On the Instantiation of ADL Operators Involving Arbitrary First-Order Formulas

Jana Koehler
Schindler Lifts Ltd.
CH-6031 Ebikon
Switzerland
jana_koehler@ch.schindler.com

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

Abstract

The generation of the set of all ground actions for a given set of ADL operators, which are allowed to have conditional effects and preconditions that can be represented using arbitrary first-order formulas is a complex process which heavily influences the performance of any planner or pre-planning analysis method.

The paper describes a sophisticated instantiation procedure that determines so-called *inertia* in a given problem representation and uses them to perform simplifications of formulas during the instantiation process. As a result, many inapplicable actions are detected and ruled out from the domain representation yielding a much smaller search space for the planner.

Introduction

A planning system that handles a more expressive language than STRIPS requires sophisticated algorithmic solutions to quite a number of problems, which have nothing to do with the actual search process for a plan. One of these problems concerns the computation of the set of *actions* as all ground instances of a given set of operators.

The aim of the instantiation process is to generate all those ground instances of the planning operators that are applicable in some legal world state. This means, that the precondition of the operator should be satisfiable and its effects should be consistent. On one hand, a naive instantiation procedure that simply expands logical quantifiers and enumerates all possible instantiations of operator parameters will quickly render even simple planning problems unsolvable. On the other hand, a rather sophisticated instantiation procedure can rule out many actions, which will never be applicable in any reachable world state or that would—if applied—yield an inconsistent state. It should also return the most simple syntactic representation of preconditions and effects.

Many planning systems do generate the complete set of actions before planning actually starts. They use this set either for the encoding of the domain in other representation formalisms such as SAT (Kautz & Selman 1996) or for the derivation of useful information that can help during planning, e.g., distance heuristics (Geffner 1999), symmetries (Fox & Long 1999), relevant actions (Nebel, Dimopoulos, & Koehler 1997), and goal orderings (Koehler 1998).

When using the PDDL language (McDermott & others 1998) to represent ADL operators (Pednault 1989), quite complex descriptions of preconditions and effects are possible:

- arbitrary function-symbol free first-order logic formulas represent preconditions,
- conditional effects have the form *(when antecedent consequent)* where the antecedent can be an arbitrary precondition and the consequent is a conjunction of literals, i.e., an atom either occurs positively or negatively in it. A conditional effect can also be universally quantified.

Given such an operator, the instantiation has to replace all occurring variables, which are either quantified or occur as parameters of the operator, by those type-consistent constants, which have been declared in the planning problem. In order to replace all variables by constants, the instantiation process proceeds in three phases:

1. The expansion of universal and existential quantifiers occurring in the first-order formulas representing preconditions or antecedents of conditional effects eliminates most of the quantified variables,
2. The expansion of universally quantified conditional effects eliminates the remaining quantified variables,
3. the instantiation of operator parameters eliminates the variable parameters.[1]

In each phase, the following *atomic instantiation task* occurs:

> *Given a variable $?x$, a constant c and an atomic formula p, determine the resulting instantiation $p[?x/c]$.*

[1] In IPP the assumption is made that different operator parameters are instantiated with different constants, i.e., the planner never generates actions like *move(a,a)*.

Copyright © 2000, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

This is a trivial problem per se. But after having determined $p[?x/c]$ one can sometimes simplify this atomic formula to FALSE or TRUE, which in turn often leads to a further simplification of the operator representation. The paper addresses exactly this problem. We describe what kind of *atomic simplifications* are performed in IPP (Koehler *et al.* 1997) under which conditions, how this process can be efficiently implemented and how it affects the search space of the planning system. The techniques have successfully been used in the 1998 AIPS planning competition where IPP demonstrated a convincing performance across a variety of STRIPS and ADL domains.

The paper is organized as follows: First, we give an overview of the three phases of the instantiation process. Then we define the notion of *inertia* predicates and describe how the knowledge about inertia is used to perform *atomic simplifications*. We prove their soundness and describe how the underlying tests can be efficiently implemented. In the second part of the paper we define how atomic simplifications can be propagated over the operator description to further simplify the operator representation. We show how unary inertia relations can be encoded as *types* to speed up the instantiation process. Finally, the impact of the instantiation process on the search space of IPP is demonstrated.

Overview over the Instantiation Process

After having parsed the domain and problem file into some appropriate data structure, a basic preprocessing step renames all variables in the logical formulas and assigns unique names to them. For example, the formula $\psi(?x) \wedge \forall ?x\, \varphi(?x)$ is equivalently transformed into $\psi(?x_1) \wedge \forall ?x_2\, \varphi(?x_2)$. Then code tables are generated, which map strings to unique numbers, i.e., we obtain one number for each predicate name, variable name, and constant name. Internally, all subsequently described operations work over trees of numbers representing the formulas.

Figure 1 shows the precondition of the **remove** operator from the *assembly* domain (McDermott 1998) with the quantifiers in frames and the underlined *requires* predicate, which will be used throughout this paper to illustrate the instantiation process. This predicate has two arguments, the first one *?whole* being an operator parameter of type *assembly* and the second one *?res* being a universally quantified variable of type *resource*.

The schematic tree-like representation of this first-order formula is shown in Figure 2. The leaves of the tree contain the atomic formulas. IPP's instantiation process traverses the tree top-down and expands quantifiers one after the other, i.e., it reaches the first quantifier *forall (?res - resource)* and extracts the variable *?res* together with its type *resource*. From the problem file, IPP knows all constants of this type. These are now used to instantiate *?res*.

The process considers all constants one after the other. For each constant, a copy of the subtree representing the quantified formula is generated. In the

```
:action remove
:parameters (?part ?whole - assembly)
:vars (?res - resource)
:precondition
(and ( forall (?res - resource)
            (imply (requires ?whole ?res)
                   (committed ?res ?whole)))
     (incorporated ?part ?whole)
     (or (and (transient-part ?part ?whole)
              ( forall (?prev - assembly)
              (imply (remove-order ?prev ?part ?whole)
                     (incorporated ?prev ?whole))))
         (and (part-of ?part ?whole)
         (not ( exists (?prev - assembly)
              (and (assemble-order ?prev ?part ?whole)
              (incorporated ?prev ?whole)))))))
```

Figure 1: Precondition of the **remove** operator from the *assembly* domain.

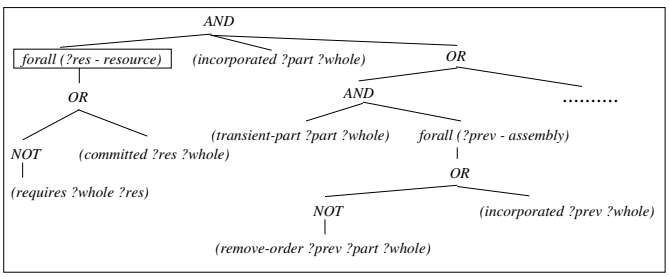

Figure 2: Tree representation of ADL formulas. Note that formulas of the form $\varphi \rightarrow \psi$ have been replaced with the equivalent $\neg\varphi \vee \psi$ already during the parsing process.

leaves of this tree, all occurrences of *?res* are replaced by the selected constant. As we will see below, this can lead to so-called *atomic simplifications*, which replace an atomic formula by either TRUE or FALSE. In turn, the whole tree can sometimes be simplified to TRUE or FALSE, which yields a dramatic reduction in the size of the formula.

In the case of a universal quantifier, the resulting trees are joined by an AND. In the case of an existential quantifier, the trees are joined by an OR. Figure 3 illustrates the result of the process.

The expansion of quantified conditional effects proceeds in a similar way. Figure 4 shows the tree representation of the **move** operator from the *briefcase* domain, whose conditional effect contains the quantifier prefix *forall (?x - object)*. The copied trees will now also contain *when* nodes, i.e., numerous partially instantiated copies of the conditional effect are generated.

The process for instantiating the parameters of an operator fits into the same scheme. In each step, it takes a variable parameter together with the set of type-consistent constants. For each of these constants, a

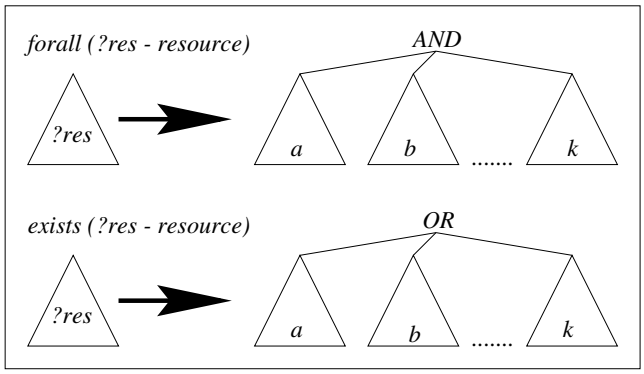

Figure 3: Copies of trees generated during the expansion of quantifiers. Obviously, if one of the subtrees resulting from the expansion of a universal (existential) quantifier can be simplified to FALSE (TRUE), then the whole formula can be simplified to FALSE (TRUE).

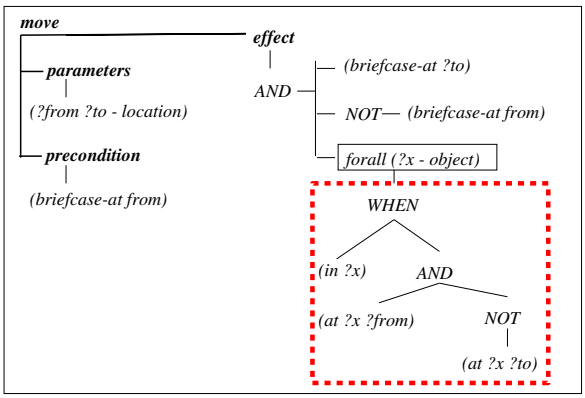

Figure 4: Tree-Representation of an operator with a quantified conditional effect. Expanding the quantifier *forall (?x - object)* results in copies of the tree starting in the *when* node.

copy of the tree representing the operator is generated, and each occurrence of the parameter in this tree is replaced with the constant. Then, the operator tree is simplified. If, for example, it's precondition simplifies to FALSE, the whole partially instantiated operator can be skipped and removed from the domain. After all parameters have been instantiated, each tree represents a ground instance (i.e., an action) of the operator.

Identification of Inertia and their Use during the Instantiation

The tree-copying process takes a variable $?x$ and a constant c as input and traverses the subformula represented in the tree. Whenever it reaches an atomic formula p, it gets replaced with $p[?x/c]$. In many situations, it is worthwhile to invest some more effort at this point and have a closer look at the result of the instantiation. Under certain conditions, namely if p represents an *inertia* relation, one can determine that $p[?x/c]$ must either always be TRUE or FALSE. This can even be the case if $p[?x/c]$ is not yet fully instantiated. Let us consider an example from the *assembly* domain. The object declaration introduces a list of objects followed by their types:

:objects doodad valve frob sprocket socket plug - assembly
charger voltmeter battery - resource

The specification of the initial state contains the following instances of the *requires* relation:

:init (requires frob charger) (requires sprocket charger)
(requires socket voltmeter) (requires doodad voltmeter)
(requires plug voltmeter)

Given the number of declared constants for the two types, the *requires* relation can be instantiated with $6 \times 3 = 18$ different type-consistent tuples, where 5 of them occur in the initial state.

The expansion of the first universal quantifier that is shown framed in Figure 1 generates three copies of the formula tree, each containing either the partially instantiated atom *(requires ?whole voltmeter)*, *(requires ?whole charger)*, or *(requires ?whole battery)*. Two observations can be made:

- If *(requires ?whole ?res) never occurs as a positive effect* of any operator then the only instances of this predicate, which can hold in any state, are those that are specified in the initial state. This, for example, implies that *(requires ?whole battery)* can never hold and is therefore equivalent to FALSE.
- If *(requires ?whole ?res) never occurs as a negative effect* of any operator then the only instances that can be FALSE in any state are those that are *not* contained in the initial state. Now, if the initial state contained all possible ground instances of, say, *(requires ?whole voltmeter)*, then this partially instantiated predicate could be replaced by TRUE. All of its instances would be initially true and thus persist in all reachable states.

In the following, we will formalize these ideas and give a precise notion of *inertia*.

Inertia Relations

IPP proceeds over the domain and problem description and collects all used relation names. For each relation it checks if it satisfies one of the following definitions:

Definition 1 *A relation is a* positive inertia *iff it does not occur positively in an unconditional effect or the consequent of a conditional effect of an operator.*

Definition 2 *A relation is a* negative inertia *iff it does not occur negatively in an unconditional effect or the consequent of a conditional effect of an operator.*

Relations, which are positive as well as negative inertia, are simply called *inertia*. Relations, which are neither positive nor negative inertia, are called *fluents*. The detection of inertia and fluents is easy because in

ADL, effects are restricted to conjunctions of literals. Furthermore, this information can be obtained with a single pass over the domain description, which takes almost no time at all. In the assembly domain, the status of all relations can be inferred as shown in Figure 5.

predicate name	pos. effect	neg. effect	status
available	yes	yes	fluent
requires	no	no	inertia
part-of	no	no	inertia
transient-part	no	no	inertia
assemble-order	no	no	inertia
remove-order	no	no	inertia
complete	yes	no	neg. inertia
committed	yes	yes	fluent
incorporated	yes	yes	fluent

Figure 5: Inertia Relations in the Assembly Domain.

Atomic Simplifications

In order to decide if an inertia can be replaced by TRUE or FALSE one needs to determine and count all type-consistent ground instances of an inertia predicate p that match a partially instantiated occurrence of p.

Definition 3 *Let τ be some* type name.

$$dom(\tau) = \{c_1, \ldots, c_m\}$$

denotes the domain *of τ, i.e., the set of constants having type τ.*

In PDDL, each constant is either explicitly declared as being of a particular type or it has the default type *object*. The same applies to all operator parameters or quantified variables. Each predicate must be explicitly declared together with its arguments, for which type names can be given or the default type is assumed.

Definition 4 *Let p be a predicate of arity n. Let $\vec{a} = (a_1, \ldots, a_n)$ be the argument vector of some partially instantiated occurrence of p where each a_i is either a constant or variable. With*

$$V(\vec{a}) = \{i \in \{1, \ldots, n\} \mid a_i \textit{ is a variable}\}$$

we denote the positions in $\vec{a}$ that are occupied by variables.

Definition 5 *Let p be a predicate and let $\vec{a}$ be the argument vector of some partially instantiated occurrence of p. Let τ_i be the type name of position i in predicate p. Then*

$$\mathsf{MAX}(p\ \vec{a}) = \prod_{i \in V(\vec{a})} |dom(\tau_i)|$$

denotes the number of all *possible type-consistent ground instances of p that unify with the argument vector $\vec{a}$. In contrast,*

$$\mathsf{N}(p\ \vec{a}) = |\{(p\ \vec{c}) \in \mathcal{I} \mid (p\ \vec{c}) \textit{ unifies with } (p\ \vec{a})\}|$$

denotes the number of unifying ground instances of p that are contained in the initial state $\mathcal{I}$. *Obviously,* $\mathsf{N}(p\ \vec{a}) \leq \mathsf{MAX}(p\ \vec{a})$ *holds.*

It is worthwhile noticing here that IPP will remove all variables or parameters that have an *empty type*. Therefore, we have $|dom(\tau_i)| \neq 0$ for each position i of any partially instantiated occurrence of the predicate p. Thus, for any $(p\ \vec{a})$ we have $\mathsf{MAX}(p\ \vec{a}) \neq 0$. As an example, let us consider $(p\ \vec{a})$ =*(requires ?whole voltmeter)*, for which one obtains

$V(\vec{a})$ = {1} /* only one variable argument */
τ_1 = *assembly*
dom(*assembly*) = {doodad, valve, frob, sprocket, socket, plug}

MAX(*requires ?whole voltmeter*) = 6
/* 6 objects can instantiate *?whole* */
N(*requires ?whole voltmeter*) = 3
/* 3 instances in the initial state contain *voltmeter* */

A partially instantiated atomic formula can be simplified to TRUE or FALSE if it satisfies one of the conditions defined below.

Definition 6 *Let $(p\ \vec{a})$ be some partially instantiated atomic formula constructed during the instantiation process.*

If *p is a positive inertia and* $\mathsf{N}(p\ \vec{a}) = 0$
then *$(p\ \vec{a})$ is simplified to* FALSE.

If *p is a negative inertia and* $\mathsf{N}(p\ \vec{a}) = \mathsf{MAX}(p\ \vec{a})$
then *$(p\ \vec{a})$ is simplified to* TRUE.

In all other cases $(p\ \vec{a})$ cannot (yet) be simplified and remains in the formula tree as it is.

From the treatment of empty types, we know that $\mathsf{MAX}(p\ \vec{a}) \neq 0$ holds for $(p\ \vec{a})$. Therefore, obviously at most one of the above tests can succeed. For example, *(requires ?whole battery)* is a positive inertia. It can be simplified to FALSE because no *requires* instance from the initial state matches the argument vector (*?whole*, *battery*), i.e., N*(requires ?whole battery)*= 0 and the first test succeeds.

That an atomic formula can sometimes be simplified to TRUE is best seen in the case when it is fully instantiated. Take, for example, *(requires plug voltmeter)*. This fact occurs in the initial state, so N*(requires plug voltmeter)*= $1 \neq 0$ and the first test fails. However, MAX*(requires plug voltmeter)*= $\prod_{i \in \emptyset} |dom(\tau_i)| = 1$ and the second test succeeds. This reflects that *(requires plug voltmeter)* is initially TRUE and will never be made FALSE because *requires* is a negative inertia.

Theorem 1 (Soundness of Simplifications) *Given a planning domain and problem, if $(p\ \vec{a})$ is simplified to*
(1) FALSE, *then no state s which is reachable from the initial state satisfies any type-consistent ground instance of $(p\ \vec{a})$.*

(2) TRUE, *then any state s which is reachable from the initial state satisfies all type-consistent ground instances of* $(p\ \vec{a})$.

All proofs can be found in (Koehler & Hoffmann 1999).

Atomic simplification requires to determine the number $\mathsf{N}(p\ \vec{a})$ of all those ground tuples in the initial state that unify with a given argument vector of arbitrary length, containing variables or constants at arbitrary positions. Using a naive solution, this means to perform a single pass over the initial state $\mathcal{I}$, testing for each fact if it unifies with $(p\ \vec{a})$. Obviously, the time complexity is $\Theta(|\mathcal{I}| * n_{max})$ where n_{max} denotes the maximum arity of the predicates. The test for atomic simplification has to be done for every leaf of every tree that is ever generated during the instantiation process. The number of these leaves is likely to be enormous, so there is a strong need for a highly efficient method to find $\mathsf{N}(p\ \vec{a})$. In the following, such a method is described, which allows to retrieve the number of matching initial facts in time linear in the *length of* $\vec{a}$, i.e., in the arity of the predicates, $O(n_{max})$.

Efficient Implementation of Atomic Simplifications

In principle, the idea behind the implementation is as simple as this: Before instantiation starts, perform a single pass over the initial state and create tables in which the occurring tuples are documented. Then later determine the proper table entry for $(p\ \vec{a})$ and look up the correct value of $\mathsf{N}(p\ \vec{a})$. What makes the process complicated is that we have to deal with *partially instantiated* argument vectors $\vec{a}$.

Let us consider the *requires* predicate as an example. For its argument vector of length 2, four cases can occur:

(1) Both arguments are variables and thus $\vec{a} = (?x_1, ?x_2)$. One needs to determine the total number of occurrences of *requires* (with arbitrary arguments) in the initial state.

(2) The first argument is instantiated, but the second argument is a variable and thus $\vec{a} = (c_1, ?x_2)$. We need the number of occurrences of *requires* where the first argument is c_1.

(3) Only the second argument is instantiated and $\vec{a} = (?x_1, c_2)$. We need to count the occurrences with c_2 at the second position.

(4) Both arguments are instantiated and $\vec{a} = (c_1, c_2)$. The question is whether the initial state contains *(requires* c_1 c_2*)*.

For each of these four cases, a separate table is constructed. The table entries are computed from the initial state. The dimension of each table corresponds to the number of instantiated positions of the argument vector. In Case (1), the table is therefore 0-dimensional and simply consists of an integer counting the number of *requires* facts in the initial state. For Cases (2) and (3), a 1-dimensional table is needed, with one entry for each object that is type-consistent with the instantiated argument. For each of these objects, the corresponding entry counts the number of times that *requires* occurred in the initial state instantiated with that object. In Case (4), a 2-dimensional table is constructed. Its entries are indexed by all pairs of type-consistent objects that can instantiate the *requires* predicate. For each such pair, the entry is set to 1 iff *requires* occurred in the initial state instantiated with that pair. All tables are shown in Figure 6.

$\emptyset$

5

$\{1\}$

doodad	1
valve	0
frob	1
sprocket	1
socket	1
plug	1

$\{2\}$

charge	2
voltmeter	3
battery	0

$\{1,2\}$

	charger	voltmeter	battery
doodad	0	1	0
valve	0	0	0
frob	1	0	0
sprocket	1	0	0
socket	0	1	0
plug	0	1	0

Figure 6: The tables to represent those facts from $\mathcal{I}$ that match a given argument vector for the *requires* predicate. Given the set of instantiated positions as $\emptyset$, $\{1\}$, $\{2\}$ or $\{1,2\}$, the corresponding tables are shown from left to right and down.

Let p be a predicate of arity n. For each subset $C \subseteq \{1,\ldots,n\}$, a table $\mathsf{T}(p\ C)$ has to be constructed. The table is $|C|$-dimensional and lists one entry $\mathsf{T}(p\ C)(\vec{c})$ for each type-consistent tuple $\vec{c}$ of constants that can possibly instantiate p at exactly the positions in C. All entries are initially set to zero.

Note that although the number of tables is exponential in the arity of the predicates, planning domain representations rarely use predicates with more than 3 or 4 arguments.

Definition 7 *Let $\vec{a} = (a_1,\ldots,a_n)$ be an argument vector of size n, each a_i being either a variable or a constant. Let $C = \{i_1,\ldots,i_k\}$ be a set of possible positions, i.e., $C \subseteq \{1,\ldots,n\}$, where the $i_1,\ldots i_k$ are ordered increasingly. With*

$$\vec{a}|_C := (a_{i_1},\ldots,a_{i_k})$$

we denote the restriction *of $\vec{a}$ to the positions in C.*

Intuitively, the restriction of a vector to some set $C \subseteq \{1,\ldots,n\}$ is obtained by simply skipping all those

positions that are not in C, but preserving the order of the arguments.

Now for each ground atom $(p\ \vec{c})$ that occurs in the initial state, the following is done:

for all sets $C \subseteq \{1, \ldots, n\}$ **do**
  increment $\mathsf{T}(p\ C)(\vec{c}|_C)$ (1)
endfor

Performing process (1), we count for all instances of p in the initial state how often each combination of constants occurs for arbitrary sets C of positions.

Definition 8 *Let p be a predicate of arity n. Let $\vec{a} = (a_1, \ldots, a_n)$ be the argument vector of some partially instantiated occurrence of p. With*

$$C(\vec{a}) := \{i \in \{1, \ldots, n\} \mid a_i \text{ is a constant}\}$$

we denote the positions where $\vec{a}$ is instantiated.

During instantiation, given a partially instantiated predicate $(p\ \vec{a})$, we determine the set $C(\vec{a})$ of positions where the argument vector is occupied by constants. The appropriate table $\mathsf{T}(p\ C(\vec{a}))$ is the one corresponding to that set. The entry in this table that we want to access is the one indexed by the constants in the current argument vector $\vec{a}$, i.e., by the restriction $\vec{a}|_{C(\vec{a})}$ of $\vec{a}$ to its constants.

Theorem 2 (Soundness of Tables)
Let the tables T be the result of performing process (1) *for each fact in the initial state. Then we have for each partially instantiated predicate $(p\ \vec{a})$:*

$$\mathsf{N}(p\ \vec{a}) = \mathsf{T}(p\ C(\vec{a}))(\vec{a}|_{C(\vec{a})}).$$

During the instantiation process it remains to find the corresponding table entry in order to determine the correct value of $\mathsf{N}(p\ \vec{a})$. Since constants are internally kept as *numbers* they can in principle be used as indices into a table. However, to directly index into the tables, one would need to define tables of *arbitrary dimension*. Instead, the implementation uses an implicit representation of the tables. The appropriate address is computed by performing a sweep over the argument vector, which takes time $O(n_{max})$. As arities are usually small, this running time is very close to constant anyway.

Ground Level Inertia

So far we have only considered the *predicates* which are never made true or false by a planning operator. These were used to constrain the instantiation process. Once the set of all actions has been determined, one can similarly define the *ground facts* that are never made true or false by one of the actions.

Definition 9 *A ground fact is a* positive ground inertia *iff it does not occur positively in an unconditional effect or the consequent of a conditional effect of an action.*

Definition 10 *A ground fact is a* negative ground inertia *iff it does not occur negatively in an unconditional effect or the consequent of a conditional effect of an action.*

An initial fact, which is a negative ground inertia, is never made FALSE and thus always satisfied in all reachable world states. It can be removed from the state description. All its occurrences in the preconditions of actions and in the antecedents of conditional effects can be simplified to TRUE.

A fact, which is a positive ground inertia and *not* contained in the initial state, is never satisfied in any reachable world state. All its occurrences in the preconditions of actions and in the antecedents of conditional effects can be simplified to FALSE.

The remaining facts are fluents that, roughly speaking, can possibly change their truth value during the planning process. They are therefore *relevant* to the representation of the planning problem.

Definition 11 *A ground fact is* relevant *iff*

1. *it is an initial fact and not a negative ground inertia, or if*
2. *it is not an initial fact and not a positive ground inertia.*

Using the table which corresponds to the fully instantiated case of the process described in the previous section, one can find all relevant facts by performing a single sweep over the initial state and the effects of all actions.

The simplified actions and the set of all relevant facts are then used by IPP to generate a bitvector representation for all states and actions, where each relevant fact corresponds to a position in a bitvector.

Simplification of Operator Representations

As we have already mentioned in the beginning, the instantiation process creates copies of trees representing formulas and operators. These trees can be simplified if one of their subtrees has been simplified to TRUE or FALSE, which can result from the atomic simplifications performed during the instantiation process. As soon as such an atomic simplification has replaced an atomic formula by TRUE or FALSE, the subsequently described non-atomic simplification operations are performed.[2]

$\neg$TRUE	$\equiv$	FALSE	$\neg$FALSE	$\equiv$	TRUE
TRUE $\wedge\ \varphi$	$\equiv$	φ	$\varphi \wedge \varphi$	$\equiv$	φ
FALSE $\wedge\ \varphi$	$\equiv$	FALSE	$\varphi \vee \varphi$	$\equiv$	φ
TRUE $\vee\ \varphi$	$\equiv$	TRUE	$\varphi \wedge \neg\varphi$	$\equiv$	FALSE
FALSE $\vee\ \varphi$	$\equiv$	φ	$\varphi \vee \neg\varphi$	$\equiv$	TRUE

Figure 7: Implemented Simplifications for First-Order Formulas.

[2] They are also performed once directly after having parsed the domain and problem file.

The first-order formulas which represent the preconditions of operators and the antecedents of conditional effects are simplified based on the well-known tautologies as shown in Figure 7. Besides this, IPP implements the following simplifications:

1. If a quantified variable does not occur in the quantified formula, the quantifier is removed, i.e., $\forall\, ?x\ \varphi(?y)$ is simplified to $\varphi(?y)$.[3]
2. If a quantified variable $?x$ has an *unknown type*, which has not been declared in the :types field of the domain file or if it has an *empty type*, for which no constant has been declared in the problem file, then the quantified formula is replaced by TRUE in the case of a universal quantifier and by FALSE in the case of an existential quantifier.
3. An equality between two identical variable names, $?x = ?x$, is simplified to TRUE. An equality between two identical constants, $c_1 = c_1$, is also simplified to TRUE. If the constants are different, i.e., $c_1 = c_2$, the equality is simplified to FALSE. In a fully instantiated formula, all occurrences of equalities have been replaced by TRUE or FALSE.

The simplification of first-order formulas can reduce a whole precondition, antecedent or consequent to TRUE or FALSE. In this case, the operator description can be simplified:

1. If the antecedent of a conditional effect becomes FALSE, the conditional effect is removed from the operator. In this case, the effect is never applicable because it requires FALSE to hold, i.e., the state must be inconsistent.
2. If the antecedent of a conditional effect becomes TRUE, the conditional effect becomes unconditional.
3. If the consequent of a conditional effect becomes TRUE, the conditional effect is removed because it does not lead to any state transition.
4. If the precondition or the unconditional effect of an operator becomes FALSE, the whole operator is removed from the domain.
5. If an operator has only TRUE as its unconditional effect and no conditional effects, then the whole operator is removed.[4]

In the final set of actions, to which no simplifications can be applied anymore, all unconditional effects are merged into a single conjunction of literals and all conditional effects with identical antecedents are merged into a single conditional effect.

[3] Unused quantified variables will usually not appear in the initial domain description. They can, however, appear as a result of atomic simplifications.

[4] Removing effects or whole operators can possibly turn fluents into inertia, i.e., one could repeat the whole analysis procedure again. However, such a phenomenon was not observed in any planning domain and therefore it seems not worth to invest the effort into such a fixpoint computation.

IPP also implements various syntax checks that help to develop proper domain representations:

1. An operator is removed (a warning is issued, but planning continues) if an operator parameter is declared using an unknown or empty type.
2. A parameter is removed from the operator description (a warning is issued, but the operator remains in the set) if it is declared, but nowhere used in the preconditions or effects.[5]

IPP aborts the instantiation process if it encounters one of the following situations:

1. A predicate symbol is overloaded. PDDL requires the declaration of predicates, their arity and the types of their arguments. When parsing the domain and problem files, IPP verifies that all occurrences of a predicate meet the declaration.
2. An equality statement occurs in an unconditional effect or in the consequent of a conditional effect.
3. An equality statement has less or more than two arguments.
4. A variable occurs that is neither declared as a parameter nor bound by a quantifier.
5. A constant occurs that has not been declared in the problem file.

Encoding Unary Inertia as Types

Many domains, in particular *all* STRIPS domains used in the 1998 AIPS planning competition contain *unary inertia*. These are predicates of arity one, which satisfy Definitions 1 and 2 and thus do not occur in any of the effects. In other words, the set of constants c that can ever (and will always) satisfy $(p\ c)$ is exactly the set of constants occurring as the arguments of the instances of p in the initial state.

Obviously, this set can be seen as the encoding of *type* information because the single variable argument of p can only be instantiated with one of these constants if we want to obtain a possibly satisfiable atomic formula. As a matter of fact, in the STRIPS domains from the planning competition, all unary inertia where intended to provide implicit type information, as there are no explicit types given in classical STRIPS, see Figure 8 for an example.

One can easily make this implicit type information explicit and remove all unary inertia from the domain description. The previously described instantiation process that identifies and simplifies inertia will also achieve the desired simplification of unary inertia, because they are simply a special case wrt. the length of the argument vector. However, doing it this way, the algorithm repeatedly generates copies of formula trees, only to find out that it can remove them immediately afterwards because they use the "wrong objects"

[5] Just like unused quantifiers, this can also happen as a result of simplifications.

in some unary inertia. For example, when instantiating the set of actions for the problem `strips-log-x-9` from the *logistics* domain used in the competition, 55088 actions are generated for which the instantiation procedure needs 527 seconds.

```
:action load-truck
:parameters (?obj ?truck ?loc)
:precondition (and (obj ?obj) (truck ?truck)
                   (location ?loc) (at ?truck ?loc)
                   (at ?obj ?loc))
:effect (and (not (at ?obj ?loc)) (in ?obj ?truck))
```

Figure 8: The **load-truck** operator from the *logistics* domain. Note the untyped parameters and the underlined unary inertia predicates that implicitly encode the type information.

Consequently, there is the need for a further optimization of the instantiation process, which can be achieved through a separate treatment of unary inertia. The optimization, which is described in detail in this section, encodes all the unary inertia *obj*, *city*, *truck*, *airplane*, *location* and *airport* directly as types, which restrict the instantiation possibilities for the arguments of the operators. Running time for this example decreases down to 63 seconds.[6]

We now give a precise notion of how implicit type information can be made explicit. First, for each unary inertia predicate p the new type symbol τ_p for *the type corresponding to* p is introduced.

Definition 12 *Let p be an inertia predicate of arity 1. The type τ_p* corresponding to *p is defined as the type whose domain comprises all constants c for which $(p\ c)$ holds in the initial state $\mathcal{I}$:*

$$dom(\tau_p) = \{c \mid (p\ c) \in \mathcal{I}\}$$

New types can be constructed from other types by intersecting or subtracting from each other the corresponding sets of constants.

Definition 13 *Let τ_1 and τ_2 be type names. Then $\tau_1 \cap \tau_2$ and $\tau_1 \setminus \tau_2$ are new type names. Their domains are defined as:*

$$dom(\tau_1 \cap \tau_2) = dom(\tau_1) \cap dom(\tau_2)$$
$$dom(\tau_1 \setminus \tau_2) = dom(\tau_1) \setminus dom(\tau_2)$$

After having extracted all types τ_p for unary inertia p from the initial state, the type structure of the domain representation is refined with the types τ_p and types that can be constructed from them.

[6]The instantiation procedure implemented in IPP 3.3 that has been used in the competition is still a bit faster: It needs only 52 seconds for this example. However, this procedure uses a specialized algorithm which is only capable of handling *conjunctive* preconditions, and it generates a total of 62261 actions because no test for ground inertia is performed.

Definition 14 *Let o be some operator and $?x$ be one of its parameters. Let p be a unary inertia. If $(p\ ?x)$ occurs in the preconditions of o or in the antecedent of one of its conditional effects, o is replaced by two new operators $o1$ and $o2$:*

- *In $o1$, the type τ that has been declared for $?x$ is restricted to $\tau \cap \tau_p$ and all occurrences of $(p\ ?x)$ are replaced with* TRUE.
- *In $o2$, the type τ that has been declared for $?x$ is restricted to $\tau \setminus \tau_p$ and all occurrences of $(p\ ?x)$ are replaced with* FALSE.

Similarly, quantified formulas in preconditions or antecedents of conditional effects are replaced.

Definition 15 *Let $\varphi = \forall\ ?x : \tau\ \psi$ be some universally quantified formula containing a unary inertia p with argument $?x$ of type τ. The formula φ is replaced with φ' defined as*

$$\varphi' = \begin{array}{l} \forall\ ?x : \tau \cap \tau_p\ \ \psi[(p\ ?x)/\mathsf{TRUE}] \wedge \\ \forall\ ?x : \tau \setminus \tau_p\ \ \psi[(p\ ?x)/\mathsf{FALSE}] \end{array}$$

Let $\varphi = \exists\ ?x : \tau\ \psi$ be some existentially quantified formula containing a unary inertia p with argument $?x$. Then φ is replaced with φ'

$$\varphi' = \begin{array}{l} \exists\ ?x : \tau \cap \tau_p\ \ \psi[(p\ ?x)/\mathsf{TRUE}] \vee \\ \exists\ ?x : \tau \setminus \tau_p\ \ \psi[(p\ ?x)/\mathsf{FALSE}] \end{array}$$

In the definition above, $\psi[(p\ \ ?x)/\mathsf{TRUE}]$ and $\psi[(p\ ?x)/\mathsf{FALSE}]$ denote the formulas, which are obtained from ψ if all occurrences of $(p\ ?x)$ have been replaced with TRUE and FALSE, resp.

The soundness of the replacements follows from the observation that under the restriction $\tau \cap \tau_p$ the atomic formula $(p\ c)$ is always TRUE because only constants c are considered which are also in $dom(\tau_p)$. Under the restriction $\tau \setminus \tau_p$ only constants $c \in dom(\tau)$ are considered that are *not* members of $dom(\tau_p)$ and thus $(p\ c)$ is always FALSE.

We formally state the soundness of the replacements for universally quantified formulas. Again, the proof can be found in (Koehler & Hoffmann 1999).

Theorem 3 (Soundness of Type Encodings) *Let p be a unary inertia predicate. Let $\varphi = \forall\ ?x : \tau\ \psi$ be a formula with $(p\ ?x)$ being a subformula of ψ. Let φ' be the formula φ gets replaced with according to Definition 15. Then, for any state s that is reachable from the initial state holds*

$$s \models \varphi \ \Leftrightarrow\ s \models \varphi'$$

The soundness of the type encoding follows from the fact that the modified operator set with the newly introduced types has exactly the set of ground instances, which is generated by the instantiation procedure using inertia and performing atomic simplifications that we described in the previous section. The soundness of the replacement of existentially quantified formulas follows

with similar arguments as in the universally quantified case.

As an example, let us consider the operator from Figure 8 again. As there is no explicitly defined type for any of the three parameters they are assigned the default type *object*. When examining the first parameter *?obj*, IPP finds that it is used in the unary inertia predicate *obj*. Therefore, it generates two copies of the operator, restricts the parameter types according to Definition 14, and performs the corresponding atomic simplification of the unary inertia. The result is shown in Figure 9.

In the first operator, the atom TRUE can obviously be removed from the conjunction, which leads to a simplified precondition. As *all* constants in the STRIPS *logistics* problems are defined to be of type *object*, the domain $dom(object \cap \tau_{obj}) = dom(\tau_{obj})$ comprises exactly those constants c for which $(obj\ c)$ is contained in the initial state.

:action **load-truck**$^{(1)}$
:parameters *(?obj - object* $\cap\ \tau_{obj}$ *?truck ?loc)*
:precondition *(and (* TRUE *)*
(truck ?truck) (location ?loc)
(at ?truck ?loc) (at ?obj ?loc))

:action **load-truck**$^{(2)}$
:parameters *(?obj - object* $\setminus \tau_{obj}$ *?truck ?loc)*
:precondition *(and (* FALSE *)*
(truck ?truck) (location ?loc)
(at ?truck ?loc) (at ?obj ?loc))

Figure 9: Parameters and preconditions of the two new **load-truck** operators, which result from the encoding of the unary inertia predicate *obj* as a type.

In the second operator, the first atomic precondition has been replaced by FALSE as no constant in $dom(object \setminus \tau_{obj})$ can satisfy $(obj\ c)$. Thus, the whole precondition of this operator simplifies to FALSE and it can be removed from the operator set as it will never be applicable. Note that in the case of arbitrary first-order preconditions one cannot usually expect that operators can be removed immediately just after they have been generated.

Repeating this process for the other two parameters, always the second copy is removed immediately after it has been generated and thus IPP obtains the final representation of the **load-truck** operator, which is shown in Figure 10.

The encoding of unary inertia as types is one possibility of how type information can be used to reduce the search space of a planner. TIM (Fox & Long 1999) implements additional sophisticated type analysis methods, but currently limited to STRIPS. The extension of this work to ADL and its combination with the instantiation method that is described in this paper remains a subject of future work.

:action **load-truck**
:parameters *(?obj - object* $\cap\ \tau_{obj}$
?truck - object $\cap\ \tau_{truck}$
?loc - object $\cap\ \tau_{location}$*)*
:precondition *(at ?truck ?loc) (at ?obj ?loc))*
:effect *(and (not (at ?obj ?loc)) (in ?obj ?truck))*

Figure 10: The new operator **load-truck**, which results from the encoding of all unary inertia as types and which replaces the original operator representation. This operator is identical with the one that is used in the typed version of this domain.

Empirical Results

Many examples could be presented, which nicely illustrate the benefits of an instantiation procedure that takes inertia into consideration. For example, in the *movie* domain used in the planning competition, 5 operators are declared to get snacks: *get-chips*, *get-dip*, *get-pop*, *get-cheese*, *get-crackers*. Each of them has a similar description, of which we only exemplify the *get-chips* operator:

get-chips
:parameters *(?x - chips)*
:precondition
:effect *(have-chips).*

One observes that the parameter $?x$ is not used anywhere in the operator description. If for example, 9 different constants are declared for each kind of snack, one obtains 9 ground instances of each operator, which are all identical and spam the search space of the planner. In all *movie* problems, the goals are reachable at time step 1, but a plan can only be extracted at time step 2, i.e., a permutation of all actions at time step 1 is performed by the complete search algorithm. Not very surprisingly, this takes almost 3 s in IPP 3.3 on a Sun Ultra 1/170 because 250973 actions must be tried before a solution is found. In contrast to this, when detecting the unused parameter, only one instance is generated for each operator, which dramatically reduces the search space down to 29 actions and thus a plan is found in only 0.06 s.

In the *assembly* domain, operators can be dramatically simplified because they contain so many inertia. For example, the complex precondition shown in Figure 1 uses 7 different predicates, but 5 of them are inertia. This means that each precondition must simplify to a formula only mentioning the fluents *incorporated* and *committed*. For many actions, the precondition reduces to a single atomic formula using only the *incorporated* predicate. IPP 4.0 is thus able to solve some *assembly* problems, while previous versions failed already during the instantiation, see Figure 11 for selected results.

Column 2 shows the number of generated actions using the instantiation process with inertia compared to the number of all possible actions using naive enumeration. The search space is measured in the number of

problem	actions	cpu sec.	search space
assem-x-1	114/760	1742.43	64 673 043
assem-x-2	84/882	18.03	848 829
assem-x-3	190/1248	0.83	108
assem-x-6	118/1800	46342.80	1 283 078 957

Figure 11: Performance of IPP on *assembly* problems on a Sun Ultra 1/170.

actions IPP tries until it finds a plan. The solution plans involve between 31 and 38 actions, but require only between 10 and 18 time steps, i.e., they involve quite some parallelism. Several other problems from this domain can be proven as unsolvable.

The determination of *ground* inertia helps IPP to discover information that it would not be able to find if only inertia *predicates* were analyzed. An interesting example of this behavior occurs in the *tower of Hanoi* domain. Given the operator

move*(?disc,?from,?to: disc)*
:precondition *(and (smaller ?to ?disc)*
(on ?disc ?from) (clear ?disc) (clear ?to))
:effect *(and (clear(?from) (on ?disc ?to)*
(not (on ?disc ?from)) (not (clear ?to))

which describes a legal move of discs, one notices that only a smaller disc can be moved onto a larger disc. IPP discovers that *smaller* is an inertia predicate and only generates the appropriate actions. But the action set also contains moves, which take a disc *from* a smaller disc and put it on a smaller disc. Indeed, the operator description says nothing about the relationship between the disc *?from* and the moving disc *?disc*, i.e., a move that takes a disc from a *smaller* disc and puts it on another smaller disc seems to be a legal action.

When performing the analysis of inertia on the ground level, IPP is able to find out that such moves are impossible. It detects that all instances of *(on ?disc ?from)*, where *?disc* is larger than *?from* are never made true by any action, i.e., they are positive inertia, and they do not hold in the initial state. Thus, these facts are unsatisfiable and all preconditions using them can be simplified to FALSE. Since all actions with FALSE as a precondition are removed from the action set, a further reduction of the size of the planning graph is achieved. For example, in the case of 3 discs, 10 out of 48 actions are eliminated. In the case of 8 discs, 140 out of 468 actions are removed. A search space of only 295.535 actions results and the plan of 255 steps is found in only 16 seconds.

Conclusion

The generation of the set of all ground actions for a given set of expressive ADL operators is a complex process which heavily influences the performance of any planner or pre-planning analysis method. The implementation comprises more than 5000 lines of C code and is available from the IPP webpage at http:www.informatik.uni-freiburg.de/~ koehler/ipp.html in the release of IPP 4.0. We hope that the instantiation procedure will become a useful part of reusable code that helps other researcher teams in setting up their own planners more quickly and without dealing with the burden of reimplementing the same preprocessing procedures again and again.

References

Fox, M., and Long, D. 1999. The detection and exploitation of symmetry in planning problems. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence*, 956–961. Morgan Kaufmann, San Francisco, CA.

Geffner, H. 1999. HSP: A heuristic search planner. web documentation.

Kautz, H., and Selman, B. 1996. Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proceedings of the 14th National Conference of the American Association for Artificial Intelligence*, 1194–1201. AAAI Press.

Koehler, J., and Hoffmann, J. 1999. Handling of inertia in a planning system. Technical Report 122, Albert-Ludwigs-University.
available at http://www.informatik.uni-freiburg.de/~ koehler/ipp.html.

Koehler, J.; Nebel, B.; Hoffmann, J.; and Dimopoulos, Y. 1997. Extending planning graphs to an ADL subset. In Steel (1997), 273–285.

Koehler, J. 1998. Solving complex planning tasks through extraction of subproblems. In Allen, J., ed., *Proceedings of the 4th International Conference on Artificial Intelligence Planning Systems*, 62–69. AAAI Press, Menlo Park.

McDermott, D., et al. 1998. *The PDDL Planning Domain Definition Language*. The AIPS-98 Planning Competition Comitee.

McDermott, D. 1998. Planning competition benchmark problems. web documentation, http://www.cs.yale.edu/users/mcdermott.hmtl.

Nebel, B.; Dimopoulos, Y.; and Koehler, J. 1997. Ignoring irrelevant facts and operators in plan generation. In Steel (1997), 338–350.

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

Steel, S., ed. 1997. *Proceedings of the 4th European Conference on Planning*, volume 1348 of *LNAI*. Springer.