

Improving Heuristics for Planning as Search in Belief Space

Piergiorgio Bertoli and Alessandro Cimatti

IRST - Istituto Trentino di Cultura, 38050 Povo, Trento, Italy

{bertoli,cimatti}@irst.itc.it

Abstract

Search in the space of beliefs has been proposed as a convenient framework for tackling planning under uncertainty. Significant improvements have been recently achieved, especially thanks to the use of symbolic model checking techniques such as Binary Decision Diagrams. However, the problem is extremely complex, and the heuristics available so far are unable to provide enough guidance for an informed search.

In this paper we tackle the problem of defining effective heuristics for driving the search in belief space. The basic intuition is that the “degree of knowledge” associated with the belief states reached by partial plans must be explicitly taken into account when deciding the search direction. We propose a way of ranking belief states depending on their degree of knowledge with respect to a given set of boolean functions. This allows us to define a planning algorithm based on the identification and solution of suitable “knowledge subgoals”, that are used as intermediate steps during the search. The solution of knowledge subgoals is based on the identification of “knowledge acquisition conditions”, i.e. subsets of the state space from where it is possible to perform knowledge acquisition actions. We show the effectiveness of the proposed ideas by observing substantial improvements in the conformant planning algorithms of MBP.

Introduction

Planning in nondeterministic domains, i.e. with uncertainty in the initial condition, partial observability, and uncertain action effects, is increasingly being recognized as an important research area. In practical domains, the ability to deal with contingencies during the planning process is often required, for instance in order to anticipate unpredictable events and counter possible failures. Significant relationships are also being recognized with the fields of controller synthesis, diagnosis (Thiebaux & Cordier 2001), and Automated Test Pattern Generation (Cimatti, Roveri, & Bertoli 2001). Planning in nondeterministic domains is also one of the most challenging problems, proved to be significantly harder than classical planning (De Giacomo & Vardi 1999). Intuitively, this is due to the fact that a plan can be associated with possibly many traces, and in order to guarantee goal achievement they all must be taken into account.

Copyright © 2002, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Several approaches to planning in nondeterministic domains have been recently tackled (Pryor & Collins 1996; Kabanza, Barbeau, & St-Denis 1997; Weld, Anderson, & Smith 1998; Cimatti, Roveri, & Traverso 1998a; 1998b; Rintanen 1999). One of the most successful research directions recasts planning in nondeterministic domains as search in the space of beliefs (Bonet & Geffner 2000). A condition of uncertainty is represented as a *belief state*, i.e. a set collecting all possible states in which the domain could be. Since the belief space is exponential in the number of states in the domain, the problem turns out to be extremely complex. The use of symbolic techniques, based on Binary Decision Diagrams, allows to stretch the scalability in conformant planning, both for breadth-first search (Cimatti & Roveri 2000) and for heuristic search (Bertoli, Cimatti, & Roveri 2001b), and in conditional planning (Bertoli *et al.* 2001b; Bertoli, Cimatti, & Roveri 2001a).

The most evident limitation of these approaches is that, even in trivial cases, the algorithms may get lost in the search space. This often results in a blow-up in the computational resources needed to solve the problem, and in highly sub-optimal solutions. The main reason for this is in the heuristic functions used to drive the search, that do not provide enough information. In this paper, we tackle the problem of defining more effective heuristic functions for searching the space of beliefs. The basic intuition is that the *degree of knowledge* associated with a certain belief state must be explicitly taken into account. In particular, we propose a way of ranking belief states depending on their degree of knowledge with respect to a given set of boolean functions. We define a new conformant planning algorithm, based on the identification and solution of suitable *knowledge subgoals*, used as intermediate steps to drive the search. The solution of knowledge subgoals is based on the identification of *knowledge acquisition conditions*, i.e. subsets of the state space from where it is possible to perform knowledge acquisition actions. The effectiveness of the proposed approach is demonstrated by a thorough experimental evaluation. In some very small instances of the simpler problems, the new algorithm may be slightly less efficient due to the overhead of reasoning about knowledge subgoals and knowledge acquisition conditions. However, the significance of such an overhead decreases with the growth of the dimension of the problem. Moreover, in any other problems we tackled, the

new algorithm exhibits improved efficiency, and always produces plans of high quality.

This paper is structured as follows. First we provide some background on conformant planning as search in the belief space. Then, we discuss the ideas underlying the definition of more effective selection functions, and we describe the algorithm for conformant planning based on such ideas. In the following section, we discuss the implementation in the setting of Planning via Symbolic Model Checking, and we present an extensive experimental evaluation. In the final section, we draw some conclusions and we discuss some related and future work.

Conformant Planning as Search in Belief Space

We consider nondeterministic planning domains represented as finite state automata. This general model can be used for providing a semantics to languages such as PDDL and its nondeterministic extension, and AR (Giunchiglia, Kartha, & Lifschitz 1997), where parallel/concurrent actions have preconditions, conditional and uncertain effects.

Definition 1 (Planning Domain) A Planning Domain is a 4-tuple $\mathcal{D} = (\mathcal{P}, \mathcal{S}, \mathcal{A}, \mathcal{R})$, where \mathcal{P} is the (finite) set of atomic propositions, $\mathcal{S} \subseteq \text{Pow}(\mathcal{P})$ is the set of states, \mathcal{A} is the (finite) set of actions, and $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{A} \times \mathcal{S}$ is the transition relation.

Intuitively, a proposition is in a state if and only if it holds in that state. In the following we assume that a planning domain \mathcal{D} is given. We use s , s' and s'' to denote states of \mathcal{D} , and α to denote actions. $\mathcal{R}(s, \alpha, s')$ holds iff when executing the action α in the state s the state s' is a possible outcome. We say that an action α is applicable in s iff there is at least one state s' such that $\mathcal{R}(s, \alpha, s')$ holds. We say that an action α is deterministic in s iff there is a unique state s' such that $\mathcal{R}(s, \alpha, s')$ holds. An action α has an uncertain outcome in s if there are at least two distinct states s' and s'' such that $\mathcal{R}(s, \alpha, s')$ and $\mathcal{R}(s, \alpha, s'')$ hold.

Consider for instance the simple navigation domain outlined in the rightmost square in Figure 1 where a robot can move in the four directions a 4x4 square. The example is very simple, and is used here only for explanatory purposes. The positions in the square can be represented by means of two fluents x and y , both ranging from 0 to 3. For this simple domain, the propositions are equalities between the variables and their possible values (e.g. $x = 0$). The black location in (1,2) is a hole that, if entered, will cause the robot to crash. This means, for instance, the GoWest action is not applicable in (0,2). When moving toward a wall, the robot does not move. For instance, if the robot performs a GoSouth action in location (0,0), it will remain in (0,0). The location in (2,1) is a slippery spot, that will make the robot move unpredictably sideways when performing an action in it. This introduces nondeterministic action effects. For instance, the effect of performing a GoWest action starting from state (2,1) may result in any of the states in $\{(3, 0), (3, 1), (3, 2)\}$.

In this paper we consider plans to be sequences of actions, i.e. elements of \mathcal{A}^* . We use ϵ for the 0-length plan, π and

ρ to denote plans, and $\pi; \rho$ for concatenation. Conformant planning is the problem of finding a plan that, if executed in any initial state in $\mathcal{I} \subseteq \mathcal{S}$, takes the domain into a set of states $\mathcal{G} \subseteq \mathcal{S}$, regardless of nondeterministic action effects. Following (Bonet & Geffner 2000), we model this problem as search in the space of *belief states*. A belief state is a set of states, intuitively expressing a condition of uncertainty, by collecting together all the states which are equally possible. The execution of actions is lifted from states to belief states by the following definition.

Definition 2 (Plan Applicability, Execution) An action α is applicable in a belief state $\emptyset \neq Bs \subseteq \mathcal{S}$ iff α is applicable in every state in Bs . If α is applicable in Bs , its execution $\text{Exec}[\alpha](Bs)$ is the set $\{s' \mid \mathcal{R}(s, \alpha, s'), \text{ with } s \in Bs\}$. The execution of a plan π in a belief state, written $\text{Exec}[\pi](Bs)$, is defined as follows.

$$\begin{aligned} \text{Exec}[\epsilon](Bs) &\doteq Bs \\ \text{Exec}[\pi](\emptyset) &\doteq \emptyset \\ \text{Exec}[\alpha; \pi](Bs) &\doteq \emptyset, \text{ if } \alpha \text{ is not applicable in } Bs \\ \text{Exec}[\alpha; \pi](Bs) &\doteq \text{Exec}[\pi](\text{Exec}[\alpha](Bs)), \text{ otherwise} \end{aligned}$$

We say that a plan is applicable in a belief state when its execution is not empty. For a conformant planning problem, solutions are applicable plans, for which all the final states must be goal states. Formally, the problem is defined as follows.

Definition 3 (Conformant Planning Problem) Let $\emptyset \neq \mathcal{I} \subseteq \mathcal{S}$ and $\emptyset \neq \mathcal{G} \subseteq \mathcal{S}$. The plan π is a conformant solution to the problem $\langle \mathcal{I}, \mathcal{G} \rangle$ iff $\emptyset \neq \text{Exec}[\pi](\mathcal{I}) \subseteq \mathcal{G}$.

In the example navigation domain, the problem is for the robot to go from the initial belief state $\{(0, 1), (0, 2)\}$ to the goal belief state $\{(3, 1), (3, 2)\}$. A possible solution to this problem is the plan (GoSouth; three times GoEast; GoSouth; GoNorth): the associated traces are depicted in the middle row in Figure 1. (The hole and the slippery spot are reported only in the first square, while in the following ones only the belief state is reported.) The reader can check that the associated plan is conformant: for all the possible associated traces, the plan never violates the applicability conditions of the attempted actions, and ends in a goal state. Notice that the first action is GoSouth in order to make sure that the preconditions of the following GoWest are met in all the states of the belief state. Notice also the last GoEast action: intuitively, the uncertainty in the position of the robot is increased by the nondeterministic effect of the action. The following GoSouth action is needed to regain some knowledge: the robot will result in (3,3) if it executes GoSouth both in (3,2) and in (3,3).

The problem of *finding* a conformant plan in this setting amounts to searching the belief space, and looking for a path from the initial belief state to the goal belief state. Figure 1 outlines (a subset of) the search space for the example, constructed forward, starting from the initial set of states toward the goal. An expansion step consists in trying all the applicable movement actions. Notice that different plans can result in the same belief state, and that cycles are possible. Notice also that, even for this simple example, several different plans are possible, since the three leftmost belief states are




Figure 2: Different ways of searching the belief space

portant to be able to express the heuristic functions in non-boolean terms, in order not to lose the structure of the problem. A variable is known to have a value in a belief state if this is the case in all the states in the belief state. Boolean functions are treated similarly. The known range is the set of all possible values of a variable in a belief state. Notice how the cardinality of the known range gives a rough but valuable information on how uncertain we are on the variable. For instance, in the above case, the cardinality of the known range of the y variable in the initial belief state is 2, while in the goal belief state it is 1.

A more informed view of a boolean function in a belief state is the following.

Definition 5 (Truth Percentage) *The truth percentage of ϕ in S , written $|\mathcal{K}|_{\top}(\phi, S)$, is defined as $\frac{|\{s \in S | s = \phi\}|}{|S|}$*

The notion of truth percentage gives an idea of how close a belief state is to the truth of a boolean function. Again, this is clearly an abstract view, since there may be correlations that are lost. For instance, the two boolean functions $x = 0$, and $y = 0$ have the same truth percentage in the belief states $\{(0, 0), (0, 1), (1, 0), (1, 1)\}$ and $\{(0, 1), (1, 0)\}$. In some cases, truth percentage can be effectively used to reach a certain belief state, such as the goal, by “hill-climbing”, i.e. trying to greedily explore belief states of increasing truth percentage.

However, simple hill climbing on truth percentage is not enough in problems with more complicate structure. Therefore, we exploit the idea that a certain knowledge is necessary to reach the goal. Similar conditions, that we call “knowledge subgoals”, can be identified by taking an abstract view of the belief space. For instance, in the case of navigation in the empty room, it is possible to detect that there is no way in which the goal could be reached without passing through a belief state where the known range of y has cardinality 1. This kind of information, resulting in knowledge subgoals such as “the known range of x must have cardinality 1”, can be acquired by simple comparison of the known range of domain variables in the initial and goal belief states.

Our approach exploits this idea, and tries to acquire the necessary knowledge before trying to reach the goal. In order to solve knowledge subgoals, we preprocess the domain and identify “knowledge acquisition conditions”, i.e. belief

states where it is possible to execute an action that guarantees the solution of the knowledge subgoal. The motivation is that a knowledge acquisition condition may be easier to reach by greedy search than the associated knowledge subgoal. Formally, a knowledge acquisition condition is defined as follows.

Definition 6 (Knowledge acquisition condition) *The belief state Bs is a knowledge acquisition condition for the boolean function ϕ iff there exists an action such that $\mathcal{K}_{\top}(\phi, Exec[\alpha](Bs))$.*

For instance, the belief state defined by $x \in \{0, 1\}$ and $y \in \{0..7\}$ (corresponding to the width-two rectangle close to the left wall) is a knowledge acquisition condition for the function $x = 0$, since the application of a GoWest action in it guarantees that $x = 0$ is known in the resulting belief state. These notions introduce a way to partition the search in the belief space. Clearly, in this framework we have to face the problem of guessing the right knowledge subgoals, and of identifying whether reasonable knowledge acquisition conditions exists. In practice, we restrict knowledge subgoals to boolean functions of the form $x = v$. We discover the knowledge acquisition conditions associated with a knowledge subgoal by means of the conformant preimage primitive. The conformant preimage is a backward computation that, given a belief state Bs , returns all pairs $\langle \alpha_i . Bs_i \rangle$ such that $Exec[\alpha_i](Bs_i) \subseteq Bs$. From this analysis, sufficient conditions for detecting the unsolvability of a problem based on knowledge degree considerations can be derived. For instance, if we get rid of the walls to create a “circular” navigation domain, the problem of Figure 2 is clearly unsolvable. There is no way to reduce the cardinality of the known range of the variables, and it is therefore impossible to reach a more “informed” portion of the search space.

The Conformant Planning Algorithm

We devised an algorithm for conformant planning based on the ideas outlined in previous section. Figure 3 outlines the Finite State Machine of the control flow of the algorithm. The algorithm is entered by initializing the target knowledge, i.e. a set of knowledge goals that are estimated to be convenient intermediate steps to solve the given problem $\langle \mathcal{I} . \mathcal{G} \rangle$. If the initial belief state \mathcal{I} is such $\mathcal{K}_{\top}(\phi_i, \mathcal{I})$ for




Figure 3: The high-level control flow of the algorithm

each knowledge subgoal ϕ_i in the target knowledge, then the algorithm enters the Ground Search mode, and greedily searches toward the final goal \mathcal{G} . Otherwise, the algorithm enters the Knowledge Acquisition mode, where it tries to reach the knowledge subgoals. In order to do so, the search considers a set of active knowledge subgoals, maintained in a subgoal agenda. The search greedily directs toward the knowledge acquisition conditions associated with the subgoals in the agenda, and then applies the suitable actions to achieve those subgoals. Once the Knowledge Acquisition mode terminates, the algorithm enters in the Ground Search mode, and tries to reach \mathcal{G} . Upon failure in the Ground Search mode, the target knowledge is refined, thus requiring the dynamic update of the knowledge subgoal agenda. When it is impossible to refine the target knowledge, and the list of open belief states is empty (i.e. all the state space has been covered), then the algorithm returns with failure. As other algorithms based on search in the beliefs space, the algorithm described above always terminates, returning a conformant plan if the problem is solvable, and failure otherwise. The conformance of plans is guaranteed by construction, while the termination follows from the complete storage of the traversed portion of search space, and on the finiteness of the belief space. To the best of our knowledge, the explicit identification of a dynamic knowledge agenda, and the revision of a target knowledge, are unique to this algorithm.

The greedy search strategy, applied in the Ground Search mode (while trying to reach the final goal) and in the Knowledge Acquisition mode (while trying to reach the knowledge acquisition conditions), is a forward expansion of the search space, similar to that described in (Bertoli, Cimatti, & Roveri 2001b). Basically, the currently selected belief state Bs is expanded by applying every action α_i that is applicable in it. The expansion returns the belief states of the form $Exec[\alpha](Bs)$, from which the belief states that have been previously generated are eliminated. Each belief state is also associated with a conformant plan for reaching it from \mathcal{I} . The newly generated belief states are merged into the open

list of belief states that deserve further expansion. The list of open belief states is ordered according to the following criteria (applied in decreasing priority):

- maintain previously reached knowledge subgoals;
- minimize the estimated distance to the local goal currently being pursued (i.e. the final goal in Ground Search mode, or a knowledge subgoal in Knowledge Acquisition mode);
- maximize the truth percentage of the local goal currently being pursued;
- in Knowledge Acquisition mode, try not to increase the distance from the final goal.

In order to reduce the cost of frontier analysis, but also to give more penetration to the search, more recent belief states are privileged: “old” sections of the frontier are considered only if the newly generated belief states all decrease current scoring function.

The implementation of the algorithm is based on this high-level control flow, together with a suitable definition of distances and intersections. In order to achieve efficiency, a number of additional refinements have been realized.

- (a) In order to reduce the cost of identifying knowledge subgoals, we only consider a fixed set of simple knowledge subgoals, corresponding to the values for fluent variables of the problem. Though simple, for many domains this seems to be a reasonable choice.
- (b) In order to reduce the cost of identifying knowledge acquisition conditions we adopt a lazy approach. Knowledge acquisition conditions are only searched for when a search at the knowledge level has been triggered. Moreover, we perform an approximation by associating each knowledge subgoal with the union of the knowledge acquisition conditions associated with it (rather than to each knowledge acquisition condition related to an action). Again, this simplification yields to significant speed ups without significantly worsening the level of information of the search.
- (c) In order to reduce the cost of computing the estimate of distances (from a belief state to either the ground goal, or a knowledge acquisition condition), we adopted three ideas. First, we defined the distance $D(Bs, Bs')$ as the sum of distances $D_i(Bs, Bs')$, each based on the projection of the domain over the i -th fluent. Second, we defined $D_i(Bs, Bs')$ as the maximal distance between any state in Bs and any state in Bs' . Statewise distance is computed by approximated reachability. Third, we enforced laziness: distances are only computed if no immediate increasing of the truth percentage of the current local subgoal is possible. The first two ideas, coupled with the choice of knowledge subgoals presented in (a), allow to independently handle more than one knowledge subgoals at once, and to handle a simplified knowledge subgoal agenda representation. The third idea significantly reduces overheads in several cases.

Implementation and Experimental Results

We implemented the algorithm described in previous section in MBP (Bertoli *et al.* 2001a). MBP (Model Based Planner) is a general system for planning in nondeterministic domains, able to tackle conditional planning problems under the hypotheses of full and partial observability, conformant planning, and planning for extended goals. MBP is built on top of the NuSMV model checker (Cimatti *et al.* 1998), and is based on symbolic model checking techniques, Binary Decision Diagrams (Bryant 1992) in particular, to compactly represent and efficiently explore the domain description in form of a Finite State Automaton. Binary decision diagrams (BDDs) are a canonical form for the representation of boolean formulae, and can be used to represent sets of states and their transformations. The implementation of the conformant planning algorithms is based on the main contribution of (Bertoli, Cimatti, & Roveri 2001b), i.e. the integration of BDD-based techniques within a heuristic search framework. Basically, each belief state is represented by a BDD. Visited BDDs are stored in a hash table, and the approach takes full advantage of the ability of the CUDD (Somenzi 1997) package to compactly represent large numbers of BDDs. Furthermore, the expansion and pruning of belief states is carried out by means of logical transformations (corresponding to the application of boolean connectives and quantifications) that avoids the explicit enumeration of possible actions.

This approach provides an effective set of primitives that we also use in the implementation of the algorithm described in this paper. Further details, not reported here for lack of space, can be found in (Cimatti & Roveri 2000) and in (Bertoli, Cimatti, & Roveri 2001b). Here we focus only on the computation primitives that are more closely related to the algorithm described in this paper. In particular, the computation of the heuristic functions is based on BDD-based transformation. It is possible to evaluate $\mathcal{K}_\top(\phi, Bs)$ simply by checking whether the BDD representing Bs entails the BDD representing ϕ . Since the BDD for ϕ is the set of states satisfying ϕ , we are checking if the set of states in Bs is included in the set of states that satisfies ϕ . The truth percentage $|\mathcal{K}_\top(\phi, Bs)|$ is also computed by means of BDD manipulations. $Bs \cap \phi$ is computed as the conjunction of two BDDs. To compute $|Bs|$ and $|Bs \cap \phi|$, we use the primitive that counts the models of (the propositional formula represented by) a given BDD, the so-called “CountMinterms”. Although this operation may appear to be complex, the computation is efficiently carried out by means of a traversal on the structure of the BDD, and is quite feasible in practice. Finally, we estimate of the distance of a belief state Bs from a target belief state by means of a backward reachability analysis, that proceeds breadth-first from the target and computes the sets of states (layers) at increasing distance until Bs is completely included. Since this approach is in general inefficient, it is approximated by means of a relaxed reachability, where the transition relation of the domain is considered variable-wise, abstracting away the interactions among variables. The set of approximate transition relations is also obtained by means of projection operations.

In the following we call KACMBP (Knowledge Acquisi-

tion CMBP) the implementation of the algorithm described in this paper. We evaluated our approach by comparing it with the approach presented in (Bertoli, Cimatti, & Roveri 2001b), in the following referred to as HSCP (Heuristic Conformant MBP). Basically, HSCP performs a backward search, from the goal to the initial state, with a pure heuristic (where the cost of reaching the belief state under analysis is disregarded) that drives the exploration toward belief states of higher cardinality. We limit the comparison to these two systems, that share the same (efficient) symbolic machinery to isolate the differences in the heuristic component. A somewhat indirect comparison of KACMBP with other relevant systems such as GPT, CGP (Smith & Weld 1998), QBFPLAN (Rintanen 1999), and the breadth-first conformant planning algorithm of MBP (Cimatti & Roveri 2000), can be derived from (Cimatti & Roveri 2000) and (Bertoli, Cimatti, & Roveri 2001b), where such systems are tested against (and outperformed by) HSCP.

We carried out an extensive experimental evaluation, by covering the standard problems presented so far in the literature to evaluate conformant planners (i.e. BT, RING, SQUARE and CUBE), and by defining some new ones. (For lack of space, we refer to (Cimatti & Roveri 2000) for the description of the standard problems.) All the experiments were run on an Pentium II 300MHz with 512Mb of memory running Linux, fixing a memory limit of 128Mb and a CPU time limit of 1 hour. For each problem instances, we report both the CPU time (in seconds) required to find a solution for the given planning problem, and the length of the generated plan. Notice that the time scale is logarithmic. The automaton construction times are not reported, being exactly the same in the two approaches. For all the examples reported in this paper, however, they never require more than 0.5 seconds.

The results of the comparison on the classical domains are depicted in Figure 4. Each plot refers to a problem class. The darker track corresponds to (the computation times and plan lengths for) HSCP, while the lighter to KACMBP (resp. Bwd and H-Fwd in the figures). We see that KACMBP either behaves very much alike that of HSCP, or it scales much better. In particular, the algorithms behave very similarly for highly constrained problems, such as the RING and the corner and face version of the CUBE domain. In such cases, the simple best-first, greedy approach of HSCP results in a well-driven search which produces an optimal plan. However, already for those slightly more complex problems where knowledge acquisition must be either interleaved with or followed by some kind of ground-level search, such as BMT or the center version of the CUBE, KACMBP shows a better scale up in time.

We also considered several variations of a “circular” navigation domain (i.e. a cube with no walls), where no trivial way of refining the knowledge is available. First, we considered a simple repositioning problem where no knowledge acquisition is needed. The difference in performance results from the more directed ground level search of KACMBP.


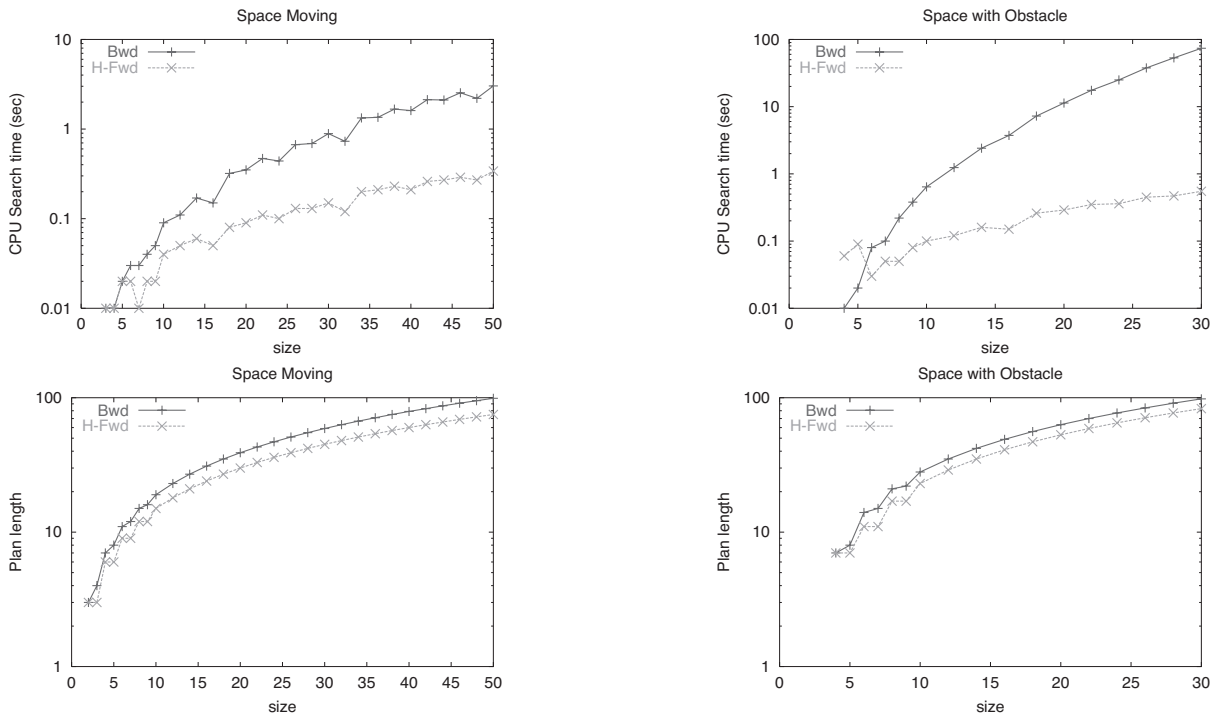
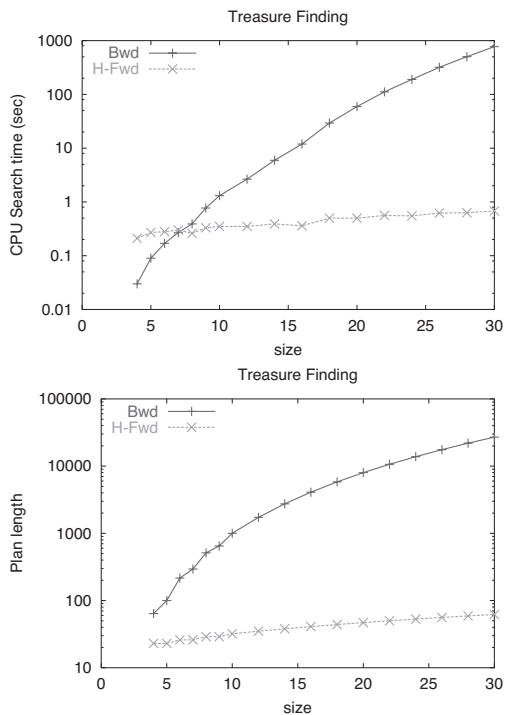


Figure 4: Experimental results on the BT, SQUARE, CUBE and RING domains



The table below reports the performance of the systems on a more difficult repositioning problem (Space with Obstacle), where only a column is available in the domain for repositioning the robot. For this problem, both knowledge acquisition and ground-level search turn out to be necessary.

Finally, in the Treasure Finding problem, the goal is to get hold of a treasure when incomplete information is available on the position. Notice how the simple structural heuristic of HSCP is unable to drive the search in the right direction.



To summarize, in the new examples KACMBP greatly outperforms HSCP in terms of search time. Moreover, in all the cases, the plans produced by KACMBP closely approximate optimal plans, whereas HSCP may end up producing very dumb plans.

This experimental evaluation is somewhat preliminary. Furthermore, much work has still to be done for tackling even more complex problems, since some of the abstractions used in the identification of the knowledge subgoals may not be as effective for more complex problems. In particular, KACMBP limits the choice of a possible knowledge subgoal to the cardinality of the known range of every fluent of the problem. This might not always be a good choice — as previously pointed out, the cardinality of the known range may not always be a sufficient index of knowledge upon a fluent. For instance, knowing that x is in $\{2, 3, 4\}$ is not the same as knowing that it is in $\{2, 4, 6\}$, if the problem is moving to an odd or even location. Furthermore, the approximate computation of distances (by component-wise projection upon fluents) may be too coarse in many cases. An example is that of navigation domains where obstacles have irregular shapes. In these cases, the low-level search may get close to the (knowledge or ground) goals without reaching them, thus giving up when success is actually near. A possible solution to this could consist of adding enforced-hill climbing techniques to the low-level search.

Further Related Work

Conformant planning has been recently devoted a significant amount of attention. The first efficient conformant planner was CGP (Smith & Weld 1998). CGP is based on the construction of a planning graph for each initial state, and tries to propagate constraints among these planning graphs to show the conformance of the plan. The main strength of

CGP is the ability to exploit the parallelism of the domain, when available. The main weakness appears to be the enumerative nature of the algorithms (for instance, in the CUBE domain of size 10, CGP would try to generate 10^3 planning graphs). QBFPLAN (Rintanen 1999) generalizes the SAT planning approach to deal with nondeterministic domains. QBFPLAN reduces the problem of finding a conformant solution of length k to the satisfiability of a Quantified Boolean Formula (QBF), and iteratively increases k until a solution is found. The strength of QBFPLAN is the logical flavor of the approach, that does not suffer from the enumerative nature of CGP. However, it is difficult for QBFPLAN to scale up to large problems, possibly for the limited maturity of the technology for QBF satisfiability. In (Cimatti & Roveri 2000), these two approaches and GPT (Bonet & Geffner 2000) are compared with (and outperformed by) the first symbolic model checking approach to conformant planning, that is bound to a breadth-first search style. More recent approaches to conformant planning are Open World planning, presented in (Filzi, Pirri, & Reiter 2000), that is based on the computation of prime implicates, and the DLV^k system, described in (Eiter *et al.* 2001), based on disjunctive Datalog. Though interesting, these approaches appear to be far in performance from the first, less advanced implementations of conformant planning based on symbolic model checking techniques.

Similarly to QBFPLAN, CPLAN (Castellini, Giunchiglia, & Tacchella 2001) extends the planning via SAT approach to the case of conformant planning. Rather than generating QBF encodings, however, CPLAN is based on standard SAT procedures. For a given plan length k , CPLAN iteratively finds a possible plan by solving a SAT problem, and tests whether it is conformant by solving another SAT problem, until a plan is found or all plans of that length have been proved not to be conformant. If no solution has been found, the length of the plan is then increased. In a sense, CPLAN implements a specialized QBF solver, tightly constrained by the features of the conformant planning problem. Compared to our approach, CPLAN is unable to decide the existence of a conformant solution of arbitrary length. We were unable to carry out an experimental comparison, since no CPLAN executable is available for this purpose. In (Castellini, Giunchiglia, & Tacchella 2001), CPLAN is reported to outperform the breadth-first version of CMBP. The analysis in previous section shows that the algorithm presented in this paper solves such problems orders of magnitude faster. Furthermore, the experimental analysis reported in (Castellini, Giunchiglia, & Tacchella 2001) disregards a large subset of the test set presented in (Cimatti & Roveri 2000). In our opinion, the examples that are not considered are likely to be extremely hard for CPLAN. For instance, in the empty room example, there is a very large number of plans that would be generated (and then would have to be excluded in the validation phase), before the minimal k allowing for a conformant plan is reached.

In (Kurien, Nayak, & Smith 2001), the new, *fragment-based* approach to conformant planning is presented. Basically, the idea is to find a plan that works for one world (an initial state), and then extending it to a full conformant

plan. The technology used is based on SAT, and the reported results, though preliminary, are very promising. A very interesting aspect of (Kurien, Nayak, & Smith 2001) is the relation of the approach with the issues of planning with faulty and partially observable systems.

Similarities between conformant planning and other fields are also being recognized. For instance, a typical problem in the field of Automated Test Pattern Generation for hardware circuits is finding a synchronization sequence, i.e. a sequence of input that takes a sequential circuit into a known state from a condition of total uncertainty (Pixley, Jeong, & Hachtel 1992; Rho, Somenzi, & Pixley 1993). Basically, this is the problem of finding a conformant plan for a set of possible goal belief states of cardinality one. A preliminary analysis of how such problems can be effectively tackled as search in the belief space is reported in (Cimatti, Roveri, & Bertoli 2001).

Conclusions and Future Work

In this paper, we have addressed the problem of heuristic search in belief space. This problem is relevant for conformant planning (Bonet & Geffner 2000), and also for planning under partial observability (Bertoli *et al.* 2001b). We show that the heuristics driving search in the space of beliefs can not be a simple combination of the standard heuristics used in classical planning, but need to take into account the fact that search is carried out in the space of beliefs. We provide the basic notion of “degree of knowledge”, and we show how such notions can be taken into account to drive the search and dynamically devise the need of knowledge acquisition. Furthermore, we show how to solve “knowledge subgoals” by defining the notion of “knowledge acquisition condition”. We implemented a conformant planning algorithm based on these ideas. The implementation in MBP shows that it improves significantly the quality of the solutions found. At the same time, it often reduces the portion state-space that is explored, leading in certain cases to dramatic speed-ups.

We will continue this research along the following lines. First, we will try to identify admissibility criteria for the heuristic functions, and conditions on the available knowledge to conclude the solvability of a problem. Then, we will integrate the planning algorithm with a front-end able to process PDDL (and its nondeterministic extensions, that are currently being standardized), and extend the efficient computations mechanisms for the heuristic estimates to the nondeterministic case. We will also explore the middle ground between the heuristic-symbolic search presented in this paper, and the breadth-first techniques of (Cimatti & Roveri 2000), to implement enforced hill-climbing techniques. Finally, the notions presented in this paper appear to be generic enough to be liftable to the case of partial observability, by including the structure of observation in addition to the actions in the treatment of knowledge acquisition conditions. Thus, we will apply the ideas to the case of planning under partial observability, where attempts at information sensing must be carefully planned in order to tackle the explosion of the search.

References

- Bertoli, P.; Cimatti, A.; Pistore, M.; Roveri, M.; and Traverso, P. 2001a. MBP: a Model Based Planner. In *Proc. of the IJCAI'01 Workshop on Planning under Uncertainty and Incomplete Information*.
- Bertoli, P.; Cimatti, A.; Roveri, M.; and Traverso, P. 2001b. Planning in Nondeterministic Domains under Partial Observability via Symbolic Model Checking. In *Proc. 7th International Joint Conference on Artificial Intelligence (IJCAI-01)*. AAAI Press.
- Bertoli, P.; Cimatti, A.; and Roveri, M. 2001a. Conditional Planning under Partial Observability as Heuristic-Symbolic Search in Belief Space. In A. Cesta, D. B., ed., *Proceedings ECP'01*.
- Bertoli, P.; Cimatti, A.; and Roveri, M. 2001b. Heuristic Search + Symbolic Model Checking = Efficient Conformant Planning. In *Proc. 7th International Joint Conference on Artificial Intelligence (IJCAI-01)*. AAAI Press.
- Bonet, B., and Geffner, H. 2000. Planning with Incomplete Information as Heuristic Search in Belief Space. In Chien, S.; Kambhampati, S.; and Knoblock, C., eds., *5th International Conference on Artificial Intelligence Planning and Scheduling*, 52–61. AAAI-Press.
- Bryant, R. E. 1992. Symbolic Boolean manipulation with ordered binary-decision diagrams. *ACM Computing Surveys* 24(3):293–318.
- Castellini, C.; Giunchiglia, E.; and Tacchella, A. 2001. Improvements to SAT-based Conformant Planning. In A. Cesta, D. B., ed., *Proceedings ECP'01*.
- Cimatti, A., and Roveri, M. 2000. Conformant Planning via Symbolic Model Checking. *Journal of Artificial Intelligence Research (JAIR)* 13:305–338.
- Cimatti, A.; Clarke, E.; Giunchiglia, F.; and Roveri, M. 1998. NUSMV: a reimplementation of SMV. In *Proceeding of the International Workshop on Software Tools for Technology Transfer (STTT-98)*, 25–31. Aalborg, Denmark: BRICS Notes Series, NS-98-4. Also IRST-Technical Report 9801-06, Trento, Italy.
- Cimatti, A.; Roveri, M.; and Bertoli, P. 2001. Searching Powerset Automata by Combining Explicit-State and Symbolic Model Checking. In *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction of Systems*, 313–327. Springer.
- Cimatti, A.; Roveri, M.; and Traverso, P. 1998a. Automatic OBDD-based Generation of Universal Plans in Non-Deterministic Domains. In *Proceeding of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*. Madison, Wisconsin: AAAI-Press. Also IRST-Technical Report 9801-10, Trento, Italy.
- Cimatti, A.; Roveri, M.; and Traverso, P. 1998b. Strong Planning in Non-Deterministic Domains via Model Checking. In *Proceeding of the Fourth International Conference on Artificial Intelligence Planning Systems (AIPS-98)*. Carnegie Mellon University, Pittsburgh, USA: AAAI-Press.

- De Giacomo, G., and Vardi, M. 1999. Automata-Theoretic Approach to Planning for Temporally Extended Goals. In Biundo, S., ed., *Proceeding of the Fifth European Conference on Planning*, Lecture Notes in Artificial Intelligence. Durham, United Kingdom: Springer-Verlag.
- Eiter, T.; Faber, W.; Leone, N.; Pfeifer, G.; and Polleres, A. 2001. The DLV^k Planning System. In *Proceedings of the IJCAI-01 Workshop on Planning with Uncertainty and Incomplete Information*.
- Filzi, A.; Pirri, F.; and Reiter, R. 2000. Open world planning in the situation calculus. In *Proceedings of Seventeenth National Conference on Artificial Intelligence (AAAI'00)*. Austin, Texas: AAAI Press.
- Giunchiglia, E.; Kartha, G. N.; and Lifschitz, V. 1997. Representing action: Indeterminacy and ramifications. *Artificial Intelligence* 95(2):409–438.
- Hoffmann, J., and Nebel, B. 2001. The FF planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research* 14:253–302.
- Kabanza, F.; Barbeau, M.; and St-Denis, R. 1997. Planning control rules for reactive agents. *Artificial Intelligence* 95(1):67–113.
- Kurien, J.; Nayak, P. P.; and Smith, D. 2001. Fragment-based Conformant Planning. Unpublished document.
- Pixley, C.; Jeong, S.-W.; and Hachtel, G. D. 1992. Exact calculation of synchronization sequences based on binary decision diagrams. In *Proceedings of the 29th Conference on Design Automation*, 620–623. Los Alamitos, CA, USA: IEEE Computer Society Press.
- Pryor, L., and Collins, G. 1996. Planning for Contingency: a Decision Based Approach. *J. of Artificial Intelligence Research* 4:81–120.
- Rho, J.-K.; Somenzi, F.; and Pixley, C. 1993. Minimum length synchronizing sequences of finite state machine. In IEEE, A.-S., ed., *Proceedings of the 30th ACM/IEEE Design Automation Conference*, 463–468. Dallas, TX: ACM Press.
- Rintanen, J. 1999. Constructing conditional plans by a theorem-prover. *Journal of Artificial Intelligence Research* 10:323–352.
- Smith, D. E., and Weld, D. S. 1998. Conformant graphplan. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98) and of the 10th Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*, 889–896. Menlo Park: AAAI Press.
- Somenzi, F. 1997. CUDD: CU Decision Diagram package — release 2.1.2. Department of Electrical and Computer Engineering — University of Colorado at Boulder.
- Thiebaux, S., and Cordier, M. 2001. Supply restoration in power distribution systems – a benchmark for planning under uncertainty. In A. Cesta, D. B., ed., *Proceedings ECP'01*.
- Weld, D. S.; Anderson, C. R.; and Smith, D. E. 1998. Extending graphplan to handle uncertainty and sensing actions. In *Proceedings of the 15th National Conference on*