action. They either need b to stand on another block, or to stand on the table, respectively. All actions that can achieve these facts need $holding(b)$ to be TRUE in the first place, and are thus excluded from the reduced action set.

After finishing the fixpoint computation, the planner tests $p\mathcal{A}(on(a,b), \mathcal{O}^*)$, where $\mathcal{O}^*$ contains all actions except those that delete $on(b,c)$ and those that use $clear(c)$ or $holding(b)$ as a precondition. It finds that the action **stack(a,b)** adds $on(a,b)$. The preconditions of this action are $holding(a)$ and $clear(b)$. These conditions are added by the actions **pickup(a)** and **unstack(a,b)**, respectively, which are both contained in $\mathcal{O}^*$: neither of them needs c to be clear or b to be in the gripper. Thus, the test finds that in fact, $on(a,b)$ is possibly achievable using the actions in $\mathcal{O}^*$, and no ordering is derived, i.e., $on(a,b) \not\leq_h on(b,c)$ follows.

Now, the other way round, $on(b,c) \leq_h on(a,b)$ is tested. The initial value for $\mathsf{F}_{DA}^{on(a,b)}$ is obtained from the single action **stack(a,b)** as

$$\mathsf{F}_{DA}^{on(a,b)} = \{clear(b), holding(a)\}$$

Again, the fixpoint computation does not cause any changes, resulting in $\mathsf{F}^* = \{clear(b), holding(a)\}$. The process now tests whether $p\mathcal{A}(on(b,c), \mathcal{O}^*)$ holds, where O^* contains all actions except those that delete $on(a,b)$ and those that use $clear(b)$ or $holding(a)$ as a precondition. The only action that can add $on(b,c)$ is **stack(b,c)**. This action needs as preconditions the facts $holding(b)$ and $clear(c)$. The process now finds that a crucial condition for achieving the first fact is violated: Each action that can achieve $holding(b)$ has $clear(b)$ as a precondition, because b must be clear first before the gripper can hold it. Since $clear(b)$ is an element of F^*, none of the actions achieving $holding(b)$ is contained in $\mathcal{O}^*$. Consequently, the test for $p\mathcal{A}(on(b,c), \mathcal{O}^*)$ fails and we obtain the ordering $on(b,c) \leq_h on(a,b)$. This makes sense as the gripper cannot grasp b and stack it onto c anymore, once $on(a,b)$ is achieved.

3.3 On the Non-Existence of Forced Goal Orderings

So far, we have introduced two easily computable ordering relations $\leq_h$ and $\leq_e$ that both approximate the reasonable goal ordering $\leq_r$. We will now explain the role of forced goal orderings $\leq_f$ and argue why we do not invest any effort in trying to find those. The results in this section are a bit more general than necessary at this point as we want to make use of them later on in the paper.

At a first glance, the early recognition of forced orderings appears to be useful for a planning system, because it can prevent a planner from "running" into a state from which the goals cannot be reached anymore. In this paper in particular, we are going to propose an *incremental* approach to planning. In a nutshell, this means that the goals are achieved one by one in a certain ordering as will be described in detail in Section 4. If the planner fails to recognize an ordering $B \leq_f A$, it might incorrectly order the goals and try to achieve A before B, which will render the problem unsolvable. As a consequence, the incremental approach can loose completeness in the presence of forced orderings. We will see in the following, however, that almost all "natural" domains do not contain forced orderings anyway.

First, we introduce the more general notion of a *deadlock* in a planning problem.

Definition 14 (Deadlock) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem. A reachable state s is called a* deadlock *iff there is no sequence of actions that leads from s to the goal, i.e., iff $s = Result(\mathcal{I}, \mathcal{P}^{\mathcal{O}})$ and $\neg\exists\, \mathcal{P}'^{\mathcal{O}} :\ \mathcal{G} \subseteq Result(s, \mathcal{P}'^{\mathcal{O}})$.*

Non-trivial forced goal orderings imply the existence of deadlocks (remember that an ordering $B \leq_f A$ or $B \leq_r A$ is called *trivial* iff there is no state $s_{(A,\neg B)}$ at all).

Lemma 5 *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let $A, B \in \mathcal{G}$ be two atomic goals. If there is a non-trivial forced ordering $B \leq_f A$ between A and B, then there exists a deadlock state s in the problem.*

Proof: Recalling Definition 9 and assuming non-triviality of $\leq_f$, we know that there is at least one state $s_{(A,\neg B)}$ where A is made TRUE, but B is still FALSE. From Definition 7, we know that there is no plan in any such state that achieves B, in particular, it is not possible to achieve *all* goals starting out from $s_{(A,\neg B)}$. Thus, the state $s := s_{(A,\neg B)}$ must be a deadlock. ■

We will now investigate deadlocks in more detail and show that most "natural" problems do not contain them. With Lemma 5, we then also know that such domains do not contain non-trivial forced goal orderings either—so there is no point in trying to find them. We do not care about trivial goal orderings. Such orderings force any reasonable planning algorithm to consider the goals in the correct order.

The existence of deadlocks depends on structural properties of a planning problem: There must be action sequences, which, once executed, lead into states from which the goals can not be reached anymore. These sequences must have undesired effects, which cannot be inverted by any other sequence of actions in $\mathcal{O}$. Changing perspective, one obtains a hint on how a sufficient condition for the *non-existence* of deadlocks might be defined. Say we have a planning problem where the effects of each action sequence in the domain can be inverted by executing a certain other sequence of actions. In such an *invertible* planning problem, it is in particular possible to get back to the *initial state* from each reachable state. Therefore, if such a problem is solvable, then it does not contain deadlocks: From any state, one can reach all goals by going back to the initial state first, and execute an arbitrary solution thereafter. We will now formally define the notion of invertible planning problems, and turn the above argumentation into a formal proof.

Definition 15 (Invertible Planning Problem) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be a planning problem, and let s denote the states that are reachable from $\mathcal{I}$ with actions from $\mathcal{O}$. The problem is called* invertible *if and only if*

$$\forall\, s :\ \forall\, \mathcal{P}^{\mathcal{O}} :\ \exists\, \overline{\mathcal{P}}^{\mathcal{O}} :\ Result(Result(s, \mathcal{P}^{\mathcal{O}}), \overline{\mathcal{P}}^{\mathcal{O}}) = s$$

Theorem 6 *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ be an invertible planning problem, for which a solution exists. Then $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ does not contain any deadlocks.*

Proof: Let $s = Result(\mathcal{I}, \mathcal{P}_s^{\mathcal{O}})$ be an arbitrary reachable state. As the problem is invertible, we know that there is a sequence of actions $\overline{\mathcal{P}}_s^{\mathcal{O}}$ for which $Result(s, \overline{\mathcal{P}}_s^{\mathcal{O}}) = \mathcal{I}$ holds.

As the problem is solvable, we have a solution plan $\mathcal{P}^{\mathcal{O}}$ starting from $\mathcal{I}$ and achieving $\mathcal{G} \subseteq Result(\mathcal{I}, \mathcal{P}^{\mathcal{O}})$. Together, we obtain $\mathcal{G} \subseteq Result(Result(s, \overline{\mathcal{P}}_s^{\mathcal{O}}), \mathcal{P}^{\mathcal{O}})$. Therefore, the concatenation of $\overline{\mathcal{P}}_s^{\mathcal{O}}$ and $\mathcal{P}^{\mathcal{O}}$ is a solution plan executable in s and consequently, s is no deadlock. ■

We now know that invertible planning problems, if solvable, do not contain deadlocks and consequently, they do not contain (non-trivial) forced goal orderings. What we will see next is that, as a matter of fact, most more "natural" planning problems *are* invertible. We arrive at a sufficient condition for invertibility through the notion of *inverse actions*.

Definition 16 (Inverse Action) *Given an action set $\mathcal{O}$ containing an action o of the form $pre(o) \Rightarrow add(o)\ \ del(o)$. An action $\overline{o} \in \mathcal{O}$ is called* inverse *to o if and only if $\overline{o}$ has the form $pre(\overline{o}) \Rightarrow add(\overline{o})\ \ del(\overline{o})$ and satisfies the following conditions*

1. $pre(\overline{o}) \subseteq pre(o) \cup add(o) \setminus del(o)$
2. $add(\overline{o}) = del(o)$
3. $del(\overline{o}) = add(o)$

Two inverse actions have inverse effects and establish each others preconditions.

Lemma 6 *Let s be a state and o be an action, which is applicable in s and for which $del(o) \subseteq pre(o)$ holds. An action $\overline{o}$ that is inverse to o in the sense of Definition 16 is applicable in $Result(s, \langle o \rangle)$ and $Result(Result(s, \langle o \rangle), \langle \overline{o} \rangle) = s$ holds.*

Proof: As o is applicable in s, we have $pre(o) \subseteq s$. The atoms in $add(o)$ are added, and the atoms in $del(o)$ are removed from s, so altogether we have

$$Result(s, \langle o \rangle) \supseteq (pre(o) \cup add(o)) \setminus del(o) \supseteq pre(\overline{o})$$

Thus, $\overline{o}$ is applicable in $Result(s, \langle o \rangle)$.

Furthermore, we have $Result(s, \langle o \rangle) = s \cup add(o) \setminus del(o)$ and with that

$$\begin{aligned}
&Result(Result(s, \langle o \rangle), \langle \overline{o} \rangle) \\
&\qquad = Result(s \cup add(o) \setminus del(o), \langle \overline{o} \rangle) \\
&\qquad = (s \cup add(o) \setminus del(o)) \cup add_{\overline{o}} \setminus del_{\overline{o}} \\
&\qquad = (s \cup add(o) \setminus del(o)) \cup del(o) \setminus add(o) && \text{(cf. Definition 16)} \\
&\qquad = s \cup add(o) \setminus add(o) && \text{(because } del(o) \subseteq pre(o) \subseteq s\text{)} \\
&\qquad = s
\end{aligned}$$

■

If we have an inverse action to each action o in a planning problem, then the problem is invertible.

Theorem 7 *Given a planning problem $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ with the set of ground actions $\mathcal{O}$, where for each action $del(o) \subseteq pre(o)$ holds. If there is an inverse action $\overline{o} \in \mathcal{O}$ for each action $o \in \mathcal{O}$, then the problem is invertible.*

Proof: Let s be a reachable state, and let $\mathcal{P}^{\mathcal{O}} = \langle o_1, \ldots o_n \rangle$ be a sequence of actions. We need to show the existence of a sequence $\overline{\mathcal{P}}^{\mathcal{O}}$ for which

$$Result(Result(s, \mathcal{P}^{\mathcal{O}}), \overline{\mathcal{P}}^{\mathcal{O}}) = s \qquad (***)$$

holds. We define $\overline{\mathcal{P}}^{\mathcal{O}} := \langle \overline{o_n}, \ldots, \overline{o_1} \rangle$, and prove $(***)$ per induction over n.

$\underline{n = 0}$: Here, we have $\mathcal{P}^{\mathcal{O}} = \overline{\mathcal{P}}^{\mathcal{O}} = \langle\rangle$, and $Result(Result(s, \langle\rangle), \langle\rangle) = s$ is obvious.

$\underline{n \to n+1}$: Now $\mathcal{P}^{\mathcal{O}} = \langle o_1, \ldots, o_n, o_{n+1} \rangle$. From the induction hypothesis we know that $Result(Result(s, \langle o_1, \ldots, o_n \rangle), \langle \overline{o_n}, \ldots, \overline{o_1} \rangle) = s$. To make the following a bit more readable, let s' denote $s' := Result(s, \langle o_1, \ldots, o_n \rangle)$. We have

$$\begin{aligned}
&Result(Result(s, \langle o_1, \ldots, o_{n+1} \rangle), \langle \overline{o_{n+1}}, \ldots, \overline{o_1} \rangle) \\
&\quad = Result(Result(s', \langle o_{n+1} \rangle), \langle \overline{o_{n+1}}, \ldots, \overline{o_1} \rangle) \\
&\quad = Result(Result(Result(s', \langle o_{n+1} \rangle), \langle \overline{o_{n+1}} \rangle), \langle \overline{o_n}, \ldots, \overline{o_1} \rangle) \\
&\quad = Result(s', \langle \overline{o_n}, \ldots, \overline{o_1} \rangle) && \text{(cf. Lemma 6 on } s' \text{ and } o_n\text{)} \\
&\quad = s && \text{(per induction)}
\end{aligned}$$

■

Altogether, we know now that invertible problems, if solvable, do not contain forced orderings. We also know that problems, where there is an inverse action to each action in $\mathcal{O}$, are invertible following Theorem 7. Theorem 7 requires $del(o) \subseteq pre(o)$ to hold for each action o. We will see that both conditions, inclusion of the delete list in the precondition list, and existence of inverse actions, hold in almost all benchmark domains.

Concerning the first condition that actions only delete facts they require as preconditions, one finds this phenomenon in *all* domains that are commonly used by the planning community, at least in those that are known to the authors. It is just something that seems to hold in any reasonable logical problem formulation. Some authors even postulate it as an assumption for their algorithms to work, cf. [FL98].

Concerning the existence of inverse actions, one usually finds them in all planning domains which attempt to model realistic (although simplified) problems. For example in the *blocksworld*, we have **stack** and **unstack**. Similarly in all domains that deal with logistics problems, for example *logistics*, *mystery*, *rocket*, *gripper* etc., one can always find inverse pairs of actions. Sometimes, two different ground instances of the same operator schema yield an inverse pair. For example, in *gripper*, the two ground instances

move(roomA, roomB)
at-robby(roomA) $\Rightarrow$ *ADD at-robby(roomB) DEL at-robby(roomA).*

and

move(roomB, roomA)
at-robby(roomB) $\Rightarrow$ *ADD at-robby(roomA) DEL at-robby(roomB).*

of the **move(?from,?to)** operator schema invert each other. Also, in *towers of hanoi*, where there is only the single **move** operator schema, an inverse instance can be found for each ground instance of the schema.

Only very rarely, non-invertible actions can be found in domains. If they occur, their role in the domain is often quite limited as for example the operators **cuss** and **inflate** in Russel's Tyreworld.

cuss
$\Rightarrow$ *DEL annoyed().*

inflate(?x:wheel)
have(pump) not-inflated(?x) intact(?x) $\Rightarrow$ *ADD inflated(?x) DEL not-inflated(?x).*

Obviously, there is not much point in defining something like a **decuss** or a **deflate** operator. More formally speaking, none of the ground actions to these operators destroys a goal or a precondition of any other action in the domain. Therefore, it does not matter that their effects cannot be inverted. In particular, no forced goal ordering can be derived wrt. these actions. [4]

The importance of inverse actions in real-world domains has also been discussed in [WN97], which describes the planner BURTON that controls the Cassini spacecraft.

In contrast to these natural domains, artificial domains such as those for example used in [BW94] almost never contain inverse actions. Consequently, in these domains plenty of forced goal orderings could be discovered and used by a planner to avoid deadlock situations. But, given the observation that natural domains are usually deadlock-free, forced goal orderings will in most cases fail in providing a reasonable problem decomposition.

3.4 An Extension of Goal Orderings to ADL Actions

The orderings which have been introduced so far can be easily extended to deal with ground ADL [Ped89] actions having conditional effects and using negation instead of delete lists. Such actions have the following syntactic structure:

$$
\begin{aligned}
o: \quad & \phi_0(o) = \mathrm{pre}_0(o) \Rightarrow \mathrm{eff}_0^+(o), \mathrm{eff}_0^-(o) \\
& \phi_1(o) = \mathrm{pre}_1(o) \Rightarrow \mathrm{eff}_1^+(o), \mathrm{eff}_1^-(o) \\
& \vdots \\
& \phi_n(o) = \mathrm{pre}_n(o) \Rightarrow \mathrm{eff}_n^+(o), \mathrm{eff}_n^-(o)
\end{aligned}
$$

All unconditional elements of the action are summarized in $\phi_0(o)$: The precondition of the action is denoted with $\mathrm{pre}_0(o)$, and its unconditional positive and negative effects

[4] The **cuss** operator, by the way, is the only one known to the authors that deletes a fact it is not using as a precondition. It is also the only one we know that could be removed from the domain description without changing anything.

with $\text{eff}_0^+(o)$ and $\text{eff}_0^-(o)$, respectively. Each conditional effect $\phi_i(o)$ consists of an effect condition (antecedent) $\text{pre}_i(o)$, and the positive and negative effects $\text{eff}_i^+(o)$ and $\text{eff}_i^-(o)$. Additionally, we denote with $\Phi(o)$ the set of all unconditional and conditional effects, i.e., $\Phi(o) = \{\phi_0(o), \phi_1(o), \ldots, \phi_n(o)\}$.

The computation of $\leq_e$ immediately carries over to ADL actions when an extension of planning graphs is used, which can handle conditional effects, e.g., IPP [KNHD97] or SGP [AW98]. One simply takes the set of exclusive facts that is returned by these systems to determine the set F_{GP}^A. The test from Definition 10, which decides whether there is an ordering $B \leq_e A$ of two atomic goals A and B, is extended to ADL as follows.

Definition 17 (Ordering $\leq_e$ for ADL) *Let $(\mathcal{O}, \mathcal{I}, \mathcal{G} \supseteq \{A, B\})$ be a planning problem. Let F_{GP}^A be the* False *set for A. The ordering $B \leq_e A$ holds if and only if*

$$\forall\, o \in \mathcal{O},\ \phi_i(o) \in \Phi(o) : B \in \mathit{eff}_i^+(o) \wedge A \notin D_i(o)\ \Rightarrow (\mathit{pre}_i(o) \cup \mathit{pre}_0(o)) \cap \mathsf{F}_{GP}^A \neq \emptyset$$

Here, $D_i(o)$ denotes all negative effects that are implied by the conditions of $\phi_i(o)$.

$$D_i(o) := \begin{cases} \mathit{eff}_0^-(o) \cup \bigcup_{\mathit{pre}_j(o) \subseteq \mathit{pre}_i(o)} \mathit{eff}_j^-(o) & i \neq 0 \\ \mathit{eff}_0^-(o) & i = 0 \end{cases}$$

Thus, B is ordered before A if all (unconditional or conditional) effects that add B either imply an effect that deletes A, or need conditions that cannot be made TRUE together with A. Note that an effect ϕ_i requires all the conditions in $\mathit{pre}_i(o) \cup \mathit{pre}_0(o)$ to be satisfied, which is impossible in any state where A holds because of the non-empty intersection with F_{GP}^A.

The computation of $\leq_h$ requires a little more adaptation effort. In order to obtain the set F_{DA}^A, we now need to investigate the conditional effects as well. For each action that has A as a conditional or unconditional effect, we determine which atoms are negated by it, no matter which effect is used to achieve A. We obtain these atoms by intersecting the appropriate sets $D_i(o)$.

$$D(o) := \bigcap_{A \in \text{eff}_i^+(o)} D_i(o)$$

These are exactly the facts that are always deleted by o when achieving A, no matter which effect we use. The intersection of the sets $D(o)$ for all actions o yields the desired set F_{DA}^A. Let us consider the following small example to clarify the computation.

$$\begin{array}{lcl} \phi_0(o) = \{U\} & \Rightarrow & \{W\}\ \{\neg X\}; \\ \phi_1(o) = \{V, W\}\ \{\neg U\} & \Rightarrow & \{A\}\ \{\neg X\}; \\ \phi_2(o) = \{W\} & \Rightarrow & \{U\}\ \{\neg Y\} \end{array}$$

We obtain $D_1(o) = \{\neg X\} \cup \{\neg Y\} = \{\neg X, \neg Y\}$, because the precondition of $\phi_2(o)$ is implied by the first conditional effect $\phi_1(o)$. As $\phi_1(o)$ is the only effect that can achieve A, we get $D(o) = D_2(o) = \{\neg X, \neg Y\}$. We obtain a smaller set $D(o)$, if we add A as an unconditional positive effect of the action.

$$\begin{array}{lll} \phi_0(o) = \{U\} & \Rightarrow & \{W, A\}\ \{\neg X\}; \\ \phi_1(o) = \{V, W\}\ \{\neg U\} & \Rightarrow & \{A\}\ \{\neg X\}; \\ \phi_2(o) = \{W\} & \Rightarrow & \{U\}\ \{\neg Y\} \end{array}$$

In this case, we need to intersect the sets $D_0(o) = \{\neg X\}$ and $D_1(o) = \{\neg X, \neg Y\}$, yielding $D(o) = \{\neg X\}$. This reflects the fact that, when achieving A via the unconditional effect of o, only X gets removed from the state.

The fixpoint computation requires to adapt the computation of $\mathcal{O}^*$. First, we repeat the same steps as in the case of simple STRIPS actions and consider the unconditional negative effects and the intersection of the preconditions with the False set:

$$\mathcal{O}^* := \mathcal{O} \setminus \{o \mid A \in \mathrm{eff}_0^-(o) \vee \mathsf{F}_{DA}^A \cap pre_0(o) \neq \emptyset\}$$

Then, we additionally remove from each action the conditional effects that either delete A or have an impossible effect condition.

$$\mathcal{O}^* := red(\mathcal{O}^*) = \{red(o) | o \in \mathcal{O}^*\}$$

Here, red is a function $red(o) : o \mapsto o'$ such that

$$\Phi(o') = \Phi(o) \setminus \{\phi_k(o) \mid A \in \mathrm{eff}_k^-(o) \vee \mathrm{pre}_k(o) \cap \mathsf{F}_{DA}^A \neq \emptyset\}$$

Finally, we need to redefine Definition 12, which expresses the conditions under which a fact is believed to be *possibly achievable* given a certain set of operators $\mathcal{O}$.

Definition 18 (Possibly Achievable Atoms for ADL) *An atom p is possibly achievable given an action set $\mathcal{O}$ (written $p\mathcal{A}(p, \mathcal{O})$) if and only if*

$$\begin{array}{l} \exists\, o \in \mathcal{O},\ \phi_i \in \Phi(o) : p \in \mathit{eff}_i^+(o) \wedge \\ \qquad \forall\, p' \in (pre_i(o) \cup pre_0(o)) :\ \exists\, o' \in \mathcal{O},\ \phi_{i'} \in \Phi(o') : p' \in \mathit{eff}_{i'}^+(o') \end{array}$$

holds, i.e., there is a positive effect for p and all of its conditions and preconditions can be made true by other effects in the reduced action set.

The process, which decides whether an atomic goal B is heuristically ordered before another goal A (i.e., whether $B \leq_h A$ holds) proceeds in exactly the same way as described in Section 3.2: The False set F_{DA}^A for A is reduced by the fixpoint computation, which remains unchanged, but employs the updated routines for computing $\mathcal{O}^*$ and for deciding $p\mathcal{A}(f, \mathcal{O}^*)$. As a result, B is ordered before A ($B \leq_h A$) if and only if it is not possibly achievable $p\mathcal{A}(B, \mathcal{O}^*)$ using the action set that results from the fixpoint.

4 The Use of Goal Orderings During Planning

After having determined the ordering relations that hold between pairs of atomic goals from a given goal set, the question is how to make use of them during planning. Several proposals have been made in the literature, see Section 6 for a detailed discussion. In this paper, we propose a novel approach that extracts an explicit ordering between subsets of the goal set—called the *goal agenda*. The planner, in our case IPP, is then run incrementally on the planning subproblems represented in the agenda.

4.1 The Goal Agenda

The first step one has to take for computing the goal agenda is to perform the *goal analysis*. During goal analysis, each pair $A, B \in \mathcal{G}$ of atomic goals must be examined in order to find out whether an ordering relation $A \leq B$, or $B \leq A$, or both, or none holds between them. For the ordering relation $\leq$, an arbitrary definition can be used. In our experiments, the relation $\leq$ was always either $\leq_e$ or $\leq_h$.

After having determined all ordering relations that hold between atomic goals, we want to split the goal set into smaller sets based on these relations, and we want to order the smaller sets, also based on these relations. More specifically, our goal is to have a sequence of goal sets $G_1, \ldots, G_n$ with

$$\bigcup_{i=1}^{n} G_i = \mathcal{G}$$

and

$$G_i \cap G_j = \emptyset$$

for $i \neq j, 1 \leq i, j \leq n$. We also want the sequence of goal sets to respect the ordering relations that have been derived between atomic goals. To make this explicit, we first introduce a simple representation of the atomic orderings that have been derived: the *goal graph* G.

$$G := (V, E)$$

where

$$V := \mathcal{G}$$

and

$$E := \{(A, B) \in \mathcal{G} \times \mathcal{G} \mid A \leq B\}$$

Now, the desired properties that the sequence of goal sets should possess can be easily stated:

- Goals A, B that lie on a cycle in G belong to the same set $A, B \in G_i$.
- If G contains a path from a goal A to a goal B, but not vice versa, then A is ordered before B, i.e., $A \in G_i, B \in G_j, i < j$.

These are the only properties that appear to be reasonable for a goal-set sequence respecting the atomic orderings. We will now introduce a simple algorithmic method that does produce a sequence of goal sets which meets these requirements.

First of all, the transitive closure of G is computed. This can be done in at most cubic time in the size of the goal set [War62]. Then, for each node A in the transitive closure, the ingoing edges A_{in} and outgoing edges A_{out} are counted. All disconnected nodes with $A_{in} = A_{out} = 0$ are moved into a separate set of goals *G-sep* containing now those atomic goals for which no $\leq$ relation was derived. For all other nodes A, their *degree*

$d(A) = A_{in} - A_{out}$ is determined as the difference between the number of ingoing edges and the number of outgoing edges. Nodes with identical degree are merged into one set. The sets are then ordered by increasing degree and yield our desired sequence of goal sets. The only problem remaining is the set *G-sep*. If it is non-empty, it is not clear in which place to put it.

Let us consider a small example of the process. Figure 3 depicts on the left the goal graph, which results from the goal set $\mathcal{G} = \{A, B, C, D, E\}$ and the ordering relations $A \leq B, B \leq C$ and $B \leq D$, and its transitive closure on the right.

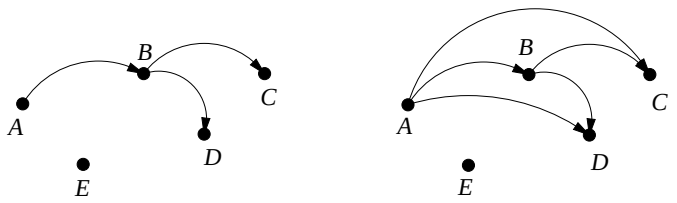

Figure 3: On the left, the goal graph depicting the $\leq$ relations between the atomic subgoals. On the right, the transitive closure of this graph.

In Figure 4, the number of in- and outgoing edges of each goal, the corresponding degrees, and resulting goal-set sequence are shown.

Figure 4: On the left, the number of in- and outgoing edges for each node. On the right, the degree of the nodes and the merged sets of goals having same degree. The node E becomes a member of the *G-sep* set and remains unordered.

It is not difficult to verify that the resulting goal sequence respects the atomic goal orderings:

- Nodes occurring on a cycle in a graph have isomorphic in- and outgoing edges in the transitive closure of that graph. Therefore, they also have the same degree and get merged into the same set G_i.
- Say we have a graph, where there is a path from A to B, but not vice versa. Then, in the transitive closure of that graph, we will have an edge from A to each node that B has a path to, and additionally the edge from A to B, i.e., $A_{out} > B_{out}$ follows. Similarly, we have an ingoing edge to B for each node that has a path to A, and additionally, the edge from A to B, which gives us $B_{in} > A_{in}$. Altogether, $d(A) = A_{in} - A_{out} < B_{in} - A_{out} < B_{in} - B_{out} = d(B)$ and thus, the degree of A is less than the degree of B, and A gets ordered before B as required.

Note that nothing is said in this argumentation about the set of unordered goals, *G-sep*. This set could, in principle, be inserted anywhere in the sequence with the resulting sequence still respecting the atomic orderings. A possible heuristic may use this goal set as the first in the sequence, because apparently there is no problem to reach all other

goals after the goals in this set have been achieved. Another heuristic could put this set at the end as there is neither a problem to reach this goal set from all other goals. We have decided to deal with the problem in a more sophisticated way by trying to derive an ordering relation between *G-sep* and other goal sets G_i that have already been derived. In order to do so, we need to extend our definitions of goal orderings to *sets* of goals.

4.2 Extension of Goal Orderings to Goal Sets

Given a set of atomic goals, it has always been a problem which of the exponentially many subsets should be compared with each other in order to derive a reasonable goal ordering between goal sets. A consideration of all possible subsets is out of question, because it will result in an exponential overhead. The partial goal agenda that we have obtained so far offers one possible answer. It suggests taking the set *G-sep* and trying to order it with respect to the goal sets emerging from the goal graph.

Given a planning problem $(\mathcal{O}, I, \mathcal{G})$ and two subsets of atomic goals $\{A_1, \ldots, A_n\} \subseteq \mathcal{G}$ and $\{B_1, \ldots, B_k\} \subseteq \mathcal{G}$, the definition of $\leq_e$ and $\leq_h$ for sets of atomic goals is straightforward. For the sake of simplicity, we consider only STRIPS actions here. The definitions can be directly extended to ADL.

To define an ordering $\leq_E$, which extends $\leq_e$ to sets, we begin by defining a set $F_{GP}^{\{A_1,\ldots,A_n\}}$ of all atoms, which are exclusive of at least one atomic goal A_i in the planning graph generated for $(\mathcal{O}, I, \mathcal{G})$:

$$F_{GP}^{\{A_1,\ldots,A_n\}} := \{p \mid p \text{ is exclusive of at least one } A_i \text{ after the graph has leveled off }\}$$

The set $\mathcal{O}_{\{A_1,\ldots,A_n\}}$ is obtained accordingly by removing from $\mathcal{O}$ all actions that delete at least one of the A_i, i.e., $\mathcal{O}_{\{A_1,\ldots,A_n\}} = \{o \in \mathcal{O} \mid \forall\, i :\ A_i \notin del(o)\}$.

Definition 19 (Ordering $\leq_E$ over Goal Sets) *Let $(\mathcal{O}, I, \mathcal{G})$ be a planning problem with $\{A_1, \ldots, A_n\} \subseteq \mathcal{G}$ and $\{B_1, \ldots, B_k\} \subseteq \mathcal{G}$. Let $\mathsf{F}_{GP}^{\{A_1,\ldots,A_n\}}$ be the* False *set for $\{A_1, \ldots, A_n\}$. The ordering $\{B_1, \ldots, B_k\} \leq_E \{A_1, \ldots A_n\}$ holds if and only if*

$$\exists\, j \in \{1, \ldots, k\} : \forall\, o \in \mathcal{O}_{\{A_1,\ldots,A_n\}} : B_j \in add(o) \Rightarrow pre(o) \cap \mathsf{F}_{GP}^{\{A_1,\ldots,A_n\}} \neq \emptyset.$$

In a similar way, $\leq_h$ can be extended to $\leq_H$. For each A_i, the sets $\mathsf{F}_{DA}^{A_i}$ are determined based on Equation (3). The set $\mathsf{F}_{DA}^{\{A_1,\ldots,A_n\}}$ is simply the union over the individual sets:

$$\mathsf{F}_{DA}^{\{A_1,\ldots,A_n\}} := \bigcup_i \mathsf{F}_{DA}^{A_i} \tag{4}$$

Then the fixpoint computation is entered with

$$\mathcal{O}^* := \mathcal{O} \setminus \{o \in \mathcal{O} \mid \exists\, i : A_i \in del(o) \vee \mathsf{F}_{DA}^{\{A_1,\ldots,A_n\}} \cap pre(o) \neq \emptyset\} \tag{5}$$

The recomputation of $\mathcal{O}^*$ in each iteration of the fixpoint algorithm from Figure 1 is done in a similar manner. Apart from this, the algorithm remains unchanged.

Definition 20 (Ordering $\leq_H$) *Let $(\mathcal{O}, I, \mathcal{G})$ be a planning problem with $\{A_1, \ldots, A_n\} \subseteq \mathcal{G}$ and $\{B_1, \ldots, B_k\} \subseteq \mathcal{G}$. Let $\mathcal{O}^*$ be the set of actions that is obtained by performing the fixpoint computation shown in Figure 1, modified to handle* sets *of facts as defined in Equations (4) and (5). The ordering $\{B_1, \ldots, B_k\} \leq_H \{A_1, \ldots, A_n\}$ holds if and only if*

$$\exists\, j\ \neg p\mathcal{A}(B_j, \mathcal{O}^*)$$

All given goal sets then undergo goal analysis, i.e., each pair of them is checked for an ordering relation $\leq_E$ or $\leq_H$. Each derived relation defines an edge in a graph with the subgoal sets as nodes. The transitive closure is determined as before, and the degree of each node is computed. If the graph contains no disconnected nodes, then a total ordering over subsets of goals results by ordering the nodes based on their degree. This ordering defines the goal agenda. In the case of disconnected nodes, we default to the heuristic of adding the corresponding goals to the last goal set in the agenda.

4.3 An Incremental Planning Algorithm

Given a planning problem $(\mathcal{O}, \mathcal{I}, \mathcal{G})$, let us assume that a goal agenda $G_1, G_2, \ldots, G_k$ with k entries has been returned by the analysis. Each entry contains a subset $G_i \subseteq \mathcal{G}$. The basic idea for an incremental planning algorithm is now to first feed the planner with the original initial state $\mathcal{I}_1 := \mathcal{I}$ and the goals $\mathcal{G}_1 := G_1$, then execute the solution plan $\mathcal{P}$ in $\mathcal{I}$, yielding the new initial state $\mathcal{I}_2 = \mathit{Result}(\mathcal{I}_1, \mathcal{P})$. Then, a new planning problem is initialized as $(\mathcal{O}, \mathcal{I}_2, \mathcal{G}_2)$. After solving this problem, we want the goals in G_2 to be TRUE, but we also want the goals in $\mathcal{G}_1$ to remain TRUE, so we set $\mathcal{G}_2 := \mathcal{G}_1 \cup G_2$. The process continues similarly, which yields a sequence of incrementally growing goal sets for the planner, namely

$$\mathcal{G}_i := \bigcup_{j=1}^{i} G_j$$

In a little more detail, the incremental planning algorithm we implemented for IPP works as follows. First, IPP is called on the problem $(\mathcal{O}, \mathcal{I}, \mathcal{G}_1)$ and returns the plan $\mathcal{P}_1$, which achieves the subgoal set $\mathcal{G}_1$. $\mathcal{P}_1$ is a sequence of parallel sets of actions, which is returned by IPP similarly to GRAPHPLAN. Given this plan, the resulting state $R(\mathcal{I}, \mathcal{P}_1) = \mathcal{I}_2$ is computed based on the operational semantics of the planning actions.[5] In the case of a set of STRIPS actions, one simply adds all ADD effects to and deletes all DEL effects from a state description in order to obtain the resulting state, according to the definition of the function *Result* in Definition 2. For STRIPS, the *Result* function coincides directly with the R function. In the case of a set of parallel ADL actions, one needs to consider all possible linearizations of the parallel action set and has to deal with the conditional effects separately. For each linearization, a different resulting state can be obtained, but each of them will satisfy the goals. To obtain the new initial state $\mathcal{I}_2$, one takes the intersection of the resulting states for each possible linearization of the actions in a parallel set. This means to compute $n!$ linearizations for a parallel action set of n actions in each time step. Since n is usually small (more than 5 or 6 ADL actions per time step are very rare), the practical costs for this computation are neglectible.

[5] See [KNHD97] for the exact definition, which we do not want to repeat here.

This way, given a solution to a subproblem $(\mathcal{O}, \mathcal{I}_i, \mathcal{G}_i)$, one calculates the new initial state $\mathcal{I}_{i+1}$ and runs the planner on the subsequent planning problem $(\mathcal{O}, \mathcal{I}_{i+1}, \mathcal{G}_{i+1})$ until the planning problem $(\mathcal{O}, \mathcal{I}_k, \mathcal{G}_k)$ is solved.

The plan solving the original planning problem $(\mathcal{O}, \mathcal{I}, \mathcal{G})$ is obtained by taking the sequence of subplans $\mathcal{P}_1, \mathcal{P}_2, \ldots, \mathcal{P}_k$. One could argue that planning for increasing goal sets can lead to highly non-optimal plans. But IPP still uses the "no-ops first" strategy to achieve goals, which was originally introduced in the GRAPHPLAN system [BF97]. Employing this strategy, the GRAPHPLAN algorithm, in short, first tries to achieve goals by simply keeping them TRUE, if possible. Since all goals $\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_i$ are already satisfied in the initial state $\mathcal{I}_{i+1}$ starting from which the planner tries to achieve $\mathcal{G}_{i+1}$, this strategy ensures that these goals are only destroyed and re-established if no solution can be found otherwise. Now, the no-ops first strategy is only a GRAPHPLAN feature, but anyway a reasonable planning strategy should, if many goals are already contained in the initial state, try to simply keep as much of them TRUE as possible.

The soundness of this incremental planning algorithm is obvious because $\mathcal{G}_k = \mathcal{G}$, so we have a sequence of sound subplans yielding a state transition from the initial state $\mathcal{I}$ to a state satisfying $\mathcal{G}$.

The completeness of the approach is less obvious and holds only if the planner cannot make wrong decisions before finally reaching the goals. More specifically, the approach is complete on domains that do not contain *deadlocks* as they were introduced in Definition 14.

Theorem 8 *Given a solvable planning problem $(\mathcal{O}, I, \mathcal{G})$, and a goal agenda $\mathcal{G}_1, \mathcal{G}_2, \ldots \mathcal{G}_k$ with $\mathcal{G}_i \subseteq \mathcal{G}_{i+1}$ and $\mathcal{G}_k = \mathcal{G}$. Running any complete planner in the incremental manner described above will yield a solution if the problem does not contain any deadlocks.*

Proof: Let us assume the planner does not find a solution in step i of the incremental algorithm, i.e., no solution is found for the subproblem $(\mathcal{O}, \mathcal{I}_i, \mathcal{G}_i)$. As the planner is assumed to be complete on each subproblem, this implies unsolvability of $(\mathcal{O}, \mathcal{I}_i, \mathcal{G}_i)$. If this problem is not solvable, then neither is the problem $(\mathcal{O}, \mathcal{I}_i, \mathcal{G})$ solvable, since $\mathcal{G}_i \subseteq \mathcal{G}$ holds. Therefore, the goals cannot be reached from $\mathcal{I}_i$. Furthermore, $\mathcal{I}_i$ is a reachable state—it was reached by executing the partial solution plans $\mathcal{P}_1, \ldots, \mathcal{P}_{i-1}$ in the initial state. Consequently, $\mathcal{I}_i$ must be a deadlock state in the sense of Definition 14, which is a contradiction. ∎

This result states the practical feasibility of our approach: As we have already shown in Section 3.3, almost all benchmark domains that intend to model realistic problems do contain inverse actions, are therefore invertible (Theorem 7), and are with that also deadlock-free (Theorem 6). Thus, with Theorem 8, our approach preserves completeness in these domains.

However in the general case, completeness cannot be guaranteed. The following example illustrates a situation where the assumption $s_A \not\models B$ (cf. the derivation of the ordering $\leq_h$ in Section 3) is wrong and forces a goal ordering under which no plan can be found anymore although the problem is solvable.

Given the initial state $\{A, C\}$ and the goals $\{B, F\}$, the planner has the following set of ground STRIPS actions :

op1: $\{A\} \Rightarrow$ ADD $\{B\}$ DEL $\{C\}$

op2: $\{C\} \Rightarrow$ ADD $\{D\}$
op3: $\{D\} \Rightarrow$ ADD $\{E\}$
op4: $\{E\} \Rightarrow$ ADD $\{F\}$

The analysis will return an ordering $B \leq_h F$ because B is only added by **op1**, but its precondition A is not an effect of any of the other actions. Thus it concludes that A is not reachable from a state in which F holds. But in this example, A holds in all states. The assumption $s_F \not\models A$ as made by the test $p\mathcal{A}(B, \mathcal{O}^*)$, is wrong. Thus, B can be reached after F. On the other hand, $F \leq_r B$ holds, we even have a forced ordering $F \leq_f B$. But when testing for $F \leq_h B$, this ordering remains undetected, because our method does not discover that the precondition E of **op4** is not achievable from the state in which B holds: we obtain $\mathsf{F}^B_{DA} = \{C\}$, which excludes **op2** from $\mathcal{O}^*$, but **op3** and **op4** remain in the set of usable actions. Thus, **op4** is considered a legal achiever of F, and **op3** is considered a legal achiever for its precondition E. We could only detect the right ordering if we regressed over the action chain **op4**, **op3**, **op2** and found out that, with C being in the F set of B, all these actions must be excluded from $\mathcal{O}^*$.

Consequently, the goal agenda $\{B\}, \{E\}$ is fed into the planner, which solves the first subproblem using **op1**, but then fails in achieving E from the state $\{A, B\}$ since there is no inverse action to **op1** and C cannot be re-established in any other way.

5 Empirical Results

We implemented both methods to approximate $\leq_r$ as a so-called Goal Agenda Manager (GAM) for the IPP planning system [KNHD97]. GAM is activated after the set of ground actions has been determined and either uses $\leq_e$ or $\leq_h$ to approximate the reasonable goal ordering. Then it calls the IPP planning algorithm on each entry from the goal agenda and outputs the solution plan as the concatenation of the solution plans that have been found for each entry in the agenda.[6]

The empirical evaluation that we performed uses the IPP domain collection, which contains 48 domains with more than 500 planning problems. Out of these domains, we were able to derive goal ordering information in 10 domains. These domains indeed pose constraints on the ordering in which a planner has to a achieve a set of goals. In all other domains, where no goal orderings could be derived, we found that either only a single goal has to be achieved, for example in the *manhattan*, *movie*, *molgen*, and *montlake* domains or the goals can be achieved in any order, as for example in the *logistics*, *gripper*, and *ferry* domains. We found no domain, in which a natural goal ordering exists, but our method failed in detecting it. As a matter of fact, looking at a goal ordering that seems to be natural, one usually finds that the ordering is reasonable in the sense of Definition 8. For example, all goal orderings in the *blocksworld*, *hanoi*, and *bulldozer* domains are reasonable orderings. Our method finds almost all of the reasonable orderings, which indicates that both approximation techniques $\leq_e$ and $\leq_h$ are appropriate for detecting ordering information.

In the following, we will first compare the $\leq_e$ and $\leq_h$ techniques in terms of running time and number of goal agenda entries generated. Then we take a closer look at the

[6]The source code of GAM and the collection of domains from which we draw the subsequent examples can be downloaded from http://www.informatik.uni-freiburg.de/~ koehler/ipp/gam.html. All experiments have been performed on a SPARC 1/170.

agendas that are generated in selected domains and investigate how they influence the performance of the IPP planning system. The exact definition of all domains can be downloaded from the IPP webpage, we just give the name of the domain and the name of the particular planning problem as well as the number of (ground) actions a domain contains, because this parameter nicely characterizes the size of a domain and with that usually the difficulty to handle it.

In all examples, the times shown to compute the goal agenda do not contain the effort to parse and instantiate the operators, i.e., to compute the set of actions. Times for parsing and instantiation are not listed explicitly, because they are, on the test examples used here, usually very close to Zero and do not influence the performance of the planner in a significant way.

5.1 Comparison of $\leq_h$ and $\leq_e$

We begin our comparison with a summary of results that we obtained in different representational variants of the *blocksworld*. The *bw_large_a* to *bw_large_d* examples originate from the SATPLAN test suite [KS96] to which we added the larger examples *bw_large_e* to *bw_large_g*. The *parcplan* example comes from [EKR96] and uses multiple grippers and limited space on the table. The *stack_n* examples use the GRAPHPLAN *blocksworld* representation and simply require to stack n blocks on each other, which are all on the table in the initial state.

problem	#actions	#agenda entries	CPU($\leq_e$)	CPU($\leq_h$)
bw_large_a	162	1	0.69	0.07
bw_large_b	242	5	1.45	0.11
bw_large_c	450	7	4.85	0.22
bw_large_d	722	11	14.18	0.35
bw_large_e	722	11	12.95	0.35
bw_large_f	1250	6	44.93	0.58
bw_large_g	1800	9	97.11	0.88
parcplan	1960	4	25.84	1.47
stack_20	800	19	6.91	0.36
stack_40	3200	39	160.00	1.74
stack_60	7200	59	840.42	4.85
stack_80	12800	79	-	11.38

Figure 5: Comparison of $\leq_e$ and $\leq_h$ on *blocksworld* problems. #actions shows the number of actions in the set $\mathcal{O}$, from which the planner tries to construct a plan. #agenda entries says how many goal subsets have been detected and ordered by GAM. Column 4 and 5 display the CPU time that is required by both methods to compute the agenda when provided with the set $\mathcal{O}$. A dash will always mean that IPP ran out of memory on a 1 Gbyte machine.

The two methods return exactly the same ordering relations across all *blocksworld* problems. But as Figure 5 confirms, the computation of $\leq_e$ based on planning graphs is much more time-consuming. It hits the computational border when a domain contains

more than 10000 actions. The computation of $\leq_h$ is much faster and also scales to larger action sets.

Figure 6 and Figure 7 show the results for the other domains, in which our method is able to detect reasonable orderings. Figure 6 lists the domains, in which both methods return the same goal agendas. The *tyreworld*, *hanoi*, and *fridgeworld* domains originate from UCPOP [PW92], while the *link-repeat* domain can be found in [VB94]. The performance results coincide with those shown in Figure 5.

domain	problem	#actions	#agenda entries	CPU($\leq_e$)	CPU($\leq_h$)
tyreworld	fixit1	26	6	0.05	0.01
	fixit2	59	6	0.20	0.03
	fixit3	108	6	0.45	0.06
	fixit4	173	6	0.84	0.10
	fixit5	254	6	1.56	0.15
	fixit10	899	6	16.29	0.64
hanoi	hanoi3	48	3	0.05	0.02
	hanoi4	90	4	0.10	0.04
	hanoi5	150	5	0.19	0.08
	hanoi6	231	6	0.35	0.12
	hanoi7	336	7	0.63	0.19
fridgeworld	fridge	779	2	0.77	0.55
link-repeat	link10	31	2	0.19	0.01
	link30	31	2	0.21	0.01

Figure 6: Comparison of $\leq_e$ and $\leq_h$ on those benchmark domains, in which they return identical agendas.

domain	problem	#actions	#agenda entries	CPU($\leq_e$)	CPU($\leq_h$)
bulldozer	bull	61	2/1	0.09	0.03
glassworld	glass1	26	2/1	0.02	0.01
	glass2	114	2/1	0.19	0.09
	glass3	122	2/1	0.22	0.09
shoppingworld	shop	81	2/1	0.07	0.02
scheduling	sched6	104	1/4	01.0	0.12
woodshop	wood1	15	1/3	0.03	0.01
	wood2	15	6/5	0.03	0.01
	wood3	43	6/5	0.14	0.06

Figure 7: Domains in which $\leq_e$ and $\leq_h$ return different goal agendas, which we give in the form n_1/n_2. The number before the slash says how many entries are contained in the agenda computed by $\leq_e$, the number following the slash says how many entries are contained in the agenda computed by $\leq_h$. #agenda entries=1 means that the agenda contains only a single entry, namely the original goal set, and no ordering was derived.

Figure 7 shows the same picture in terms of runtime performance, but in these domains different agendas are returned by $\leq_e$ and $\leq_h$.

The *woodshop* and *scheduling* domains contain actions with conditional effects, while the other domains only use STRIPS operators. The computation of $\leq_e$ fails to derive goal orderings for all *scheduling world* problems (of which we only display the largest problem *sched6*) and for the *wood1* problem. The explanation for this behavior can be found in the different treatment of conditional effects by both methods. IPP does only find a very limited form of *mutex* relations between conditional effects when building the planning graphs. A goal, which is achieved with a conditional effect, will not very often be exclusive to a large number of other facts in the graph. Thus, the F sets are very small or sometimes even empty and consequently, only very few actions can be excluded when performing the reachability analysis and thus, reasonable orderings may remain undetected. Direct analysis investigates the conditional effects in more detail and is therefore able to derive much larger F sets.

The behavior of the $\leq_h$ method in the STRIPS domains *bulldozer*, *glassworld*, and *shopping world* is caused by the same phenomenon. In these domains, one can derive much larger F sets using planning graphs and in turn these sets exclude more actions. Since direct analysis finds smaller or empty F sets, it finds also less $\leq_h$ relations. The *woodshop* domain shows that the results can differ within the same domain, but depending on the specific planning problem. The problem *wood2* contains all goals from the problem *wood1* plus one additional goal. The presence of this additional goal changes the *mutex* relations in the planning graph and yields two additional F sets, based on which more $\leq_e$ relations between the atomic subgoals in *wood2* can be derived. The problem *wood3* contains additional objects and many more goals, which can also be successfully ordered.

In the following experiments, we decided to use solely the heuristic ordering $\leq_h$ because the computation of $\leq_h$ is less costly than the computation of $\leq_e$ in all cases, yielding comparable agendas in most cases. In the three domains that we investigate more closely, namely the *blocksworld*, *tyreworld* and *hanoi* domains, the agendas derived by both methods are, in fact, exactly the same.

5.2 Influence of Goal Orderings on the Performance of IPP and Interaction with RIFO

In this section, we analyze the influence of the goal agenda on the performance of IPP and combine it with another domain analysis method, called RIFO [NDK97]. RIFO is a family of heuristics that enables IPP to exclude irrelevant actions and initial facts from a planning problem. It can be very effectively combined with GAM, because if IPP plans incrementally for only a subset of goals from the original goal set, it is very likely that also only a subset of the relevant actions is needed to find a plan. More precisely, we obtain one subproblem for each entry in the agenda, and, for each such subproblem, use RIFO for preprocessing before planning with IPP. In this configuration, GAM reduces the search space for IPP by decreasing the number of subgoals the planner has to achieve at each moment, while RIFO reduces the search space dramatically by selecting only those actions that are relevant for these goal subsets.

5.2.1 The Blocksworld

Figure 8 illustrates the *parcplan* problem [EKR96] in detail. Seven robot arms can be used to order 10 blocks into 3 stacks on 5 possible positions on the table. The goal agenda

derived by IPP orders the blocks into horizontal layers:

1: on-table(21, t2) ∧ on-table(11, t1)
2: on-table(31, t3) ∧ on(22, 21) ∧ on(12, 11)
3: on(32, 31) ∧ on(13, 12) ∧ on(23, 22)
4: on(14, 13) ∧ on(24, 23)

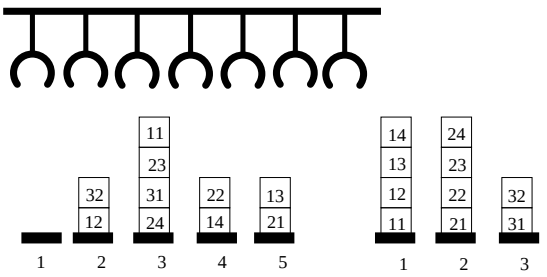

Figure 8: The *parcplan* problem with limited space on the table, seven robot arms, and several stacks.

The optimal plan of 20 actions solving the problem is found by IPP in 14 s, where it spends one second on computing the goal agenda, almost 13 seconds to build the planning graphs, but only 0.01 second to search for a plan. Only 70 actions have to be tried to find the solution. Without the goal analysis, IPP needs approx. 47 s and searches 52893 actions in more than 26 seconds.

RIFO [NDK97] fails in detecting a subset of relevant actions when the original goal set has to be considered, but it succeeds in selecting relevant actions for the subproblems stated in the agenda. It reduces runtime down to less than 8 s with 1 s again spent on the goal agenda, almost 6 s spent on the removal of irrelevant actions and initial facts, less than 1 s spent on building the planning graphs. As previously, almost no time is spent on planning.

Figure 9 shows IPP on the SATPLAN *blocksworld* examples from [KS96], the *bw_large.e* example taken from [DNK97], and two very large examples *bw_large.f* (containing 25 blocks and requiring to build 6 stacks in the goal state) and *bw_large.g* with 30 blocks/8 stacks.

IPP 3.3 without GAM can only solve the *bw_large.a* and *bw_large.b* problems. The *bw_large.a* problem is solved in approx. 2 s. For *bw_large.b*, the optimal solution of 18 actions is found in approx. 129 s. Using a goal agenda, the plans become slightly longer, but performance is increasing dramatically. Plan length is growing because some blocks are accidentally put in positions where they cut off goals that are still ahead in the agenda and thus, additional actions need to be added to the plan to remove blocks from wrong positions. A further speed-up is possible when RIFO is additionally used, because it reduces the size of planning graphs dramatically. Finally, goals that belong to the same subset in the agenda can be linearized based on the heuristic assumption that if the analysis found no reasonable goal orderings then the goals are achievable in any order. With this option, the problems are solved almost instantly.[7]

[7]The reader may wonder at this point why we use linearization of agenda entries only as an extra option and do not investigate it further. The reason is that linearization does have negative side effects in most domains that we investigated. For example, it yields much longer plans in the *logistics* domain

SATPLAN	# actions	plan length	+G	+G+R	+G+R+L
bw_large.a	162	12	0.74	0.58	0.34
bw_large.b	242	22	0.86	0.55	0.52
bw_large.c	450	48	7.34	2.42	2.58
bw_large.d	722	54	11.62	3.74	3.81
bw_large.e	722	52	11.14	3.99	3.97
bw_large.f	1250	90	-	-	16.01
bw_large.g	1800	84	-	117.56	28.71

Figure 9: Performance on the extended SATPLAN *blocksworld* test suite. The second column shows the plan length, i.e., the number of actions contained in the plan. +G means that IPP is using the goal agenda manager, +G+R means IPP uses the goal agenda manager and RIFO, +G+R+L means that subgoals from the same set in the agenda are arbitrarily linearized. All runtimes cover the whole planning process starting with parsing the operator and domain file, performing the GAM and RIFO analysis (if active), and then searching the graph until a plan is found.

Figure 10 shows IPP on the *stack_n* problems. IPP without any domain analysis can handle up to 12 blocks in less than 5 minutes, but for 13 blocks more than 15 minutes are needed. Using GAM, 40 blocks can be stacked in less than 5 minutes. Using GAM and RIFO, the 5 minutes limit is extended to 80 blocks, while *stack100* is solved in 11.5 min where 11.3 min are spent for both analysis methods and only 0.2 min are needed for building planning graphs and the search for a plan.

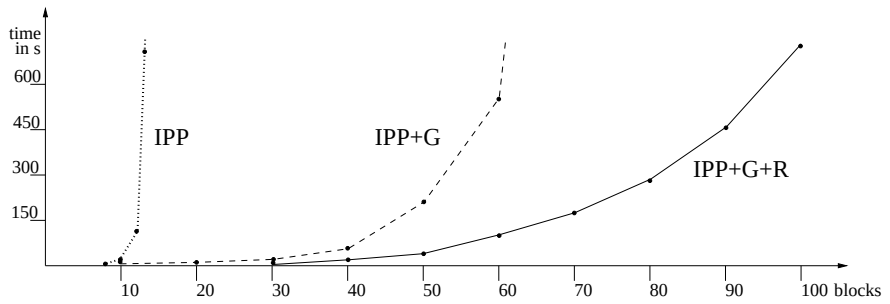

Figure 10: IPP 3.3 on a simple, but huge stacking problem.

Figure 11 shows the sharing of the overall problem-solving time between GAM, RIFO and the IPP search algorithm on *blocksworld* problems. Similar results are obtained in the *tyreworld*. GAM takes between 3 and 16 %, RIFO takes between 75 and 96 %, and the search effort is reduced down to approx. 1 %. The overall problem solving time is clearly determined by RIFO, while the search effort becomes a marginal factor in the determination of performance. This indicates that a further speed-up is possible when improving the performance of GAM and RIFO. It also indicates that even the hardest planning problems can become easy if they are structured and decomposed in the right way.

and all its variants. When linearizing the single entry that the agenda for a *logistics* problem contains, all packages get transported to their goal position one by one. Of course, this takes much more planning steps than simultaneously transporting packages with coinciding destinations.

problem	# actions	GAM	RIFO	search algorithm
stack_20	800	0.31 = 16 %	1.44 = 75 %	0.13 = 7 %
stack_40	3200	1.57 = 7 %	18.77 = 90 %	0.51 = 2 %
stack_60	7200	4.40 = 4 %	93.10 = 94 %	1.15 = 1 %
stack_80	12800	9.60 = 3 %	283.60 = 96 %	2.33 = 1 %
parcplan	1960	0.86 = 12 %	5.52 = 76 %	0.83 = 11 %

Figure 11: Distribution of problem-solving time on *blocksworld* examples between GAM, RIFO, and the search algorithm, which comprises the time to build and search the planning graph. The remaining fraction of total problem-solving time, which is not shown in the table, is spent on parsing and instantiating the operators.

5.2.2 The Tyreworld

The tyreworld problem, originally formulated by Stuart Russell, asks a planner to find out how to replace a flat tire. It is easily solved by IPP within a few milliseconds. The problem becomes much harder if the number of flat tires is increasing as Figure 12 shows.

Tires	# actions	IPP	+G+R	+G+R+L	Search Space
1	26	0.10 (12/19)	0.15 (14/19)	0.16 (17/19)	1298/88
2	59	17.47 (18/30)	0.41 (24/32)	0.32 (30/34)	1290182/210
3	108	-	2.87 (32/44)	0.63 (41/46)	-/366
4	173	-	-	1.12 (52/60)	-/565
5	254	-	-	1.93 (63/73)	-/807
6	353	-	-	3.42 (73/85)	-/1092
7	464	-	-	4.81 (84/98)	-/1420
8	593	-	-	8.07 (95/121)	-/1791
9	738	-	-	11.27 (106/124)	-/ 2205
10	899	-	-	16.80 (118/136)	-/2662

Figure 12: IPP in the Tyreworld. The numbers in parentheses show the time steps, followed by the number of actions in the generated plan. The last column compares the search spaces. The number before the slash shows the "number of actions tried" parameter for the plain IPP planning algorithm, while the number following the slash shows the "number of actions tried" for IPP using GAM, RIFO, and the linearization of entries in the agenda. A dash means that the "number of actions tried" is unknown because IPP failed in solving the corresponding planning problem.

IPP is only able to solve the problem for 1 and 2 tires. Using GAM and RIFO, 3 tires can be handled. Solution length under GAM is slightly increasing, which is caused by superfluous *jack-up* and *jack-down* actions. In short, this is explained as follows. Each wheel needs to be mounted on its hub, which is expressed by an $on(?r, ?h)$ goal. To mount a wheel, its hub must be jacked up. After mounting, the nuts are done up. Then, the hub needs to be jacked down again, in order to tighten the nuts, achieving a $tight(?n, ?h)$ goal. Now, GAM puts all of the *on* goals one entry ahead of the *tight* goals. Thus, solving the entry containing the *on* goals, each hub is jacked up, the wheel is put on, and the hub

is immediatly jacked down again in order to replace the next wheel. Afterwards, solving the *tight* goals, each hub must be jacked up—and down—one more time for doing up the nuts. Solving the problem in this manner, the planner inserts one superfluous *jack-up*, and one superfluous *jack-down* action for each wheel.[8] In the case of 3 tires, the following goal subsets are identified and ordered:

1: inflated(r3), inflated(r2), inflated(r1)
2: on(r3, hub3), on(r1, hub1), on(r2, hub2)
3: tight(n2, hub2), tight(n3, hub3), tight(n1, hub1)
4: in(w3, boot), in(pump, boot), in(w1, boot), in(w2, boot)
5: in(jack, boot)
6: in(wrench, boot)
7: closed(boot)

The hardest subproblem in the agenda is to achieve the $on(r_i, hub_i)$ goals in entry 2, i.e., to mount inflated spare wheels on the various hubs. Trying to generate a maximum parallelized plan is impossible for IPP for more than 3 tires. But since the goals are completely independent of each other, any linearization of them will perfectly work. [9] Thus using +G+R+L in the case of 10 tires, only 2662 actions need to be tried until a plan of 136 actions is found, which takes 0.08 s. GAM requires 0.55 s, RIFO requires 14.42 s, 1.74 s are needed to generate the planning graphs, and 0.08 s are spent to compute the initial states for all subproblems. The remaining 0.02 s are needed for parsing and instantiating.

5.2.3 The Tower of Hanoi

A surprising result is obtained in the *tower of hanoi* domain. In this domain, a stack of discs has to be moved from one peg to a third peg with an auxiliary second peg between them, but never a larger disc can be put onto a smaller disc. In the case of three discs $d1$, $d2$, $d3$ of increasing size, the goals are stated as *on(d3,peg3), on(d2,d3), on(d1,d2)*. GAM returns the following agenda, which correctly reflects the ordering that the largest disc needs to be put in its goal position first.

1: on(d3,peg3)

[8] More precisely, superfluous *jack-up* and *jack-down* actions are inserted for all but one wheel, namely the wheel that is last mounted when solving the *on* goals. After mounting this wheel, all *on* goals are achieved, and the planner proceeds to the next agenda entry with this wheel still being jacked up. Then, trying to achieve the *tight* goals, IPP recognizes that the shortest plan (in terms of the number of parallel steps) results when the nuts are first done up on the hub that is already jacked up. Thus, this hub is only jacked up *one* time, achieving the corresponding *on* goal, and jacked down again one time, before achieving its *tight* goal.

[9] The plans become slightly longer due to the way that the *tight* goals are achieved when using the -L option. We noticed earlier that for one wheel (the one that is last mounted when solving the *on* goals) no superfluous *jack-up* and *jack-down* actions need to be inserted into the plan. Linearizing the agenda entries, superfluous *jack-up* and *jack-down* actions must most likely be inserted for *all* wheels, yielding plans that are two steps longer. The reason for that is that any *tight* goal might be the first in the linearization. Most likely, this is not the *tight* goal corresponding to the hub that is still jacked up, so the planner needs to insert one superfluous *jack-down* action here. Later, it must jack this hub up again, yielding another superfluous action.

discs	#actions	IPP	IPP +G	UCPOP	UCPOP on subproblems
2	21	0.02	0.02	0.12 (27)	0.06 (17) + 0.02 (6)
3	48	0.08	0.07	8.00 (2291)	0.18 (48) + 0.06 (13) + 0.01 (6)
4	90	0.33	0.25	-	-
5	150	1.57	3.10	-	-
6	231	9.71	88.45	-	-
7	336	69.44	2339.94	-	-

Figure 13: Runtimes of IPP with and without the goal agenda on *hanoi* problems compared to UCPOP without agenda and UCPOP on the agenda subproblems.

2: on(d2,d3)
3: on(d1,d2)

The goal agenda leads to a partition into subproblems that corresponds to the recursive formulation of the problem solving algorithm, i.e., to solve the problem for n discs, the planner first has to solve the problem for $n-1$ discs, etc. For the first entry, a plan of 4 actions (time steps 0 to 3 below) is generated, which achieves the goal *on(d3,peg3)*.[10] Then a plan of 2 actions (time steps 4 and 5) achieves the goals *on(d3,peg3)* and *on(d2,d3)* with *on(d3,peg3)* holding already in the initial state. Finally, a one-step plan (time step 6) is generated that moves the third disc with the other two discs being already in the goal position.

```
time step 0:  move(d1,d2,peg3)
time step 1:  move(d2,d3,peg2)
time step 2:  move(d1,peg3,d2)
time step 3:  move(d3,peg1,peg3)
time step 4:  move(d1,d2,peg1)
time step 5:  move(d2,peg2,d3)
time step 6:  move(d1,peg1,d2)
```

Surprisingly, IPP is not able to benefit from this information, but runtime of IPP using GAM is exploding dramatically for increasing number of discs, see Figure 13. We are not able to provide an explanation for this phenomenon, but the division into subproblems causes a much larger search space for the planner although the same solution plans result.[11] RIFO cannot improve on the situation because it selects all actions as relevant.

The *tower of hanoi* domain is the only one we found where IPP gets slowed down when using GAM. We do currently not see a way of how one can tell in advance whether IPP will gain an advantage from using GAM or not. However, the phenomenon seems to be specific to IPP. We simulated the information that is provided by GAM in UCPOP and obtained a quite different picture. The fifth column in Figure 13 shows the runtime of UCPOP with the number of explored partial plans in parentheses. UCPOP can only solve the problem for 2 and 3 discs. In the last column of the figure, we show the runtime and number of explored partial plans, which result when UCPOP is run on the subproblems

[10] A *move* action takes as first argument the disc to be moved, as second the disc from which it is moved, and as third argument the disc or peg to which it is moved.

[11] See also Section 6 for a discussion of related complexity results.

that result from the agenda. These are exactly the same which IPP has to solve. In the case of three discs, UCPOP first has to move the large disc into its goal position, second it has to move the middle disc on top of the large disc being already in the goal position. Finally, it has to move the smallest disc on the middle disc, with both the middle and the large disc being in the goal position. The performance improves significantly. Instead of taking 8 s and exploring 2291 partial plans, UCPOP only takes 0.18+0.06+0.01=0.25 s and explores only 48+13+6=67 plans. Unfortunately, any problems or subproblems with more than 3 discs remain beyond the performance of UCPOP. Knoblock [Kno94] also reports an improvement in performance for the Prodigy planner [FV94] when it is using the abstraction hierarchy generated for this domain by the ALPINE module.

6 Summary and Comparison to Related Work

Many related approaches have been developed to provide a planner with the ability to decompose a planning problem by giving it any kind of goal ordering information. Subsequently, we discuss the most important of them and review our own work in the light of these approaches.

Our method introduces a preprocessing approach, which derives a total ordering for subsets of goals by performing a static, heuristic analysis of the planning problem at hand. The approach works for domains described with STRIPS or ADL operators and is based on polynomial-time algorithms. The purpose of this method is to provide a planner with search control, i.e., we opt at deriving a goal selection order and then incrementally call the planner on the totally ordered subsets of goals.

The method preserves the soundness of the planning system, but the completeness only in the case that the planning domain does not contain deadlocks. We argue that real-world domains will usually possess this property, which is also supported by the discussion in [WN97].

The computation of $\leq_h$ and $\leq_e$ requires only polynomial time, but both methods are incomplete in the sense that they will not detect all reasonable goal orderings in the general case. The complexity of deciding on the existence of *forced* and *reasonable* goal orderings has been proven to be PSPACE-hard in Section 2 and therefore, trading completeness for efficiency seems to be an acceptable solution. Our complexity results relate to those found in [Byl92] who proves the PSPACE-completeness of serial decomposability [Kor87]. Given a set of subgoals, *serial decomposability* means that previously satisfied subgoals do not need to be violated later in the solution path, i.e., once a subgoal has been achieved, it remains valid until the goal is reached. The purpose of our method is to derive constraints that make those orderings explicit under which no serial decomposability of a set of goals can be found, i.e., we consider the complementary problem, which is also reflected in our complexity proofs.

In many cases, we found that the goal agenda manager can significantly improve the performance of the IPP planning system, but we found at least one domain, namely the *tower of Hanoi*, where a dramatic decrease in performance can be observed although IPP still generates the optimal plan when processing the ordered goals from the agenda. Such a behavior has been predicted in [BJ95], who showed that planning with abstraction hierarchies can be exponentially less efficient in the worst case.

The idea to analyze the effects and preconditions of operators and to derive ordering

constraints based on the interaction of operators can also be found in a variety of approaches. While we analyze harmful interactions of operators in our method by studying the delete effects, the approaches described in [DS77, Kor85, Kno94] concentrate on the positive interactions between operators. The successful matching of effects to preconditions forms the basis to learn *macro-operators* in [DS77, Kor85].

The ALPINE system described in [Kno94] learns abstraction hierarchies for the Prodigy planner [FV94]. The approach is based on an ordering of the preconditions and the effects of each operator, i.e., all effects of an operator must be in the same abstraction hierarchy and its preconditions must be placed at the same or a lower level than its effects. This introduces an ordering between the possible subgoals in a domain, which is orthogonal to the ordering we compute: In ALPINE, a subgoal A is ordered before a subgoal B if A *enables* B, i.e., A must be *possibly* achieved first in order to achieve B. Our method orders A before B if A cannot be achieved without *necessarily* destroying B. The result of ALPINE and GAM are a set of binary constraints. In the case of ALPINE, the constraints are computed between all atoms in a domain, while GAM restricts the analysis to the goals only. Both approaches represent the binary constraints in a graph structure. ALPINE merges atomic goals together if they belong to a strongly connected component in the graph. GAM merges sets of goals together if they have identical degree. Then they both compute a topological sorting of the sets that is consistent with the constraints. The resulting goal orderings can be quite similar as the examples from [Kno94] demonstrate, but a detailed empirical comparison of when both approaches yield similar or differing results is beyond the scope of this paper.[12]

An analysis, which is quite similar to ALPINE, but which is performed in the framework of HTN planning, describes [THN98]. The approach analyzes the *external* conditions of methods, which cannot be achieved when decomposing the method further. This means, such conditions have to be established by the decomposition of those methods, which *precede* the method using this external condition. Two strategies to determine the decomposition order of methods are defined and empirically compared. Here lies the main difference to the other approaches described so far: Instead of trying to automatically construct the decomposition orderings, they are predefined and fixed for all domains and problems.

Harmful interactions among operators are studied in [SP93] and [Etz93]. A *threat* of an operator o to a precondition p as defined in [SP93] occurs if there is an instantiation of o such that its effects are inconsistent with p. The knowledge about threats is used to control a plan-space planner. In contrast to a state-space planner such as IPP, computing an explicit ordering of goals does not prevent the presence of threats in a partial plan because the order in which the goals are processed does not determine the order in which actions occur in the plan. The notion of forced and reasonable goal orderings is not comparable to that of a threat because a threat still has the potential of being resolved by adding binding or ordering constraints to the plans. In contrast to this, a forced or reasonable goal ordering persists under all bindings and enforces a specific ordering of the subgoals.

Given a planning problem, STATIC [Etz93] computes a backchaining tree from the goals

[12] Abstraction hierarchies are more general than the goal orderings we compute. They cannot only serve for the purpose of providing a planner with goal ordering information, but also allow to generate plans at different levels of refinement. Two other approaches generating abstraction hierarchies based on numerical *criticality* values are [Sac74, BGSW96].

in the form of an AND/OR graph, which it subsequently analyses for the occurrence of goal interactions that will *necessarily* occur. This analysis is much more complicated than ours, because STATIC has to deal with uninstantiated operators and axioms, which describe properties of legal states. The result of the analysis are goal ordering rules, which order goals if certain conditions are satisfied in a state. This is the main difference to GAM, which generates explicit goal orderings independently of a specific state. It does not need to extract conditions that a specific state has to satisfy because it considers the generic state $s_{(A,\neg B)}$ in the analysis, which represents all states satisfying A, but not B. As GAM, STATIC is incomplete in the sense that it cannot detect all existing goal interactions. The problem for GAM is that deciding reasonable orderings is PSPACE-hard, as we have proven in this paper. The problem for STATIC is that it has to compute the *necessary* effects of an operator in a given state. As [Etz93] conjectures and [NB94] prove, this problem is computationally intractable and therefore, any polynomial-time analysis method must be incomplete.

Last, but not least there have been quite a number of approaches in the late Eighties, which focused directly on subgoal orderings. These fall into two categories: The approaches described in [DC89, HH89] focus on the detection of conflicts caused by goal interdependencies to guide a partial order planner during search. We do not investigate these approaches in more detail here, as they do not extract explicit goal orderings as a preprocess to planning, like we do. The works described in [IC87, CI89, JR90] implement preprocessing approaches, which perform a structural analysis of the planning task to determine an appropriate goal ordering before planning starts. [IC87] compute a relation $\prec$ between pairs of goals, which—roughly speaking—orders a goal A after a goal B if B must be achieved before A can be achieved. Their formalism is rather complicated and the theoretical properties of the relation are not investigated. In [CI89], the approach is extended such that *sets* of goals can be ordered with respect to each other. The exact properties of the formalism remain unclear. In [JR90], a graph-theoretical approach is described that generates a graph with all atoms from a given domain description as nodes and draws an arc between a node A and a node B if an operator exists that takes A as precondition and has B as an effect. When assuming that all operators have inverse counterparts, identifying connected components in the graph is proposed as a means to order goals. The approach is unlikely to scale to the size of state spaces today's planners consider and it is also completely outdated in terms of terminology.

Finally, one can wonder how the reasonable and forced goal orderings relate to others defined in the literature. There is only one attempt of which we know where an ordering relation is explicitly defined and its properties are studied, see [HMAW99]. In this paper, the notion of *necessary* goal orderings is introduced, which must be true in all *minimal solution plans* [Kam95].[13] The approach uses extended operator graphs based on [SP93] and orders a goal based on three criteria called *goal subsumption*, *goal clobbering*, and *precondition violation*. Goal subsumption $A < B$ holds if every solution plan achieving a goal B in a state s also achieves a goal A in a state s' preceding s, and no plan achieving one of the goals in $\mathcal{G} \setminus \{A\}$ deletes A. Goal clobbering holds if any solution plan for A deletes B and thus, $A < B$. Precondition violation holds if any solution for B results in a deadlock from which A cannot be reached anymore, i.e., again $A < B$. A composite

[13] A plan is minimal if it contains no subplan that is also a solution plan. We remark that minimality does not mean that only shortest plans having the least number of actions are considered. In fact, minimal plans can be highly non-optimal as long as no action is truly superfluous.

criterion is defined that tests all three criteria simultaneously.[14] A goal A is *necessarily ordered* before B if it satisfies the composite criterion.

We remark that *precondition violation* seems to be equivalent to the *forced* orderings we introduced, while *goal clobbering* appears to be similar to our *reasonable* orderings. It is not possible for us to verify this conjecture as the authors of [HMAW99] do not give exact formal definitions. We have nothing similar to *goal subsumption* and we argue that this criterion will be rarely satisfied in natural problems: if a goal A is achieved by every solution for a goal B anyway, then the goal A can be removed from the goal set without changing the planning task.

The authors report that they are able to detect *necessary* orderings in the artificial domains D_iS_i, cf. [BW94], but fail in natural domains such as the *blocksworld* or the *tyreworld*. The reason for this seems to be that their operator graphs do not represent all possible instantiations of operator schemes. As the authors claim, this makes operator graph analysis very efficient. However, the heuristic ordering $\leq_h$ that we introduced in this paper also takes almost no computation time, and succeeds in finding the goal orderings in natural domains.

7 Outlook

Two promising avenues for future research are as follows. First, one can imagine that goal ordering information is also used *during* the search process, i.e., by not only ordering the original goal set, but also other goals that emerge during search. The major challenge seems to balance the effort on computing the goal ordering information with the savings that can result for the search process. One can easily imagine that ordering all goal sets that are ever generated can become a quite costly investment without yielding a major benefit for the planner.

Second, the refinement of the goal agenda with additional subgoals is another interesting future line of work. A first investigation using so-called *intermediate goals* (these are facts that the planner must make TRUE before it can achieve an original goal) has been explored inside GAM and the results are reported in [KH98]. Earlier work addressing the task of learning intermediate goals can be found in [RK89], but this problem has not been in the focus of AI planning research since then.

Currently, we are exploring a third line of work, which addresses the interaction of GAM with a planning system. We have seen that GAM preserves correctness of a planner, and that it preserves completeness, at least on deadlock-free planning problems. We have also seen, however, that solution plans using GAM can get longer, i.e., GAM does *not* preserve optimality of a planner. Now, recent times have seen some success of planning systems that do not search for optimal plans anyway, like the HSP system (first described in [BLG97]) or GRT [RV99]. Both of these are heuristic search planners, searching forward in state space with non-admissible (but informative) heuristics.

It seems to be an interesting piece of work to find out how this kind of planning algorithms interact with GAM. In addition to the smaller problems they are facing when using the goal agenda, their heuristics will be influenced: both use some technique for estimating the goal distance from a state. With the goal agenda, the goals for which they

[14] Here, the authors are not very precise what they mean with this. We argue that this means that two goals are ordered if they satisfy at least one of the criteria.

have to estimate this distance will be different at each stage of the planning process. We are at the moment developing our own heuristic search planning system, and have already been able to draw great benefits from the goal agenda guidance information.

References

[AAA96] *Proceedings of the 14th National Conference of the American Association for Artificial Intelligence*. AAAI Press, 1996.

[All98] J. Allen, editor. *Proceedings of the 4th International Conference on Artificial Intelligence Planning Systems*. AAAI Press, Menlo Park, 1998.

[AW98] C. Anderson and D. Weld. Conditional effects in Graphplan. In Allen [All98], pages 44–53.

[BF95] A. Blum and M. Furst. Fast planning through planning graph analysis. In IJCAI-95 [IJC95], pages 1636–1642.

[BF97] A. Blum and M. Furst. Fast planning through planning graph analysis. *Artificial Intelligence*, 90(1–2):279–298, 1997.

[BGSW96] A. Bundy, F. Giunchiglia, R. Sebastiani, and T. Walsh. Computing abstraction hierarchies by numerical simulation. In AAAI-96 [AAA96], pages 523–529.

[BJ95] C. Bäckström and P. Jonsson. Planning with abstraction hierarchies can be exponentially less efficient. In IJCAI-95 [IJC95], pages 1599–1604.

[BLG97] B. Bonet, G. Loerincs, and H. Geffner. A robust and fast action selection mechanism for planning. In *Proceedings of the 14th National Conference of the American Association for Artificial Intelligence*, pages 714–719, 1997.

[BW94] A. Barrett and D. Weld. Partial-order planning: Evaluating possible efficiency gains. *Artificial Intelligence*, 67:71–112, 1994.

[Byl92] T. Bylander. Complexity results for serial decomposability. In *Proceedings of the 10th National Conference of the American Association for Artificial Intelligence*, pages 729–734, San Jose, CA, July 1992. MIT Press.

[Byl94] T. Bylander. The computational complexity of propositional STRIPS planning. *Artificial Intelligence*, 69:165–204, 1994.

[Cha87] D. Chapman. Planning for conjunctive goals. *AIJ*, 32(3):333–377, 1987.

[CI89] J. Cheng and K. Irani. Ordering problem subgoals. In IJCAI-89 [IJC89], pages 931–936.

[DC89] M. Drummond and K. Currie. Goal ordering in partially ordered plans. In IJCAI-89 [IJC89], pages 960–965.

[DNK97] Y. Dimopoulos, B. Nebel, and J. Koehler. Encoding planning problems in non-monotonic logic programs. In Steel [Ste97], pages 169–181.

[DS77] C. Dawsson and L. Siklossy. The role of preprocessing in problem solving systems. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence*, pages 465–471, Cambridge, MA, August 1977.

[EKR96] A. El-Kholy and Barry Richards. Temporal and resource reasoning in planning: the parcPLAN approch. In W. Wahlster, editor, *Proceedings of the 12th European Conference on Artificial Intelligence*, pages 614–618. John Wiley & Sons, Chichester, New York, 1996.

[Etz93] O. Etzioni. Acquiring search-control knowledge via static analysis. *Artificial Intelligence*, 62:255–301, 1993.

[FL98] M. Fox and D. Long. The automatic inference of state invariants in TIM. *Journal of Artificial Intelligence Research*, 9:367–421, 1998.

[FL99] M. Fox and D. Long. Efficient implementation of the plan graph in STAN. *Journal of Artificial Intelligence Research*, 10:87–115, 1999.

[FV94] E. Fink and M. Veloso. Prodigy planning algorithm. Technical Report CMU-94-123, Carnegie Mellon University, 1994.

[HH89] J. Hertzberg and A. Horz. Towards a theory of conflict detection and resolution in nonlinear plans. In IJCAI-89 [IJC89], pages 937–942.

[HMAW99] J. Hüllem, H. Munoz-Avila, and F. Weberskirch. Extracting goal orderings to improve partial-order and Graphplan-based planning. Technical report, University of Kaiserslautern, 1999.

[IC87] K. Irani and J. Cheng. Subgoal ordering and goal augmentation for heuristic problem solving. In D. McDermott, editor, *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 1018–1024, Milan, Italy, August 1987. Morgan Kaufmann.

[IJC89] *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, MI, August 1989. Morgan Kaufmann.

[IJC95] *Proceedings of the 14th International Joint Conference on Artificial Intelligence*. Morgan Kaufmann, San Francisco, CA, 1995.

[JR90] D. Joslin and J. Roach. A theoretical analysis of conjunctive-goal problems. *Artificial Intelligence*, 41:97–106, 1990.

[Kam95] S. Kambhampati. Admissible pruning strategies based on plan minimality for plan-space planning. In IJCAI-95 [IJC95], pages 1627–1633.

[KH98] J. Koehler and J. Hoffmann. Planning with goal agendas. Technical Report 110, University of Freiburg, 1998. available at http://www.informatik.uni-freiburg.de/~ koehler/ipp.html.

[KNHD97] J. Koehler, B. Nebel, J. Hoffmann, and Y. Dimopoulos. Extending planning graphs to an ADL subset. In Steel [Ste97], pages 273–285.

[Kno94] C. Knoblock. Automatically generating abstraction for planning. *Artificial Intelligence*, 68(2):243–302, 1994.

[Kor85] R. Korf. Macro-operators: A weak method for learning. *Artificial Intelligence*, 26:35–77, 1985.

[Kor87] R. Korf. Planning as search: A quantitative approach. *Artificial Intelligence*, 33:65–88, 1987.

[KS96] H. Kautz and B. Selman. Pushing the envelope: Planning, propositional logic, and stochastic search. In AAAI-96 [AAA96], pages 1194–1201.

[NB94] B. Nebel and C. Bäckström. On the computational complexity of temporal projection, planning, and plan validation. *Journal of Artificial Intelligence*, 66(1), 1994.

[NDK97] B. Nebel, Y. Dimopoulos, and J. Koehler. Ignoring irrelevant facts and operators in plan generation. In Steel [Ste97], pages 338–350.

[Ped89] E. Pednault. ADL: Exploring the middle ground between STRIPS and the Situation Calculus. In R. Brachman, H.J. Levesque, and R. Reiter, editors, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, pages 324–332, Toronto, Canada, 1989. Morgan Kaufmann.

[PW92] J. Penberthy and D. Weld. UCPOP: A sound, complete, partial order planner for ADL. In B. Nebel, W. Swartout, and C. Rich, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, pages 103–113. Morgan Kaufmann, San Mateo, 1992.

[RK89] D. Ruby and D. Kibler. Learning subgoal sequences for planning. In IJCAI-89 [IJC89], pages 609–615.

[RN95] S. Russel and P. Norvig. *Artificial Intelligence - A modern Approach*. Prentice Hall, 1995.

[RV99] I. Refanidis and I. Vlahavas. Grt: A domain independent heuristic for strips worlds based on greedy regression tables. In *Proceedings of the 5th European Conference on Planning*, pages 346–358, 1999.

[Sac74] E. Sacerdoti. Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 5:115–135, 1974.

[SP93] D. Smith and M. Peot. Postponing threats in partial-order planning. In *Proceedings of the 11th National Conference of the American Association for Artificial Intelligence*, pages 500–506. AAAI Press, MIT Press, 1993.

[Ste97] S. Steel, editor. *Proceedings of the 4th European Conference on Planning*, volume 1348 of *LNAI*. Springer, 1997.

[THN98] R. Tsuneto, J. Hendler, and D. S. Nau. Analyzing external conditions to improve the efficiency of htn planning. In Allen [All98], pages 913–920.

[VB94] M. Veloso and J. Blythe. Linkability: Examining causal link commitments in partial-order planning. In K. Hammond, editor, *Proceedings of the 2nd International Conference on Artificial Intelligence Planning Systems*, pages 170–175. AAAI Press, Menlo Park, 1994.

[War62] J. Warshall. A theorem on boolean matrices. *Journal of the ACM*, 9(1):11–12, 1962.

[WN97] B. Williams and R. Nayak. A reactive planner for a model-based executive. In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, pages 1178–1185. Morgan Kaufmann, San Francisco, CA, 1997.