Simple Causal Minimizations for Temporal Persistence and Projection

Brian A. Haugh
Martin Marietta Laboratories
1450 South Rolling Road
Baltimore, Maryland 21227

Abstract

Formalizing temporal persistence and solving the temporal projection problem within traditional non-monotonic logics is shown possible through two different approaches, neither of which requires any special minimization techniques. Minimizing *potential causes* is shown to yield a type of temporal persistence that is useful for the temporal projection problem, although it differs significantly from the ordinary conception of temporal persistence. A conception of *determined causes* is then developed whose minimization does yield the results preferred by ordinary temporal persistence. Finally, previous approaches to formalizing temporal persistence using chronological minimizations are shown inadequate for certain classes of scenarios, which causal minimizations formalize correctly.

I. Introduction

A. Temporal Persistence

Temporal persistence of facts (i.e., their presumed continuation through time in the absence of contrary information) was introduced to the AI community by Drew McDermott as an important part of formalizing planning (McDermott, 1982). It contributes to a solution of the "frame problem" (i.e., the determination of what facts will continue to hold after the occurrence of some sequence of events) in automatic planning systems (McCarthy and Hayes, 1969). Applying McDermott's original conception of persistence, referred to here as "the ordinary conception of temporal persistence," enables the deduction that all previously holding facts continue to hold (for their persistence period) unless explicit causal rules (or other provable facts) entail their cessation. Other, non-standard, conceptions of the conditions under which facts are presumed to persist are possible, and will be shown to be preferable under certain circumstances.

B. Temporal Projection

Temporal persistence may also contribute to solutions of a narrower problem, the "temporal projection problem," which has been recently described as:

> "... given an initial description of the world (some facts that are true), the occurrence of some events, and some notion of causality (that an event occurring can cause a fact to become true), what facts are true once all the events have occurred?" (Hanks and McDermott, 1986, p. 330)

While this definition could be more explicit, it appears to define the temporal projection problem as narrower than the frame problem since (in the context of its use) it seems intended to admit to an initial description only facts that hold in some initial state, excluding facts that hold prior or subsequent to that state. Ordinary temporal persistence has been proposed as a basis for the solution to both these problems. Thus, we examine various alternatives for formalizing temporal persistence to determine whether they accurately model its ordinary conception or provide a reasonable basis for solution to either of these two problems.

C. Difficulties In Logical Formulation

Early AI papers on temporal persistence[1] expressed optimism that its non-monotonic features would eventually be adequately modeled by some version of a non-monotonic logic. It was recently realized (Hanks & McDermott, 1985) that formalizing this concept was more difficult than it appeared, due to unacceptable models associated with certain obvious formulations. Furthermore, the intended interpretation of temporal persistence seemed relatively obvious and procedural implementations had already been developed. Thus, it was argued that ordinary non-monotonic logics (e.g., the NML of McDermott and Doyle, 1980; the circumscription of McCarthy, 1980, 1986; or the default logic of Reiter, 1980) were inherently incapable of formalizing an important type of ordinary inference, casting doubt on the suitability of logic as a foundation for continuing work in artificial intelligence (Hanks and McDermott, 1985, 1986). In response to these arguments we describe how useful notions of temporal persistence can be formalized by minimizing certain causal relations using ordinary non-monotonic logics.

II. Causal Minimizations

A. Informal Scenario Description

The simple shooting scenario presented in Hanks and McDermott (1986) is used to illustrate how different proposed logical formulations achieve the temporal persistence properties important to solving the problems of temporal projection. In this <u>idealized</u> scenario, a gun is loaded and, after a brief wait, is fired at someone. Furthermore, there is a causal rule asserting that if the gun is shot while loaded, someone will die. Temporal persistence is required to ensure that the gun remains loaded and the intended result obtains.

B. Situation-Calculus Formalization

A situation-calculus type formalism (McCarthy, 1968) is used here, although it differs in two respects from typical logics of this type. First, causal relations are represented using causal predicates (after McDermott, 1982, and Allen, 1984). Second, we use result relations (e.g., Result(Event1, S1, S2)) rather than result functions (e.g., S2 = Result(Event1, S1)), since relations are easier to restrict than functions when minimizations are being performed. Causal relations also seem required for

[1] For example, McDermott (1982), p.122.

supporting the intended causal minimizations, since truth functional representations of causality will hold in many circumstances in which they do not express causal relations. The causal predicate used is of the form *Causes(precondition1, cause1, effect1)*, and is informally interpreted as: when *precondition1* holds in a state, then any event of type *cause1* occurring in that state causes the effect *effect1* to hold in the situation that results.

Scenario Axioms

The particular known facts are represented in this formulation by the axioms:

A1) T(Alive, S0)

A2) Result(Load, S0, S1)

A3) Result(Wait, S1, S2)

A4) Result(Shoot, S2, S3)

The general causal relations are expressed by:

A5) Causes(True, Load, Loaded)

A6) Causes(Loaded, Shoot, not(Alive))

where

A7) $\neg$(Load = Wait v Load = Shoot v Wait = Shoot v Loaded = True v Loaded = not(Alive))

General Temporal Causal Axioms

A set of axioms is required to define the basic temporal and causal relations. The condition that a fact has just ceased being true (clipped) is defined:

T1) Clipped(f, s) <=> $(\exists c,s')$ [Result(c, s', s) & T(f, s') & $\neg$T(f,s)]

A prior fact become false (changed), is defined:

T2) Changed(f, s) <=> $(\exists s')$ [Before(s',s) & T(f, s') & $\neg$T(f, s)]

where

T3) Before(s', s) <=> $(\exists c,s')$ [Result(c, s', s) v $(\exists s'')$ (Before(s', s'') & Result(c, s'', s)]

T4) Before(s', s) => $\neg$Before(s, s')

One special feature of this formulation is that it requires clippings of facts to have immediate explanations in terms of applicable causal laws[2], expressed by:

T5) Clipped(f, s) => $(\exists p,c,s')$ [Causes(p, c, not(f)) & T(p, s') & Result(c, s', s)]

where

T6) T(not(f), s) => $\neg$T(f, s)

T7) f = not(not(f))

T8) T(True, s)

The consequence that the effect holds whenever the cause occurs and the preconditions of a causal relation hold is expressed by:

T9) Causes(preconditions, cause, effect) => $(\forall s,s')$ [(T(preconditions, s) & s' = Result(cause, s)) => T(effect, s')]

C. Minimizing Potential Causes

1. Distinguishing Preferred Models Causally

Examination of the alternative models of the shooting scenario reveals several significant distinctions between them apart from the fact that clipping is minimized in the preferred models. Most notably, there is no known cause for the gun becoming unloaded, nor is any causal law identified that could explain how "Loaded" might come to be clipped. In contrast, the clipping of "Alive" is covered by a causal axiom and the required cause is known to occur. Preferred models, then, contain fewer potential causal explanations in them, i.e., some sort of minimization of causal relations or potential causes is indicated to ensure the preferred conclusions in all minimal models.

[2]As suggested by (Hayes, 1971).

2. Applied to the Shooting Scenario

The condition of "being covered by a causal law whose cause occurs" is one promising candidate for minimization. The existence of a "potential cause" may be defined to capture this condition, as follows:

M1) Potential_cause(p,c,e,s) <=> [Causes(p,c,e) & $(\exists s')$ Result(c, s', s)]

Informally, a "Potential_cause(p,c,e,s)" statement may be understood as asserting that state s results from a potential cause c occurring in the previous state, where c's occurrence would cause effect e to hold in situation s if precondition p is met. Under this definition, the only potential causes in the preferred model of the shooting scenario are the shooting, resulting in the death in S3, and the loading, resulting in the loaded gun in S1 (i.e., Potential_cause(True, Load, Loaded, S1) and Potential_cause(Loaded, Shoot, not(Alive), S3)). These are potential causes in all models, even those in which the gun is unloaded and the shooting remains merely a potential cause, so that "Alive" is not actually clipped. Undesired models will not be minimal in potential causes since they will require an additional potential cause (Potential_cause(P1, C1, Not(Loaded), S2)) to explain the clipping of "Loaded" (since our axioms require every clipping to have a cause). Thus, this conception of minimization is adequate to distinguish the preferred model in this shooting scenario, yielding the desired conclusion that T(Dead, S3) holds in all minimal models.[3] However, this solution does not accurately model the ordinary notion of temporal persistence in many other situations.

3. Inequivalence to Temporal Persistence

The results expected of ordinary temporal persistence cannot be obtained by minimizing potential causes when minor modifications are made to the shooting scenario to include a potential causal explanation for the gun being unloaded. For example, suppose someone tried to unload the gun, and would have, if he had known how, while waiting; that is, we add the following axiom:

A8) Cause(Knows, Wait, not(Loaded))

Here we assume that there is no information about the truth of the precondition ("Knows") in causal relation A8, so that we cannot conclude that the unloading was successfully performed.

This modified set of axioms admits models in which "T(Knows, S1)" is false, and in which "Loaded" persists through S2 and S3. Such models are the preferred ones by the ordinary notion of temporal persistence, which requires that provable facts (e.g., "Loaded") persist when possible (barring conflicts with other persistences). Even

[3]After completing this analysis, we learned of a similar approach to causal minimization developed by Vladimir Lifschitz (Lifschitz, 1987) that appears to achieve the same results in the examined cases. Lifschitz's approach also includes a computationally efficient method for performing the required minimizations as well as other types of default causal reasoning, although it does not appear as readily extensible to more expressive temporal formalisms.

though there is some conflict of persistences in this case, it is clear that the ordinary notion of persistence requires the persistence of "Loaded" in this case just as it did in the original. If no attempt were made to shoot the gun, then there would be no conflict of persistences, and "Loaded" unquestionably ought to persist when the ordinary conception is in force. Whether or not someone attempts to fire the gun subsequently should have no bearing on whether "Loaded" persisted previously, an intuition well illustrated by the original scenario. Furthermore, this persistence result is obtained by application of chronological minimization of clipping (Kautz, 1986, Lifschitz, 1986), the most successful previous approach to formalizing temporal persistence.

Minimizing potential causes, however, does not lead to the results required by temporal persistence in these cases. Whether or not the gun is actually unloaded, Potential_cause(Knows, Wait, not(Loaded), S2) will be true; hence, minimizing potential causes cannot distinguish the preferred model, and will not accurately model the ordinary notion of temporal persistence. However, we must ask whether the ordinary notion of persistence is the preferred inference procedure in such circumstances.

4. Minimizing Potential Causes Preferred

Suppose that a potential assassin knew that an attempt would be made to unload the gun, but had no idea whether the attempt would be successful. Surely it would be foolhardy to proceed with the expectation of a loaded gun, simply because there was no definite proof of its unloading. Knowledge of the occurrence of a potential cause for the clipping of a fact provides reasonable grounds for doubting its persistence. It seems more reasonable in such circumstances to consider the persistence of potentially clipped facts as uncertain, and (in planning contexts) to make plans for both alternatives that ensure the desired effects (e.g., checking the gun after the unloading attempt). This strategy corresponds to minimization of potential causes, and entails that after the attempted unloading, a loaded condition can no longer be deduced, only the disjunction "T(Loaded, S2) v T(not(Loaded), S2)" can be derived.

Thus, minimization of potential causes actually provides a better model of commonsense temporal reasoning in the cases examined than does ordinary temporal persistence or the chronological minimization of clipping. Whether this advantage persists in arbitrary temporal projection scenarios requires further investigation. In any case, this inquiry has created a new perspective on how persistence may best be modeled by planning systems. Since ordinary temporal persistence might still be useful, we have continued to pursue its formalization.

D. Minimizing Determined Causes

1. Initial Conception of Determined Causes

To use causal minimizations for ordinary temporal persistence requires a causal concept that is more discriminating than that of potential cause. We observe that in the revised scenario, the preferred clipping (of "Alive") differs in that its precondition ("Loaded") must either hold or have changed previously. In contrast, the precondition ("Knows") for clipping "Loaded" need not have changed if it is false, since it may have always been false. Thus, we may distinguish the preferred model by minimizing causes whose preconditions could not have always been false. In all models, such "determined causes" either have true preconditions or their preconditions have previously changed to false, i.e.:

M2) Determined_cause(p,c,e,s) <=> [Causes(p,c,e) & (∃s') [Result(c, s', s) & (T(p, s') v Changes(p,s'))]]

Such causes are "determined" in the sense that the axioms of the system taken together with ordinary temporal persistence will determine the truth value of their preconditions, unlike merely potential causes, whose preconditions may be of indeterminate truth value. Minimizing determined causes will then favor as minimal the models of the latest scenario preferred by ordinary persistence, since the attempt at unloading will not be a determined cause in all models. More formally, the wffs *Determined_cause(True, Load, Loaded, S0)* and *Determined_cause(Loaded, Shoot, not(Alive), S3)* are true in all models of the scenario, while *Determined_cause(P1, Wait, not(Loaded), S1)* is false in some models in which the attempt at unloading is unsuccessful. Thus, all models minimal in determined causes are ones in which the gun remains loaded, as required by temporal persistence.

The method by which minimizing determined causes works in other cases is interesting to observe. Whenever something qualifies as a determined cause by virtue of a change in the truth value of its precondition, models in which the precondition retains its last provable value are preferred, since any further change would require an additional cause. Thus, determined causes in such cases will not be effective in the minimal models unless their preconditions are presumed to be true assuming ordinary persistence, as intended.

2. Determined Causes in Causal Chains

This conception of determined cause needs modification to account for causal chains, in which the precondition of one causal relation is a result of another causal relation. For example, a simple causal chain can be described:

A1') T(A, S0) & ¬T(B,S0) & ¬T(C, S0)
A2') Result(Wait, S0, S1)

A3') Result(E1, S1, S2)
A4') Result(E2, S2, S3)
A5') Causes(A, E1, B)
A6') Causes(B, E2, C)

In such situations, ordinary temporal persistence prescribes that A persists through the "Wait," enabling E1 to cause B, which in turn enables E2 to cause C. However, event E2 will not be always be a determined cause of C under our latest definition, since there are models in which A is clipped and C remains false without changing. This results in indifference between models in which A is clipped and A persists, since when A is clipped, there is an extra determined cause for that clipping, and when A persists, there is an extra determined cause for the change in C. Thus, the definition of determined cause must be modified to also apply when the precondition of a cause is the result of a determined cause or its persistence, making it recursive, as follows:

M3) Determined_cause(p,c,e,s) <=> [Causes(p,c,e) & (∃s')[Result(c, s',s) & (T(p, s') v Changes(p,s') v (∃s'',p1,c1)(Before(s'',s') & Determined_cause(p1,c1,p,s'')))]]

This new definition will now handle our chaining example properly: since E2 will be a determined cause of C in all models, any models with A clipped will not be minimal in determined causes; hence, A is not clipped. The recursive nature of the new definition also ensures that the consequences of any length chain of determined causes will also be determined causes.

All such determined causes will be potential causes. But, when there is no information bearing on the truth of the preconditions of a potential cause, it will not be a determined cause, allowing the preference of models favoring the persistence of known facts over the truth of unknown preconditions. Thus, we have isolated a conception of determined cause whose ordinary minimization yields the results entailed by the ordinary conception of temporal persistence in all the varied situations considered so far. However, because of the many ways in which different persistences may interact and conflict, the adequacy of this conception must be considered provisional upon further investigations.

III. Circumscription Proofs

To illustrate our assertion that the causal minimizations discussed could be achieved in any ordinary non-monotonic logic, we here sketch an approach to performing these minimizations using McCarthy's circumscription techniques (McCarthy 1986). Circumscription works by supplementing a theory with a set of circumscription axioms that entail the minimization of an identified predicate's extension. Variable circumscription (Perlis and Minker, 1986) is a simplified form of formula circumscription (McCarthy, 1986) in which the circumscription axioms are specified by a schema of the following form:

$$[\mathbf{A}[Z_0,\ldots,Z_n] \;\&\; (x)(Z_0x \Rightarrow P_0x)] \Rightarrow (y)(P_0y \Rightarrow Z_0y)$$

where **A** is the original theory specified as a conjunction of all its axioms; $\mathbf{A}[P_0,\ldots,P_n]$ is the same conjunction of axioms identified as a function of certain predicates that appear in them; P_0 is the original predicate to be minimized; $P_1,\ldots,P_n$ are other predicates in the theory **A** that are allowed to vary along with the minimization predicate P_0; and $\mathbf{A}[Z_0,\ldots,Z_n]$ is the result of substituting the formulas $Z_0,\ldots,Z_n$ for the original predicates $P_0,\ldots,P_n$ in the theory **A**.

Proving the intended results of minimizing our causal relations (e.g., *Potential_cause*) using circumscription requires choosing other, intimately connected, relations ($P_1,\ldots,P_n$) to vary, so that the original theory under appropriate substitutions, $\mathbf{A}[Z_0,\ldots,Z_n]$, can be proven. We review how such a proof proceeds for circumscribing *Potential_cause* in the original shooting scenario. When the circumscription axiom schema for this case is instantiated, choosing *T, Result, Causes, Clipped, Changed,* and *Before* as the $P_1,\ldots,P_6$ to vary, we get:

$$\{[\mathbf{A}[Z_0,\ldots,Z_6] \;\& \;(\forall p,c,e,s)(Z_0(p,c,e,s) \Rightarrow \text{Potential\_cause}(p,c,e,s)] \Rightarrow (\forall p',c',e',s')(\text{Potential\_cause}(p',c',e',s') \Rightarrow Z_0(p',c',e',s'))\}$$

The theory **A** in our original scenario is the conjunction of our axioms A1-A7 and T1-T9. The substitution, Z_0, for *Potential_cause(p,c,e,s)*, should be a formula that uniquely identifies the minimal set of potential causes in this situation, as follows:

$$[\,(p = \text{True} \;\&\; c = \text{Load} \;\&\; e = \text{Loaded} \;\&\; s = \text{S1}) \;\text{v}\; (p = \text{Loaded} \;\&\; c = \text{Shoot} \;\&\; e = \text{not(Alive)} \;\&\; s = \text{S3})]$$

Any quadruple $<p,c,e,s>$ satisfying this Z_0 obviously satisfies the *Potential_cause* predicate in this scenario. The substitutions for the other varying predicates should also specify their minimal extension in the preferred models. For T(f,s) we may substitute the Z_1:

$$[(f=\text{Alive} \;\&\; (s=\text{S0 v } s=\text{S1 v } s=\text{S2})) \text{ v } (f=\text{Loaded} \;\&\; (s = \text{S0 v } s=\text{S1 v } s=\text{S2 v } s=\text{S3})) \text{ v } (f=\text{not(Alive)} \;\&\; s=\text{S3}) \text{ v } (f=\text{True} \;\&\; (s=\text{S0 v } s=\text{S1 v } s=\text{S2 v } s=\text{S3}))\,]$$

For Result(c,s,s') we many substitute the Z_2:

$$[(c = \text{Load} \;\&\; s = \text{S0} \;\&\; s' = \text{S1}) \text{ v } (c = \text{Wait} \;\&\; s = \text{S1} \;\&\; s' = \text{S2}) \text{ v } (c = \text{Shoot} \;\&\; s = \text{S2} \;\&\; s' = \text{S3})\,]$$

For Causes(p,c,e) we may substitute the Z_3:

$$[(p = \text{True} \;\&\; c = \text{Load} \;\&\; e = \text{Loaded}) \text{ v } (p = \text{Loaded} \;\&\; c = \text{Shoot} \;\&\; e = \text{not(Alive)})\,]$$

For Clipped(f,s) we may substitute the Z_4:

$$[\,f = \text{Alive} \;\&\; s = \text{S3}\,]$$

For Changed(f,s) we may substitute the Z_5:

$$[\,f = \text{Alive} \;\&\; s = \text{S3}\,]$$

For Before(s, s') we may substitute the Z_6:

$$[(s = \text{S0} \;\&\; (s' = \text{S1 v } s' = \text{S2 v } s' = \text{S3})) \text{ v } (s = \text{S1} \;\&\; (s' = \text{S2 v } s' = \text{S3})) \text{ v } (s = \text{S2} \;\&\; s' = \text{S3})]$$

These substitutions define an interpretation of their predicates which can be proven to satisfy the original axioms, although there is not space here to complete the proof. When this is done, it is possible to prove the antecedent of the circumscription axiom, leading to the result that the potential causes specified by Z_0 are the only ones that exist, which then allows proving that the loaded state of the gun persists, and the shooting is effective as intended.

Although our knowledge of the preferred result assisted the derivation in this example, the circumscriptive axiom schema provides the basis for such proofs whether or not one knows how to choose the most useful substitutions for the varying predicates. Still, our example illustrates the validity of criticisms of circumscription (e.g., Hanks and McDermott, 1986) which emphasize the absence of any efficient general procedure for applying it. Whether an efficient method can be developed for applying circumscriptions to such causal minimizations remains a topic for further research. However, we have, nevertheless, demonstrated that such ordinary non-monotonic techniques can achieve the results desired for the temporal projection problem. Finally, the prospects for efficient non-monotonic techniques for causal minimizations have been brightened by recent independent development (Lifschitz, 1987) of such techniques for a broad class of cases.

IV. Chronological Minimization

A. Background

A chronological minimization of a time-indexed predicate is roughly defined as one that prefers admitting an instance that occurs later in time over one that occurs earlier in time. Chronological minimization (either of clipping or of knowledge) has been widely advocated for formalizing temporal persistence (Kautz, 1986; Lifschitz, 1986; Shoham, 1986). However, while it has been found to yield the desired results in temporal projection problems similar to our original shooting example, we have discovered that when there is incomplete knowledge of the initial state, it may not lead to the conclusions preferred by commonsense. Now it will be argued that, in certain cases, such chronological minimization cannot fully formalize the ordinary conception of temporal persistence either, so that its domain of application is further restricted. Although this argument addresses chronological minimization of clipping -- because it has been more clearly formulated than such a minimization of knowledge -- there is good reason to believe that chronological minimization of knowledge will also suffer from similar limitations.

Chronological minimization is rather suspect as a general basis for formalizing persistence. Why, in general, if there is a choice between two different changes occurring, should temporal persistence prefer the later change to the earlier one? Of course, given a choice between an earlier clipping of a fact with no potential explanation (e.g., of "Loaded") and a later clipping of one with a potential explanation (e.g., of "Alive"), then the later clipping is preferred. But, the basis for decision here is not simply a preference for later clippings over earlier ones, but a preference for clippings that have "known" potential explanations over those that do not. If causal minimizations were the actual basis for such decisions, then we would not expect there to be a basis for choice when competing clippings were not distinguished causally, regardless of whether they were distinguished chronologically. And, indeed, this is what we find.

B. Inadequacy for Temporal Persistence

In situations where neither of two alternatively required clippings has a potential explanation, neither commonsense nor the ordinary notion of temporal persistence provides a basis for choosing between them. A general situation of this type is characterized by the following axioms:

A1'') T(A, S1) & T(B, S1)
A2'') Result(E1, S1, S2)
A3'') Result(E2, S2, S3)
A4'') T(not(A), S2) v T(not(B), S3)

In such a situation, the principle of presuming facts to persist when possible provides no basis for choosing between models in which A or B is clipped, and any faithful logic of persistence should leave both of these possibilities open.

Chronologically minimizing clipping in models of axioms A1'' - A4'' will favor models in which A persists, since the clipping of B occurs later, and thus, does not capture the ordinary notion of persistence in such cases. Minimizing potential or determined causes, however, will remain indifferent between these choices in accord with the ordinary notion of persistence.

Chronological minimization of clipping is, thus, inadequate for formalizing the ordinary sense of temporal persistence in general. However, it might still be considered adequate for applications to the frame problem or to the more narrowly understood temporal projection problem, since our formulation of the above counterexample is not characteristic of such problem descriptions.

C. Inadequacy for Frame Problem

Our counterexample schema above is not an example of a temporal projection problem because it specifies indefinite factual information about times other than the initial state. However, situations fitting our schema may easily arise in ordinary planning contexts, which include no specific information about the future. For example, our axioms (A1'' - A4'') might be specified in a context in which the current state is S4, which was the result of some event in S3. It might be known in S4 that either A or B was clipped earlier (as indicated), due to some other fact known to hold in S4 that is incompatible with both A and B persisting. Since such incomplete knowledge of the present and the past is characteristic of most real-world planning domains, chronological minimization will not provide the results preferred by ordinary temporal persistence (or commonsense temporal projection, for that matter) in many realistic planning situations; hence, it does not provide the basis for a general solution to the frame problem.

D. Inadequacy for Temporal Projection

Forced choice clipping situations like those just described can arise even within scenarios fitting the narrowly conceived problem of temporal projection. Even if the initial state is fully specified, certain sets of general causal relations may lead to a forced choice of clipping one of two facts, neither of which has any potential causal explanation. Consider situations in which two facts, A and B, are initially known to be true, but the continued persistence of both in circumstances where certain events occur would lead to contradictory causal results:

A1''') T(A, S0) & T(B, S0) & ¬T(C, S0)
 & ¬T(D1, S0) & ¬T(D2, S0)
A2''') Result(Wait, S0, S1)
A3''') Result(E1, S1, S2)
A4''') Result(E2, S2, S3)
A5''') Causes(A, E1, C)
A6''') Causes(B, E2, D1)
A7''') Causes(C, E2, D2)
A8''') T(D1, s) => ¬T(D2, s)

Here, the persistence of A would cause C to hold in S1 which would cause D1 to hold in S3, while the persistence of B would cause D2 to hold in S3. Thus, if A and B were both to persist, D1 and D2 would both hold in S3, which is impossible by A8'''. Hence, the preconditions to one of the general causal relations A6''' and A7''' must be false in S2. Since temporal persistence provides no basis for choice here, only the disjunction (¬T(B, S2) v ¬T(C, S2) follows. If C is false in S2, then the precondition, A, to the causal law supporting it, must also be false (i.e., ¬T(A, S1)). Thus, temporal persistence supports only the conclusion that either A is clipped in S1 or B is clipped in S2. Since chronological minimization favors the later clipping of B (for no good reason), it does not model temporal persistence in such temporal projection situations.

This temporal projection schema may be instantiated for an autonomous vehicle, Robbie, in a factory domain using the following informal definitions:

A =df John can lock out Robbie's forward gears.
B =df Robbie's reverse gears are locked out.
C =df Robbie's forward gears are locked out.
D1 =df Robbie is observed moving forward.
D2 =df Robbie is observed moving in reverse.
Wait =df Robbie waits.
E1 =df John tries to lock out Robbie's forward gears.
E2 =df Robbie moves.

The causal laws are thus informally interpreted:

A5''') If John is able to lock out Robbie's forward gears, and attempts to do so, he will succeed.

A6''') If Robbie's forward gears are locked out, its movement will cause it to be observed moving in reverse.

A7''') If Robbie's reverse gears are locked out, its movement will cause it to be observed moving forward.

In the initial conditions, reverse gears are locked out, forward gears are not locked out, John is able to lock the forward gears, and no robot movement is observed. Thus, we would ordinarily expect the robot to be unable to move in S2, except that we are assuming that movement occurs. Hence, one of the gears must not be locked, but there is no basis for deciding which one. If reverse is unlocked, then in accord with persistence, we assume that its previously locked state was clipped in S2. If forward is unlocked, then the precondition (John's ability to lock it) for its presumptive cause (John's attempt) must be clipped in S1. Chronological minimization of clipping clearly entails that the reverse gear must have been unlocked, while minimizing determined causes (or potential causes) allows either alternative since neither is a determined cause in all models.

This quite ordinary sort of situation uncovers a problematic assumption in the chronological minimization approach. It assumes that one can start with an initial state and a planned set of events and sweep forward in time allowing all facts to persist until the result of some event clips them. However, this model is clearly inadequate for handling conflicting results; temporal backtracking must be allowed in order to accommodate the adjustments required by such conflicts.

Although causal minimizations are, thus, superior to chronological minimization of clipping in this type of scenario, other conflicting result scenarios might conceivably create problems for our approach as well. This remains an area for further investigation.

V. Further Research

Much work remains to extend the simple formulation presented here to handle variable length persistences, more complex facts as preconditions and effects, multiple simultaneous events of the same type, densely ordered time, durations of states and events, event-event causation and default causal generalizations. Our continuing work (Haugh, 1987) is aimed at a temporal causal logic supporting all these capabilities.

VI. Summary

The conclusions reached here may be summarized as follows:

1) Contrary to Hanks and McDermott (1985, 1986), minimizations using ordinary non-monotonic logics can handle the temporal projection problem.
2) Minimizing potential causes leads to the results preferred by commonsense in all the examined instances of the temporal projection problem.
3) Minimizing determined causes models the ordinary notion of temporal persistence better than chronological minimization of clipping.
4) Ordinary temporal persistence does not yield the conclusions preferred by commonsense in certain cases of incomplete initial knowledge.
5) Chronological minimization of clipping does not provide an adequate basis for solving the temporal projection problem nor for modeling temporal persistence in a variety of cases.
6) Therefore, ordinary minimizations of potential and determined causes better formalize temporal projection and temporal persistence, respectively, than the chronological minimizations previously advocated.

Acknowledgements

I would like to thank Donald Perlis for helpful discussions, encouragement in pursuing issues in temporal reasoning, and for pointing out Lifschitz's independent work in causal minimization. Special thanks to Steve Barash for assistance in developing the robot domain counterexample. I also thank Vladimir Lifschitz, Yukiko Sekine, and Frederick Suppe for helpful comments.

References

Allen, J. 1984. "Towards a General Theory of Action and Time." *Artificial Intelligence*, 24 (2): 123-154.

Hanks, S. and McDermott, D. 1985. *Temporal Reasoning and Default Logics.* YALEU/CSD/RR #430. New Haven: Yale University Department of Computer Science

Hanks, S. and McDermott, D. 1986. "Default Reasoning, Nonmonotonic Logics, and the Frame Problem," *AAAI-86.* Palo Alto: Morgan Kaufmann Publishers, Inc., pp. 328-333.

Haugh, B. 1987. *Non-Monotonic Formalisms for Commonsense Temporal-Causal Reasoning.* Ph.D. Dissertation, Department of Philosophy, University of Maryland, College Park. (in preparation).

Hayes, P. 1971. "A Logic of Actions." *In* D. Michie and B. Meltzer, (eds.), *Machine Intelligence 6.* Edinburgh: Edinburgh University Press.

Kautz, H.A. 1986. "The Logic of Persistence." *AAAI-86.* Palo Alto: Morgan Kaufmann Publishers, Inc., pp. 401-405.

Lifschitz, V. 1986. "Pointwise Circumscription: Preliminary Report." *AAAI-86.* Palo Alto: Morgan Kaufmann Publishers, Inc., pp. 406-410.

Lifschitz, V. 1987. "Formal Theories of Action" *In* Brown, F. (ed.), *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop.* Los Altos: Morgan Kaufman Publishers, Inc.

McCarthy, J. 1968. "Programs with Common Sense." In Minsky, M. (ed.), *Semantic Information Processing.* Cambridge: The MIT Press, pp. 403-418.

McCarthy, J. 1980. "Circumscription - A Form of Non-Monotonic Reasoning." *Artificial Intelligence,* 13:27-39.

McCarthy, J. 1986. "Aplications of Circumscription to Formalizing Common-Sense Knowledge." *Artificial Intelligence*, 28:89-116.

McCarthy, J. and Hayes, P. 1969. "Some Philosophical Problems from the Standpoint of Artificial Intelligence." *In* D. Michie and B. Meltzer, (eds.), *Machine Intelligence 4.* Edinburgh: Edinburgh University Press.

McDermott, D. 1982. "A Temporal Logic for Reasoning about Processes and Plans." *Cognitive Science*, 6: 101-155.

McDermott, D. and Doyle, J. 1980. "Non-Monotonic Logic I." *Artificial Intelligence*, 13:41-72.

Perlis, D. and Minker, J. 1986. "Completeness Results for Circumscription." Aritificial Intelligence, 28:29-42.

Reiter, R. 1980. "A Logic for Default Reasoning." *Artificial Intelligence* 13:81-132.

Shoham, Y. 1986. "Chronological Ignorance: Time, Nonmonotonicity, Necessity and Causal Theories." *AAAI-86.* Palo Alto: Morgan Kaufmann Publishers Inc. pp. 389-393.