

DISCOPLAN: an Efficient On-line System for Computing Planning Domain Invariants*

Alfonso Gerevini¹ Lenhart Schubert²

¹ Dipartimento di Elettronica per l'Automazione, Università di Brescia
Via Branze 38, 25123 Brescia, Italy. E-mail: gerevini@ing.unibs.it

² Department of Computer Science, University of Rochester
Rochester, NY 14627-0226. E-mail: schubert@cs.rochester.edu

Abstract

DISCOPLAN is an efficient system for discovering state invariants in planning domains with conditional effects. Among the types of invariants found are implicative constraints relating a fluent predication to a fluent or static predication (with allowance for static supplementary conditions), single-valuedness constraints, exclusiveness constraints, and several others. The algorithms used are polynomial-time for any fixed bound on the number of literals in an invariant. Some combinations of constraints are found by simultaneous induction, and the methods can be iterated by expanding operators using previously found invariants. The invariants found by DISCOPLAN have been shown to enable large performance gains in SAT planners, and they can also be helpful in planning domain development and debugging.

Introduction

State invariants (or *state constraints*) in planning are properties of objects or relationships among objects that hold in all states reachable from the initial state. For example, a familiar invariant in a blocks world is the property that if one block is on another, the latter is not clear. In our terminology, this is an *implicative* constraint. Another example is that a block can be on at most one other block; this is a *single-valuedness* constraint (sv-constraint).

A point that has become widely recognized in the planning community (and that we amplify in what follows) is that knowledge of state invariants is important for efficient planning. However, such knowledge cannot in general be assumed to be available *a priori* in a given planning domain. Rather, planning domains are generally considered fully specified once a set of operators with well-defined preconditions and effects has been supplied, along with an initial state. This is defensible since state invariants are implicit in the specification of the operators and initial state; i.e., under a STRIPS assumption the only properties and relationships that change when an operator is applied are those spelled out in the effects of the operator. So a separate specification of what remains unchanged when operators are applied would be logically redundant. However, it is far from obvious from inspection of a given set of planning operators and an initial state what the invariants of the domain are. The goal of our research has been to formulate automatic, efficient methods for inferring the most important such invariants, and to implement these methods in our DISCOPLAN system.

*The on-line DISCOPLAN system can be accessed at <http://prometeo.ing.unibs.it/discoplan>. DISCOPLAN is written in Common Lisp.

The importance of state invariants for efficient planning is that they can be used to radically restrict the search space. This is so for any approach to planning that involves explicit or implicit exploration of incompletely specified possible states of the world, as is the case for deductive planning, regression planning, bidirectional planning, and planning by incremental constraint satisfaction (in particular, SAT-based planning).

In our work we have focused on SAT-based planners. These implicitly search a space of state sequences, constrained by disjunctions of ground literals. Their performance depends critically on the invariants added (as ground instances) to the mix of disjunctions, and intuitively this is because state invariants constrain the alternative states that are possible at each time step under consideration. Some results showing the dramatic improvements in the performance of SAT-based planners like SATPLAN [8] and MEDIC [2] obtainable through the use of automatically inferred invariants are included in [5].

DISCOPLAN finds a variety of different types of constraints, including static (type) constraints (most importantly, supertype/subtype and exclusion relations among static monadic predicates – ones unaffected by any operator), and predicate domain constraints (sets of possible argument tuples corresponding to each predicate in the domain, after 0, 1, ..., t actions have occurred). But the majority of its algorithms are devoted to the discovery of state invariants, using a *hypothesize-and-test* paradigm. All the algorithms instantiating this paradigm are applicable to sets of operators conforming with UCPOP or PDDL syntax [11, 7], allowing for *when*-clauses (conditional effects) but not disjunctive or universally quantified conditions. We will be referring to the unconditional part of an operator as its *primary when-clause*. The allowance for conditional effects is a major distinction of DISCOPLAN from related systems.

Very briefly, the hypothesize-and-test paradigm consists of hypothesizing invariants Γ of some particular syntactic type, such as implicative constraints ($\text{IMPLIES } \phi \psi$) where ϕ and ψ are literals that may contain universal variables, augmenting these hypotheses with potential supplementary static conditions, and then testing them against all *when*-clauses of all operators and against the given initial conditions. In the testing phase, minimal sets Σ of supplementary conditions are found, up to sets of some limited size (e.g., 2 or 3) that suffice to ensure that $\Sigma \Rightarrow \Gamma$ holds in all states reachable from the initial state. The hypothetical invariants Γ of a particular type are chosen by inspecting the preconditions and effects of particular operators, to find conditions that appear to become or remain true when

certain kinds of effects occur. The idea is to choose the constituents of Γ in such a way that a proof by induction of the invariance of Γ will be at least locally enabled. In this way large numbers of syntactically possible invariants are eliminated from consideration. The testing phase can be viewed as an automated inductive proof attempt (with addition of supplementary conditions as needed to allow the proof to succeed). An important point is that Γ may actually consist of multiple hypotheses that can be proved to be invariants by simultaneous induction. Typically, such multiple hypotheses consist of an implicative hypothesis ($\text{IMPLIES } \phi \psi$) along with sv-hypotheses corresponding to argument positions in ϕ and ψ occupied by universal variables occurring in only one of ϕ, ψ . The point is important since the invariance of the individual formulas in such cases cannot be proved in isolation. Our various hypothesize-and-test algorithms have been proved to yield correct invariants, and run in polynomial time for fixed bound on the number of supplementary conditions Σ added to Γ .

In a little more detail, the hypothesize-and-test algorithms conform with the following structure (iterating over all possible candidate constraints Γ found in the first step).

1. Hypothesize a constraint Γ based on co-occurrences of literals in a *when*-clause w of an operator and in the corresponding primary *when*-clause w_1 (if different). For example, effects ϕ and ψ might lead to an implicative hypothesis ($\text{IMPLIES } \phi \psi$), and possibly sv-hypotheses about the predicates involved.
2. Add a set of candidate supplementary conditions $\{\sigma_1, \dots, \sigma_n\}$, consisting of the static preconditions of w and w_1 and if $w \neq w_1$, the negations of static preconditions of other *when*-clauses (except ones that unify with static preconditions of w or w_1 or their negations).
3. Test hypothesis Γ relative to each *when*-clause of each operator, using the relevant *verification conditions*; for each apparent violation of Γ find the corresponding possible “excuses” for the violation. An excuse is a set of provisos $\{\sigma'_1, \dots, \sigma'_m\}$, chosen from the candidate supplementary conditions, that weaken the hypothesis sufficiently to maintain its truth. If a violation has no excuses, abandon the hypothesis Γ , otherwise record the set of possible excuses of the violation on a global list.
4. Find all minimal subsets (up to a given size, e.g., 3) of $\{\sigma_1, \dots, \sigma_n\}$ that “cover” all apparent violations of Γ ; a subset of $\{\sigma_1, \dots, \sigma_n\}$ covers an apparent violation of Γ if it contains all elements of at least one “excuse” for that violation;
5. Check hypothesis ($\Gamma \sigma'_1 \dots \sigma'_m$) (i.e., the original hypothesis together with added provisos) for each of the minimal subsets $\{\sigma'_1, \dots, \sigma'_m\}$ of $\{\sigma_1, \dots, \sigma_n\}$ found in the previous step for truth in the initial conditions of the problem being solved; return the variant hypotheses that pass this test as the verified hypotheses.

The *verification conditions* referred to in step 3 depend on the form of Γ , and are designed to ensure that if Γ together with specified supplementary conditions holds in a given state, it also holds in every possible successor state. For example, in the case of a simple implicative constraint ($\text{IMPLIES } \phi \psi$) together with a set of static supplementary conditions, the verification conditions say (roughly) that any operator effect matching ϕ

```
(define (operator Put)
:parameters (?x ?y ?z)
:precondition (and (on ?x ?z) (clear ?x)
                  (neq ?x Table) (neq ?y ?z) (neq ?x ?y))
:effect (and (when (eq ?y Table)
                 (and (on ?x ?y) (clear ?z) (not (on ?x ?z))))
            (when (and (neq ?y Table) (clear ?y))
                 (and (on ?x ?y) (clear ?z) (not (on ?x ?z))
                     (not (clear ?y))))) )
```

Figure 1: A Formalization of the blocks world.

must be accompanied by an effect or persistent precondition matching ψ , or else the preconditions must entail the falsity of a supplementary condition; and similarly for the contrapositive, ($\text{IMPLIES } \neg\psi \neg\phi$). (The conditions are actually slightly more complicated because of the allowance for conditional effects.)

Types of DISCOPLAN Invariants

The input of DISCOPLAN is a domain description consisting of the specification of an initial state and a set of extended STRIPS operators which may involve conditional effects, negated preconditions, constants, typed and untyped parameters (Figure 1 gives a very simple formalization of the blocks world containing some of these features). In the following we describe the types of invariants that are discovered by the current version of DISCOPLAN (for a more detailed description the reader is referred to [5, 6]).

Predicate Domain Constraints. Predicate domain constraints are sets of possible argument tuples corresponding to each predicate in the domain after 0, 1, ..., t actions have occurred. These constraints are computed using a simplified version of the planning graph for the given problem [1].

Static Predicates and Static Constraints. Static constraints are invariants involving type-predicates, i.e., static monadic predicates that occur positively in the initial state – ones unaffected by any operator. Static constraints consist of a (possibly empty) set of objects for each type-predicate and a list of supertype, subtype, and incompatible relationships between type-predicates.

Simple Implicative Constraints. Simple Implicative Constraints are constraints of form $((\phi \Rightarrow \psi) \sigma_1 \dots \sigma_k)$, where ϕ, ψ , and $\sigma_1, \dots, \sigma_k$ are function-free *literals*, i.e., negated or unnegated atomic formulas whose arguments are constants or variables. Such constraints are to be interpreted as saying “In every state, for all values of the variables, if ϕ then ψ , *provided that* σ_1, \dots , and σ_k ”. We assume that the variables occurring in ϕ include all those occurring in ψ and in the supplementary conditions $\sigma_1, \dots, \sigma_k$. The predicate in ϕ is a fluent predicate, while ψ may be fluent or static. However, if ϕ contains variables that do not occur in ψ , then ψ is required to be “upward monotonic”, in the sense that no instances of it can become false ($\neg\psi$ does not unify with effect of any operator; this is certainly true if ψ is static). Finally, we require $\sigma_1, \dots, \sigma_k$ to be static. The following is an example of this type of constraint in the blocks world stating that the table cannot be on any block: $\forall x, y \text{ ON}(x, y) \Rightarrow \text{NEQ}(x, \text{TABLE})$.

Single-valuedness Constraints. An sv-constraint states that the value of a certain predicate argument is unique for any given values of the remaining arguments. An example of an sv-constraint is the following blocks-world constraint stating that any object can be ON at most one other object:

$$\forall x, y, z. (ON(x, y) \wedge ON(x, z)) \Rightarrow y = z.$$

Implicative Constraints + Single-Valuedness Constraints.

These invariants are formed by an implicative constraint and a set of sv-constraints that are simultaneously discovered by DISCOPLAN. We distinguish two cases which require different verification conditions: the case of subsumed variables and the case of non-subsumed variables. The blocks-world constraint

$$((IMPLIES (ON ?*X ?Y) (NOT (CLEAR ?Y))) (NEQ ?Y TABLE))$$

is an example of a combined implicative and sv-constraint for the first case. In general, the implicative constraints we are considering here have as their antecedent a positive literal that contains at least one “starred” variable not occurring in the consequent, and zero or more “unstarred” variables occurring in the consequent. The stars indicate that for all values of the unstarred variables, the antecedent holds for at most one tuple of values of the starred variables.

In the second case we have implications in which both antecedent and consequent contain variables not contained in the other. All such variables are “starred”, while the shared variables are unstarred. An example is the following constraint from the *Logistics* domain:

$$((IMPLIES (AT ?X ?*Y) (NOT (IN ?X ?*Z))) (OBJECT ?X)).$$

This is an *exclusive* state constraint, i.e., it states that no object can simultaneously be AT something and IN something (and in addition an object can be AT no more than one thing, and IN no more than one thing).

Antisymmetry Constraints. Antisymmetry constraints are particular implicative constraints of the form

$$((IMPLIES (P t_1 t_2) (NOT (P t_2 t_1))) \sigma_1 \sigma_2 \dots \sigma_n),$$

where t_1 and t_2 can be constants or universally quantified variables, and $\sigma_1, \dots, \sigma_n$ are supplementary conditions whose variables are a subset of $\{t_1, t_2\}$. An example of an antisymmetry constraint in the blocks world is

$$\forall x, y. ON(x, y) \Rightarrow \neg ON(y, x),$$

i.e., if one object is on another, then the second is not on the first.

OR and XOR Constraints. OR and XOR-constraints are state constraints of the form

$$(((X]OR \phi \psi) \sigma_1 \sigma_2 \dots \sigma_n),$$

where ϕ and ψ are positive fluent literals, such that non-shared variables are existentially quantified, while shared variables are universally quantified, and where the variables in $\sigma_1, \sigma_2, \dots, \sigma_n$ can only be variables shared by ϕ and ψ . An example of an XOR-constraint in the logistics domain is

$$((XOR (AT ?X ?Y) (IN ?X ?Z)) (OBJECT ?X)),$$

stating that in any reachable state, any object is either at some place or in something.

Strict Single-Valuedness and n-Valuedness Constraints.

This type of invariant is a generalization of sv-constraints. A *nv*-constraint states that a certain predicate can be bound to at most n arguments for any given values of the remaining arguments. A strict *nv*-constraint states that a certain predicate is bound to exactly n arguments for any given values of the remaining arguments. An example of a strict *nv*-constraint with $n = 1$ in the blocks world is the invariant stating that any block is on exactly one thing (either another block or the table).

Using “Expanded Operators” to Infer Further Constraints. DISCOPLAN’s package includes routines

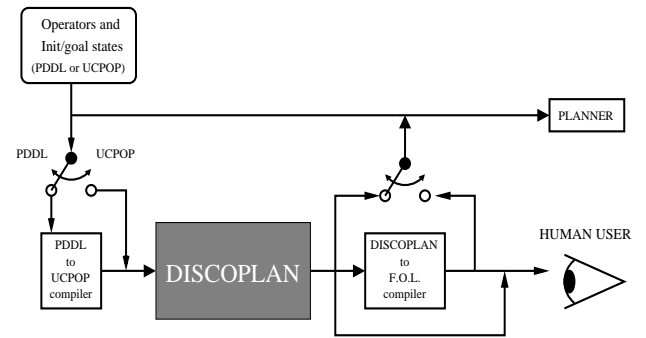


Figure 2: General scheme of DISCOPLAN’s input/output

for expanding an operator with a set of given invariants. The operator expansion consists of enriching the operator description with additional preconditions and effects that are entailed by the given invariants. By using expanded operators DISCOPLAN may infer new invariants, which can be used to expand the operators again. This can be iterated until no new constraints are inferred.

Constraints with Exceptions. Some hypotheses are rejected by DISCOPLAN only because they are not verified against the initial state. For example, consider the blocks world formalization of Figure 1, where we have just one operator in which the parameters are not typed. If we have the simple initial state $((ON A TABLE) (ON C A) (CLEAR C) (ON B TABLE) (CLEAR B))$, DISCOPLAN discovers $(ON ?X ?*Y)$, which then becomes a hypothesis with a strict sv-constraint $((ON ?X ?Y !1)$ in DISCOPLAN format). But $(ON ?X ?Y !1)$ is not confirmed because the test against the initial state fails. This is because the object TABLE is on nothing in the initial state. In order to deal with these *exceptions*, we have recently weakened the test against the initial state, so that a hypothesis can be verified by restricting the domain of certain variables. In our example $(ON ?X ?Y !1)$ can be satisfied in the initial state, provided that $?X$ is not instantiated to TABLE. Hence, DISCOPLAN weakens the hypothesis by excluding TABLE from the domain of $?X$, and derives $((ON ?X ?Y !1) (NOT (MEMBER ?X (TABLE))))$.

These exceptions are computed during the test against the initial state by keeping track of unifiers that assign anomalous tuples of values to the unconstrained variables, i.e., tuples for which strict single-valuedness is violated (e.g., TABLE/?X in the previous example), and weakening the constraint by excluding these values from the domains of the relevant variables.

This analysis of the initial state is also used to derive additional supplementary conditions rescuing hypotheses that were rejected because a required simultaneous sv-constraint was not satisfied in the initial state (while all the other required verification conditions were satisfied). For example, if in the logistics domain the initial state of a problem contains the facts $(AT ORANGES MIAMI)$, $(AT ORANGES ORLANDO)$ and $(OBJECT ORANGES)$, then the invariants

$$((IMPLIES (AT ?X ?*Y) (NOT (IN ?X ?*Z))) (OBJECT ?X))$$

$$((IMPLIES (AT ?X ?*Y) (NOT (IN ?X ?*Z))) (OBJECT ?X) (NOT (MEMBER ?X (ORANGES))))¹$$

¹The complete output for these examples can be seen by running DISCOPLAN *on-line* at the web site <http://prometeo.ing.unibs.it/discoplan>.

Interacting with DISCOPLAN on-line

The general input/output scheme of DISCOPLAN is depicted in Figure 2. The input domain and problem descriptions can be specified using the syntax of either UCPOP or PDDL. Since the core functions of DISCOPLAN assume UCPOP descriptions, when the input is specified using PDDL, it is automatically translated into a UCPOP set of operators.

The output of DISCOPLAN can be given as input to either a planner that can exploit this information, or to a domain developer, as an aid to domain specification and debugging. The syntax of the output can be either FOL or the compact format using implicit quantification and “starred” variables as in the previous sections. The compactness of the starred-variable format is due to the fact that it allows an implicative or exclusive constraint to be augmented with simultaneously discovered sv-constraints merely by starring some variables, rather than adding explicit FOL formulas. The FOL description of the state constraints is obtained by a postprocessing step translating the constraints computed in DISCOPLAN format into FOL.

DISCOPLAN on-line is a version of the system that can be remotely run through any web browser. In particular, from the “test and demo” page of the web site of DISCOPLAN the user can run DISCOPLAN either on a set of predefined domains and problems, or on any other domain and problem that is supplied by the user from her/his local machine (see Figure 3). Before running the system, the user can set some parameters, such as the style of the output, the maximum number of supplementary conditions an invariant can have, the automatic computation of the operator parameter domains using techniques described in [4], etc. Finally, the user can inspect the domain and problem selected.

Related Work and Conclusions

We have sketched how many natural types of state invariants in planning domains with conditional effects can be efficiently inferred, and have described the implementation of our techniques in the DISCOPLAN system. The invariants inferred include predicate domain constraints, relations among static type predicates, implicative constraints, strict and non-strict sv-constraints, combinations of implicative and sv-constraints (where these cannot be inferred in isolation), and OR and XOR constraints. All invariants are found by algorithms that are polynomial-time for any fixed bound on the number of literals in an invariant, and the algorithms can be iterated to find additional invariants after expanding operators using previously found invariants. The outputs can be presented as FOL formulas or in a concise format with implicit universal quantification and “starred” variables indicating single-valuedness. The automatically derived invariants have been shown to radically boost the performance of SAT planners, and are also potentially useful for other planning styles, and as a help in domain analysis and debugging.

Other approaches for the automatic inference of state invariants have been proposed including [10, 9, 3, 13, 14, 12], but to the best of our knowledge the only other implemented system that is available is Fox and Long’s TIM. A major difference between the DISCOPLAN and these approaches is that DISCOPLAN can process domains specified using a more expressive planning language. In particular, TIM does not handle opera-

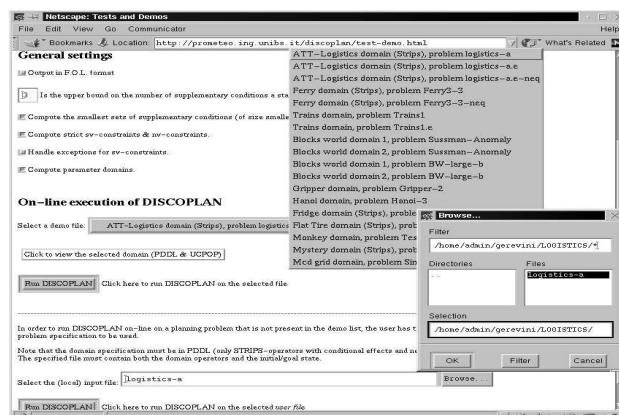


Figure 3: Test and Demo page of DISCOPLAN on-line

tors with conditional effects and negated preconditions. Moreover, DISCOPLAN infers some types of constraints that are not inferred by TIM, such as antisymmetry constraints, XOR-constraints and some implicative constraints involving variable binding constraints or predicates without parameters.² On the other hand, some of TIM’s “state membership invariants” and “uniqueness invariants” are not inferred by the currently implemented version of DISCOPLAN.

It remains unclear how important the “omissions” in each system, relative to the other, are for planning and domain analysis purposes. In any case a reasonable strategy at this time, for builders of planning systems that can benefit from state invariants, would be to combine the invariants found by TIM and DISCOPLAN.

We have developed some further algorithms for inferring invariants, beyond those implemented in DISCOPLAN. The most general of these is an algorithm for inferring n -ary disjunctions of fluent literals, together with sv-constraints and static supplementary conditions, for n not limited to 2 (as at present). This algorithm is a candidate for future implementation.

References

- [1] A. Blum and M.L. Furst. Fast planning through planning graph analysis. In *Proc. of IJCAI-95*, pp. 1636–1642. 1995.
- [2] M.D. Ernst, T.D. Millstein, and D.S. Weld. Automatic SAT-compilation of planning problems. In *Proc. of IJCAI-97*, pp. 1169–1176. 1997.
- [3] M. Fox and D. Long. The automatic inference of state invariants in TIM. *JAIR*, 9:367–421, 1998.
- [4] A. Gerevini and L. Schubert. Accelerating Partial-Order Planners: Some Techniques for Effective Search Control and Pruning. *JAIR*, 5:95–137, Sept. 1996.
- [5] A. Gerevini and L. Schubert. Inferring state constraints for domain-independent planning. In *Proc. of AAAI-98*, pp. 905–912. 1998.
- [6] A. Gerevini and L. Schubert. Inferring state constraints in DISCOPLAN: Some new results. In *Proc. of AAAI-00*, pp. 761–767. 2000.
- [7] M. Ghallab, A. Howe, G. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL – planning domain definition language. Available at <http://cs-www.cs.yale.edu/homes/dwm/>.
- [8] H.A. Kautz and B. Selman. Pushing the envelope: Planning, propositional logic, and stochastic search. In *Proc. of AAAI-96*, 1996.
- [9] G. Kelleher. Determining general consequences of sets of actions. Technical Report TR CMS.14.96, Liverpool Moores University, 1996.
- [10] J. Kelleher and A. Cohn. Automatically synthesizing domain constraints from operator descriptions. In *Proc. of ECAI-92*, pp. 653–655. 1992.
- [11] J.S. Penberthy and D.S. Weld. UCPOP: A sound, complete, partial order planner for ADL. In *Proc. of KR’92*, pp. 103–114. 1992.
- [12] U. Scholz. Extracting state constraints from PDDL-like planning domains. In *Working Notes of the AIPS00 Workshop on Analysing and Exploiting Domain Knowledge for Efficient Planning*, pp. 43–48. 2000.
- [13] J. Rintanen. A planning algorithm not based on directional search. In *Proc. of KR’98*, pp. 617–624. 1998.
- [14] J. Rintanen. An iterative algorithm for synthesizing invariants. In *Proc. of AAAI-00*, pp. 806–811. 2000.

²Examples of these constraints are: ((IMPLIES (ON ?X ?Y) (NEQ ?X ?Y))) in the blocks works, and ((IMPLIES (HASBANANAS) (HASKNIFE))) in the Monkey domain.