

On the Nature of Modal Truth Criteria in Planning*

Subbarao Kambhampati

Department of Computer Science and Engineering
Arizona State University
Tempe, AZ 85287-5406
Email: rao@asu.edu

Dana S. Nau

Department of Computer Science,
Institute for Systems Research,
and Institute for Advanced Computer Studies
University of Maryland
College Park, MD 20742
Email: nau@cs.umd.edu

Abstract

Chapman's paper, "Planning for Conjunctive Goals," has been widely acknowledged for its contribution toward understanding the nature of nonlinear (partial-order) planning, and it has been one of the bases of later work by others---but it is not free of problems. This paper addresses some problems involving modal truth and the Modal Truth Criterion (MTC). Our results are as follows:

1. Even though modal duality is a fundamental axiom of classical modal logics, it does not hold for modal truth in Chapman's plans; i.e., "necessarily p " is not equivalent to "not possibly $\neg p$."
2. Although the MTC for necessary truth is correct, the MTC for possible truth is incorrect: it provides necessary but *insufficient* conditions for ensuring possible truth. Furthermore, even though necessary truth can be determined in polynomial time, possible truth is NP-hard.
3. If we rewrite the MTC to talk about modal *conditional* truth (i.e., modal truth conditional on executability) rather than modal truth, then both the MTC for necessary conditional truth and the MTC for possible conditional truth are correct; and both can be computed in polynomial time.

1 Introduction

Chapman's paper, "Planning for Conjunctive Goals," [2] has been widely acknowledged as an important step towards formalizing partial-order planning, and it has been one of the bases of later work by others (for example, [5, 7, 9, 12, 14]). Unfortunately, however, Chapman's work is not free of problems, and this has led to confusion about the meaning of his results. Previous papers [5, 9, 14] have pointed out several of these problems.

One of the fundamental concepts used by Chapman is the idea of modal truth in plans. We will discuss the details of this concept later---but a simple version of it is that if P is a partially-ordered, partially-instantiated plan and p is a ground literal, then p is *necessarily* (or *possibly*) true in P 's

*The authors' names are given in alphabetical order. Kambhampati's research is supported in part by an NSF Research Initiation Award IRI-9210997, and ARPA/Rome Laboratory planning initiative under grant F30602-93-C-0039. Nau's research supported in part by NSF Grants IRI-9306580 and NSFD CDR-88003012.

final situation if for every (or some) totally-ordered ground instance P' of P , p is true after executing P' . Chapman's Modal Truth Criterion (MTC) purports to give necessary and sufficient conditions for ensuring that p is necessarily or possibly true. As we describe below, this paper addresses several problems with modal truth and the MTC.

Chapman explicitly states and proves the MTC for necessary truth, and claims that by modal duality (i.e., the equivalence of "necessarily p " and "not possibly $\neg p$ "), the MTC for possible truth is obtained via a simple rewording of the MTC for necessary truth. But in this paper, we show that although modal duality is a fundamental axiom of classical modal logics, it does *not* hold for modal truth in Chapman's plans.¹ This has several consequences:

1. The MTC for possible truth is not completely correct: it provides necessary but insufficient conditions for ensuring possible truth. Furthermore, although necessary truth in plans can be computed in polynomial time as pointed out by Chapman, the same is not true for possible truth. Instead, the problem of computing possible truth in plans is NP-hard.²
2. We can define a concept called *modal conditional truth*, which is similar to modal truth but does not require that a plan be executable as modal truth does. Necessary conditional truth and possible conditional truth *are* duals of each other, and both can be computed in polynomial time. Furthermore, if we rewrite the MTC to talk about modal conditional truth rather than modal truth, then both the MTC for necessary conditional truth *and* the MTC for possible conditional truth are correct.

This paper is organized as follows. Section 2 contains basic definitions, and clarifications/corrections of some of Chapman's terminology. Section 4 presents results about

¹Although Chapman does not explicitly state that his usage is consistent with modal logics, it seems clear to us that this is what he had in mind. In particular, Chapman explicitly appeals to modal duality in his proof of the MTC [2, p. 368].

²If modal duality held, then both necessary truth and possible would be at similar levels of complexity: either both would be polynomial, or one would be NP-hard and the other co-NP-hard. Section 5.2 discusses some formulations of planning in which this occurs.

modal duality, the complexity of modal truth, and the modal truth criterion, and compares and contrasts these results with Chapman’s claims, as well as with other related work. Section 6 contains concluding remarks. Proofs of all the theorms stated in this paper can be found in [8].

2 Basics

The planning language \mathcal{L} is any function-free first-order language. Since \mathcal{L} is function-free, every term is either a variable symbol or a constant symbol, and thus every ground term is a constant symbol. We follow the usual convention of defining an *atom* to be a predicate symbol followed a list of terms, a *literal* to be an atom or its negation, and a *proposition* to be a 0-ary atom. Thus, what Chapman calls a *proposition*, we call a *literal*.

A *state* is any finite collection of ground atoms of \mathcal{L} . If a state s contains a ground atom p , then p is true in s and $\neg p$ is false in s ; otherwise p is false in s and $\neg p$ is true in s .

If T is a finite set of terms, then a *codesignation constraint* on T is a syntactic expression of the form ‘ $t \approx u$ ’ or ‘ $t \not\approx u$ ’, where $t, u \in T$. Let D be a set of codesignation constraints on T , and θ be a *ground substitution* over T (i.e., a substitution that assigns a ground term to each variable in T). Then θ *satisfies* D if $t\theta = u\theta$ for every syntactic expression ‘ $t \approx u$ ’ in D , and $t\theta \neq u\theta$ for every syntactic expression ‘ $t \not\approx u$ ’ in D . D is *consistent* if there is at least one ground substitution θ that satisfies D . If $t\theta = u\theta$ for every θ that satisfies D , then t *codesignates* with u .

A *step* is a triple $a = (\text{name}(a), \text{pre}(a), \text{post}(a))$, where $\text{name}(a)$ is a constant symbol called a ’s *name*, and $\text{pre}(a)$ and $\text{post}(a)$ are collections of literals called a ’s *preconditions* and *postconditions*. A *plan* is a 4-tuple $P = (s_0, A, D, O)$, where s_0 is a state called P ’s *initial state*, A is a set of steps, D is a set of codesignation constraints on the terms of P (i.e., the terms in s_0 and A), and O is a set of ordering constraints on the steps of A . P is *complete* if there is a unique total ordering $a_1 \prec a_2 \prec \dots \prec a_n$ over A that satisfies O , and a unique ground substitution θ over the terms of P that satisfies D . (Note that a complete plan need not necessarily be executable). Suppose that P is complete, and let k be the largest integer $\leq n$ for which there are states s_1, s_2, \dots, s_k such that for $1 \leq i \leq k$, s_{i-1} satisfies a_i ’s preconditions, and s_i is the state produced by performing the step a_i in the state s_{i-1} . Then for $1 \leq i \leq k$, a_i is *executable* in the *input state* s_{i-1} , *producing the output state* s_i . If $k = n$, then P is *executable*, and it *produces the final state* s_n .

A plan $P' = (s'_0, A', D', O')$ is a *constraintment* of a plan $P = (s_0, A, D, O)$ if $s'_0 = s_0$, $A' = A$, $O \subseteq O'$, and $D \subseteq D'$. A *completion* of P is any constraintment of P that is complete.³ P is *consistent* if it has at least one completion; otherwise P is *inconsistent*.

³Chapman’s definition of a completion does not make it entirely clear whether a completion of P should include only the steps in P , or allow other steps to be added. However, other statements in his paper make it clear that he means for a completion to include only the steps in P , so this is how we and most others (e.g., [9, 14]) use the term.

3 Situations, Truth, and Modality

One of the basic concepts in Chapman’s planning framework is the idea of a *situation*. To avoid some problems with Chapman’s definitions, we define plans using STRIPS-style states of the world, and then define situations in terms of states. The intent of our definitions is that if a plan is complete and can be executed at least far enough to reach the situation s , then s corresponds to some state t that arises while executing the plan; and what is true and false in s is precisely what is true and false in t . Otherwise, *nothing* is true or false in s although certain things may be *conditionally* true or false (as defined below). These ideas are formalized below.

If P is a plan, then associated with each step a of P are two symbols $\text{in}(a)$ and $\text{out}(a)$, called a ’s *input* and *output* situations. Associated with P are symbols init and fin called the *initial* and *final* situations of P . All of these symbols must be distinct. Whenever $a \prec b$, we will also say that $x \prec y$, where x may be a or $\text{in}(a)$ or $\text{out}(a)$, and y may be b or $\text{in}(b)$ or $\text{out}(b)$.

We now define what is *true* and *false* in a situation of a complete plan. Let P be a complete plan, and p be a ground literal. Then p is true in init if p is true in P ’s initial state, and p is true in fin if p is true in P ’s final state. If a is an executable step of P , then p is true in $\text{in}(a_i)$ (or $\text{out}(a_i)$) if p is true in a ’s input state (or output state, respectively). A ground literal p is false in a situation s iff $\neg p$ is true in s . Note that if P is not executable, then the law of the excluded middle does not apply, for p will be neither true nor false in P ’s final situation.

As a consequence of the above definitions, it follows that p is true in s (which we write symbolically as $\mathcal{M}(p, s)$) iff the following three conditions are satisfied:

Establishment: either p codesignates with a postcondition of some step $a \prec s$, or else $p \in s_0$.

Nondeletion: for all steps b between a (or s_0) and s , no postcondition of b codesignates with $\neg p$.

Executability: every step that precedes s is executable.

A closely related concept is *conditional truth*, which is like ordinary truth except that it does not require executability: p is *conditionally true* in s (which we write symbolically as $\mathcal{C}(p, s)$) iff the establishment and nondeletion conditions hold.

We defined truth and conditional truth only for complete plans, because for incomplete plans, what is true or conditionally true will vary depending on which completion we choose. In incomplete plans, we instead need to talk about *modal truth*, which Chapman defines as follows [2, p. 336]:

I will say ‘‘*necessarily p*’’ if p is true of all completions of an incomplete plan, and ‘‘*possibly p*’’ if p is true of some completion.

Above, Chapman apparently means p to be nearly any statement about a plan: examples in his paper include not only statements about specific literals and situations in the plan, but also statements about the entire plan (e.g., the statement [2, p. 341] that a plan ‘‘necessarily solves the

problem’’). However, unless we place some restrictions on the nature of p , this has some dubious results---for example, if P is an incomplete plan, then all completions of P are complete, and therefore P itself is necessarily complete. Therefore, for the formal results in the paper, we will use ‘‘necessarily’’ and ‘‘possibly’’ only in the following cases (although we will sometimes use them informally in a broader sense). If p is an atom, P is a plan, and s is a situation in P , then:

- p is *necessarily* (or *possibly*) true in s (written $\Box\mathcal{M}(p, s)$ and $\Diamond\mathcal{M}(p, s)$, respectively) iff $\mathcal{M}(p, s)$ in every (or some) completion of P ;
- p is *necessarily* (or *possibly*) conditionally true in s (written $\Box\mathcal{C}(p, s)$ and $\Diamond\mathcal{C}(p, s)$, respectively) iff $\mathcal{C}(p, s)$ in every (or some) completion of P .

We now define the following decision problems (where P is a plan and p is a ground literal):

NECESSARY TRUTH: given p and P , is p necessarily true in P ’s final situation fin ?

POSSIBLE TRUTH: given p and P , is p possibly true in P ’s final situation fin ?

NECESSARY CONDITIONAL TRUTH: given p and P , is p necessarily conditionally true in P ’s final situation fin ?

POSSIBLE CONDITIONAL TRUTH: given p and P , is p possibly conditionally true in P ’s final situation fin ?

4 Duality, and Complexity of Modal Truth

Given the definitions of modal truth and modal conditional truth above, it is easy to see that a literal p is necessarily true in the final situation fin of a plan P if and only if (1) p is necessarily conditionally true in fin , and (2) for every action a of the plan and every precondition p_a of a , p_a is necessarily conditionally true in the situation $\text{in}(a)$. Thus,⁴

$$\Box\mathcal{M}(p, \text{fin}) \equiv \Box \left[\mathcal{C}(p, \text{fin}) \wedge \bigwedge_{\forall a \in P, \forall p_a \in \text{pre}(a)} \mathcal{C}(p_a, \text{in}(a)) \right], \quad (1)$$

⁴We could consider generalizing Eq. 1 to apply to situations $s \neq \text{fin}$, by replacing the condition ‘‘ $\forall a \in P$ ’’ with the condition ‘‘ $\forall a \in S$,’’ where S is the set of all actions that precede s in at least one completion of P . However, such a generalized version of Eq. 1 would not always hold, as illustrated by the following counterexample (due to Backstrom [1]). Let P be a plan with three actions a, b, c , such that $a \prec b, a \prec c, \text{pre}(a) = \emptyset, \text{post}(a) = \{\neg p\}, \text{pre}(b) = \emptyset, \text{post}(b) = \{p\}, \text{pre}(c) = \{\neg p\},$ and $\text{post}(c) = \emptyset$. Then $\Box\mathcal{M}(p, \text{Out}(b))$, but it is not true that $\Box[\mathcal{C}(p, \text{Out}(b)) \wedge [\bigwedge\{\mathcal{C}(p_d, \text{in}(d)) : d \in S \& p_d \in \text{pre}(d)\}]]$. To see this, note that P has two completions, one executable and one non-executable. c precedes b in one completion (the executable one) and thus $c \in S$. However for c ’s precondition ($\neg p$), $\mathcal{C}(\neg p, \text{in}(c))$ fails in the other (non-executable) completion of P . The main reason for this is that the set of steps that precede $\text{Out}(b)$ is different in different completions -- $\{a, b\}$ in one, and $\{a, b, c\}$ in the other. Thus, the correct way of generalizing Eq. 1 will involve doing the inner conjunction with S ranging over each of these values, and disjoining all the resulting conjunctions.

Now, since modal necessity commutes over conjunctions (i.e., $\Box(p \wedge q) \equiv \Box(p) \wedge \Box(q)$), we can write Eq. 1 as

$$\Box\mathcal{M}(p, \text{fin}) \equiv \left[\Box\mathcal{C}(p, \text{fin}) \wedge \left[\bigwedge_{\forall a \in P, \forall p_a \in \text{pre}(a)} \Box\mathcal{C}(p_a, \text{in}(a)) \right] \right]. \quad (2)$$

Thus computing whether p is necessarily true in fin involves computing whether p is necessarily conditionally true in fin , as well as computing the necessary conditional truth of all preconditions of all steps preceding fin . As noted in Chapman, computing the necessary conditional truth of a literal in a situation (which involves checking whether the MTC’s establishment and declobbering clauses are consistent with the plan’s ordering and codesignation/non-codesignation constraints) can be done in time polynomial ($O(n^3)$) in the plan length. Thus, since the total number of preconditions in a plan is of the order of number of actions in the plan, computing whether p is necessarily true can also be done in polynomial time. Coming to the case of possible truth, we have

$$\Diamond\mathcal{M}(p, \text{fin}) \equiv \Diamond \left[\mathcal{C}(p, \text{fin}) \wedge \left[\bigwedge_{\forall a \in P, \forall p_a \in \text{pre}(a)} \mathcal{C}(p_a, \text{in}(a)) \right] \right]. \quad (3)$$

But possible truth does not commute over conjunctions (i.e., in general, $\Diamond(p \wedge q) \not\equiv \Diamond(p) \wedge \Diamond(q)$), so there is no way to simplify Eq. 3 into component tests of computing possible conditional truth of individual literals. Thus, even though possible conditional truth in fin and necessary conditional truth in fin are duals of each other (i.e., $\Diamond\mathcal{C}(p, \text{fin}) \equiv \neg\Box\neg\mathcal{C}(p, \text{fin})$), possible truth in fin and necessary truth in fin are *not* duals of each other. More specifically:

Theorem 1 *There is a ground literal p , a plan P , with the final situation fin such that $\Box\mathcal{M}(p, \text{fin}) \not\equiv \neg\Diamond\neg\mathcal{M}(p, \text{fin})$.*

Thus, unlike necessary conditional truth and possible conditional truth, necessary truth and possible truth do not obey the modal duality that is obeyed by all classical modal logics [3, p. 62], and thus do not define a well-formed modal logic. It is easy to understand why this is so. The semantics of modal logics are based on Kripke structures (a.k.a. possible worlds). In this formulation, if p is a ground literal, then for every possible world, p must either be true or false in that world. For partially ordered plans, one might expect that each completion of the plan would give rise to a possible world. However, the modal truth of p in a situation of a plan requires that the plan’s actions be executable in order to produce that situation. Thus, if a completion is not executable, then truth of p is not defined in the corresponding possible world.⁵

⁵Although TWEAK plans cannot be modeled using the semantics of classical modal logics, they can be modeled in a variant of modal logics, called *first order dynamic logic* [13]. Dynamic

Given a ground literal p and a plan P , p is possibly true in P 's final situation if and only if there is an executable completion of P that produces a final state in which p is true, and this happens iff it is not the case that every executable completion of P produces a final state in which $\neg p$ is true. Thus, POSSIBLE TRUTH is the dual of the following problem:

PARTIAL TRUTH: given a ground literal p and a plan P , does every executable completion of P produce a final state in which p is true?⁶

Lemma 1 PARTIAL TRUTH is NP-hard.

PARTIAL TRUTH is a weaker condition than both NECESSARY TRUTH and NECESSARY CONDITIONAL TRUTH. There are some cases (one occurs in the proof of Lemma 1) in which every executable completion of P produces a final state in which p is true, but p is neither necessarily true nor necessarily conditionally true in P 's final situation.

Another way of understanding the problem with simplifying Eq. 3 is to note that if p is possibly conditionally true and that all the preconditions of the preceding actions are possibly conditionally true, this only implies that each of them is *individually* true in at least one completion---and this condition is *necessary* but *insufficient* for ensuring possible truth. We could check possible truth by checking to see whether all these conditions are *collectively* true in at least one completion of the plan, but since the number of completions of a plan is exponential in the number of actions of the plan, this would take exponential time. Furthermore, the following theorem shows that unless P=NP, there is no polynomial-time approach for solving this problem.

Theorem 2 POSSIBLE TRUTH is NP-hard.

Thus, NECESSARY TRUTH and POSSIBLE TRUTH have different levels of complexity. If modal duality held, then this would not be so, for each would be reducible to the other's complement via an equivalence of the form $\diamond \mathcal{M}(p, \text{fin}) \equiv \neg \square \neg \mathcal{M}(p, \text{fin})$. Thus it would follow [6, p. 29] that either POSSIBLE TRUTH would be polynomial like NECESSARY TRUTH, or else NECESSARY TRUTH would be co-NP-hard. In Section 5.2, we discuss some planning situations where this occurs.

5 Comparison with Other Work

5.1 The Modal Truth Criterion

Chapman states the MTC as follows [2, p. 340]:

logic, which has been used to provide semantics for programs and plans, provides a clean way to separate executability/termination conditions from goal satisfaction conditions. More about this in Section 5.2.

⁶PARTIAL TRUTH corresponds closely to the notion of partial correctness, which was studied in connection with dynamic-logic-based modeling of computer programming languages [11, 13].

Modal Truth Criterion. A [literal] p is necessarily true in a situation s iff two conditions hold:⁷ there is a situation t equal or necessarily previous to s in which p is necessarily asserted; and for every step C possibly before s and every [literal] q possibly codesignating with p which C denies, there is a step W necessarily between C and s which asserts r , a [literal] such that r and p codesignate whenever p and q codesignate. The criterion for possible truth is exactly analogous, with all the modalities switched (read ‘‘necessary’’ for ‘‘possible’’ and vice versa).

If we take these words literally, then the definition of modal truth tells us that the plan must be modally executable. This is consistent with Chapman's definition of a situation [2], from which it follows that a step's output situation (and hence what is true in that situation) is only defined if the step can be executed. However, a careful look at Chapman's proof of necessity and sufficiency of his MTC reveals that his proof deals with necessary *conditional* truth rather than necessary truth.⁸ In proving that any literal with an establisher and no clobberer must be necessarily true, Chapman's proof refers to white-knight steps for every potential clobberer, [2, p. 370], without checking that the white knights are in fact executable.⁹

For the ‘‘necessary truth’’ version of the MTC, this does not affect the validity of Chapman's proof, since executability occurs naturally as a consequence of applying necessary conditional truth recursively to prerequisites of all preceding steps. The same, however, cannot be guaranteed for possible truth, since modal possibility does not commute over conjunctions---and thus Chapman's proof cannot be extended to possible truth. In particular, the following theorem shows that the ‘‘possible truth’’ version of the MTC sometimes fails:

Theorem 3 There is a plan P and a ground literal p such that in P 's final situation, p is not possibly true but the MTC concludes otherwise.

The above discussion suggests an alternative interpretation of the MTC that sidesteps the problem: drop the executability requirement, and interpret the MTC as a statement about modal *conditional* truth rather than modal truth. This alternative interpretation is not as far-fetched as it might sound. To see this, note that Chapman defines the notion of truth of a literal in a situation as follows [2, p. 338]:

A [literal] is true in a situation if it codesignates with a [literal] that is a member of the situation. A step asserts a [literal] in its output situation if the [literal] codesignates with a postcondition of the step.

⁷The second of these conditions is the ‘‘white-knight decllobbering clause’’ that we refer to elsewhere.

⁸Had Chapman explicitly noted this use of modal conditional truth in his proof, we believe he would have noticed the non-duality of necessary and possible truths.

⁹Note that in Chapman's terminology, the establisher is a *situation*, while clobberers and white knights are *steps*.

Here, there is no explicit requirement that the step be executable. This suggests that the MTC does not require that P be modally executable, and thus suggests that Chapman was talking about modal conditional truth. This interpretation is also consistent with his “nondeterministic achievement procedure” [2, Fig. 7], where to make a literal necessarily true in a situation, he only ensures establishment and declobbering without explicitly stating that the establisher needs to be executable. (As explained above, for the case of necessary truth, executability follows from making every prerequisite of every action necessarily conditionally true.)

The “conditional truth” interpretation of MTC gives a quasi-local flavor to planning, by separating the process of ensuring local establishment and declobbering from the process of ensuring executability, with the understanding that if all preconditions are necessarily established and declobbered, then the whole plan itself will be executable and correct. In fact, some latter rewrites of the MTC (e.g. [14, 9]) use this interpretation to eliminate the notion of situations entirely, and state MTC solely in terms of steps (operators) and their preconditions and postconditions.

Although a truth criterion for modal conditional truth does have utility in plan generation, it is of limited utility in projecting plans or partially ordered events. As mentioned in Section 3, the latter are more naturally related to modal truth.

5.2 Modal Duality and Universal Executability

In Section 4, we observed that the main reason why necessary truth and possible truth are not duals in TWEAK-style plans is that such plans can contain unexecutable completions. Thus, one way to achieve duality between necessary truth and possible truth is to restrict our attention to plans whose completions are always executable. One way to guarantee that plans will always be executable is to restrict the actions to have no preconditions, i.e., to consider only those plans P such that $\text{pre}(a) = \emptyset$ for every step a of P .

This approach is clearly too restrictive, since it precludes modeling actions with any form of preconditions. But if we relax the restrictions of TWEAK-style action representation, there is a more reasonable way to guarantee universal executability: let an action a be executable even if its preconditions are not satisfied. If the preconditions are satisfied, then a will produce its postconditions; otherwise, a will simply have no effects.¹⁰ For plans that contain only this type of actions, possible truth and necessary truth are duals of one another, computation of possible truth is NP-hard, and computation of necessary truth is co-NP-hard. As discussed below, this approach has been used in different forms by several different researchers.

¹⁰While seemingly unintuitive, this relaxation is in fact very much consistent with the original formalization of actions in situational calculus [3]. In this formalism, actions are modeled as situation-transformers, with the transformation given by the `Result` function, which takes an action and a situation as the arguments. Having universally executable steps corresponds to having the `Result` be a total rather than a partial function.

To our knowledge, the above approach was first used in Rosenchein’s work [13] on providing semantics to plans based on first-order propositional dynamic logic. Rosenchein restricts the use of conditionals in PDL to guarantee that the plan terminates irrespective of which branch of the conditional it takes.

A very similar idea is used in Dean and Boddy’s work on temporal projection [4]. In Dean and Boddy’s formulation, a partially ordered set of events A is projectible even when a rule’s preconditions don’t hold (in which case the rule simply has no effect). Hence in their formalism, determining possible truth and necessary truth are duals, and both are NP-hard.

Chapman [2, p. 371] uses universally executable actions (he calls them conditional steps) in proving his intractability theorem for actions containing conditional effects. A plan composed entirely of such steps will always be executable, leading to the same results as in Dean and Boddy’s formalism.

Since Chapman’s intractability theorem is based on planning operators that have conditional effects, it has been natural for planning researchers to interpret it to mean that the conditionality of these operators is what causes necessary truth to be intractable. However, this interpretation is misleading. The intractability result depends just as much on the universal executability of Chapman’s conditional steps as it does on their conditionality. Here’s why:

Consider an incomplete plan P composed of ordinary “unconditional” steps as defined in Section 2, and let a be a step of P such that $\text{pre}(a)$ $\text{post}(a)$ contain an unbound variable x . Then for the purposes of both planning and temporal projection, a has conditional effects: its effects will be different in different completions of P , depending on what we bind x to. However, computing necessary truth in such plans is still polynomial. Since Chapman’s planning language has an infinite number of constant symbols, it follows that in the plan P we can *always* find a binding for x that makes a unexecutable. As a consequence, P will always have at least one unexecutable completion. Hence, determining necessary truth is trivial: nothing will be necessarily true in P ’s final situation.

Now, suppose we restrict our planning language \mathcal{L} to contain only finitely many constant symbols (and thus only finitely many ground terms, since \mathcal{L} is function-free). Then there will be some plans in which a is executable for every binding of x . In this case, as the following theorem shows, checking necessary truth will be co-NP-hard, even with unconditional steps.

Theorem 4 *If the language \mathcal{L} contains only finitely many constant symbols, then NECESSARY TRUTH is co-NP-hard.*

Notice that this result is related to Chapman’s observation [2, p. 356] that restricting the range of a variable to a finite set will defeat the MTC, and make constraint computations NP-complete.

Finally, a recent investigation by Nebel and Backstrom [10] on the computational complexity of plan-validation and temporal projection has yielded results related to those

presented in this paper. While our investigation is initially motivated by the apparent lack of modal duality in Chapman's MTC, Nebel and Backstrom's work is motivated by the apparent asymmetry in the complexity of plan validation through modal truth criterion, and temporal projection (c.f. [4]). Rather than interpret MTC in terms of modal conditional truth, and use that to explain the asymmetry in the possible and necessary truth, as we have done in this paper, Nebel and Backstrom choose to restrict applicability of MTC only for plans whose completions are all executable (they term this property *coherence*). Another difference with their research is that they concentrate on ground (variable-less) plans, while we look at the more general variablized plans. We believe that the results in this paper complement theirs and together provide a coherent interpretation of the role of modal truth criteria in planning.

6 Concluding Remarks

In this paper, we have presented the following results about modal truth and the modal truth criterion:

1. Contrary to Chapman's statement, the principle of modal duality that is obeyed by all classical modal logics is not obeyed in TWEAK-style plans. The lack of duality between necessary truth and possible truth is related to (a) the fact that modal truth of a literal in a situation of a plan requires that the plan's actions be executable in order to produce that situation, and (b) the asymmetry in the way necessary conditional truth and possible conditional truth commute over conjunctions: $\Box(p \wedge q) \equiv \Box(p) \wedge \Box(q)$ while $\Diamond(p \wedge q) \not\equiv \Diamond(p) \wedge \Diamond(q)$. To achieve modal duality, one needs universally executable plans.
2. Even though necessary truth in plans can be determined in polynomial time as stated by Chapman, the same statement does not hold for possible truth. Instead, the problem of determining possible truth in plans is NP-hard.¹¹
3. As stated by Chapman, the MTC is correct only as a criterion for necessary truth (not as a criterion for possible truth). However, if we reinterpret it as a criterion for modal *conditional* truth (i.e., modal truth conditional on plan executability), then it is correct as a criterion for both necessary conditional truth and possible conditional truth.

Because of the wide impact of Chapman's paper, it is important to correct any misimpressions that may result from it. We hope readers will find this paper useful for that purpose. Finally, while we concentrated on clarifying the nature of modal truth criterion, there have also been several misimpressions regarding its *role* in plan generation. In the extended version of this paper [8], we also address these confusions.

¹¹Checking possible truth has several applications in plan projection [4] as well as plan generalization [9].

Acknowledgement

We appreciate the helpful comments of Christer Backstrom and Bernhard Nebel.

References

- [1] C. Backstrom. Personal communication, 1993.
- [2] D. Chapman. Planning for conjunctive goals. *Artificial Intelligence*, 32:333--379, 1987.
- [3] E. Davis. *Representations of Commonsense Knowledge* Morgan Kaufmann Publishers, Inc. San Mateo, California, USA, 94403.
- [4] T. Dean and M. Boddy. Reasoning about partially ordered events. *Artificial Intelligence*, 36:375-399, 1988.
- [5] K. Erol, D. Nau, and V. S. Subrahmanian. When is planning decidable? In *Proc. First Intl. Conference on AI Planning Systems*, pp. 222--227, June 1992.
- [6] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness* W.H. Freeman and Company, New York, 1979.
- [7] S. Hanks and D. S. Weld. Systematic adaptation for case-based planning. In *Proc. First Internat. Conf. AI Planning Systems*, pp. 96--105, June 1992.
- [8] S. Kambhampati and D.S. Nau. On the Nature and Role of Modal Truth Criteria in Planning *Artificial Intelligence*, 1994 To appear. (Available as Tech. Report. ISR-TR-93-30, Inst. for Systems Research, University of Maryland, March, 1993. Anonymous ftp: enws318.eas.asu.edu:pub/rao/role-mtc-planning.ps)
- [9] S. Kambhampati and S. Kedar. A unified framework for explanation-based generalization of partially ordered and partially instantiated plans. *Artificial Intelligence*, Vol. 67, No. 2, June 1994.
- [10] B. Nebel and C. Backstrom. On the computational complexity of temporal projection, planning and plan validation. *Artificial Intelligence*, Vol. 56, No. 1, 1994.
- [11] V. Pratt. Semantical considerations on Floyd-Hoare Logic. In *Proc. 17th FOCS*, 109-121.
- [12] M. A. Peot. Conditional nonlinear planning. In *Proc. First International Conference on AI Planning Systems*, pp. 189--197, 1992.
- [13] S. Rosenchein. Plan Synthesis: A logical perspective. In *Proc. IJCAI-81*, pp. 331-337, 1981.
- [14] Q. Yang and J. D. Tenenber. Abtweak: Abstracting a nonlinear, least commitment planner. In *AAAI-90*, pp. 204--209, 1990.