

On the complexity of case-based planning

P. LIBERATORE

Università di Roma “La Sapienza”, Italy

(Received: June 2005)

We analyze the computational complexity of problems related to case-based planning: planning when a plan for a similar instance is known, and planning from a library of plans. We prove that planning from a single case has the same complexity than generative planning (i.e. planning ‘from scratch’); using an extended definition of cases, complexity is reduced if the domain stored in the case is similar to the one to search plans for. Planning from a library of cases is shown to have the same complexity. In both cases, the complexity of planning remains, in the worst case, PSPACE-complete.

Keywords: automated planning, case-based planning, computational complexity.

1 Introduction

Case-based reasoning (Kolodner, 1993, Aamodt and Plaza, 1994, Watson, 1997) is a problem solving methodology based on using a library of solutions for similar problems, i.e. a library of ‘cases’ with their respective solutions. Roughly speaking, case-based planning consists into storing generated plans and using them for finding new plans (Hammond, 1989, Bergmann *et al.*, 1998, Spalazzi, 2001). In practice, what is stored is not only a specific problem with a specific solution, but also some additional information that is considered useful to the aim of solving new problems, e.g. information about how the plan has been derived (Veloso, 1994), why it works (Kambhampati, 1990b, Hanks and Weld, 1995), when it would not work (Ihrig and Kambhampati, 1997), etc. Different case-based planners differ on how they store cases, which cases they retrieve when the solution of a new problem is needed, how they adapt a solution to a new problem, whether they use one or more cases for building a new solution, etc. The survey papers by Bergmann *et al.* (1998) and by Spalazzi (2001) give a detailed introduction to case-based planning.

In this paper, we study the computational complexity of case-based planning, i.e. we characterize the complexity of case-based planning using the theory of the NP-completeness (Johnson, 1990). What makes this analysis different from the other results on the complexity of planning (Bylander, 1991, Nebel *et al.*, 1997, Nebel and Bäckström, 1994, Bäckström and Nebel,

1995, Bäckström, 1992, Bäckström and Nebel, 1992, Bäckström and Nebel, 1993, Bäckström and Jonsson, 1995). is that case-based planning is not actually a *problem*, but rather a *family* of solutions for a problem. In fact, the theory of computational complexity can only characterize the complexity of problems, not of family of solutions. On the other hand, case-based planning always requires solving some specific subproblems, such as, for example, the adaptation of a plan to a different problem. In this paper, we study the complexity of some problems that are related to case-based planning.

The first problem we consider is that of finding a plan given another plan that works on a slightly different domain. This problem formalizes the task of plan adaptation that case-based planners have to perform. While the formalization of this problem disregards many aspect of plan adaptation in practice (e.g. more plans for similar cases may be available), it nevertheless contains the main requirement of plan adaption: building a new plan given an old one. This problem has been already analyzed by Nebel and Koehler (1993, 1995) under the constraint of minimal plan change: the new plan should be as similar to the old one as possible; plan adaptation in this case is called *conservative*. There are some scenarios in which this constraint is reasonable (e.g. we schedule the actions of the old plan before starting planning, and not executing them results in an additional cost); while conservatism is sometimes obtained as a byproduct of algorithms that work by changing an old plan (Kambhampati, 1990b), it is usually not a requirement. On the contrary, not enforcing minimal change has been considered a viable alternative for escaping the high complexity of conservative plan adaptation (Au *et al.*, 2002). The absence of the constraint of conservatism is also evident in the *generative case-based* planning approach, which differs from the *transformational case-based* approach outlined above in that the new plan is generated from scratch (rather than from the old plan), and the case is used to guide the choices made during the search (Bergmann *et al.*, 1998).

The main problem of analyzing a solution technique, rather than a specific problem, is that the implemented solutions may greatly differ to each other. While the abstract idea of reusing old plans is part of all case-based planners, even giving a common definition of ‘an old plan’ is not easy. This may be simply a specific plan, but can also be an abstract plan, or a partially ordered plan, and may include additional information (Kambhampati, 1990b, Veloso, 1994, Ihrig and Kambhampati, 1997). Moreover, the plan adaptation problem may not be tackled at all: planning by derivational analogy (Carbonell, 1986, Veloso and Carbonell, 1993) uses the traces of the search done for finding a plan, rather than the plan itself; generative case-based planners (Bergmann *et al.*, 1998) build a plan from scratch by using information from more than one case, rather than adapting a specific plan. In the opinion of the author of the present paper, the problem of plan adaptation in its simpler formalization is

a good starting point for the computational analysis of plan adaptation, even if it is not the problem that is faced in practice.

What makes giving a formalization that is at the same time general and precise a difficult task is that plan adaptation cannot always be separated from the other steps of case-based planning. For example, the plan to adapt is not chosen arbitrarily in the library of plans; it is chosen because it is expected to be adaptable to the new situation; moreover, more cases may be selected, leading to more than one ‘starting point’ for the search of the new plan. Being impossible to precisely formalize the whole case-based planning process and remain general enough, we consider the problem of case-based planning in a very general way, in which the only assumption that is made is that a case, or a library of cases, is a data structure of polynomial size.

The problems that are therefore considered are: planning from a specific known plan, and planning from a ‘generalized’ case or a library of cases. The first problem formalizes the basic plan adaptation step; the second problem formalizes the fact that cases usually contain additional information other than the plan itself; the third problem is a formalization of the whole case-based planning process.

In the analysis of these problems, we do not enforce the constraint of minimal change. The resulting freedom in the search for plans may simplify the algorithms (Au *et al.*, 2002), but makes the computational analysis more complicated. Indeed, the theory of NP-completeness formalizes problems that can be expressed in a scheme like the following one:

INSTANCE: a planning domain;
QUESTION: is there a plan for the domain?

(for some reason not clear to the author, the words ‘instance’ and ‘question’ are often written in small caps font.) The fact that this is a decision problem (the solution can only be yes or not) is not a big restriction, as search problems (where the solution can be a more complex data structure) can be usually reduced to a number of decision problems. What makes this formalization restrictive in the setting of plan adaptation is that there is no slot in the scheme where to place the old plan. Indeed, the old plan is not exactly part of the instance, as its presence does not change the definition of the problem. As well, it is not part of the question. The plan adaptation problem would be better formalized by a list of three parts:

INSTANCE: a planning domain;
HINT: a plan for a similar domain;
QUESTION: is there a plan for the new domain?

This is still a decision problem (its solution is either yes or not). However, its

definition contains a new part, the *hint*, which is not necessary for answering the question, but can be useful nevertheless. In other words, the existence of a plan only depends on the domain, not on the hint. The difference between the hint and the instance makes the analysis of plan adaptation carried on in this paper different from that of conservative plan adaptation: in the latter problem, the plan for the similar domain is part of the instance, as the new plan to be found depends on the old one (Nebel and Koehler, 1993). In other words, the instance of conservative plan adaptation contains a plan for a domain, and the question is whether there exists a similar plan for the new domain. Roughly speaking, an hint is something we could disregard while looking for a solution, while the old plan cannot in the case of conservative plan adaptation.

The problems we analyze in this paper are the problem of adaptation from a specific plan, from a plan plus other information, and the problem of planning with a library of cases. In all these cases, we are given an instance of the planning problem and an hint, which can be a plan, a plan plus other information, or a plan library. The technical results of this paper is that plan adaptation of a single specific plan can be as hard as generative planning (i.e. planning ‘from scratch’ disregarding the hint completely); planning from a more general definition of cases may be easier than generative planning in some cases, but is as hard in general. The same results hold for case-based planning, i.e. planning from a library of cases.

2 Preliminaries

The planning problems analyzed in this paper are formalized using the propositional STRIPS formalism. A STRIPS instance (or a STRIPS domain) is a 4-tuple $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$, where \mathcal{P} is the set of conditions, \mathcal{O} is the set of operators, \mathcal{I} is the initial state, and \mathcal{G} is the goal. The *conditions* are facts that can be true or false in the world of interest. A *state* S is a set of conditions, and represents the state of the world in a certain time point. The conditions in S are those representing facts that are true in the world, while those not in S represent facts currently false.

The *initial state* is a state, thus a set of conditions. The *goal* is specified by giving a set of conditions that should be achieved, and another set specifying which conditions should not be made true. Thus, a goal \mathcal{G} is a pair $\langle \mathcal{M}, \mathcal{N} \rangle$, where \mathcal{M} is the set of conditions that should be made true and \mathcal{N} is the set of conditions that should be made false.

The *operators* are actions that can be performed to achieve the goal. Each operator is a 4-tuple $\langle \phi, \eta, \alpha, \beta \rangle$, where ϕ , η , α , and β are sets of conditions. When executed, such an operator makes the conditions in α true, and those in β false, but only if the conditions in ϕ are currently true and those in η are

currently false. The conditions in ϕ and η are called the positive and negative *preconditions* of the operator. The conditions in α and β are called the positive and negative *effects* or *postconditions* of the operator.

Given an instance of a STRIPS planning domain $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$, we define a plan for it as a sequence of operators that, when executed in sequence from the initial state, lead to a state where all the conditions in \mathcal{M} are true and all those in \mathcal{N} are false. More details about the definition of STRIPS can be found in other papers (Fikes and Nilsson, 1971, Bylander, 1991).

A *planning case* is a pair $\langle D_o, P_o \rangle$, where D_o is a planning instance and P_o is a plan for it. Plan adaptation is formalized as follows: we are given a plan case and a specific domain D that contains the same conditions and actions of the domain D_o . The problem is that of finding a plan for D . This is a simplification of the plan adaptation step of case-based planners, but contains all main components: the case with a solution and a new instance to be solved. Somehow, we are studying this problem in isolation from the general case-based planning process, as for example we disregard the fact that the case is typically chosen in such a way P_o is expected to be useful for finding a plan for D . We then consider the problem of plan adaptation in the assumption that the case is not only a pair $\langle D_o, P_o \rangle$ but may also contain other information; we abstract over the kind of information that is stored, so that our results are independent on whether we use derivational traces, abstractions, previous failures, etc. We prove that planning adaptation is feasible, in this case, provided that the domain of the case and the current domain are similar.

The second problem we consider is whether a library of plans can improve the efficiency of finding a plan for a new domain. A simple (and somehow simplistic) definition of a plan library is a set of cases $\{\langle D_i, P_i \rangle \mid 1 \leq i \leq m\}$. We consider this definition of plan library as a special case of some importance, but the result we prove holds for any kind of plan library. This generalization is necessary because a case library usually contains much more information than just a set of planning instances with their relative solutions (e.g. plans may be abstract, they may be partially ordered, the library may contains indexes using for example description logics (Koehler, 1996), the cases themselves may be abstract and stored hierarchically, etc.) By assuming that the library of cases is an arbitrary data structure, our results carry on to all specific classes of plan libraries. We remark that the complexity-theoretic results for this generalized definition of case library are more related to the generative case-based planning than to the transformational approach, in that the plan library is only assumed to be used when it is useful in the process of searching for a new plan, and is not necessarily used to provide a starting point of the search process.

3 Results

The problem of deciding whether there exists a plan for a STRIPS instance is denoted as PLANSAT, and is known to be PSPACE-complete (Bylander, 1991, Bylander, 1994). If a planning case $\langle D_o, P_o \rangle$ is also known, the problem cannot become more difficult, as we can simply disregard the case and find the plan using D only. Essential to this trivial result is that we do not enforce adaptation to be conservative, as motivated in the Introduction. The problem of plan adaptation is therefore in PSPACE. While a problem cannot be made harder by the presence of an hint, it may become easier. The following theorem shows that this is not the case for plan adaptation.

THEOREM 3.1 *Deciding whether there exists a plan for a STRIPS instance D , given a case $\langle D_o, P_o \rangle$, is PSPACE-complete, even if D_o and D only differ from one condition of the initial state.*

Proof. The problem is in PSPACE, as we can find the a plan for D disregarding the case $\langle D_o, P_o \rangle$.

We prove that the problem of plan adaptation is PSPACE-hard by showing a reduction from PLANSAT. In other words, we show that, given an instance of PLANSAT $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$, there exists $\langle D_o, P_o \rangle$ and D such that P_o is a plan for D_o , the domains D_o and D only differ for one condition of the initial state, and D has a plan if and only if $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$ has a plan.

The domains D_o and D are obtained from $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$ by adding a single condition a , a single action $e = \langle \{a\}, \emptyset, \mathcal{M}, \mathcal{N} \rangle$, where $\mathcal{G} = \langle \mathcal{M}, \mathcal{N} \rangle$ is the goal of the original instance, and possibly modifying the initial state:

$$\begin{aligned} D_o &= \langle \mathcal{P} \cup \{a\}, \mathcal{O} \cup \{e\}, \mathcal{I} \cup \{a\}, \mathcal{G} \rangle \\ D &= \langle \mathcal{P} \cup \{a\}, \mathcal{O} \cup \{e\}, \mathcal{I}, \mathcal{G} \rangle \end{aligned}$$

The plan P_o that is part of the case is $\langle e \rangle$. It can be easily verified that P_o is a plan for D_o (its preconditions are verified in the initial state of D_o , and its consequences are exactly the goal.) As a result, $\langle D_o, P_o \rangle$ and D constitute a valid plan adaptation instance.

A case-based planner, while looking for a plan for D , would first check whether P_o works in D and then try to adapt it if it does not. In this case, P_o is not executable at all in D , as the precondition a of e is not true in the initial state of D . Moreover, the plan P_o is completely useless for solving D , as it requires a to be true, while a is not initially true in D and cannot even be made true because no action in $\mathcal{O} \cup \{e\}$ makes it true. As a result, a can be removed altogether from D , along with e that is never executable. After this change, D becomes $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$, which is exactly the original instance. We can therefore conclude that D has a plan if and only if the original instance has a

plan, thus proving the PSPACE-hardness of the plan adaptation problem. \square

This theorem proves that checking the existence of a case $\langle D_o, P_o \rangle$ does not simplify the check of existence of a plan for a domain D that is similar to D_o . Namely, the problem remains PSPACE-complete. We remark that the hardness part of this theorem holds *even* if the instances are similar, not only in this case. This result is therefore relevant even in those settings in which the choice of the case to reuse is not driven by the similarity of the domain to the one to be solved (Kambhampati, 1990a).

Analyzing the proof, we observe that the reduction used for proving the hardness part can be described as: ‘given a planning instance $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$, there exists a case $\langle D_o, P_o \rangle$ and a domain D such that...’. The important part is that we have chosen the specific plan P_o that is part of the case. In a way, this theorem proves that *there exists* a plan for D_o that does not simplify the plan existence problem for D , i.e. some plans do not help.

A natural question is now: given that some plans do not help, may it be that some other plans do? In other words, it may be that a sensible choice of a plan P_o for D_o is useful for solving the problem for D . We give a negative answer to this question.

THEOREM 3.2 *The problem of checking the existence of a plan for D , given a case $\langle D_o, P_o \rangle$, where P_o is any plan of D_o , is PSPACE-complete even if D and D_o only differ from one condition of the initial state.*

Proof. The proof is based on the idea of reducing the PLANSAT problem to the plan adaptation problem in which P_o is the only plan of D_o . By proving that P_o is not useful for solving PLANSAT for D , we prove that no plan is, in general, useful.

The proof of Theorem 3.1 has the only problem that $P_o = \langle e \rangle$ may not be the only plan of D_o . We therefore modify the domain D_o to make it so. Given a PLANSAT instance $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$, the domains D and D_o are still based on the set of conditions $\mathcal{P} \cup \{a\}$, where a is a new condition. The operators are $\mathcal{O}' \cup \{e\}$, where:

$$\begin{aligned} \mathcal{O}' &= \{ \langle \alpha, \beta \cup \{a\}, \gamma, \delta \rangle \mid \langle \alpha, \beta, \gamma, \delta \rangle \in \mathcal{O} \} \\ e &= \langle \{a\}, \emptyset, \mathcal{M}, \mathcal{N} \rangle \end{aligned}$$

In words, we add a as a negative precondition of each operator in \mathcal{O} . The new action e is defined as in the proof of the previous theorem. The domains D and D_o are defined as:

$$\begin{aligned} D_o &= \langle \mathcal{P} \cup \{a\}, \mathcal{O}', \mathcal{I} \cup \{a\}, \mathcal{G} \rangle \\ D &= \langle \mathcal{P} \cup \{a\}, \mathcal{O}', \mathcal{I}, \mathcal{G} \rangle \end{aligned}$$

The plan P_o is $\langle e \rangle$. This time, P_o is the only plan for D_o , as all other actions have a as a negative precondition, while a is initially true and no action makes it false. On the other hand, the action e is not executable in D , as a is initially false and no action makes it true. As a result, we can remove a and e from D , and this makes it identical to the original instance. As a result, D as a plan if and only if $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$ has a plan. This proves that the problem of plan adaptation remains PSPACE-hard regardless of how the plan in the case $\langle D_o, P_o \rangle$ is chosen. \square

This theorem proves that no plan for D_o can be useful in finding a plan for D , even if D and D_o only differ for a single condition of the initial state. This result, however, only holds when we assume that a case is exactly a pair $\langle D_o, P_o \rangle$, i.e. a domain and a plan for it. As explained in the Introduction, case-based planners usually record with plans other information such as derivational traces, abstractions of the plan, information that anticipates when the plan may fail, etc. This additional information can simplify the plan adaptation process. Since case-based planners greatly differ on what information is recorded, we simply assume that the case is recorded as a pair $\langle D_o, A_o \rangle$, where A_o is a polynomial-size data structure. We can prove that such cases may be of help if the new domain to be solved is similar to the one in a case.

The theorem showing this fact, as well as the following one, requires a digression into the topic of problem preprocessing. Problems such as plan adaptation can be formalized as set of pairs $\langle D_o, D \rangle$, where the question is: given D_o , is there any A_o such that solving the question of plan existence of D can be done in polynomial time given $\langle D_o, A_o \rangle$? In terms of problem preprocessing, we are given $\langle D_o, D \rangle$ and ask whether preprocessing D_o can result in a data structure A_o that makes the PLANSAT problem on D easy to solve, i.e. polynomial-time. Problems for which preprocessing lowers the complexity to P form the class $\sim\text{P}$, a.k.a. comp-P, a.k.a. compilable to P. The following lemma shows that the problem of plan adaptation is compilable to P.

LEMMA 3.3 *The problem of checking the plan existence of D is in $\sim\text{P}$, given $\langle D_o, D \rangle$, where D_o is a domain that only differs from D for a constant number of conditions in the initial state and goal.*

Proof. Cadoli *et al.* (2002) proved that all problems for which a polynomial number of varying parts (in this case, the number of possible D) that correspond to the same fixed part (in this case, D_o) is a constant, then the problem is compilable to P. This is actually the case, as D and D_o only differ for a constant number c of conditions in the initial states or goal; as a result, for any D_o there are only $|\mathcal{P}|^c$ possible D . Since c is a constant, this function $|\mathcal{P}|^c$ is polynomial in the size of D_o and D . As a result, the problem can be compiled to P, i.e. it is in $\sim\text{P}$. \square

The following is an easy corollary of the above lemma.

THEOREM 3.4 *For every domain D_o there exists a data structure A_o such that $\langle D_o, A_o \rangle$ allows for solving the problem of plan existence of a domain D that only differs from D_o for a constant number of conditions in the initial state or the goal in polynomial time.*

Proof. By the above lemma, the fixed part D_o of the problem $\langle D_o, D \rangle$ can be preprocessed in such a way the result of this phase is polynomial in size and allows for solving the problem of plan existence for D in polynomial time, if D and D_o only differ for a constant number of conditions of the initial state and goal. As a result, for any given D_o there is a polynomial-size data structure A_o such that $\langle D_o, A_o \rangle$ makes polynomial the problem PLANSAT for every D that only differ to D_o only for a constant number of conditions of the initial state or goal. \square

We remark that this result is theoretical, in that it abstracts over the possible ‘extended representations’ of cases. How this result apply to the various specific representations of cases is an open question.

The question of whether an extended representation of cases, or a whole library of cases, is of help in finding a plan for a new domain that can differ from the ones that have already been analyzed for a non-constant number of conditions can be given a negative answer by showing that the problem of solving PLANSAT on $\langle D_o, D \rangle$ is not compilable to P. We indeed prove that this problem is $\|\rightsquigarrow$ PSPACE-complete. This is proved by means of the following lemma.

LEMMA 3.5 *For every operator o over \mathcal{P} , the STRIPS instance $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$ has a plan if and only if the instance $\langle \mathcal{P} \cup \{y\}, \mathcal{O} \cup \{o\}, \mathcal{I}', \mathcal{G} \rangle$ has a plan, where y is a new condition not in \mathcal{P} , the operator o' is o with the addition of y as a positive precondition, and \mathcal{I}' is \mathcal{I} or $\mathcal{I} \cup \{y\}$ depending on whether $o \in \mathcal{O}$ or not.*

Proof. Since y is a new condition not in \mathcal{P} , and is only mentioned as a precondition of o' , no operator change its value. As a result, it is true if and only if it is true in the initial state. Moreover, y is positive in the initial state if and only if $o \in \mathcal{O}$. As a result, o' is equivalent to o if $o \in \mathcal{O}$, and cannot be executed if $o \notin \mathcal{O}$. Since the addition of y and o' and the possible removal of o are the only changes from the first instance to the second, the property of plan existence does not change. \square

This lemma looks like a trivial property, but in fact it says something interesting about the complexity of the PLANSAT problem: regardless of whether $o \in \mathcal{O}$, we can let o' to be in \mathcal{O} , and use a variable y to encode whether o in \mathcal{O} or not. In other words, we are making an element of the set of operators fixed,

as o' is added to the set of operators regardless of whether $o \in \mathcal{O}$. Since the original problem can be reduced to this new one in which the set of operators has a fixed part, the complexity of the new problem is at least as high as that of the original problem. By iterating this procedure, we can make the set of operators completely fixed. As a result, the complexity of the problem remains the same even if the set of operators is fixed.

The formal proof uses a sufficient condition for proving that a problem is non compilable to P called representative equivalence.

LEMMA 3.6 *The problem of determining a plan for D , given $\langle D_o, D \rangle$, where D_o and D contain the same conditions and operators, is $\|\rightsquigarrow$ PSPACE-complete.*

Proof. Membership follows from the fact that PLANSAT is in PSPACE, and that any problem in PSPACE is also in $\|\rightsquigarrow$ PSPACE.

Hardness: we show a reduction from the PLANSAT problem when operators are restricted to have two preconditions and two postconditions. The problem of plan existence has been proved PSPACE-hard even under this restriction (Bylander, 1994).

In order for using the condition of representative equivalence (Liberatore, 2001), we need first to show three functions for the PLANSAT problem: a classification, representative, and extension function. We make the following choice:

$$\begin{aligned} \text{Class}(\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle) &= |\mathcal{P}| \\ \text{Repr}(n) &= \langle \{x_1, \dots, x_n\}, \emptyset, \emptyset, \langle \emptyset, \emptyset \rangle \rangle \\ \text{Exte}(\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle, m) &= \langle \mathcal{P} \cup \{x_{|\mathcal{P}|+1}, \dots, x_m\}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle \end{aligned}$$

In words, the class of an instance is the number of its conditions. The representative of the class n is the instance that has n conditions, no operators, and empty initial state and goal (this instance has the empty plan $\langle \rangle$ as the initial state satisfies the goal, but this is irrelevant.) Extending an instance is obtained by simply adding conditions that are not then contained in any operator nor in the initial state or goal. Technically, these new conditions are false in the initial state and not required to have any specific value in the goal. As a result, $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$ has a plan if and only if $\text{Exte}(\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle, m)$ has a plan, as required for the extension function.

We now show a reduction that satisfies the condition of representative equivalence. By Lemma 3.5, we can replace \mathcal{O} with a set of operators \mathcal{O}' which contains a fixed operator o' . By iterating this reduction for *each possible operator* of four conditions over \mathcal{P} , we end up with an instance $\langle \mathcal{P} \cup Y, \mathcal{O}_{\mathcal{P}}, \mathcal{I}'', \mathcal{G} \rangle$, where $\mathcal{O}_{\mathcal{P}}$ is obtained from the set of all possible operators of four conditions

over \mathcal{P} , and Y is a set of new conditions in correspondence with the operators of $\mathcal{O}_{\mathcal{P}}$.

Here we exploit the restriction on the number of preconditions and postconditions. The number of possible pairs of preconditions is given by $n(n-1)$; since each precondition can be either positive or negative, we have exactly four combinations. As a result, the number of possible preconditions of an operator are $4n(n-1)$, assuming that no operator contains the same precondition both positively and negatively. For the same reason, there are exactly $4n(n-1)$ postconditions. As a result, there are exactly $16n^2(n-1)^2$ possible operators over n conditions. Therefore, the size of $\langle \mathcal{P} \cup Y, \mathcal{O}_{\mathcal{P}}, \mathcal{I}'', \mathcal{G} \rangle$ is polynomially larger than the size of $\langle \mathcal{P}, \mathcal{O}, \mathcal{I}, \mathcal{G} \rangle$.

The instances that result from implementing this reduction from two instances that have the same conditions are the same. By renaming all conditions to $\{x_1, \dots, x_n\}$, the results of the reduction is the same if the two instances have the same number of conditions. As a result, this reduction satisfies the condition of representative equivalence, which is sufficient to show that the problem that is reduced to (PLANSAT) is C-hard for any class C for which the problem that is reduced from (PLANSAT) again). As a result, the PLANSAT problem is $\|\mapsto$ PSPACE-hard. \square

This proof, based on Lemma 3.5, has an intuitive construction: we can progressively make \mathcal{O} fixed while maintaining the property of existence of plans. As a result, the PLANSAT problem can be reduced to the PLANSAT problem where \mathcal{O} is fixed, i.e. when \mathcal{O} can be preprocessed. Since making the set of conditions \mathcal{P} fixed is only a matter of condition renaming, this reduction shows that the complexity of the problem does not change even if \mathcal{P} and \mathcal{O} are fixed, i.e. the fixed part of the instance is actually fixed. This method can be used also for other problems. However, giving a general formulation is made difficult by the fact that ‘making a part fixed’ depend on the specific problem under analysis.

The result of $\|\mapsto$ PSPACE-hardness of PLANSAT implies that the knowledge of a case does not reduce the complexity of plan existence, even if the case can contain arbitrary data besides the plan.

THEOREM 3.7 *The problem of checking the plan existence for D does not become polynomial even if a case $\langle D_o, A_o \rangle$ is known, where D_o has the same conditions and operators of D , and A_o is an arbitrary polynomial-size data structure depending only on D_o , unless the polynomial hierarchy collapses.*

Proof. Assume, on the contrary, that for every D_o there exists a polynomial-sized data structure A_o (i.e. the ‘extended plan’) that makes solving PLANSAT on D a polynomial task. If this were the case, it would be possible to preprocess D_o obtaining A_o , as the preprocessing phase is not constrained in any way but that its result must be of polynomial size. Since $\langle D_o, A_o \rangle$ allows for solving

PLANSAT on D in polynomial time, we have that the problem of plan existence on $\langle D_o, D \rangle$ is in $\sim\text{P}$. Since the same problem is also $\|\rightsquigarrow\text{PSPACE}$ -hard, it follows that $\|\rightsquigarrow\text{PSPACE} \subseteq \sim\text{P}$, which implies that $\text{PSPACE}/\text{poly} \subseteq \text{P}/\text{poly}$ thanks to a result by Cadoli *et al.* (2002), which in turns implies that $\Sigma_2^P \cap \Pi_2^P = \text{PSPACE}$ thanks to a result by Karp and Lipton (1980). In other words, the polynomial hierarchy collapses to its second level. \square

This result can also be extended to the case in which a whole ‘library of cases’ $\{\langle D_i, P_i \mid 1 \leq i \leq m \rangle\}$ is given, provided that all domains D_i have the same conditions and operators of the current domain D , and the whole library is of polynomial size.

THEOREM 3.8 *The problem of checking the existence of plans for D cannot be solved in polynomial time even if a library of plans $\{\langle D_i, P_i \mid 1 \leq i \leq m \rangle\}$ of polynomial size is given, where all D_i have the same conditions and operators of D , unless the polynomial hierarchy collapses.*

Proof. This is only a consequence of the above theorem: assume that $A_o = \{\langle D_i, P_i \mid 1 \leq i \leq m \rangle\}$; since $\langle D_o, A_o \rangle$ does not make PLANSAT on D polynomial-time, then $\{\langle D_i, P_i \mid 1 \leq i \leq m \rangle\}$ does not make it polynomial-time either. \square

As it is clear, *some* library of plans simplify the problem: for example, if the library contains a case $\langle D_i, P_i \rangle$ where $D = D_i$, a plan for D is simply P_i . The theorem indeed proves that such a simplification is not possible in general. As it is also clear from the proof, using an extended definition of cases, in which not only plans are recorded, is not useful either. This result requires some discussion, which will be given in the next section.

4 Conclusions

In this paper, we have proved two kinds of results: first, adapting a specific plan to a new domain is as hard as planning from scratch; second, a case or a library of cases composed of domains, plans, and additional information, can simplify the problem of planning in some cases, but remains hard in general. These results are interesting because they are *formally*, and not only *empirically*, proved.

The first kind of results are somehow not surprising, as it has already been observed that even changing a single condition of the initial state or goal may make a plan completely useless. Our results are simply formal proofs of this observation.

The second kind of results are, in the author’s opinion, more interesting. The fact that ‘extended’ cases sometimes lower the complexity of finding plans

formally validates the trend in case-based planning of storing complex information in cases, rather than simple plans. Nevertheless, this formal result only holds in a very general settings, in which no assumption, besides polynomiality of space, is made about what is stored in the case. How this result extends to specific form of information used in case-based planners, such as abstract plans or derivational traces, is an open question.

The negative results about the complexity of planning from an extended plan or a library of plans are the most interesting ones of this paper. First, they formally prove that case-based planning can be as hard as generative planning; while this phenomena has been empirically observed (Bergmann *et al.*, 1998), it was not yet proved that it is intrinsic to the problem and not related to the specific implementations. Second, compared with the result on ‘small changes’, it proves that there exists a trade-off of efficiency and size of library in case-based planning. Again, what is interesting in this result is that it has formally proved, not only empirically observed.

Let us compare the results of this paper with similar work in the literature. The computational complexity of planning has been deeply investigated by several authors (Bylander, 1991, Bylander, 1994, Nebel *et al.*, 1997, Nebel and Bäckström, 1994, Bäckström and Nebel, 1995, Bäckström, 1992, Bäckström and Nebel, 1992, Bäckström and Nebel, 1993, Bäckström and Jonsson, 1995). All these works are on planning from first principles, i.e. are about planning given only a planning domain. More related to the present work are that by Nebel and Koehler (1993, 1995), who analyzed the complexity of conservative plan adaptation. Their complexity results are about planning when the constraint to be similar to another known plan is enforced. While there are some scenarios where this constraint is important, it seems not to be enforced often in case-based planning. Some work on plan adaptation has also been done by Bylander (1996), who has shown that, probabilistically, plan modification is simpler than plan generation provided that the domain of the case is similar to the one to be solved.

As it has already been noticed in the introduction, the problem of plan adaptation and that of planning with a library of cases is of a different kind of most problems that are studied in computational complexity, in that the data we are given include an ‘hint’ that can very well be neglected if necessary. Such problem format include other problems, such as case-based reasoning in general. The complexity of other forms of case-based reasoning is an interesting problem which is however out of the scope of this paper.

References

A. Aamodt and E. Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. *AI Communications—The Euro-*

pean Journal on Artificial Intelligence, 7:39–59, 1994.

T.-C. Au, H. Muñoz-Avila, and D. Nau. On the complexity of plan adaptation by derivational analogy in a universal classical planning framework. In *Sixth European Conference on Advances in Case-Based Reasoning (ECCBR 2002)*, pages 13–27, 2002.

C. Bäckström and P. Jonsson. Planning with abstraction hierarchies can be exponentially less efficient. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 1599–1605, 1995.

C. Bäckström and B. Nebel. On the computational complexity of planning and story understanding. In *Proceedings of the Tenth European Conference on Artificial Intelligence (ECAI'92)*, pages 349–353, 1992.

C. Bäckström and B. Nebel. Complexity results for SAS+ planning. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 1430–1435, 1993.

C. Bäckström and B. Nebel. Complexity results for SAS+ planning. *Computational Intelligence*, 11:625–656, 1995.

C. Bäckström. Equivalence and tractability results for SAS+ planning. In *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR'92)*, pages 126–137, 1992.

R. Bergmann, H. Muñoz-Avila, M. Veloso, and E. Melis. CBR applied to planning. In M. Lenz, B. Bartsch-Spörl, H.-D. Burkhard, and S. Wess, editors, *Case-based reasoning technology, from foundations to applications*. Springer, 1998.

T. Bylander. Complexity results for planning. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI'91)*, pages 274–279, 1991.

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

T. Bylander. A probabilistic analysis of propositional STRIPS planning. *Artificial Intelligence*, 81:241–271, 1996.

M. Cadoli, F. Donini, P. Liberatore, and M. Schaerf. Preprocessing of intractable problems. *Information and Computation*, 176(2):89–120, 2002.

- J. Carbonell. Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In R. Michalski, J. Carbonell, and T. Mitchell, editors, *Machine Learning: An Artificial Intelligence Approach*. Morgan Kaufmann, Los Altos, 1986.
- R. Fikes and N. Nilsson. STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.
- K. Hammond. *Case-based Planning: Viewing Planning as a Memory Task*. Academic Press, 1989.
- S. Hanks and D. Weld. A domain-independent algorithm for plan adaptation. *Journal of Artificial Intelligence Research*, 2:319–360, 1995.
- L. Ihrig and S. Kambhampati. Storing and indexing plan derivations through explanation-based analysis of retrieval failures. *Journal of Artificial Intelligence Research*, 7:161–198, 1997.
- D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2, pages 67–161. Elsevier Science Publishers (North-Holland), Amsterdam, 1990.
- S. Kambhampati. Mapping and retrieval during plan reuse: a validation structure based approach. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI'90)*, pages 170–175, 1990.
- S. Kambhampati. A theory of plan modification. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI'90)*, pages 176–182, 1990.
- R. M. Karp and R. J. Lipton. Some connections between non-uniform and uniform complexity classes. In *Proceedings of the Twelfth ACM Symposium on Theory of Computing (STOC'80)*, pages 302–309, 1980.
- J. Koehler. Planning from second principles. *Artificial Intelligence*, 87:145–186, 1996.
- J. Kolodner. *Case-Based Reasoning*. Morgan Kaufmann, Los Altos, 1993.
- P. Liberatore. Monotonic reductions, representative equivalence, and compilation of intractable problems. *Journal of the ACM*, 48(6):1091–1125, 2001.
- B. Nebel and C. Bäckström. On the computational complexity of temporal projection, planning, and plan validation. *Artificial Intelligence*, 66(1):125–160, 1994.

- B. Nebel and J. Koehler. Plan modification versus plan generation: a complexity-theoretic perspective. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'93)*, pages 1436–1440, 1993.
- B. Nebel and J. Koehler. Plan reuse versus plan generation: A theoretical and empirical analysis. *Artificial Intelligence*, 76:427–454, 1995.
- B. Nebel, Y. Dimopoulos, and J. Koehler. Ignoring irrelevant facts and operators in plan generation. In *Proceedings of the Fourth European Conference on Planning (ECP'97)*, pages 338–350, 1997.
- L. Spalazzi. A survey on case-based planning. *Artificial Intelligence Review*, 16:3–36, 2001.
- M. Veloso and J. Carbonell. Derivational analogy in PRODIGY: Automating case acquisition, storage, and utilization. *Machine Learning*, 10(3):249–278, 1993.
- M. Veloso. Flexible strategy learning: Analogical replay of problem solving episodes. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI'94)*, pages 595–600, 1994.
- I. Watson. *Applying Case-Based Reasoning: Techniques for Enterprise Systems*. Morgan Kaufmann, Los Altos, 1997.