On the Limits of Planning over Belief States under Strict Uncertainty

Sebastian Sardina
Dept. of Comp. Sci.
RMIT University
Melbourne, AUSTRLIA
ssardina@cs.rmit.edu.au

Giuseppe De Giacomo
Dipart. di Infor. e Sist.
Università di Roma "La Sapienza"
Roma, ITALY
degiacomo@dis.uniroma1.it

Yves Lespérance
Dept. of Comp. Sci. & Eng.
York University
Toronto, CANADA
lesperan@cs.yorku.ca

Hector J. Levesque
Dept. of Comp. Sci.
University of Toronto
Toronto, CANADA
hector@cs.toronto.edu

Abstract

A recent trend in planning with incomplete information is to model the actions of a planning problem as nondeterministic transitions over the belief states of a planner, and to search for a plan that terminates in a desired goal state no matter how these transitions turn out. We show that this view of planning is fundamentally limited. Any plan that is successful by this criteria has an upper bound on the number of actions it can execute. Specifically, the account will not work when iterative plans are needed. We also show that by modifying the definition slightly, we obtain another account of planning that does work properly even for iterative plans. Although the argument is presented in an abstract form, we illustrate the issues using a simple concrete example.

Introduction

Most work in AI planning today deals with *sequential* planning: generating a sequence of actions to achieve a goal. A smaller community is concerned with *conditional* planning where plans can be tree-like structures, and an even smaller community is concerned with *iterative* planning, where plans can be graph-like structures with loops. The motivation for the latter two cases involves planning problems where there is information about the domain that is unavailable at plan time, but can be acquired at runtime by the robot or agent executing the plan. It is the job of the plan to specify what to do depending on how this information turns out. Typically, we imagine a robot or agent armed with sensors of some sort, and the plan must specify courses of action for the various outcomes of these sensors.

Since conditional and iterative planning are done in a setting where full information about the world is not available, a recent trend has been to perform the planning over *belief states* rather than world states (e.g., (De Giacomo *et al.* 1997; Bonet & Geffner 2000; Bertoli *et al.* 2001; Petrick & Bacchus 2002; Bryce & Kambhampati 2004)).[1] Instead of formulating the problem in terms of states of the world where fluents (the changing properties of the world) hold or do not hold and actions that change these world states, we think of the problem as involving states of belief about the world, where beliefs are held or not held and actions that change these belief states.

To take an example from (Moore 1985; Bacchus & Petrick 1998), consider the effect of dipping a litmus paper into a solution. From a world viewpoint, this is a *deterministic* action: if the solution is acidic, the paper turns red; otherwise, it stays blue; either way the effect is completely determined. But from a belief viewpoint, we can think of this as a *nondeterministic* action that leads to two possible belief states: one in which we believe that the paper is red and the solution acidic, and the other in which we believe the paper is blue and the solution not acidic. Or as a variant, we might think of the dipping action as deterministic (not changing beliefs about the color of the paper at all), but introduce a sensing action to examine the color of the paper, with again two possible outcomes. Either way, the effect is the same: we have two possibilities to consider, and no information whatsoever at plan time to choose between them.[2]

Now imagine planning in a setting like this. Different systems will have different ideas about what a plan will look like and how beliefs are to be represented, but in all cases, plans must specify what to do depending on how the runtime information turns out. And what we are after is an *adequate* plan, that is, one that will achieve the goal in all cases.[3]

But what do we mean by this? This is taken to mean something like the following: a plan is adequate iff it works from the initial belief state, where a plan is considered to work from some state iff either (1) it says do nothing and the goal is believed to hold, or (2) it says to do some action, and for every possible state you can get to by doing the action, the remaining plan works in the resulting state.

In this paper, we will make these notions precise and prove that this view of planning is fundamentally limited. We will show that every plan that is adequate according to this definition has an upper bound on the number of actions it can perform. Specifically, it will not work for iterative planning, where unbounded loops may be necessary. How-

[1]This is sometimes called "planning at the knowledge level."

[2]All the actions in this paper will be considered to be deterministic in the world, but potentially nondeterministic in terms of their believed effects.

[3]For some purposes, a weaker notion of planning is considered, such as achieving the goal with some probability or achieving the goal in some cases. These notions will not be addressed here.

ever, we will also show that by slightly varying this view of planning, we obtain another account of adequacy that does work properly even for iterative plans. Our accounts will remain abstract and independent of the representations and planning algorithms that might be used.

Observe that in this paper we assume a setting of *strict uncertainty* in that the space of possibilities (possible effects of actions or sensing outcomes) is known, but the probabilities of these potential alternatives cannot be quantified. This contrasts with settings where a probability distribution over the set of possibilities is available as it is the case with Markov decision processes (MDPs/POMDPs) (e.g., (Kaelbling, Littman, & Cassandra 1998; Boutilier, Dean, & Hanks 1999)) or probabilistic planning (e.g., (Draper, Hanks, & Weld 1994; Bonet & Geffner 2000)). We shall claim our results valid only in the context of strict uncertainty and claim nothing for settings where quantitative information about the likelihood of possible outcomes of action or sensing is available. We will elaborate on this in the discussion section.

Belief-Based Planning

We believe that the limitations on planning over belief states sketched in the introduction do not depend on the details of how the planning problem is formalized. Let us assume that a belief-based planning problem $\mathcal{B} = \langle B, A, T, b_0, G\rangle$ is characterized by five entities: a set B of belief states, a finite set A of actions, a transition relation $T \subseteq B \times A \times B$, and for simplicity, one initial state $b_0 \in B$, and a set of goal states $G \subseteq B$. The interpretation is that $T(b, a, b')$ holds when action a performed in belief state b *may* lead to a successor belief state b'. An action a is said to be *possible* in b if there is at least one b' such that $T(b, a, b')$ holds, and the action is *deterministic* if there is at most one such b'. Though B is typically finite, the only constraint we really need to insist on is that, for any b and a, the set of b' such that $T(b, a, b')$ must be *finite*. Note that if we were to think of the states B as world states, we might want to allow the possibility of an agent not knowing what state it was in; but we are modeling belief states, and so assume that an agent always has enough introspection to be able to tell what belief state it is in.

Solving a planning problem means finding actions that will take us from the initial state to one of the goal states. There are very different forms of plans for different purposes. What we care about, however, is that when using a plan we can always get the next action to perform. For a sequential (or conformant) plan, we need a sequence of actions: a_1 then a_2 then ... then a_n. A conditional plan, however, can make the next action depend on the state that results from the execution of the previous one. To sidestep issues of plan structure, we simply assume that a plan P is a quadruple $P = \langle Q, q_0, \mathit{nexta}, \mathit{nexts}\rangle$ defined as follows:

- Q is a (possibly infinite) set of plan states;
- $q_0 \in Q$ is the initial plan state;
- $\mathit{nexta} \in [Q \times B \rightarrow A \cup \{\mathit{stop}\}]$, where $\mathit{nexta}(q, b)$ returns the next action to execute starting from a plan state q and a belief state b;
- $\mathit{nexts} \in [Q \times B \times A \rightarrow Q]$, where $\mathit{nexts}(q, b, a)$ returns the next state of the plan, starting from a plan state q and belief state b and performing action a.

Intuitively, a plan issues at each step, according to the function *nexta*, an action a or a special command *stop* that indicates that the plan has terminated. If an action a is issued, then the plan state evolves to a next state, according to the function *nexts*. Note that the plan state argument of function *nexta* allows us to arrive at a belief state more than once but take different actions.

Even though the state of the plan could always be seen as part of the agent's beliefs, for convenience, we keep this separate from the belief state. In addition, we assume the plan state is always known by the agent.

Finally, we insist on one constraint for a plan to be *legal*[4] with respect to a planning problem $\mathcal{B}$, namely, that it *does not prescribe impossible actions*: If $\mathit{nexta}(q, b) = a$ and $a \neq \mathit{stop}$, then for some b', $T(b, a, b')$. In what follows, we will only care about plans that are legal in this sense, and we will not mention the issue again.

Adequacy

We now present our first definition of adequacy, along the lines sketched in the introduction. We first define the set R_P of pairs $\langle q : b\rangle$ such that running P in plan state q and in belief state b is guaranteed to get to a goal state. Then, P is adequate if $\langle q_0 : b_0\rangle \in R_P$. More precisely, the definition is as follows:

Definition 1 *A plan $P = \langle Q, q_0, \mathit{nexta}, \mathit{nexts}\rangle$ is considered adequate with respect to a belief-based planning problem $\mathcal{B} = \langle B, A, T, b_0, G\rangle$ iff the pair $\langle q_0 : b_0\rangle \in R_P$, where R_P is the least set satisfying the following two properties:*

1. *If $\mathit{nexta}(q, b) = \mathit{stop}$ and $b \in G$, then $\langle q : b\rangle \in R_P$;*
2. *If $\mathit{nexta}(q, b) = a$, $\mathit{nexts}(q, b, a) = q'$, and for all b' such that $T(b, a, b')$ we have that $\langle q' : b'\rangle \in R_P$, then $\langle q : b\rangle \in R_P$.*

It is not hard to show that this definition is well formed in that there is indeed a unique least set R_P satisfying the two conditions. For the first condition, if P tells us to stop, then we must already be at the goal; for the second, if P tells us to continue with a, then no matter what belief state b' this takes us to (and there must be at least one since the plan is assumed to be legal and must not prescribe impossible actions), P must take us to the goal from there. Note that nothing stops us from having an action a that has transitions to b' and to b'', where from b' we can get to the goal in n steps, but from b'' we can only get to the goal in some different number of steps. We will now investigate how big those numbers can be.

We first define the configuration tree for a plan P:

Definition 2 *The configuration tree of a plan $P = \langle Q, q_0, \mathit{nexta}, \mathit{nexts}\rangle$ with respect to a belief-based planning problem $\mathcal{B} = \langle B, A, T, b_0, G\rangle$ is the smallest tree whose nodes are labelled with pairs $\langle q : b\rangle$ as follows:*

(i) the root of the tree is labelled $\langle q_0 : b_0\rangle$;

[4] We follow Reiter (2001) in the use of the term "legal." There is no deontic or judicial connotation.

(ii) from any node labelled $\langle q : b\rangle$, there is an edge to another node $\langle q' : b'\rangle$, for every b' such that $nexta(q, b) = a$, $nexts(q, b, a) = q'$ and $T(b, a, b')$.

A branch in the configuration tree is a, possibly infinite, sequence of nodes corresponding to a maximal path from the root node. If $\langle q : b\rangle$ is the last element in a (finite) branch, we say that the branch terminates in belief state b.

We will show an example of a configuration tree in the next section.

While plans that are adequate may get to the goal in different numbers of steps depending on the transitions taken, the following result shows that every branch of the configuration tree must eventually get to the goal:

Theorem 1 *A plan P is adequate with respect to $\mathcal{B}$ iff every branch of the configuration tree of P with respect to $\mathcal{B}$ terminates in a goal state.*

Proof: ($\Leftarrow$) We assume that every branch of the configuration tree terminates in a goal state. Let L be the set of all labels on nodes in the configuration tree. We will show that if R_P is any set that satisfies (1) and (2), then $L \subseteq R_P$. Since $\langle q_0 : b_0\rangle \in L$, it then follows that $\langle q_0 : b_0\rangle$ is an element of the least set satisfying (1) and (2), and therefore that P is adequate.

So suppose that R_P is any set that satisfies (1) and (2), and that $\langle q : b\rangle \in L$. We prove that $\langle q : b\rangle \in R_P$ by induction on the length d of the longest path from $\langle q : b\rangle$ to a leaf node (which is well defined, since each branch is finite).

1. If $d = 0$, then $\langle q : b\rangle$ is a leaf, and so $P(q, b) = stop$ and $b \in G$. Therefore, $\langle q : b\rangle \in R_P$ since R_P satisfies (1).
2. If $d > 0$, then $\langle q : b\rangle$ is not a leaf, and $nexta(q, b) = a$ for some a. Let $\langle q' : b'\rangle$ be any successor child of this node. Then, $q' = nexts(q, b, a)$ and $T(b, a, b')$ must hold. Also, node $\langle q' : b'\rangle$ has a smaller d, and by induction, $\langle q' : b'\rangle \in R_P$. Therefore, $\langle q : b\rangle \in R_P$, since R_P satisfies (2).

($\Rightarrow$) Let us assume that there is a branch of the configuration tree labelled with pairs $L = \{\langle q_0 : b_0\rangle, \langle q_1 : b_1\rangle, \langle q_2 : b_2\rangle, \ldots\}$ that does not terminate in a goal state (either it is infinite or the leaf node is not labelled $\langle q : b\rangle$ where $b \in G$). We will show that P cannot be adequate. What we will show is that if R_P is any set of pairs $\langle q : b\rangle$ that satisfies conditions (1) and (2) in the definition of adequate, then the set $R_P - L$ also satisfies (1) and (2). Since $\langle q_0 : b_0\rangle \in L$, it then follows that $\langle q_0 : b_0\rangle$ is not in the *least* set that satisfies (1) and (2), and therefore P is not adequate.

So assume that R_P is any set that satisfies (1) and (2).

1. Suppose that $nexta(q, b) = stop$ and $b \in G$. Then we have that $\langle q : b\rangle \notin L$, but $\langle q : b\rangle \in R_P$ since R_P satisfies (1). Therefore, $\langle q : b\rangle \in R_P - L$. So $R_P - L$ satisfies (1).
2. Suppose that $nexta(q, b) = a$ and that for every b' such that $T(b, a, b')$, $\langle q' : b'\rangle \in R_P - L$, with $q' = nexts(q, b, a)$. Therefore, for every b' such that $T(b, a, b')$, we have that $\langle q' : b'\rangle \in R_P$. Then $\langle q : b\rangle \in R_P$ since R_P satisfies (2). Also, by assumption, for every b' such that $T(b, a, b')$, $\langle q' : b'\rangle \notin L$. It then follows that $\langle q : b\rangle \notin L$, since each element $\langle q : b\rangle$ of L, except for when $nexta(q, b) = stop$, has a successor in L. Therefore, $\langle q : b\rangle \in R_P - L$ and $R_P - L$ satisfies (2).

This completes the proof. ∎

Next, we define whether a plan is bounded as follows:

Definition 3 *Plan P is bounded with respect to $\mathcal{B}$ if the configuration tree of P with respect to $\mathcal{B}$ is finite.*

Thus, a plan P is bounded iff there is some number n such that we cannot make more than n transitions using the actions specified by P (where n is the maximum depth of the configuration tree). Note that we can have *unbounded* plans even in a system with a single action a, a single plan state q_0, and a single belief state b: simply let $nexta(q_0, b) = a$ and $nexts(q_0, b, a) = q_0$. A less obvious case is where there are two states b_1 and b_2 where $T(b_1, a, b_1)$ and $T(b_1, a, b_2)$, and where $nexta(q_0, b_1) = a$, $nexta(q_0, b_2) = stop$, and $nexts(q_0, b_1, a) = q_0$. This is unbounded because of the infinite sequence of belief states $b_1, b_1, \ldots$, even though a transition to b_2 would allow us to terminate.

The main result of this section is the following:

Theorem 2 *If a plan P is adequate with respect to $\mathcal{B}$, then it is bounded with respect to $\mathcal{B}$.*

Proof: Suppose to the contrary that P is not bounded. Then, the configuration tree of P has infinitely many nodes. Since for any action, the transition relation T leads to only finitely many successor states, the tree is finitely branching. Therefore, by König's Lemma, the tree contains an infinite branch. Hence, at least one branch of the tree does not terminate. By Theorem 1, P is not adequate. ∎

So the notion of adequacy that we have defined, however intuitive it may appear, never holds for unbounded plans. In other words, for any adequate plan P, there is a number n such that P cannot perform more than n actions. We now turn to a simple example where unbounded plans are required.

A Problematic Example

Consider a situation in which an agent wants to cut down a tree. Assume that the agent has a primitive action *chop* to chop at the tree, and that the tree *will eventually come down* if it is hit often enough. Assume also that the agent can find out whether the tree is down by doing a (binary) sensing action *look*: a result of 1 indicates that the tree is down; a result of 0 means that the tree is still up.

In its simplest form, we can model this problem using the belief-based planning problem $\mathcal{B}_{tc} = \langle B_{tc}, \{chop, look\}, T_{tc}, b_u, \{b_d\}\rangle$, where B_{tc} has three belief states: b_u is the state where it is known that the tree is up; b_d is the state where it is known that the tree is down; and $b_?$ is the state where the status of the tree is *not* known. The transition relation T_{tc} is the following:

$$(b_u, chop, b_?),\ (b_u, look, b_u),\ (b_d, look, b_d),$$
$$(b_?, look, b_u),\ (b_?, look, b_d).$$

That is, a *chop* action is only possible if we believe the tree to be up, and the result is to move to a state where we do not

$$
\begin{tikzcd}
\langle q : b_u \rangle \arrow[r, "chop"] & \langle q : b_? \rangle \arrow[r, "look"] \arrow[d, "look"] & \langle q : b_u \rangle \arrow[r, "chop"] & \langle q : b_? \rangle \arrow[r, "look"] \arrow[d, "look"] & \cdots \\
& \langle q : b_d \rangle & & \langle q : b_d \rangle &
\end{tikzcd}
$$

Figure 1: Execution tree of program P_{tc} with respect to belief-based tree chopping planning problem $\mathcal{B}_{tc}$. A link between two nodes means that a transition is "possible" between the two corresponding configurations.

know the status of the tree; a *look* action is possible in any state, and if done when we do not know the status of the tree, we move nondeterministically to one of the states where we do know it. Initially, we are in the state b_u (i.e., the tree is known to be up), and the goal is to get to b_d (i.e., the tree is known to be down).

The most obvious "reactive" plan then is P_{tc}, which has only one state q, and its next action and next plan state functions are defined as follows:

$$
\begin{aligned}
&\textit{nexta}(q, b_d) = \textit{stop}, \\
&\textit{nexta}(q, b_u) = \textit{chop}, \\
&\textit{nexta}(q, b_?) = \textit{look}, \\
&\textit{nexts}(q, b, \textit{chop}) = q, \ \text{ for } b \in \{b_?, b_u, b_d\}, \\
&\textit{nexts}(q, b, \textit{look}) = q, \ \text{ for } b \in \{b_?, b_u, b_d\}.
\end{aligned}
$$

In other words, we chop at the tree when we know it is up, stop when we know it is down, and look otherwise. We repeat this unless we have stopped. Intuitively, this is all it takes to get the tree down no matter how many chops are needed—the tree will go down after a finite, though unbounded, number of chops. Note that the above strategy would not work if the tree were made of *metal* and chopping at it therefore had no effect (we shall deal with this variant later in the paper). In Figure 1, the configuration tree of plan P_{tc} with respect to belief-based planning problem $\mathcal{B}_{tc}$ is depicted.

Unfortunately, this plan is not adequate:

Theorem 3 *The plan P_{tc} is not adequate with respect to planning problem $\mathcal{B}_{tc}$.*

Proof: The plan is not bounded, since we can go through the states $b_u, b_?, b_u, b_?, \ldots$ indefinitely. Then, we apply Theorem 2. ∎

This result is perhaps not too surprising since nothing in our characterization $\mathcal{B}_{tc}$ captures the fact that the tree will eventually go down if it is hit often enough. To do so, we will look at belief-based planning problems from a different perspective, where they are induced by world-based ones.

World-Based Planning

Let us imagine that we start with a potentially infinite set of *world* states W. Rather than describing actions as nondeterministic transitions among belief states, we now think of actions as *deterministic* transitions among world states. We assume that actions can return a sensing result, which for simplicity, we take to be binary. In different world states, we may get different sensing results.[5] However, in this view, we still want to apply plans to belief states and not world states, since an agent may not know what world state it is in. We think of belief states as certain non-empty subsets of W. Intuitively, a belief state $b \subseteq W$ is one where the agent believes that any element of b might be the real world.

Formally, $\mathcal{W} = \langle W, A, \tau, \sigma, i, g \rangle$ is a world-based planning problem where W is a set of world states, A is a finite set of actions, $\tau \in [A \times W \to W \cup \{\perp\}]$ is a transition function, $\sigma \in [A \times W \to \{0, 1\}]$ is a sensing function, $i \subseteq W$ is a non-empty set of initial world states, and $g \subseteq W$ is a non-empty set of goal world states. Symbol $\perp$ here is used for an action that is not possible in a world state.

It is possible to induce a belief-based planning problem from a world-based planning one, as follows. As suggested above, the states of the induced belief-based planning problem will be sets of world states. Probably the most difficult task is to characterize what the resulting belief state (i.e., set of world states) is when an action is executed in some, possibly different, belief state. To that end, it is convenient to define, for any $w \in W$, any $a \in A$, and any $b \subseteq W$, the set of world states $\kappa(w, a, b)$ as follows:

$$
\begin{aligned}
&w'' \in \kappa(w, a, b) \ \text{ iff} \\
&\quad \text{for some world state } w' \in b, \\
&\quad w'' = \tau(a, w') \text{ and } \sigma(a, w') = \sigma(a, w).
\end{aligned}
$$

The definition of $\kappa(w, a, b)$ is a version of the successor state axiom for knowledge proposed by Scherl and Levesque (Scherl & Levesque 2003), where w is the real world, b is the set of worlds accessible from w, and $\kappa(w, a, b)$ is the new set of accessible worlds after doing action a in world w and belief state b. Overall, if all the world states in b agree on the sensing result for action a, there will be one resulting belief state b' (a deterministic outcome); otherwise there will be two (a nondeterministic outcome).

We can now formally define the induced belief-based planning problem from a world-based one.

Definition 4 *Let $\mathcal{W} = \langle W, A, \tau, \sigma, i, g \rangle$ be a world-based planning problem. The belief-based planning problem induced by $\mathcal{W}$ is $\langle B, A, T, b_0, G \rangle$, where*

1. *$b_0 = i$ (i.e., the initial belief state is the set of all initial world states);*
2. *$B \subseteq 2^{\mathcal{W}}$ is the least set such that: (i) $b_0 \in B$; (ii) for all $b \in B$ and $a \in A$, if $T(b, a, b')$, then $b' \in B$;*
3. *$T(b, a, b')$ holds iff (i) for all $w \in b$, $\tau(w, a) \neq \perp$, and (ii) for some $w^* \in b$, $b' = \kappa(w^*, a, b)$;*
4. *$G = \{b \mid b \in B \text{ and } b \subseteq g\}$.*

Thus, the belief states we deal with are those subsets of W that result from starting with belief state b_0 and closing under the T relation. The T relation for action a starting in a belief state b involves progressing (using τ) the world states that are elements of b, but dividing them into groups that have the same sensing result (using σ). A belief state

[5] We assume that non-sensing actions always return the same sensing result.

is a goal state in the induced problem if it is in B and is a subset of g. Note that, when building the induced T relation, we require that the action be *known* to be possible, that is, possible in all (accessible) world states (condition (i) in point (3) of Definition 4).

Let us now define a world-based version of the tree chopping problem $\mathcal{W}_{tc} = \langle W_{tc}, \{chop, look\}, \tau_{tc}, \sigma_{tc}, i_{tc}, g_{tc}\rangle$. To capture the fact that the tree will eventually go down if it is hit sufficiently often, we can model world states using $W_{tc} = \mathbb{N}$ (i.e., the natural numbers including 0), with the interpretation that a world n is one where n chops are necessary to bring the tree down—in world state 0, the tree is down. So, in every world, the tree will go down if it is hit enough times. The initial states $i_{tc} = \mathbb{N} - \{0\}$ are those where the tree is up. The goal state $g_{tc} = \{0\}$ is the single state where no further chops are needed (and so the tree is down). We can define $\sigma_{tc}(chop, n) = 0$ and $\tau_{tc}(look, n) = n$ for all n, and

$$\sigma_{tc}(look, n) = \begin{cases} 1 & \text{if } n = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$\tau_{tc}(chop, n) = \begin{cases} \bot & \text{if } n = 0 \\ n-1 & \text{otherwise} \end{cases}$$

So the *look* action reports a 1 if the tree is down and 0 otherwise, and the chop action changes the world so that one less chop is needed to fell the tree.

Despite the fact that $\mathcal{W}_{tc}$ has an infinite set of world states, the belief-based planning problem it induces has just three (belief) states:

Theorem 4 *The belief-based planning problem induced by $\mathcal{W}_{tc}$ is equivalent to $\mathcal{B}_{tc}$.*

Proof (sketch): The induced belief-based planning problem is $\mathcal{B}^{\mathcal{W}_{tc}} = \langle\{b_?, b_u, b_d\}, \{chop, look\}, T_{tc}, b_u, \{b_d\}\rangle$, where $b_? = \mathbb{N}$, $b_u = \mathbb{N} - \{0\}$, $b_d = \{0\}$, and relation T_{tc} is the one from the previous section. (See the appendix for the full proof.) ∎

So, while a world-based planning problem may look intractable in terms of the number of states it deals with, the resulting belief-based problem may be quite small.

Adequacy Reconsidered

What is the point of redefining the tree chopping example in terms of worlds? After all, given that we end up with exactly the same belief-based planning problem, the plan P_{tc} remains unbounded and therefore not adequate. The answer is that the world-based version will let us see what goes wrong in the definition of adequacy and how to fix it.

First observe that the plan P_{tc} will be adequate for any *bounded* version $\mathcal{W}^n_{tc}$ of the tree chopping problem:

Theorem 5 *Let $\mathcal{W}^n_{tc}$ be just like $\mathcal{W}_{tc}$ except that the set of world states W_n is the set $\{0, 1, \ldots, n\}$, for some fixed $n \geq 0$. Then, the plan P_{tc} is adequate for the belief-based planning problem $\mathcal{B}^{\mathcal{W}^n_{tc}}$ induced by $\mathcal{W}^n_{tc}$.*

Proof: Consider the configuration tree of P_{tc} with respect to the induced problem $\mathcal{B}^{\mathcal{W}^n_{tc}}$. Each branch will be finite, since no branch can have more than n chop actions. Also, each leaf will be labelled $\langle q_0 : \{0\}\rangle$. So, P_{tc} is adequate by Theorem 1. ∎

We point out that although $\mathcal{W}^n_{tc}$ has only finitely many world states, the belief-based planning problem it induces has $2n$ belief states, compared to only 3 belief states for the infinite $\mathcal{W}_{tc}$.

Now let us reconsider tree chopping in the general case for $\mathcal{W}_{tc}$. We start in the belief state $b_u = \mathbb{N} - \{0\}$. We do not know how many chops are needed—every $m > 0$ is a member of the belief state b_u. Then, we do a *chop* action, which subtracts one from every element of b_u, producing $b_? = \mathbb{N}$. Next, we do a *look*, one of whose results is b_d and the other is b_u. In b_d, we know that the tree is down, but in b_u, we are back to where we started. In terms of our beliefs, we have made no progress, and it appears that P_{tc} is no closer to a state where the tree is down.

But this is an illusion. While it is true that we do not know how many chops are needed, in the real world, some number of chops is sufficient. If we start in a state where n chops are really needed, action *look* will unambiguously tell us that the tree is up until we do n chops, at which point, it will unambiguously tell us that the tree is down. Once we fix the starting state, the behavior of P_{tc} is completely determined.

This is unfortunately not what happens in our definition of adequate. As we perform *chop* actions, we ask what *look* will do with respect to *any member of our current belief state*. To see this more formally, note that condition (2) in the definition of adequate (Definition 1) requires that

$\langle q' : b'\rangle \in R_P$, for every b' such that $T(b, a, b')$.

For a planning problem induced by a world-based one, this is the same as requiring that: $\langle q' : b'\rangle \in R_P$, for every b' such that for some $w \in b$, $b' = \kappa(w, a, b)$. What our definition of adequate is saying, in effect, is that at any given point, our plan must work under the assumption that sensing results may come from any element w of the current belief state b. Since, until we know the tree is down, there will always be elements of b with more chops remaining, we never see the progress happening in the real world.

To remedy this, we need to consider each world state in the initial belief state separately, and ask if the plan will work in each case. Instead of considering a set of pairs $\langle q : b\rangle$, we consider a set of triples $\langle w : q : b\rangle$ where $w \in b$ and such that running P from its state q in belief state b *and with sensing determined by* w is guaranteed to get to a goal state. More precisely, the new definition is as follows:

Definition 5 *A plan $P = \langle Q, q_0, nexta, nexts\rangle$ is considered <u>adequate′</u> with respect to a world-based planning problem $\mathcal{W} = \langle W, A, \tau, \sigma, i, g\rangle$ iff for every $w_0 \in i$, the triple $\langle w_0 : q_0 : i\rangle \in R_P$, where R_P is the least set satisfying the following:*

1. *If $nexta(q, b) = stop$ and $b \subseteq g$, then $\langle w : q : b\rangle \in R_P$;*
2. *If $nexta(q, b) = a$, $nexts(q, b, a) = q'$, and $\langle \tau(a, w) : q' : \kappa(w, a, b)\rangle \in R_P$, then $\langle w : q : b\rangle \in R_P$.*

This is exactly like the previous definition, except that sensing is now done with respect to a world state w (that changes systematically as actions are performed) rather than with respect to arbitrary elements of the current belief state b.

We can define an analogue to branches of the configuration trees for world-based problems:

Definition 6 *The* <u>*run*</u> *of a plan P with respect to $\mathcal{W}$ on initial state $w_0 \in i$ is the smallest sequence of triples $\langle w : q : b\rangle$ such that:*

1. *the initial triple of the sequence is $\langle w_0 : q_0 : i\rangle$;*
2. *if $\langle w : q : b\rangle$ is in the sequence, and $nexta(q, b) = a$ and $nexts(q, b, a) = q'$, then the next triple in the sequence is $\langle \tau(a, w) : q' : \kappa(w, a, b)\rangle$.*

Then we get an analogue to Theorem 1 with a similar proof:

Theorem 6 *A plan P is adequate$'$ with respect to $\mathcal{W}$ iff for every initial state $w_0 \in i$ the run of P with respect to $\mathcal{W}$ on state w_0 terminates in a goal state.*

Proof: See appendix. ∎

Applying this to tree chopping, we get the following:

Theorem 7 *The plan P_{tc} is adequate$'$ with respect to $\mathcal{W}_{tc}$.*

Proof: Using Theorem 6 (every run of P_{tc} with respect to $\mathcal{W}_{tc}$ terminates in a goal state). ∎

Hence, adequacy$'$ differs from adequacy in that it correctly judges the plan P_{tc} to solve the tree chopping problem. Here is the crux of the matter: whereas each of the infinitely many runs of P_{tc} is finite, the configuration tree of P_{tc} has an infinite branch that does not correspond to anything in reality.

As a result, adequacy and adequacy$'$ are *not* the same. They are very close, however, and in fact, agree on bounded plans:

Theorem 8 *Let $\mathcal{B}$ be a belief-based problem induced by some $\mathcal{W}$. Suppose a plan P is bounded with respect to $\mathcal{B}$. Then, P is adequate$'$ with respect to $\mathcal{W}$ iff P is adequate with respect to $\mathcal{B}$.*

Proof: See appendix. ∎

So the two definitions are identical in the bounded case: planning with one is the same as planning with the other. Of course, these properties of adequacy$'$ would still be true if *every* unbounded plan were adequate$'$. To show that adequacy$'$ is not too weak, we discuss a variant of the tree chopping example that is intuitively unsolvable.

Imagine that not only do we not know the number of chops needed to fell the tree, we also have the possibility of being in a world state $*$ where the tree is up and *chop* does not change the state (e.g., the tree is actually a steel lamp post). Formally, we can let $\mathcal{W}_*$ be just like $\mathcal{W}_{tc}$ except that i now includes the new world state $*$, where $\tau(a, *) = *$ and $\sigma(a, *) = 0$, for every action a. Obviously, the planning problem induced by $\mathcal{W}_*$ should not be solvable: if it turns out that the real world is $*$, nothing can be done to fell the tree. We get the following:

Theorem 9 *No plan is adequate$'$ with respect to $\mathcal{W}_*$.*

Proof: The run of any plan P on initial state $* \in b_0$ will never get to the goal state. Therefore, by Theorem 6, P is not adequate$'$. ∎

Hence, adequacy$'$ follows our intuitions once more in this case. Interestingly, from the point of view of belief-based planning, nothing has changed in this variant:

Theorem 10 *The belief-based planning problem induced by $\mathcal{W}_*$ is equivalent to the one induced by $\mathcal{W}_{tc}$.*

Proof: The belief-based planning problem induced by $\mathcal{W}_*$ is very similar to that induced by $\mathcal{W}_{tc}$, namely $\mathcal{B}^{\mathcal{W}_{tc}}$, described in the proof of Theorem 4. It also has 3 belief states, corresponding to b_u, b_d, and $b_?$, which are as in $\mathcal{B}^{\mathcal{W}_{tc}}$, except that the $*$ world is included in both b_u and $b_?$. The induced transition relation is exactly as in $\mathcal{B}^{\mathcal{W}_{tc}}$. ∎

So, with respect to belief-based planning, $\mathcal{W}_{tc}$ and $\mathcal{W}_*$ seem to be equally unsolvable (and so much the worse for belief-based planning).

To recap: for any belief-based planning problem induced by a world-based one, the adequate plans coincide with the adequate$'$ plans in the bounded case. In the unbounded case, there are no adequate plans, but some adequate$'$ plans may exist, and just in those cases where the planning problem appears to be intuitively solvable. The conclusion: adequacy$'$ is the appropriate definition of plan correctness, not adequacy.

Discussion and Conclusion

Even though superficially similar, adequate and adequate$'$ are very different. Adequacy is a condition on the *belief states* that are reachable through possible transitions of the plan being considered; adequacy$'$, in contrast, depends on what is believed about possible transitions in *world states*. Our results mean that the former is unable to recognize progress towards the goal in the general case, that is, when the solution to the planning problem cannot be bounded in advance. The latter then, though more involved, is required to correctly capture when a plan is a solution to a planning problem.

Our results also mean that planning with loops is more difficult than what one might have expected. Note that in this paper, we were not interested in developing algorithms for iterative planning, but only in investigating the conceptual limitations of planning over belief states. For practical approaches to iterative planning, we refer to (Lin & Dean 1996; Son, Baral, & McIlraith 2001; Cimatti *et al.* 2003; Levesque 2005).

In (Cimatti *et al.* 2003), three different types of solutions for planning problems in nondeterministic domains are devised, namely, *weak*, *strong*, and *strong cyclic* solutions. Weak solutions are plans that may achieve the goal, but are not guaranteed to do so; strong solutions are plans that are guaranteed to achieve the goal. In our work, we have not been concerned with weak solutions, but rather with strong “safe” ones. However, Cimatti *et al.*’s notion of strong plans does not account for (unbounded) iterative behavior. Strong cyclic plans are an alternative way to deal with unbounded domains by capturing the notion of “acceptable” iterative trial-and-error strategies. A strong cyclic plan is not required to actually reach a goal state in the unbounded cases, but only to remain on a path to a goal state. This may be advantageous if for instance, the world might change to delay the attainment of the goal. On the other hand, a strong cyclic plan may procrastinate indefinitely. Hence, adequate$'$ plans are stricter, in the sense that indefinite procrastination is disallowed. In fact, whereas our definition of adequate$'$ separates the two tree chopping scenarios $\mathcal{W}_{tc}$ and $\mathcal{W}_{*}$ correctly, Cimatti *et al.*'s account would not as it would produce the same (strong cyclic) plan with the same guarantee in both cases. Roughly speaking, this is because planning problems are cast as finite state systems and, as a result, it is not possible to express fairness directly which then must be assumed as a meta-constraint on the system.

The reader may wonder how our work relates to decision theoretic and probabilistic planning (e.g., (Kaelbling, Littman, & Cassandra 1998; Boutilier, Dean, & Hanks 1999; Draper, Hanks, & Weld 1994; Bonet & Geffner 2000)) in which the underlying setting is similar, tough somewhat richer, than the one assumed here. In those areas, belief states are often represented as probability distributions over the set of world states. It is well known that belief states comprise sufficient statistics for acting *optimally* (see (Aström 1965)). There are basically two main differences with the work presented here. First, we are concerned with checking whether a plan is (always and completely) "effective" for the given goal rather than checking whether the plan is "optimal" or "satisficing" (i.e., gets maximum expected reward or has sufficient chances of success). Still, one can imagine ways of recasting our objective in terms of utilities and probabilities. Second, and most importantly, we have assumed that the uncertainty present in the domain cannot be quantified, whereas the aforementioned approaches do assume that a probability distribution over the world states is available. When these probabilities are available, one can annotate configuration trees with transition probabilities between a node and each of its children. We suspect that, by making use of this additional information on the configuration trees, one could develop a purely belief-based account of adequacy expressive enough to deal with even unbounded problems like the tree chopping example. The intuitive reason for this is that one could make use of the available probabilities, together with the reward system, to recognize progress towards the goal. When it comes to our tree chopping example, a belief state will be a distribution over an infinite set of world states with zero reward except for only one absorbing state (where there are 0 chops remaining) with reward $r > 0$. The expected utility of plan P_{tc} will turn out to be exactly r. A fuller study of this is left for future work.

In this paper, we made assumptions on plans sufficient to guarantee two fundamental properties: (i) epistemic feasibility and (ii) determinism. A plan is *epistemically feasible* if an executor will *always* have enough information to continue the execution up to successful termination—a necessary requirement for autonomous execution of the plan. The plans used in this paper are always epistemically feasible, but if the programming constructs in our planning language are general enough, the issue of verifying that a plan is actually epistemically feasible arises (see (Sardina *et al.* 2004)). As for *determinism*, this implies that plans do not leave any leeway to the executor. This is certainly the case in the vast majority of the planning literature, and indeed we required it in this paper. However, different notions of plans can be developed in which plans are in fact nondeterministic programs. In this case, one must consider how the choice between different transitions available to the executor is made, perhaps arbitrarily by a "dumb executor" or perhaps intelligently, by an executor that plans ahead. Interestingly, natural generalizations of *adequate* and *adequate$'$* for this class of plans can be defined (see (De Giacomo *et al.* 2003)).

Thus, our results are also relevant for agent programming language design and semantics, where one may need a notion of agent program adequacy that handles unbounded iterative programs. Also, we point out that an account of agent ability (e.g., (Moore 1985; Lespérance 2003)) would necessarily have to include an account of plan adequacy in the style of what we have done here: to be able to achieve ϕ is to know of an "adequate" plan for it. We will address some of the above issues in future work.

Acknowledgments

This work was funded by the National Science and Engineering Research Council of Canada, an Ontario Graduate Scholarship, the European FET basic research project FP6-7603 Thinking Ontologies (TONES), and the Australian Research Council and Agent Oriented Software under the grant LP0560702. The authors would also like to thank the anonymous reviewers for their comments.

References

Aström, K. J. 1965. Optimal control of markov decision processes with incomplete state estimation. *Journal of Mathematical Analysis and Applications* 10:174–205.

Bacchus, F., and Petrick, R. 1998. Modeling an agent's incomplete knowledge during planning and execution. In *Principles of Knowledge Representation and Reasoning (KR)*, 432–443.

Bertoli, P.; Cimatti, A.; Roveri, M.; and Traverso, P. 2001. Planning in nondeterministic domains under partial observability via symbolic model checking. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 473–478.

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., *Proceedings of the International Conference on AI Planning & Scheduling (AIPS)*, 52–61. Breckenridge, CO: AAAI Press.

Boutilier, C.; Dean, T.; and Hanks, S. 1999. Decision-theoretic planning: Structural assumptions and compuattional leverage. *Journal of Artificial Intelligence Research* 11:1–94.

Bryce, D., and Kambhampati, S. 2004. Heuristic guidance measures for conformant planning. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*, 365–375.

Cimatti, A.; Pistore, M.; Roveri, M.; and Traverso, P. 2003. Weak, strong, and strong cyclic planning via symbolic model checking. *Artificial Intelligence Journal* 1–2(147).

De Giacomo, G.; Iocchi, L.; Nardi, D.; and Rosati, R. 1997. Planning with sensing for a mobile robot. In *Proceedings of 1997 European Conference on Planning (ECP-97)*, 156–168.

De Giacomo, G.; Lespérance, Y.; Levesque, H.; and Sardina, S. 2003. On deliberation under incomplete information and the inadequacy of entailment and consistency-based formalizations. In *Proceedings of the 1st Programming Multiagent Systems Languages, Frameworks, Techniques and Tools workshop (PROMAS-03)*.

Draper, D.; Hanks, S.; and Weld, D. S. 1994. Probabilistic planning with information gathering and contingent execution. In *Proceedings of the International Conference on AI Planning & Scheduling (AIPS)*, 31–36.

Kaelbling, L. P.; Littman, M. L.; and Cassandra, A. R. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence Journal* 101:99–134.

Lespérance, Y. 2003. On the epistemic feasibility of plans in multiagent systems specifications. *Logic Journal of the IGPL* 11(2):161–178.

Levesque, H. J. 2005. Planning with loops. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*, 509–515.

Lin, S.-H., and Dean, T. 1996. Generating optimal policies for high-level plan. In Ghallab, M., and Milani, A., eds., *New Directions in AI Planning*. Amsterdam: IOS Press. 187–200.

Moore, R. C. 1985. A formal theory of knowledge and action. In Hobbs, J. R., and Moore, R. C., eds., *Formal Theories of the Common Sense World*. 319–358.

Petrick, R., and Bacchus, F. 2002. A knowledge-based approach to planning with incomplete information and sensing. In *Proceedings of the International Conference on AI Planning & Scheduling (AIPS)*, 212–221.

Reiter, R. 2001. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press.

Sardina, S.; De Giacomo, G.; Lespérance, Y.; and Levesque, H. 2004. On the semantics of deliberation in IndiGolog – from theory to implementation. *Annals of Mathematics and Artificial Intelligence* 41(2–4):259–299. (Previous version appeared in Proc. of KR-2002).

Scherl, R., and Levesque, H. J. 2003. Knowledge, action, and the frame problem. *Artificial Intelligence Journal* 144(1–2):1–39.

Son, T.; Baral, C.; and McIlraith, S. 2001. Extending answer set planning with sequence, conditional, loop, non-deterministic choice, and procedure constructs. In *Proceedings of the AAAI Spring Symposium on Answer Set Programming: Towards Efficient and Scalable Knowledge Representation and Reasoning*, 202–209.

Proofs

Theorem 4 *The belief-based planning problem induced by $\mathcal{W}_{tc}$ is equivalent to $\mathcal{B}_{tc}$ (as defined in the previous section).*

Proof: Let us recall the world-based version of the tree chopping problem:

$$\mathcal{W}_{tc} = \langle \mathbb{N}, \{chop, look\}, \tau_{tc}, \sigma_{tc}, \mathbb{N} - \{0\}, \{0\} \rangle,$$

where $\sigma_{tc}(chop, n) = 0$ and $\tau_{tc}(look, n) = n$ for all n, and

$$\sigma_{tc}(look, n) = \begin{cases} 1 & \text{if } n = 0 \\ 0 & \text{otherwise} \end{cases}$$

$$\tau_{tc}(chop, n) = \begin{cases} \bot & \text{if } n = 0 \\ n-1 & \text{otherwise} \end{cases}$$

So, let us build inductively the induced belief-based planning problem $\mathcal{B}^{\mathcal{W}_{tc}} = \langle B^{\mathcal{W}_{tc}}, \{chop, look\}, T^{\mathcal{W}_{tc}}, \mathbb{N} - \{0\}, G^{\mathcal{W}_{tc}} \rangle$. First of all, we know that $\mathbb{N} - \{0\} \in B^{\mathcal{W}_{tc}}$. Next, we build $T^{\mathcal{W}_{tc}}$ and $B^{\mathcal{W}_{tc}}$ simultaneously by starting with the initial belief state b_u and applying the two available actions:

1. $T^{\mathcal{W}_{tc}}(\mathbb{N} - \{0\}, look, b)$ iff $b = \mathbb{N} - \{0\}$ (and $\mathbb{N} - \{0\}$ is already in $B^{\mathcal{W}_{tc}}$).
2. $T^{\mathcal{W}_{tc}}(\mathbb{N} - \{0\}, chop, b)$ iff $b = \mathbb{N}$. This is because for all $n > 0$, $\tau(n, chop) \neq \bot$ and $\kappa(n, chop, b_u) = \mathbb{N}$. At this point, we obtained a new belief state $\mathbb{N}$, and hence, we add it to the belief base of $\mathcal{B}^{\mathcal{W}_{tc}}$, that is, $\mathbb{N} \in B^{\mathcal{W}_{tc}}$.
3. $T^{\mathcal{W}_{tc}}(\mathbb{N}, look, b)$ iff $b \in \{\mathbb{N} - \{0\}, \{0\}\}$. First of all, for every $n \geq 0$, $\tau(n, look) \neq \bot$. Moreover, $\kappa(0, look, \mathbb{N}) = \{0\}$, and $\kappa(m, look, \mathbb{N}) = \mathbb{N} - \{0\}$ for $m > 0$. Thus, we have a new belief state $\{0\} \in B^{\mathcal{W}_{tc}}$.
4. $T^{\mathcal{W}_{tc}}(\mathbb{N}, chop, b)$ does not hold for any b because $\tau(0, chop) = \bot$. That is, it is not possible to do a *chop* action in belief state $\mathbb{N}$.
5. $T^{\mathcal{W}_{tc}}(\{0\}, look, b)$ iff $b = \{0\}$, since $\tau(0, look) \neq \bot$ and $\kappa(0, look, \{0\}) = \{0\}$.
6. $T^{\mathcal{W}_{tc}}(\{0\}, chop, b)$ does not hold for any b because $\tau(0, chop) = \bot$. That is, it is not possible to do a *chop* action in belief state $\{0\}$.

At this point, we have finished defining $T^{\mathcal{W}_{tc}}$ and the resulting set of belief states is $B^{\mathcal{W}_{tc}} = \{\mathbb{N}, \mathbb{N} - \{0\}, \{0\}\}$. Finally, the set of goal states for $\mathcal{B}^{\mathcal{W}_{tc}}$ is easily obtained as $G^{\mathcal{W}_{tc}} = \{\{0\}\}$.

Putting it all together, the induced belief planning problem is as follows:

$$\mathcal{B}^{\mathcal{W}_{tc}} = \langle \{\mathbb{N}, \mathbb{N} - \{0\}, \{0\}\}, \{chop, look\}, T^{\mathcal{W}_{tc}}, \mathbb{N} - \{0\}, \{\{0\}\} \rangle.$$

It is not hard to see that $\mathcal{B}^{\mathcal{W}_{tc}}$ is just a “renaming” of $\mathcal{B}_{tc}$ from the example section, by talking $b_? = \mathbb{N}$, $b_u = \mathbb{N} - \{0\}$, and $b_d = \{0\}$. ∎

Theorem 6 *A plan P is adequate$'$ with respect to $\mathcal{W}$ iff for every initial state $w_0 \in i$ the run of P with respect to $\mathcal{W}$ on state w_0 terminates in a goal state.*

Proof: The proof mirrors the one for Theorem 1.
($\Leftarrow$) We assume that every run of P terminates in a goal state. Let L be the set of labels on all the nodes of all the runs. It can be shown, using an argument like the one in Theorem 1, that if R is any set that satisfies conditions (1) and (2) in the definition of adequate$'$, then $L \subseteq R$. Since for every $w_0 \in i$, $\langle w_0 : q_0 : i\rangle \in L$, it then follows that $\langle w_0 : q_0 : i\rangle$ is an element of the least set satisfying (1) and (2), and therefore that P is adequate$'$.
($\Rightarrow$) Now let us assume that there is a run of plan P labelled with triples $L = \{\langle w_0 : q_0 : i\rangle, \langle w_1 : q_1 : b_1\rangle, \langle w_2 : q_2 : b_2\rangle, \ldots\}$, for some $w_0 \in i$, that does not terminate in a goal state (either it is infinite or the leaf node is not labelled $\langle w : q_k : b_g\rangle$ with $b_g \subseteq G$). It can be shown, using an argument like the one in Theorem 1, that if R_P is any set of triples $\langle w : q : b\rangle$ that satisfies conditions (1) and (2) in the definition of adequate$'$, then the set $R_P - L$ also satisfies (1) and (2). Since $\langle w_0 : q_0 : i\rangle \in L$, it then follows that $\langle w_0 : q_0 : i\rangle$ is not in the *least* set that satisfies (1) and (2), and therefore P is not adequate$'$. ∎

Theorem 8 *Let $\mathcal{B}$ be a belief-based planning problem induced by some $\mathcal{W}$. Suppose a plan P is bounded with respect to $\mathcal{B}$. Then P is adequate$'$ with respect to $\mathcal{W}$ iff P is adequate with respect to $\mathcal{B}$.*

Proof: ($\Leftarrow$) Suppose P is adequate (and hence, bounded). By Theorem 1, every branch of the configuration tree of P terminates in a goal state. Now consider an arbitrary run of P. If we replace any label $\langle w : q : b\rangle$ in this run by the pair $\langle q : b\rangle$, we obtain a branch of the configuration tree. Hence the run is finite and terminates in a goal state. Since this holds for any run, by Theorem 6, P is adequate$'$.

($\Rightarrow$) Suppose P is adequate$'$. By Theorem 6, every run of P terminates in a goal state. Now consider an arbitrary branch of the configuration tree of P. In general, there may be no run corresponding to this branch (*e.g.* the infinite branch for P_{tc}). But if P is bounded, we can construct a run corresponding to the finite branch as follows:

1. if the leaf of the branch is $\langle q : b\rangle$, we make the leaf of the run be $\langle w : q : b\rangle$, for some $w \in b$;
2. if we have a node $\langle q' : b'\rangle$ on the branch corresponding to $\langle w' : q' : b'\rangle$ on the run, and the predecessor node of $\langle q' : b'\rangle$ is $\langle q : b\rangle$, then we add the node $\langle w : q : b\rangle$ as the predecessor to $\langle w' : q' : b'\rangle$, where w is some element of b such that $w' = \tau(a, w)$, for which it will then follow that $b' = \kappa(w, a, b)$. (Such a w must exist since every element of b' (including w') is of the form $\tau(a, w'')$ for some $w'' \in b$.)

Since every run terminates in a goal state, it follows that the branch also terminates in a goal state. Since this holds for any branch, by Theorem 1, P is adequate. ∎