

Expressive Equivalence of Formalisms for Planning with Sensing

Jussi Rintanen

Albert-Ludwigs-Universität Freiburg, Institut für Informatik
Georges-Köhler-Allee, 79110 Freiburg im Breisgau
Germany

Abstract

There have been several proposals for expressing planning problems with different forms of uncertainty, including non-determinism and partial observability. In this paper we investigate two questions. First, the restriction to certain normal forms of operators, for example, restricting to operators in which nondeterministic choice must be outside conditional effects, or vice versa. We show that some such restrictions lead to an exponentially less succinct representation of problem instances. Second, we consider the problem of reducing certain features of formalisms for planning problem to other, more basic features. We show that compound observations can be reduced to atomic observations, sensing uncertainty can be reduced to effect uncertainty, dependence of observations on the operator last applied (special sensing actions) can be reduced to the case in which same observations are always possible. We show that these reductions are possible without significantly affecting quantitative properties of problem instances. One reduction doubles plan length, and the others do not affect plan length and only increase problem instance size slightly.

Introduction

There are many planning algorithms for partially observable planning in which the problems discussed in this work show up. These include policy construction algorithms for POMDPs (Sondik 1978; Kaelbling *et al.* 1998) and algorithms for conditional planning (Weld *et al.* 1998; Bonet and Geffner 2000; Bertoli *et al.* 2001; Rintanen 2002). These planners take input in differing input languages, and it has not in all cases been clear what is the exact relation between the problems addressed by these works. For example, Bertoli *et al.* (2001) address compound observations, that is observations the values of which are a Boolean combination of values of individual state variables, whereas Rintanen (2002) does not but claims that compound observations could be handled by an extension of the algorithms he presents.

The present work addresses planning problems that are expressed in terms of state variables and plan operators as used in most of AI planning. Some problems also show up with less succinct enumerative representations of partially

observable planning problems, but the reductions then correspond to less natural transformations of the state space. Early work on planning with partial observability and especially the algorithms for POMDPs have used explicit (flat, enumerative) representations of state spaces. Another class of succinct representations of state spaces for probabilistic planning uses Bayesian networks (Littman 1997; Boutilier *et al.* 1999).

In this framework we set out to investigate relations between observation models with different properties. For different advanced properties an observation model in planning can have, we show that these properties can be reduced to a basic model in which a fixed set of state variables are observable at all time points (assuming certain basic features in the basic language, like conditional effects and arbitrary formulae as preconditions.) Extensions to this basic model like compound observations or observations dependent on the last operator application do not bring additional expressivity or complexity on top of the basic problem. Hence, this work justifies once and for all the restriction to the basic model of partially observable planning with only atomic observations and no sensory uncertainty. The existence of this kind of reductions seems to be part of the folklore in planning but they have not earlier been formalized in more detail. The ideas underlying the reductions are simple, but in some cases small technical problems necessitate a complicated looking reduction.

The only problematic case is reduction of compound observations to atomic observations. Reducing it to the basic model seems to require an increase in plan length when no restrictions on nondeterministic effects are imposed.

We consider the following properties of formalisms for planning.

1. Compound observations means that the value of a formula can be observed without necessarily the possibility of observing the values of its subformulae.
2. By heterogeneous observability we mean that the observability of some state variables depends on the last action that has been taken. For example, certain observations are possible only after special sensing actions.
3. Sensing uncertainty means that an observation gives correct information only with a certain probability. Uncertainty can be decreased by repeating the observation.

Also, we consider a number of normal forms of operators. For example, Kushmerick et al. (1995) and Smith and Weld (1998) restrict to operators with a single nondeterministic choice between deterministic effects. We show that this is exponentially less succinct than the general form of nondeterministic effects we consider.¹ On the other hand, all effects can be transformed in polynomial time into a normal form in which several independent nondeterministic choices can be done in parallel.

The results of the work can be used in different ways. First, on theoretical work on planning under uncertainty, the results justify considering just a basic problem definition without all the bells and whistles, as the more general problem definitions can be reduced to the basic definition. Second, on more practical level, planner implementations can use the reductions to compile away more advanced features of the input language.

The structure of the paper is as follows. We first present our formal framework which consists of the definition of problem instances in planning with uncertainty. We discuss a number of normal forms of nondeterministic operators. Then we describe two reductions of compound observations to atomic observations, a reduction for making all operators have the same observability properties, and finally a reduction of sensing uncertainty to effect uncertainty. We conclude the paper by discussing related work and new research directions.

The Formal Framework

In this section we define the type of problem instances and planning problems we address in this paper. Although we explicitly consider so-called reachability goals only, the results also hold for more general types of plan quality criteria, for example based on rewards. For relations between reachability goals and rewards see work by Condon (1992).

Definition 1 *Let A be a set of state variables. A problem instance in planning is $\langle I, O, G \rangle$ where I is a formula over A describing the initial states, G is a formula over A describing the goal states, and O is a set of operators $\langle c, e, b \rangle$ where c is a formula over A describing the precondition, b is the set of observations $\langle p, \beta \rangle$ where p is success probability $0 < p \leq 1$ and β is a formula over A , and e is an effect. Effects are recursively defined as follows.*

1. a and $\neg a$ for state variables $a \in A$ are effects.
2. $e_1 \wedge \dots \wedge e_n$ is an effect if e_1, \dots, e_n are effects (the special case with $n = 0$ is the empty conjunction \top .)
3. $c \triangleright e$ is an effect if c is a formula over A and e is an effect.
4. $p_1 e_1 | \dots | p_n e_n$ is an effect if e_1, \dots, e_n for $n \geq 2$ are effects, $p_i > 0$ for all $i \in \{1, \dots, n\}$ and $\sum_{i=1}^n p_i = 1$.

If the operator does not contain $p_1 e_1 | \dots | p_n e_n$ then the operator is deterministic.

¹Here we consider reducing one operator to another. If one operator may be replaced by several and plan length and structure do not have to be preserved, the exponential size increases discussed in the connection with the normal forms can be avoided.

The observations are used in the plans that are solutions to the problem instances. After applying an operator, the respective observations can be made. Based on the values of observations β , the plan determines the course of actions to be taken. The probability associated with an observation indicates how likely is the observation going to be correct. Observation uncertainty in the POMDP context (Kaelbling et al. 1998) has been defined as a conditional probability of how likely a given observation is in a given state.

Without loss of generality we assume that a plan always starts with an operator application, and we therefore do not have to specify what is initially observable.

Definition 1 without nondeterminism and observations corresponds to languages commonly used in describing planning problems. A main difference is that languages like PDDL (Ghallab et al. 1998) allow the description of high numbers of operators by one schematic operator description. The schematic description contains variables that are instantiated with constants to obtain operators like considered in the present work. However, almost all planning algorithms are defined in terms of non-schematic operators, and at this level Definition 1 describes a language that is relevant for most of planning research.

We define $ac_a(e)$ as the sufficient and necessary condition for the effect e to change (possibly, taking into account nondeterminism) the truth-value of the state variable a . Formally,

$$\begin{aligned} ac_a(a) &= \top \\ ac_a(\neg a) &= \top \\ ac_a(a') &= \perp \text{ if } a \neq a' \\ ac_a(\neg a') &= \perp \text{ if } a \neq a' \\ ac_a(c \triangleright e) &= c \wedge ac_a(e) \\ ac_a(e_1 \wedge \dots \wedge e_n) &= ac_a(e_1) \vee \dots \vee ac_a(e_n) \\ ac_a(p_1 e_1 | \dots | p_n e_n) &= ac_a(e_1) \vee \dots \vee ac_a(e_n) \end{aligned}$$

Here \perp and \top are respectively the constants false and true.

A related definition, that will be used later, is the following. For a given effect e define a set $as(e)$ of state variables as follows.

$$\begin{aligned} as(a) &= \{a\} \\ as(\neg a) &= \{a\} \\ as(c \triangleright e) &= as(e) \\ as(e_1 \wedge \dots \wedge e_n) &= as(e_1) \cup \dots \cup as(e_n) \\ as(p_1 e_1 | \dots | p_n e_n) &= as(e_1) \cup \dots \cup as(e_n) \end{aligned}$$

This is the set of state variables possibly changed by the effect, or in other words, the set of state variables occurring in the effect not in the antecedent c of a conditional $c \triangleright e$.² We will also use this generalized to sets of effects as $as(\{e_1, \dots, e_n\}) = as(e_1) \cup \dots \cup as(e_n)$.

Allowing the same state variable to be simultaneously affected by several parts of the effect leads to problems. Clearly, effects a and $\neg a$ taking place simultaneously is not well-defined. One could allow the same effect a in two places, but this leads to technical complications that do not justify the very small additional generality. From now on we assume that all operator effects have the following property.

²Notice that $as((a \wedge \neg a) \triangleright b) = \{b\}$, that is, there is no guarantee that the state variable indeed can be changed by the operator.

Property 2 Let Φ be the strongest invariant of a problem instance $P = \langle I, O, G \rangle^3$. Then the effect e of an operator $\langle c, e, b \rangle \in O$ has the following property. Assume a state variable $a \in A$ occurs in two conjuncts e' and e'' of a conjunctive effect $\epsilon = e_1 \wedge \dots \wedge e_n$ occurring in e . Let c' be the conjunction of all the antecedents c'' of conditionals $c'' \triangleright f$ so that ϵ is a subeffect of f . Then $\Phi \cup \{c, c', ac_a(e'), ac_a(e'')\}$ must be inconsistent.

This means that in no state of the problem instance can two occurrences of a state variable simultaneously contribute in the determination of the successor state. Testing the satisfaction of Property 2 is PSPACE-hard⁴, but there are simple sufficient syntactic conditions that can be easily tested and that guarantee its fulfillment. Further, the property can often easily be satisfied by applying equivalences from Table 1.

Next we give a formal semantics for the application of an operator. Each operator assigns a probability distribution to the set of successor states of a state.

Definition 3 (Operator application) Let $\langle c, e, b \rangle$ be an operator over A . Let s be a state, that is an assignment of truth values to A . The operator is applicable in s if $s \models c$.

Recursively assign each effect e a set $[e]_s$ of pairs $\langle p, l \rangle$ where p is a probability $0 < p \leq 1$ and l is a set of literals a and $\neg a$ for $a \in A$.

1. $[a]_s = \{\langle 1.0, \{a\} \rangle\}$ and $[\neg a]_s = \{\langle 1.0, \{\neg a\} \rangle\}$ for $a \in A$.
2. $[e_1 \wedge \dots \wedge e_n]_s = \{\langle \prod_{i=1}^n p_i, \bigcup_{i=1}^n f_i \rangle \mid \langle p_1, f_1 \rangle \in [e_1]_s, \dots, \langle p_n, f_n \rangle \in [e_n]_s\}$.
3. $[c' \triangleright e]_s = [e]_s$ if $s \models c'$ and $[c' \triangleright e]_s = \{\langle 1.0, \emptyset \rangle\}$ otherwise.
4. $[p_1 e_1 \mid \dots \mid p_n e_n]_s = \{\langle p_1 \cdot p, e \rangle \mid \langle p, e \rangle \in [e_1]_s\} \cup \dots \cup \{\langle p_n \cdot p, e \rangle \mid \langle p, e \rangle \in [e_n]_s\}$

Above in (4) the union of sets is defined so that for example $\{\langle 0.2, \{a\} \rangle\} \cup \{\langle 0.2, \{\neg a\} \rangle\} = \{\langle 0.4, \{a\} \rangle\}$: same sets of changes are combined by summing their probabilities.

The successor states of s under the operator are ones that are obtained from s by making the literals in f for $\langle p, f \rangle \in [e]_s$ true and retaining the truth-values of state variables not occurring in f . The probability of a successor state is the sum of the probabilities p for $\langle p, f \rangle \in [e]_s$ that lead to it.

Each $\langle p, f \rangle$ means that with probability p the literals that become true are those in f , and hence indicate the probabilities of the possible successor states of s . For any $[e]_s = \{\langle p_1, f_1 \rangle, \dots, \langle p_n, f_n \rangle\}$ the sum of probabilities is $\sum_{i=1}^n p_i = 1.0$.

Example 4 Consider the operator $\langle a, (0.1\neg a \mid 0.9\neg b) \wedge \neg c, \emptyset \rangle$ and a state s with a, b and c true. Now $[e]_s =$

³This means that for all valuations s of state variables in A , $s \models \Phi$ if and only if s can be reached from I with a sequence of operators from O .

⁴This is because testing whether the strongest invariant is consistent with a formula (even with an atomic formula) is PSPACE-hard. This problem is equivalent to the plan existence problem of deterministic full-information (classical) planning.




Figure 1: Hierarchy of planning problems with partial observability as a Hasse diagram.

$\{\langle 0.1, \{\neg a, \neg c\} \rangle, \langle 0.9, \{\neg b, \neg c\} \rangle\}$, and the successor states respectively satisfy $\{\neg a, b, \neg c\}$ and $\{a, \neg b, \neg c\}$. ■

Properties of Observability and Sensing

The three advanced properties of sensing considered in this work are formalized as follows.

Definition 5 A problem instance in planning $\langle I, O, G \rangle$ has the properties C, H, U respectively under the following conditions.

- A problem instance has compound observations (which we denote by C) if there is operator $\langle c, e, b \rangle \in O$ with $\langle p, \beta \rangle \in b$ and $\beta \notin A$, that is, one of the possible observations is not an atomic proposition.
- A problem instance has heterogeneous observations (which we denote by H) if there are operators $\langle c, e, b \rangle \in O$ and $\langle c', e', b' \rangle \in O$ so that $b \neq b'$.
- A problem instance has sensing uncertainty (which we denote by U) if there is operator $\langle c, e, b \rangle \in O$ with $\langle p, \beta \rangle \in b$ such that $p < 1.0$.

The eight combinations of these properties are depicted in Figure 1. We will call those problem instances *basic* that have none of the properties C, H or U . In the second part of the paper we reduce problem instances CHU first to HU , then further to H and finally to *basic* instances.

Normal Forms for Effects

We introduce three normal forms for effects. These normal forms are either used for establishing the results on reductions between classes of problem instances, or they are interesting on their own right. In earlier work on algorithms for planning with uncertainty, different syntactic restrictions have been used. We show how some of the restrictions lead to an exponential blow-up in the size of operators.

Table 1 lists a number of equivalences on effects. Their proofs of correctness with Definition 3 are straightforward. An effect e is equivalent to $\top \wedge e$, and nondeterministic effects and conjunctions of effects can be arbitrarily reordered while preserving equivalence. These trivial equivalences will later be used without explicitly mentioning them, for example in the definitions of the normal forms and when applying Equivalences 7, 8 and 9.

$$\begin{aligned}
c \triangleright (e_1 \wedge \dots \wedge e_n) &\equiv (c \triangleright e_1) \wedge \dots \wedge (c \triangleright e_n) & (1) \\
c \triangleright (e' \triangleright e) &\equiv (c \wedge e') \triangleright e & (2) \\
c \triangleright (p_1 e_1 | \dots | p_n e_n) &\equiv p_1 (c \triangleright e_1) | \dots | p_n (c \triangleright e_n) & (3) \\
(c_1 \triangleright e) \wedge (c_2 \triangleright e) &\equiv (c_1 \vee c_2) \triangleright e & (4) \\
e \wedge (c \triangleright e) &\equiv e & (5) \\
e &\equiv \top \triangleright e & (6) \\
e \wedge (p_1 e_1 | \dots | p_n e_n) &\equiv p_1 (e \wedge e_1) | \dots | p_n (e \wedge e_n) & (7) \\
p_1 (p'_1 e'_1 | \dots | p'_n e'_n) | p_2 e_2 | \dots | p_n e_n &\equiv (p_1 p'_1) e'_1 | \dots | (p_1 p'_n) e'_n | p_2 e_2 | \dots | p_n e_n & (8) \\
p_1 (e' \wedge (c \triangleright e_1)) | p_2 e_2 | \dots | p_n e_n &\equiv (c \triangleright (p_1 (e' \wedge e_1) | p_2 e_2 | \dots | p_n e_n)) \wedge (\neg c \triangleright (p_1 e' | p_2 e_2 | \dots | p_n e_n)) & (9)
\end{aligned}$$

Table 1: Equivalences on effects

Example 6 An example of an effect with nesting of conditionality and nondeterminism in different ways and independent nondeterministic effects is the following.

$$(a \triangleright (0.3b | 0.7(c \wedge f))) \wedge (0.2(d \wedge e) | 0.8(b \triangleright e))$$

We use this effect for demonstrating the normal forms in Examples 8, 11 and 15. ■

Next we define the normal forms.

Unary Conditionality (UC) Normal Form

The first normal form corresponds to moving the conditionals inside so that their consequents are atomic effects. This normal form is very useful in some of the reductions we will give later. This transformation causes only a polynomial increase in the size of the effect.

Definition 7 A effect e is in unary conditionality (UC) normal form if for all $c' \triangleright e'$ in e the effect e' is either a or $\neg a$ for some $a \in A$.

Example 8 The effect from Example 6 transformed into UC normal form is the following.

$$(0.3(a \triangleright b) | 0.7((a \triangleright c) \wedge (a \triangleright f))) \wedge (0.2(d \wedge e) | 0.8(b \triangleright e))$$

Theorem 9 For every effect there is an equivalent one in UC normal form. There is one that has a size that is polynomial in the size of the effect.

Proof: By using Equivalences 1, 2 and 3 in Table 1 we can transform any effect into UC normal form. Applying each equivalence by replacing an instance of the left-hand-side by the right-hand-side reduces the number of occurrences of \wedge , \triangleright and $|$ after \triangleright . When the number is down to zero, the effect is in UC normal form.

In this transformation the conditions c in $c \triangleright e$ are copied in front of the atomic effects. Let m be the sum of the sizes of all the conditions c , and let n be the number of occurrences of atomic effects a and $\neg a$ in the effect. An upper bound on size increase is $O(nm)$, which is polynomial. □

The main application of the UC normal form is the definition of regression, that is, the computation of a formula that represents the possible predecessor states of a given set of states before applying a given operator. If atomic effect a occurs in conjuncts $c \triangleright a$ and $d \triangleright \neg a$ of a deterministic effect in UC normal form, then $c \vee (a \wedge \neg d)$ is true before applying the operator if and only if a is true after applying the operator.

Conditionality (CO) Normal Form

In the UC normal form conditionality is inside nondeterminism. It is also possible to move conditionality outside nondeterminism.

Definition 10 An effect e is in conditionality (CO) normal form if e does not contain⁵ $p_1(e_1 \wedge (c \triangleright e')) | \dots | p_n e_n$.

Example 11 The effect from Example 6 transformed into CO normal form is the following.

$$(a \triangleright (0.3b | 0.7(c \wedge f))) \wedge (b \triangleright (0.2(d \wedge e) | 0.8e)) \wedge (\neg b \triangleright (0.2(d \wedge e) | 0.8\top))$$

Theorem 12 For every effect there is an equivalent one in CO normal form.

Proof: By applications of Equivalence 9. □

However, contrary to the UC normal form, transformation into CO normal form may increase size exponentially. This reduces the usefulness of this normal form for planner input languages.

Example 13 Consider the effect $\frac{1}{2^n}(a_1 \triangleright \neg a_1) | \frac{1}{2^n}(\neg a_1 \triangleright a_1) | \dots | \frac{1}{2^n}(a_n \triangleright \neg a_n) | \frac{1}{2^n}(\neg a_n \triangleright a_n)$. When the conditional choices have to be done first, 2^n different valuations of a_1, \dots, a_n have to be considered, and for each a variable to be reversed is chosen nondeterministically. Because the variable to be reversed is not known when making the conditional choices first, every valuation of a_1, \dots, a_n has to be considered separately. The size of the effect is proportional to the number of valuations 2^n . ■

⁵Here e_1 may be \top , that is, it may be missing.

Unary Nondeterminism (1ND) Normal Form

Finally, some previous works on planning with nondeterminism, notably Kushmerick et al. (1995) and Smith and Weld (1998), restrict to operator effects with a single nondeterministic choice between deterministic effects. This is strictly more restrictive than the UC normal form.

Definition 14 An effect $\langle c, e, b \rangle$ is in unary nondeterminism (1ND) normal form if e is deterministic or of the form $p_1 e_1 | \dots | p_n e_n$ and every e_i is deterministic.

Example 15 The effect from Example 6 transformed into 1ND normal form is the following.

$$\begin{array}{l|l} (0.06((a \triangleright b) \wedge d \wedge e) & 0.24((a \triangleright b) \wedge (b \triangleright e)) \\ & 0.14((a \triangleright (c \wedge f)) \wedge d \wedge e) \\ & 0.56((a \triangleright (c \wedge f)) \wedge (b \triangleright e)) \end{array}$$

■

Theorem 16 For every effect there is an equivalent one in 1ND normal form.

Proof: By first translating into UC normal form and then applying Equivalences 7 and 8. □

Like in the CO normal form, the size of the effect might increase exponentially, and therefore one should not expect to be able to efficiently transform all nondeterministic effects into this normal form.

Example 17 Consider the effect $e = (0.5a_1 | 0.5\neg a_1) \wedge \dots \wedge (0.5a_n | 0.5\neg a_n)$. An operator $\langle c, e, b \rangle$ produces 2^n successor states for any state. A single nondeterministic choice producing the same successor states must have 2^n alternatives. ■

Notice that exponential size increase can be avoided in transformation into normal forms like 1ND if plan structure does not have to be preserved. A conjunction like that in Example 17 could be split to n operators the application of which is enforced to take place sequentially, in each case making a nondeterministic choice between two alternatives.

Definition of Plans

In this section we give a basic definition of conditional plans. We will not extensively use this definition and it is only superficially referred to in some of the proofs in the following sections, and other reasonable definitions of plans could be used just as well.

Plans are directed graphs with nodes of degree 1 labeled with operators and edges from nodes of degree ≥ 2 labeled with formulae.

Definition 18 Let $\langle I, O, G \rangle$ be a problem instance in planning. A plan is a triple $\langle N, r, l \rangle$ where

- N is a finite set of nodes,
- $r \in N$ is the initial node,

- $l : N \rightarrow (O \times N) \cup 2^{\mathcal{L} \times N}$ is a function that assigns each node an operator and a successor node $\langle o, n \rangle \in O \times N$ or a set of conditions and successor nodes $\langle \phi, n \rangle$.

The formulae ϕ are Boolean combinations of formulae β such that $\langle p, \beta \rangle$ is an observation for every operator that may be the last one applied before reaching the branch node.

The definition of plan execution should be intuitively clear. When a successor node of a branch node is chosen, the result is undefined if more than one condition is true.

Solutions to problem instances $\langle I, O, G \rangle$ are evaluated in terms of their probability of reaching a state in G when starting from a state in I .

Reduction of Compound Observations to Atomic Observations

In the rest of the work we reduce more complex forms of sensing to more basic ones. We start by eliminating compound observations. A compound observation means observing the truth-value of a formula without (necessarily) being able to observe the truth-values of its subformulae. We show how compound observations can be reduced to atomic observations, that is, observations of state variables only. These reductions are based on modifying the operators so that values of compound observations are copied to state variables that are observable.

We give two reductions.

The first reduction preserves the size and the structure of plans but it is restricted to the case in which two state variables occurring in the same compound observation may not be affected by two nondeterministic effects that occur independently. This restriction is imposed to have a polynomial-size reduction.

Condition 19 For no two state variables $a \in A$ and $a' \in A$ occurring in a compound observation β does the operator effect contain $e_1 \wedge \dots \wedge e_n$ with occurrences of $p' e' | \dots$ and $p'' e'' | \dots$ respectively in conjuncts e_i and $e_j, i \neq j$ such that a occurs in e_i and a' occurs in e_j .

This condition is trivially fulfilled by deterministic operators and by operators in 1ND normal form, and we believe that also almost all planning problems occurring in practice fulfill it.

The next example suggests that operators that do not satisfy this condition might be difficult to handle within one operator.

Example 20 Consider the effect $(0.5a_1 | 0.5\neg a_1) \wedge \dots \wedge (0.5a_n | 0.5\neg a_n)$ and an observation β in which all of a_1, \dots, a_n occur. Evaluating the new value of β may require knowing all of the n nondeterministic outcomes, and this seems possible only if this evaluation is performed within one nondeterministic effect that determines values of all a_1, \dots, a_n (this requires the exponential blow-up in the size of the effect) or within another operator that is applied after the operator in question. ■

The second reduction works for all operators but requires auxiliary operators for evaluating the compound observations and copying their values into observable state variables, one for every operator.

Both reductions are computable in polynomial time. The first does not affect plan length, and the second increases plan length by a factor of 2.

First Reduction

Let $P = \langle I, O, G \rangle$ be a problem instance. We define $C_1(P) = \langle I, O', G \rangle$ where O' is defined by replacing every operator in O with a modified one.

Let $\Theta = \{\beta_1, \dots, \beta_n\}$ be the formulas in the observations in the operators of the problem instance. We introduce new auxiliary state variables $A_\Theta = \{a_{\beta_1}, \dots, a_{\beta_n}\}$.

We first transform the operators into UC normal form, then apply Equivalence 7 to not have conjunctions of both deterministic and nondeterministic effects, and further apply Equivalences 4 and 6 so that all conjuncts of conjunctions of deterministic effects are conditionals $c \triangleright l$ and each such conjunction contains at most one $c \triangleright a$ and at most one $c' \triangleright \neg a$ for any $a \in A$. Then we replace $\langle c, e, b \rangle$ by $\langle c, fd(e), \{\langle p, a_\beta \rangle \mid \langle p, \beta \rangle \in b\} \rangle$ where

$$\begin{aligned} fd(p_1 e_1 \mid \dots \mid p_k e_k) &= p_1 fd(e_1) \mid \dots \mid p_k fd(e_k) \\ fd(e_1 \wedge \dots \wedge e_k) &= fd(e_1) \wedge \dots \wedge fd(e_k) \\ &\text{if } e_i \text{ are nondeterministic} \\ fd(e_1 \wedge \dots \wedge e_k) &= e_1 \wedge \dots \wedge e_k \wedge f(\{e_1, \dots, e_k\}) \\ &\text{otherwise} \end{aligned}$$

Here $f(S)$ copies the values of the old compound observations to the new observable state variables. Condition 19 is essential: it guarantees that on any application of the operator, if the value of a compound observation is changed by S , it is not changed by any other part of the operator effect.

$$\begin{aligned} f(S) &= \bigwedge \{f_\beta(S) \mid \beta \in \Theta, a \in as(S), a \text{ occurs in } \beta\} \\ f_\beta(S) &= (\beta[\phi_{a_1}/a_1, \dots, \phi_{a_m}/a_m] \triangleright a_\beta) \wedge \\ &\quad (\neg \beta[\phi_{a_1}/a_1, \dots, \phi_{a_m}/a_m] \triangleright \neg a_\beta) \end{aligned}$$

Here ϕ_a is the regressed value of $a \in A = \{a_1, \dots, a_m\}$, that is, a formula whose value in the predecessor state equals the value of a in the successor state. It is defined as follows.

If a is not the consequent of any conjunct of S , add $\perp \triangleright a$ to S . If $\neg a$ is not the consequent of any conjunct of S , add $\perp \triangleright \neg a$ to S . Hence there is exactly one conjunct $z \triangleright a$ and exactly one conjunct $z' \triangleright \neg a$ in S . Now we define

$$\phi_a = (a \wedge \neg z') \vee z.$$

The formula $(a \wedge \neg z') \vee z$ says that either a was true before and did not become false, or it became true, and hence it expresses the value of a after applying the operator in terms of its value before applying the operator.

In summary, we modify the effects to copy the new values of the compound observations into the new observable state variables a_β .

Example 21 Consider the effect

$$(0.5(a \wedge (g \triangleright b)) \mid 0.5(g \triangleright a)) \wedge c$$

and the observations $\beta_1 = a \wedge b \wedge e$ and $\beta_2 = c \wedge \neg d$. The deterministic conjunct c is moved inside the nondeterministic choice by using Equivalence 7.

$$0.5(a \wedge (g \triangleright b) \wedge c) \mid 0.5((g \triangleright a) \wedge c)$$

Consider $a \wedge (g \triangleright b) \wedge c$. Now $\phi_a = (a \wedge \neg \perp) \vee \top \equiv \top$, $\phi_b = (b \wedge \neg \perp) \vee g \equiv b \vee g$, and $\phi_c = (c \wedge \neg \perp) \vee \top \equiv \top$. The regressed value of β_1 is therefore $\top \wedge (b \vee g) \wedge e$, and the regressed value of β_2 is $\top \wedge \neg d$.

Consider $(g \triangleright a) \wedge c$. Now $\phi_a = (a \wedge \neg \perp) \vee g \equiv a \vee g$, $\phi_b = (b \wedge \neg \perp) \vee \perp \equiv b$, and $\phi_c = (c \wedge \neg \perp) \vee \top \equiv \top$. The regressed value of β_1 is therefore $(a \vee g) \wedge b \wedge e$, and the regressed value of β_2 is $\top \wedge \neg d$. Finally, the new effect is

$$\begin{aligned} 0.5 & (a \wedge (g \triangleright b) \wedge c \\ & \wedge (((b \vee g) \wedge e) \triangleright p_{\beta_1}) \wedge (\neg d \triangleright p_{\beta_2})) \\ \mid 0.5 & ((g \triangleright a) \wedge c \\ & \wedge (((a \vee g) \wedge b \wedge e) \triangleright p_{\beta_1}) \wedge (\neg d \triangleright p_{\beta_2})) \end{aligned}$$

Theorem 22 Let $P = \langle I, O, G \rangle$ be a problem instance that satisfies Condition 19. Then P has a solution plan of length n if and only if $C_1(P)$ has.

Proof: We only give a brief proof sketch.

Plans for P and $C_1(P)$ can be mapped to each other simply by interchanging the operators with their counterparts in P and $C_1(P)$. This is because the differences in the operators involve the new auxiliary state variables only. These are referred to in the observations the plans use, and do not affect the applicability or other effects of the operators.

That observing β can be replaced by observing p_β is because under Condition 19, the formula $\beta[\phi_{a_1}/a_1, \dots, \phi_{a_m}/a_m]$ in the definition of $f_\beta(S)$ always has the same truth-value as β in the next time point, assuming that those nondeterministic choices are made that lead that subeffect to take place.

Condition 19 guarantees that the new truth-value of β is affected only by one maximal deterministic subeffect of the operator's effect. \square

Second Reduction

This reduction works for all problem instances and is simpler but doubles the length of operator sequences. Plan length increase may substantially decrease the efficiency of some types of algorithms.

Let $P = \langle I, O, G \rangle$ be a problem instance. Let β_1, \dots, β_n be the formulas in the observations in the operators in O . We introduce new auxiliary state variables $a_{\beta_1}, \dots, a_{\beta_n}$, z , and z_i for every $i \in \{1, \dots, m\}$ where m is the number of operators in O . State variable a_β will have the same value as the respective observation β . The state variables z_i control the application of operators that evaluate the values of compound observations. Now $C_2(P) = \langle I \wedge z \wedge \neg z_1 \wedge \dots \wedge \neg z_m, O', G \rangle$ where O' consists of

$$\begin{aligned} & \langle z \wedge c, \neg z \wedge z_i \wedge e, \emptyset \rangle \\ & \langle z_i, z \wedge \neg z_i \wedge \nu, \{\langle p, a_\beta \rangle \mid \langle p, \beta \rangle \in b\} \rangle \end{aligned}$$

for every operator $\langle c, e, b \rangle \in O$ (operator's index is i). The first operator replaces the old operator and the second evaluates the observations.

The effect ν evaluates the observations β that are possible after applying $\langle c, e, b \rangle$ and assigns the values to the corresponding state variables a_β :

$$\nu = \bigwedge \{ (\beta \triangleright a_\beta) \wedge (\neg\beta \triangleright \neg a_\beta) \mid \langle p, \beta \rangle \in b \}$$

The state variable z indicates that any operator can be applied, z_i indicates that operator i has been applied and the corresponding observations are to be evaluated.

Example 23 Below we have a simple problem instance with compound observations and its translation to a problem instance with atomic observations only.

$$\begin{aligned} P &= \langle a, \{ \langle a, b \triangleright c, \{ \langle 0.5, a \vee b \rangle \} \} \rangle, c \rangle \\ C_2(P) &= \langle a \wedge z \wedge \neg z_1, \\ &\quad \{ \langle z \wedge a, \neg z \wedge z_1 \wedge (b \triangleright c) \rangle, \emptyset \}, \\ &\quad \langle z_1, z \wedge \neg z_1 \wedge ((a \vee b) \triangleright a_{a \vee b}) \\ &\quad \quad \quad \wedge (\neg(a \vee b) \triangleright \neg a_{a \vee b}), \\ &\quad \quad \quad \{ \langle 0.5, a_{a \vee b} \rangle \} \rangle, \\ &\quad c \rangle \end{aligned}$$

■

Theorem 24 *Let P be a problem instance. Then P has a solution plan if and only if $C_2(P)$ has, and for every plan for P of length n , there is a plan for $C_2(P)$ of length $2n$.*

Proof: We only give a brief proof sketch.

The reduction guarantees that the value of a_β coincides with the value of β after an operator with observation β has been applied.

Translations from plans for P to plans for $C_2(P)$ and vice versa do not require changing the structure of the plans, only one operator is interchanged with a sequence of two operators, or vice versa.

Plans for P can be translated into plans for $C_2(P)$ simply by replacing observations of β by observations of a_β and replacing every operator with the corresponding two new operators.

Plans for $C_2(P)$ can be transformed into plans for P by replacing observations a_β by β and making the converse replacement of operators. The two operator in $C_2(P)$ obtained from one operator in P always appear together, and plans cannot branch between them because the first operator does not have anything observable. □

Despite the increase in plan length, we believe that in many cases the reduction would be almost as practical as the first one. First, many types of algorithms do not suffer from the plan size increase because the new pairs of operators that double plan length always have to be applied together, and therefore no increase in the size of the search space would follow for example in algorithms that construct plans in a strict forward or backward direction. Second, the operators evaluating observations could be restricted to cases in which a problematic compound observation might have been affected, and non-problematic observations could be handled like in the reduction C_1 .

Reduction of Sensing Uncertainty to Effect Uncertainty

The POMDP model of probabilistic planning is often defined to allow sensing uncertainty (Kaelbling *et al.* 1998) so that making an observation in a given state can have any probability between 0 and 1. Sensing uncertainty may for example be caused by imperfections in sensors. In this section we show how sensing uncertainty can be reduced to nondeterminism in action effects. So, in the presence of nondeterminism no additional expressivity is obtained by including sensing uncertainty to the input language.

The reduction uses nondeterministic effects to copy sense data into observable state variables. The observations are reliable but copying introduces uncertainty about sense data.

We assume that all observations in P are atomic, as justified by the reductions in the previous section, and that the effects are in UC normal form.

Let $P = \langle I, O, G \rangle$ be a problem instance. We define $U(P) = \langle I, O', G \rangle$ where

$$O' = \{ \langle c, e' \wedge f(b), \{ \langle 1, a' \rangle \mid \langle p, a \rangle \in b \} \mid \langle c, e, b \rangle \in O \}.$$

Here a' for every observable $a \in A$ is a new state variable to which we copy the current value of a . Copying takes place correctly with probability p and incorrectly with probability $1 - p$.

Next we define f and e' . The function f handles observations that are not changed by e , and the rest of the observations are handled in e' .

Let $N = \{a_1, \dots, a_n\}$ be the state variables that are observable with some of the operators. Let $B = \{a'_1, \dots, a'_n\}$. Define

$$\begin{aligned} rnd_1(a) &= a \\ rnd_1(\neg a) &= \neg a \\ rnd_p(a) &= pa \mid (1-p)\neg a \\ rnd_p(\neg a) &= p\neg a \mid (1-p)a \end{aligned}$$

for state variables $a \in A$ and real numbers $p, 0 < p < 1$. Hence $rnd_p(l)$ sets literal l true by probability p and false by probability $1 - p$.

Now the old observable state variables $a \in A$ that are not affected by the operator are simply copied to the new observable state variables $a' \in B$. We define $f(b)$ as the conjunction of $(a \triangleright rnd_p(a')) \wedge (\neg a \triangleright rnd_p(\neg a'))$ for all $\langle p, a \rangle \in b$ such that $a \notin as(e)$.

The rest of this section describes how those observable state variables that may be affected by the operator are handled in e' . The difficulty is that a state variable may occur in several conjuncts in the effect, but copying may take place in only one to avoid violating Property 2.

Clearly, copying whenever the state variable changes is not a problem because change may take place in only part of an effect at a time. The problem is to decide where to copy when none of the possible changes, for example by effects $c \triangleright a$, takes place. There could be several places in the effect where a state variable potentially changes.

The solution to this difficulty is as follows. For every conjunction $e_1 \wedge \dots \wedge e_n$ we perform copying in the no-change case only in e_1 . When change takes place, copying is done wherever the state variable changes. In the following

we parenthesize all conjunctions $e_1 \wedge e_2 \wedge e_3 \wedge \dots \wedge e_n$ to $e_1 \wedge (e_2 \wedge (e_3 \wedge \dots \wedge e_n))$.

We define $e' = r_{\{\langle \top, a \rangle | a \in as(e), \langle p, a \rangle \in b\}}(e)$ where $r_Q(e)$ is recursively defined below. The set Q consists of elements $\langle c, a \rangle$ that describe the conditions c under which the state variable a is copied to a' .

$$\begin{aligned} r_Q(e_1) &= e_1 \wedge \bigwedge_{\langle c, a \rangle \in Q} cp_{e_1, c}(a) \\ &\quad \text{if } e_1 \text{ is deterministic} \\ r_Q(e_1 \wedge e_2) &= r_{\{\langle c \wedge \neg ac_a(e_2), a \rangle | \langle c, a \rangle \in Q\}}(e_1) \wedge \\ &\quad r_{q \cup q'}(e_2) \\ \text{where } q &= \{\langle c, a \rangle \in Q | a \in as(e_2) \setminus as(e_1)\} \\ \text{and } q' &= \{\langle \perp, a \rangle | a \in as(e_2) \cap as(e_1)\} \\ r_Q(p_1 e_1 | \dots | p_n e_n) &= p_1 r_Q(e_1) | \dots | p_n r_Q(e_n) \end{aligned}$$

In the first case, when e_1 is deterministic, the effect e_1 is extended with copying the state variables a mentioned in Q . The second case is the conjunction, the definition of which is slightly complicated because of the need to avoid violating Property 2. Any a may be copied in only one conjunct. If a variable is shared by e_1 and e_2 , then it is copied in e_1 if it cannot change in e_2 (also when it does not change in e_1), and it is copied in e_2 exactly when it can change in it. In the third case, nondeterministic choice, the same state variables – as indicated by Q – are copied in every alternative although different state variables may change in them.

Copying of observable state variables a in $cp_{e, c}(a)$ has several cases because the value that is copied to a' depends on how the value of a is changed by e .

$$\begin{aligned} cp_{e, c}(a) &= (pc_e(a) \triangleright rnd_p(a')) \wedge \\ &\quad (pc_e(\neg a) \triangleright rnd_p(\neg a')) \wedge \\ &\quad ((a \wedge c \wedge \neg pc_e(\neg a) \wedge \neg pc_e(a)) \triangleright rnd_p(a')) \wedge \\ &\quad ((\neg a \wedge c \wedge \neg pc_e(\neg a) \wedge \neg pc_e(a)) \triangleright rnd_p(\neg a')) \\ &\quad \text{if } \langle p, a \rangle \in b \\ pc_e(l) &= \begin{cases} \top & \text{if } e \text{ has conjunct } l \\ c & \text{if } e \text{ has conjunct } c \triangleright l \\ \perp & \text{otherwise} \end{cases} \end{aligned}$$

Meanings of the above notations are summarized as follows.

$r_Q(e)$ Effect e modified to copy the values of state variables $a \in A$ to the observable state variables $a' \in B$.

$cp_{e, c}(a)$ Copying the value of one observable state variable $a \in A$ to a' . A prerequisite for copying is that the variable cannot be affected anywhere outside e (the condition c).

The four conjuncts describe the four possible cases: a becomes true, a becomes false, a does not change and was true, and a does not change and was false. The prerequisite c is needed in the no-change cases to guarantee that a' is indeed changed only in one part of an operator effect, thus satisfying Property 2.

$pc_e(l)$ Condition for literal l becoming true in e .

Now, whenever one of the old observable state variables a with $\langle p, a \rangle \in b$ is changed, its value is copied to the new observable variable a' with probability p and negated with probability $1 - p$. If the state variable does not change, the old value is copied similarly, and this copying takes place in only one part of the effect to preserve Property 2.

Example 25 Consider the following operator.

$$o = \langle \top, (0.5(\phi \triangleright a) | 0.5b) \wedge (0.5(\theta \triangleright a) | 0.5c), \{\langle 0.9, a \rangle, \langle 0.8, b \rangle, \langle 0.8, c \rangle\} \rangle$$

Now we derive the new operator o' corresponding to o as follows. We do not spell out the new effect completely because even in a simple case like this it would look rather complicated.

$$\begin{aligned} Q &= \{\langle \top, a \rangle, \langle \top, b \rangle, \langle \top, c \rangle\} \\ Q' &= \{\langle \top \wedge \neg \theta, a \rangle, \langle \top \wedge \neg \perp, b \rangle, \langle \top \wedge \neg \top, c \rangle\} \\ Q'' &= \{\langle \perp, a \rangle, \langle \top, c \rangle\} \\ o' &= \langle \top, r_Q((0.5(\phi \triangleright a) | 0.5b) \wedge \\ &\quad (0.5(\theta \triangleright a) | 0.5c)), \\ &\quad \{\langle 1, a' \rangle, \langle 1, b' \rangle, \langle 1, c' \rangle\} \rangle \\ &= \langle \top, r_{Q'}(0.5(\phi \triangleright a) | 0.5b) \wedge \\ &\quad r_{Q''}(0.5(\theta \triangleright a) | 0.5c), \\ &\quad \{\langle 1, a' \rangle, \langle 1, b' \rangle, \langle 1, c' \rangle\} \rangle \\ &= \langle \top, (0.5r_{Q'}(\phi \triangleright a) | 0.5r_{Q'}(b)) \wedge \\ &\quad (0.5r_{Q''}(\theta \triangleright a) | 0.5r_{Q''}(c)), \\ &\quad \{\langle 1, a' \rangle, \langle 1, b' \rangle, \langle 1, c' \rangle\} \rangle \end{aligned}$$

Here for example $r_{Q'}(\phi \triangleright a)$ yields the following.

$$\begin{aligned} r_{Q'}(\phi \triangleright a) &= (\phi \triangleright a) \wedge cp_{\phi \triangleright a, \top \wedge \neg \theta}(a) \\ &\quad \wedge cp_{\phi \triangleright a, \top \wedge \neg \perp}(b) \wedge cp_{\phi \triangleright a, \top \wedge \neg \top}(c) \\ &\equiv (\phi \triangleright a) \wedge cp_{\phi \triangleright a, \neg \theta}(a) \\ &\quad \wedge cp_{\phi \triangleright a, \top}(b) \wedge cp_{\phi \triangleright a, \perp}(c) \\ &\equiv (\phi \triangleright a) \wedge \\ &\quad (\phi \triangleright rnd_{0.9}(a')) \wedge \\ &\quad (\perp \triangleright rnd_{0.9}(\neg a')) \wedge \\ &\quad ((a \wedge \neg \theta \wedge \neg \perp \wedge \neg \phi) \triangleright rnd_{0.9}(a')) \wedge \\ &\quad ((\neg a \wedge \neg \theta \wedge \neg \perp \wedge \neg \phi) \triangleright rnd_{0.9}(\neg a')) \\ &\quad \wedge cp_{\phi \triangleright a, \top}(b) \wedge cp_{\phi \triangleright a, \perp}(c) \\ &\equiv (\phi \triangleright a) \wedge \\ &\quad (\phi \triangleright (0.9a' | 0.1\neg a')) \wedge \\ &\quad ((a \wedge \neg \theta \wedge \neg \phi) \triangleright (0.9a' | 0.1\neg a')) \wedge \\ &\quad ((\neg a \wedge \neg \theta \wedge \neg \phi) \triangleright (0.9\neg a' | 0.1a')) \\ &\quad \wedge cp_{\phi \triangleright a, \top}(b) \wedge cp_{\phi \triangleright a, \perp}(c) \end{aligned}$$

The effects obtained from $cp_{\phi \triangleright a, \top \wedge \neg \theta}(a)$ illustrate the main technical difficulty in the reduction.

$$\begin{aligned} &(\phi \triangleright (0.9a' | 0.1\neg a')) \wedge \\ &((a \wedge \neg \theta \wedge \neg \phi) \triangleright (0.9a' | 0.1\neg a')) \wedge \\ &((\neg a \wedge \neg \theta \wedge \neg \phi) \triangleright (0.9\neg a' | 0.1a')) \end{aligned}$$

They show when the observable state variable a' gets a new value in the first alternative of the first nondeterministic choice in the effect. This is when ϕ is true and a becomes true, or when ϕ is false and a cannot be affected by the second nondeterministic choice in the operator effect because θ is false. ■

Theorem 26 Let $P = \langle I, O, G \rangle$ be a problem instance. Then P has a plan if and only if $U(P)$ has.

Proof: We only give a proof sketch.

The proof is based on lemmata that show that after applying an operator that makes a observable with probability p , the truth-value of a is assigned to a' with probability p

and otherwise the opposite truth-value is assigned to a' with probability $1 - p$.

Mappings of plans for P to $U(P)$ and back are modular: each operator is replaced by its counterpart. Plan structure is preserved because sets of observable state variables stay the same modulo replacement of a by a' or vice versa. \square

Reduction of Heterogeneous Observability to Homogeneous Observability

In this section we show how dependency of observations on the last action taken, that is special sensing actions, is not necessary, and how such problem instances can be reduced to the basic case in which observability is the same at all points of time.

Let $P = \langle I, O, G \rangle$ be a problem instance in planning. We assume that all the observations in P are atomic and have probability 1, which is justified by the reductions in the previous sections. Let $\{a_1, \dots, a_n\} \subseteq A$ be those state variables that are observable with some of the operators in O . We introduce the set $B = \{a'_1, \dots, a'_n\}$ of new state variables. Define a mapping H from P to

$$H(P) = \langle I \wedge \bigwedge_{i=1}^n (a_i \leftrightarrow a'_i), O', G \rangle$$

so that $H(P)$ has homogeneous observability. The set O' is obtained from O by modifying the operators to copy values of old observable state variables a into a' if a was originally observable after applying the operator in question. The operator might also change a and then copying also has to take place. Let

$$O' = \{ \langle c, e' \wedge f(b), \{ \langle 1, a' \rangle | a' \in B \} \rangle | \langle c, e, b \rangle \in O \}$$

where

$$f(b) = \bigwedge_{\langle 1, a \rangle \in b, a \notin as(e)} (a \triangleright a') \wedge (\neg a \triangleright \neg a').$$

We assume that all the operators are in UC normal form and Equivalences 4, 5 have been applied so that no conjunct has more than one occurrence of atomic effect a and $\neg a$ for any $a \in A$.

Similarly to handling observation uncertainty with non-deterministic copying in the previous section, also here we have the problem of satisfying Property 2 when copying the value of a to a' when a does not change. We use the same construct already used in the previous section. The only difference is the function $cp_{e,c}(a)$ which does not randomize when copying from a to a' . Define $e' = r_{\{ \langle \top, a \rangle, a \in as(e), \langle 1, p \rangle \in b \}}(e)$ where $r_Q(e)$ is defined as follows.

$$\begin{aligned} r_Q(e_1) &= e_1 \wedge \bigwedge_{\langle c, a \rangle \in Q} cp_{e_1, c}(a) \\ &\quad \text{if } e_1 \text{ is deterministic} \\ r_Q(e_1 \wedge e_2) &= r_{\{ \langle c \wedge \neg a c_a(e_2), a \rangle | \langle c, a \rangle \in Q \}}(e_1) \wedge \\ &\quad r_{q \cup q'}(e_2) \\ \text{where } q &= \{ \langle c, a \rangle \in Q | a \in as(e_2) \setminus as(e_1) \} \\ \text{and } q' &= \{ \langle \perp, a \rangle | a \in as(e_2) \cap as(e_1) \} \\ r_Q(p_1 e_1 | \dots | p_n e_n) &= p_1 r_Q(e_1) | \dots | p_n r_Q(e_n) \end{aligned}$$

$$cp_{e,c}(a) = \begin{aligned} &(pc_e(a) \triangleright a') \wedge \\ &(pc_e(\neg a) \triangleright \neg a') \wedge \\ &((a \wedge c \wedge \neg pc_e(\neg a) \wedge \neg pc_e(a)) \triangleright a') \wedge \\ &((\neg a \wedge c \wedge \neg pc_e(\neg a) \wedge \neg pc_e(a)) \triangleright \neg a') \end{aligned}$$

$$pc_e(l) = \begin{cases} \top & \text{if } e \text{ has conjunct } l \\ c & \text{if } e \text{ has conjunct } c \triangleright l \\ \perp & \text{otherwise} \end{cases}$$

Theorem 27 Let $P = \langle I, O, G \rangle$ be a problem instance. Then P has a plan if and only if $H(P)$ has.

Proof: We only give a proof sketch.

Plans of P can be transformed to plans of $H(P)$ simply by replacing the operators by their counterparts in $H(P)$ and by replacing observations a by a' .

The transformation from plans of $H(P)$ to P is similar, but there is the possibility that certain observation a' is used later than when the copying from a into a' last took place, that is, at a point of a plan which a is not observable in P . However, in those cases we can transform the plan by moving branching earlier to obtain a valid plan for P . We sketch a small example to illustrate the transformation. Let the plan consist of operator applications o_1 and o_2 followed by branching on a'_1 (corresponding to the observable variable a_1 in P). The problem might be that originally a_1 was not observable after the operator corresponding to o_2 , but only after o_1 . The transformation involves moving the branch earlier: apply o_1 , branch on the value of a'_1 , in both branches first apply o_2 and then continue with the subplans that originally followed the branch. \square

Note that this reduction could be easily combined with the one for reducing sensing uncertainty to nondeterminism: simply have the set B observable for every operator in the reduction in the previous section.

Related Work

Littman (1997) and Boutilier et al. (1999) show the equivalence of different types of representations of nondeterministic actions. In this paper we have used only one definition of nondeterministic operators that is syntactically less restricted than the definition of probabilistic STRIPS operators used by Boutilier et al. and by Littman. These definitions of operators are equally expressive when allowing reductions in which one operator may be replaced with several and in which plan length and structure do not have to be preserved. Neither Littman nor Boutilier et al. address planning with sensing and partial observability.

Expressivity of different formalisms for deterministic planning has earlier been investigated by Bäckström (1995) and Nebel (2000). The topics of interest in classical planning have been the form of preconditions (positive literals only, conjunctions of literals only, DNF, CNF, arbitrary formulae) and conditional effects. In contrast, in this work we consider extended formalisms with the possibility to express nondeterminism and observability.

The reductions presented in this work can be defined as Nebel's (2000) compilation schemes, and, with the exception of the more general reduction for compound observations which doubles the length of operator sequences, also in the ESP reductions framework of Bäckström (1995). Of course, both of these works restrict to deterministic full information planning (classical planning), so the embeddability we allude to here requires generalizations of ESP reductions and compilation schemes to conditional plans. Such generalizations are straightforward.

Conclusions

We have presented three normal forms for plan operators with conditional effects and nondeterminism, shown how arbitrary effects can be transformed into these normal forms, and discussed the size increase of these transformations.

We have also shown how three advanced features of planning with partial observability, compound observations and sensing uncertainty and heterogeneous observability, can be reduced to observability of the same set of atomic state variables at all time points. Thus no additional expressivity or complexity is obtained by these advanced features.

All the reductions are modular, that is, they involve replacing one operator by another one with a closely related behavior, or in some cases by two operators that must be used together. The modularity goes so far that the structures of the plans for the original problem instance and the ones produced by the translation are the same.

The only difficulties in the reductions are caused by conjunctions of nondeterministic effects: inside one operator, there is no way of expressing anything about the combined results of mutually independent nondeterministic effects (without the exponential size reduction of the independent effects to one nondeterministic effect.) At the level of the language for describing the operators all these problems could easily be resolved by allowing a two (or n) phase definition of operator effects, that is, sequential composition of effects. Extending planner input languages with sequential composition might therefore be beneficial.

References

C. Bäckström. Expressive equivalence of planning formalisms. *Artificial Intelligence*, 76(1–2):17–34, 1995.

P. Bertoli, A. Cimatti, M. Roveri, and P. Traverso. Planning in nondeterministic domains under partial observability via symbolic model checking. In B. Nebel, editor, *Proceedings of the 17th International Joint Conference on Artificial Intelligence*, pages 473–478. Morgan Kaufmann Publishers, 2001.

B. Bonet and H. Geffner. Planning with incomplete information as heuristic search in belief space. In S. Chien, S. Kambhampati, and C. A. Knoblock, editors, *Proceedings of the Fifth International Conference on Artificial Intelligence Planning Systems*, pages 52–61. AAAI Press, 2000.

C. Boutilier, T. Dean, and S. Hanks. Planning under uncertainty: structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.

A. Condon. The complexity of stochastic games. *Information and Computation*, 96(2):203–224, 1992.

M. Ghallab, A. Howe, C. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL - the Planning Domain Definition Language, version 1.2. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, Yale University, October 1998.

L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101(1-2):99–134, 1998.

N. Kushmerick, S. Hanks, and D. S. Weld. An algorithm for probabilistic planning. *Artificial Intelligence*, 76(1-2):239–286, 1995.

M. L. Littman. Probabilistic propositional planning: Representations and complexity. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI-97) and 9th Innovative Applications of Artificial Intelligence Conference (IAAI-97)*, pages 748–754, Menlo Park, July 1997. AAAI Press.

B. Nebel. On the compilability and expressive power of propositional planning formalisms. *Journal of Artificial Intelligence Research*, 12:271–315, 2000.

J. Rintanen. Backward plan construction under partial observability. In M. Ghallab, J. Hertzberg, and P. Traverso, editors, *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems*, pages 173–182. AAAI Press, 2002.

D. E. Smith and D. S. Weld. Conformant Graphplan. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98) and the Tenth Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*, pages 889–896. AAAI Press, 1998.

E. J. Sondik. The optimal control of partially observable Markov processes over the infinite horizon: discounted costs. *Operations Research*, 26(2):282–304, 1978.

D. S. Weld, C. R. Anderson, and D. E. Smith. Extending Graphplan to handle uncertainty and sensing actions. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98) and the Tenth Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*, pages 897–904. AAAI Press, 1998.