A Circumscriptive Calculus of Events

Murray Shanahan

Imperial College
Department of Computing,
180 Queen's Gate,
London SW7 2BZ.
England.

Tel: +44 71 589 5111 x 5076
Email: mps@doc.ic.ac.uk

Abstract

A calculus of events is presented in which domain constraints, concurrent events, and events with non-deterministic effects can be represented. The paper offers a non-monotonic solution to the frame problem for this formalism that combines two of the techniques developed for the situation calculus, namely causal and state-based minimisation. A theorem is presented which guarantees that temporal projection will not interfere with minimisation in this solution, even in domains with ramifications, concurrency, and non-determinism. Finally, the paper shows how the formalism can be extended to cope with continuous change, whilst preserving the conditions for the theorem to apply.

Introduction

The frame problem was first described by McCarthy and Hayes in the Sixties [23], and has occupied the thoughts of AI researchers ever since. In a nutshell, the problem is this: if we deploy classical logic in a straightforward way to describe the effects of actions, we have to represent explicitly which properties are not affected by each action, as well as those that are. This is a problem because the number of properties that are not affected by an action tends to be huge in all but the most trivial domains.[1] By the early Eighties, it was thought that the frame problem could be solved using the newly developed techniques of formal default reasoning (McCarthy [22]). However, Hanks and McDermott [8] demonstrated that the naïve application of these techniques could lead to counter-intuitive results. Many authors, such as Lifschitz [13] and Shoham [37], rose to the challenge of finding solutions that yielded correct conclusions with the examples that had undermined earlier attempts.[2]

By the end of the Eighties, the emphasis of research on the frame problem had shifted towards providing solutions which could deal with the features of complex domains. Baker [1], for example, addressed the issue of ramifications (which arise with domain constraints). Lin and Shoham [19], on the other hand, examined the problem of concurrent events. A complete solution to the frame problem is perhaps now within our grasp. However, a number of outstanding issues still need to be resolved, chief among which are actions with non-deterministic effects, that is actions whose precise effects are unknown, and continuous change.

Besides focussing on the features of complex domains such as concurrency, non-determinism and continuous change, current proposals for solving the frame problem are distinguished from their antecedents in another way. No longer is it considered acceptable to argue for the validity of a proposed solution using only a small number of benchmark examples, such as Hanks and McDermott's Yale Shooting scenario. Following Lifschitz [15], a proposal has to be mathematically justified for a substantial class of problems (see also Sandewall [32]).

This paper offers a predicate calculus based formalism for representing and reasoning about change, which facilitates the representation of concurrent events and events with non-deterministic effects.[3] A non-monotonic solution to the frame problem is given for this formalism, based on the idea of minimising the extensions of certain predicates using standard prioritised circumscription.[4] The solution works for examples involving ramifications (domain constraints) and events with non-deterministic effects. A mathematical result is demonstrated which suggests that the solution has very general application. Another result is presented which facilitates the construction of temporal projection algorithms. Finally, the paper shows how to extend the formalism to cope with continuous change, whilst preserving the conditions of applicability of these mathematical results.

Unlike most of the work cited in this introduction, the solution to the frame problem offered in this paper is not based on the situation calculus of McCarthy and Hayes [23]. Although some authors have offered hints and suggestions as to how continuous change could be formulated using the situation calculus (Gelfond, *et al.* [6]), it is not yet clear how this could be done. On the other hand, work already exists which extends the narrative-based event calculus of Kowalski and Sergot [11] to deal with continuous change (Shanahan [35]), and this work, taken out of the framework of logic programming and augmented with a circumscriptive approach to default reasoning, is the basis of the formalism presented here.

The ontology of Kowalski and Sergot's formalism includes event types, time points and properties. Properties are initiated and terminated by events. In their formalism, once a property has been initiated, negation-as-failure is used to assume that it persists by default until an event occurs to terminate it, and this is how the frame problem is addressed. The calculus of events presented in this paper is similar in many respects.[5] However, unlike Kowalski and Sergot's formalism, which is expressed in extended Horn clauses and uses negation-as-failure, that presented here exploits the full expressive power of first-order predicate calculus, and therefore negation-as-failure cannot be relied on to supply a solution to the frame problem. The circumscriptive solution offered here takes advantage of the insights of several previous authors, whose work I will now briefly review.

After Hanks and McDermott [8] discovered flaws in early efforts to overcome the frame problem, such as McCarthy's [22], several authors developed more robust solutions, such as Lifschitz [12], Kautz [10] and Shoham [37] who use chronological minimisation, Haugh [9] and Lifschitz [13] who use causal minimisation, and Baker [1], [2], Baker and Ginsberg [3] and Lifschitz [15] who use state-based minimisation. The last two approaches, which were developed for the situation calculus, are based on the following simple observation. The Hanks-McDermott problem does not arise if temporal projection is independent of minimisation.[6] This independence can be achieved by designing a formalism in which the frame problem can be overcome by minimising predicates whose extensions do not depend on the outcome of projection. In other words, minimisation can be separated from projection if the predicates to be minimised express timelessly true facts, rather than time-varying facts. For example, with causal minimisation, the predicates which are minimised express timelessly true facts about the effects of actions.

The difficulty with this approach of separating minimisation from projection is to ensure that the formalism has sufficient expressive power to capture rich domains, whilst preserving the principle that the only predicates to be minimised express timelessly true facts. For example, the existing work using causal minimisation does not cope adequately with ramifications, as shown by Baker [1]. State-based minimisation handles ramifications much better (Baker [1]). With state-based minimisation, the predicates which are minimised express timelessly true facts about the abnormality of certain actions with respect to change. These facts are timelessly true because they are relativised to states, whose existence and properties are independent of temporal projection. The solution to the frame problem described here uses a hybrid of causal and state-based minimisation.

The paper is organised as follows. The event calculus and the accompanying solution to the frame problem are presented in the next two sections, followed by a traditional Blocks World example of its use, without any proof of correctness. Then the mathematical properties of the formalism are investigated. Two results are developed: one which suggests that the frame problem is solved for a wide class of problems, and another to aid the construction of temporal projection algorithms. Further examples of the application of the formalism are then presented, which illustrate the use of these mathematical results. These include examples with ramifications and events with non-deterministic effects. Finally, the formalism is extended to cope with continuous change in a way which ensures that the mathematical results still apply.

1. States

I will use a many-sorted language of first-order predicate calculus with equality, including variables for time points (t, t1, t2, etc.), properties (p, p1, p2, q, q1, q2, etc.), states (s, s1, s2, etc.), truth values (v, v1, v2, etc.), and truth elements (f, f1, f2, etc.).

The domain of truth values has two members, denoted by the constants True and False. A pair $\langle p,v \rangle$ is a truth element. The first part of the formalism to be presented concerns the properties of states. A state is represented as a set of truth elements. To capture this, the predicate $\in$ is defined as follows.[7]

$$s1=s2 \leftrightarrow \forall f\ [f \in s1 \leftrightarrow f \in s2] \qquad \text{(S1)}$$

$$\forall s1,f1\ \exists s2\ \forall f2\ [f2 \in s2 \leftrightarrow [f2 \in s1 \vee f2=f1]] \qquad \text{(S2)}$$

$$\exists s\ \forall f\ [\neg\ f \in s] \qquad \text{(S3)}$$

Axiom (S1) says that two states are equal if they have the same truth elements. Axiom (S2) says that any truth element can be added to any state to give another state. Axioms (S2) and (S3) guarantee that a state exists for every combination of truth elements, and are analogous to the existence-of-situations axiom in Baker's formulation [1], [2]. The properties which hold in a given state are described by the predicate HoldsIn. We have the following axioms.

$$\mathrm{HoldsIn}(p,s) \leftarrow [\langle p,\mathrm{True} \rangle \in s \wedge \neg\ \mathrm{AbState}(s)] \qquad \text{(E1)}$$

$$\neg\ \mathrm{HoldsIn}(p,s) \leftarrow [\langle p,\mathrm{False} \rangle \in s \wedge \neg\ \mathrm{AbState}(s)] \qquad \text{(E2)}$$

Axioms (S1) to (S3) and (E1) and (E2) will be part of the formalisation of any domain. If $\langle p,\mathrm{True} \rangle \in s$ then the property p holds in state s. Conversely if $\langle p,\mathrm{False} \rangle \in s$ then the property p does not hold in state s. If $\langle p,\mathrm{True} \rangle \notin s$ and $\langle p,\mathrm{False} \rangle \notin s$ then, in the absence of further information about HoldsIn, we cannot say whether or not the property p holds in state s. However, further information of this kind may be present in the form of domain constraints expressed as extra HoldsIn formulae. Such formulae can be admitted without giving rise to contradiction because of the AbState conditions on (E1) and (E2). The predicate AbState will be minimised, making Axioms (E1) and (E2) into defaults. Abnormal states are those ruled out by domain constraints. Baker's approach adopts a similar tactic.

Although there is no overall partitioning of properties into primitive and derived, the members of a set can be thought of as the primitive properties that hold/don't hold in the corresponding state. Domain constraints in the form of extra HoldsIn formulae can then be thought of as yielding "derived" properties. The presence of such domain constraints means that an event can have complicated ramifications, which must be dealt with by any approach to the frame problem. It is because of domain constraints and ramifications that states are represented by <u>partial</u> descriptions of the properties that hold in them, and (E1) and (E2) cannot be replaced by a simple biconditional. Ramifications will be discussed in some detail in a later section.

2. A Calculus of Events

Now the main axiom of the formalism is presented in a form which is suitable for domains which involve only discrete change. Later, it will be modified to cater for the continuous case. The axiom defines the predicate State. The formula State(t,s) represents that time point t is associated with state s. Each time point is associated with a single, characterising state s,[8] such that $\langle p,\mathrm{True} \rangle \in s$ if and only if the property p was initiated by some event before t and still holds at t, and $\langle p,\mathrm{False} \rangle \in s$ if and only if p was terminated by some event before t and still does not hold at t. The axiom we require is essentially the following.

$$\begin{array}{l} State(t,s) \leftrightarrow \\ \quad \forall p[[\langle p,True\rangle \in s \leftrightarrow Initiated(p,t)] \wedge \\ \quad [\langle p,False\rangle \in s \leftrightarrow Terminated(p,t)]] \end{array}$$

In order to make the formalism easier to use, the final form of the axiom will facilitate the representation of the initial situation. But for now, let's assume the above axiom, and consider the meaning of Initiated and Terminated. The formulae Initiated(p,t) and Terminated(p,t) are not part of the language, but are just abbreviations, which are defined as follows.[9] Several more predicates are introduced here, along with variables for the new sort of event types (e, e1, e2, etc.).

$$\begin{array}{l} Initiated(p,t2) \equiv_{def} \\ \quad \exists e,t1,s[Happens(e,t1) \wedge t1 < t2 \wedge State(t1,s) \wedge \\ \quad Initiates(e,p,s) \wedge \neg\, Clipped(t1,p,t2)] \end{array}$$

$$\begin{array}{l} Terminated(p,t2) \equiv_{def} \\ \quad \exists e,t1,s[Happens(e,t1) \wedge t1 < t2 \wedge State(t1,s) \wedge \\ \quad Terminates(e,p,s) \wedge \neg\, Declipped(t1,p,t2)] \end{array}$$

Again, the formulae Clipped(t1,p,t2) and Declipped(t1,p,t2) are not part of the language, but are abbreviations, which are defined as follows.

$$\begin{array}{l} Clipped(t1,p,t3) \equiv_{def} \\ \quad \exists e,t2,s[Happens(e,t2) \wedge t1 < t2 \wedge t2 < t3 \wedge \\ \quad State(t2,s) \wedge Terminates(e,p,s)] \end{array}$$

$$\begin{array}{l} Declipped(t1,p,t3) \equiv_{def} \\ \quad \exists e,t2,s[Happens(e,t2) \wedge t1 < t2 \wedge t2 < t3 \wedge \\ \quad State(t2,s) \wedge Initiates(e,p,s)] \end{array}$$

The formula Initiates(e,p,s) represents that, in state s, the property p is initiated by an event of type e, and the formula Terminates(e,p,s) represents that, in state s, the property p is terminated by an event of type e. The "causal" predicates Initiates and Terminates will be minimised. The formula Happens(e,t) represents that an event of type e happens at time t. The formula t1 < t2 represents that time t1 is before time t2. For the discrete case, time points can be interpreted by the naturals, and I will assume that we are considering only models in which < is interpreted accordingly.[10]

To make it easy to represent the initial situation (the state at time 0), the predicate Initially is introduced. The formula Initially(p) represents that property p holds in the initial situation. Conversely, the formula Initially(Not(p)) represents that property p does not hold in the initial situation. Note that Initially will be minimised, although our knowledge of the initial situation might be incomplete. The final form of the main axiom incorporates the Initially predicate as follows. Two further abbreviations are used to make the axiom clear.

$$\begin{array}{l} State(t,s) \leftrightarrow \\ \quad \forall p[[\langle p,True\rangle \in s \leftrightarrow TrueAt(p,t)] \wedge \\ \quad [\langle p,False\rangle \in s \leftrightarrow FalseAt(p,t)]] \end{array} \qquad (E3)$$

$$TrueAt(p,t) \equiv_{def} Initiated(p,t) \vee [Initially(p) \wedge p \neq Not(p') \quad \wedge \neg Clipped(0,p,t)]$$

$$\begin{array}{l} FalseAt(p,t) \equiv_{def} \\ \quad Terminated(p,t) \vee [Initially(Not(p)) \wedge \neg Declipped(0,p,t)] \end{array}$$

It is important to note here that the state s which characterises a time t is not the set of all properties that hold (or don't hold) at t. Rather, s is a subset of those properties — the "primitive" ones. The presence of domain constraints, expressed in terms of HoldsIn, means that further "derived" properties might hold (or not hold) at t. The last axiom of the formalism defines the predicate HoldsAt, which takes into account this possibility. The formula HoldsAt(p,t) represents that property p holds at time t.

$$\text{HoldsAt}(p,t) \leftrightarrow \exists s[\text{State}(t,s) \wedge \text{HoldsIn}(p,s)] \tag{E4}$$

The HoldsAt predicate is still not a complete description of which properties hold at what times, because there may be non-deterministic actions, or an incompletely described initial situation. However, it takes into account all that is known about each time point. In the rest of the paper, the conjunction of Axioms (S1) to (S3) with (E1) to (E4) will be denoted by EC. In general, a temporal projection problem will be captured by the conjunction of EC with a conjunction of Happens and Initially formulae representing a history of events, and a conjunction of Initiates, Terminates and HoldsIn formulae representing the domain. The answer to the temporal projection problem resides in the set of HoldsAt formulae that are consequences.

The frame problem is overcome using circumscription [21]. Circumscription works by minimising the extensions of certain predicates. To minimise the extension of a predicate is to insist that it contains only those objects it is forced to contain by the formula being circumscribed. The extensions of other predicates are optionally allowed to vary to accommodate this. We write CIRC[λ ; P* ; Q*] to denote the circumscription of λ, minimising P* and allowing Q* to vary, where P* and Q* are sets of predicates. If Q* is empty, this is sometimes written CIRC[λ ; P*]. The circumscription of a formula is conventionally defined by a second-order sentence, but it can be equivalently presented in terms of minimal models. Consider two models M1 and M2. We have,

M1 $\sqsubseteq_{P*;Q*}$ M2 if

- M1 and M2 agree on the interpretation of everything except P* and Q*.
- For all p in P*, the extension of p in M1 $\subseteq$ its extension in M2.

Then, a model M of λ is *minimal* with respect to $\sqsubseteq_{P*;Q*}$ if there is no model M' of lambda such that M' $\sqsubseteq_{P*;Q*}$ M and not M $\sqsubseteq_{P*;Q*}$ M'. A formula is true in CIRC[λ ; P* ; Q*] if and only if it is true in all models of λ which are minimal with respect to $\sqsubseteq_{P*;Q*}$. This notion of circumscription can be extended to include the idea of minimising certain predicates with a higher priority than others (Lifschitz [12]). Let P1*, P2* and Q* be sets of predicates. We write CIRC[λ ; P1*>P2* ; Q*] to denote the circumscription of λ, minimising P1* with a higher priority than P2*, and allowing Q* to vary. Again, the conventional definition is via a second-order sentence, but this is equivalent to,

$$\text{CIRC}[\lambda ; P1^* ; P2^* \cup Q^*] \wedge \text{CIRC}[\lambda ; P2^* ; Q^*]$$

The exact choice of which predicates are minimised in a circumscription, the order in which they are minimised, and which predicates are allowed to vary, is called the *circumscription policy*. For a thorough review of the theory of circumscription, see Lifschitz [16].

The circumscription policy for overcoming the frame problem, representing the assumptions that the only domain constraints are the known domain constraints, that the only events which occur are those which are known to occur, and that the only effects of events are the known effects, is to minimise AbState at a high priority, and to minimise Initially, Happens, Initiates and Terminates at a lower priority, allowing HoldsAt and

State to vary. It is necessary to prioritise the minimisation of AbState in order to exclude models in which a larger than necessary extension of AbState is traded for a smaller than desired extension of Initiates or Terminates. The circumscription of a formula λ according to this policy will be written $CIRC_{ec}[\lambda]$.

3. The Blocks World

Fortunately, the foregoing machinery is mostly transparent to anyone who uses the formalism, and descriptions of domains and histories are intuitive and elegant. In the next section, the mathematical properties of the event calculus are investigated, and a result is developed which supports the claim that the frame problem has been solved for a large class of examples. But first, I will show how the formalism could be used to represent a simple version of the Blocks World. The ontology of this world includes blocks and locations. A new sort is introduced for these, with variables x, y and z.[11] An event of type Move(x,y) is an attempt to move block x to location or block y. The property On(x,y) represents that block x is at location y or on block y. The domain of time points is assumed to be the natural numbers.[12] We have the following formulae.

Initiates(Move(x,y),On(x,y),s) ←
    HoldsIn(Clear(x),s) ∧ HoldsIn(Clear(y),s) ∧ x ≠ y

Initiates(Move(x,y),Clear(z),s) ←
    HoldsIn(Clear(x),s) ∧ HoldsIn(Clear(y),s) ∧ x ≠ y ∧
    HoldsIn(On(x,z),s) ∧ y ≠ z

Terminates(Move(x,y),On(x,z),s) ←
    HoldsIn(Clear(x),s) ∧ HoldsIn(Clear(y),s) ∧ x ≠ y ∧
    HoldsIn(On(x,z),s) ∧ y ≠ z

Terminates(Move(x,y),Clear(y),s) ←
    HoldsIn(Clear(x),s) ∧ HoldsIn(Clear(y),s) ∧ x ≠ y

The key feature of these formulae is that no mention is made of the HoldsAt predicate. No mention is made of actual times at all. Instead, the HoldsIn predicate is used to give access to the properties that hold when an event occurs, by referring to the corresponding state. Because states are timeless, the formulae are timelessly true. This style of representation permits the vital separation of minimisation from temporal projection.

Instead of explicitly specifying the conditions under which an event initiates or terminates the Clear property, a domain constraint could be used. This would be achieved by writing HoldsIn formulae which related the Clear property to the On property, and which constrained every block or location to have at most one block on it.

For each domain, a set of uniqueness-of-names axioms is required for properties and event types. It might also be necessary to include other uniqueness-of-names axioms, in this case for blocks and locations. I will not explicitly list these axioms in the examples in rest of the paper, but here is the full set of uniqueness-of-names axioms for this Blocks World domain, using Lifschitz's UNA notation [13].

UNA[Move]

UNA[Clear, On]

UNA[A, B, C, X1, X2, X3]

A particular sequence of events or actions is described using Happens formulae. For example, the following formulae represent that block A is moved onto block B at time 5, then block C is moved onto location X1 at time 8.

Happens(Move(A,B),5)

Happens(Move(C,X1),8)

Concurrent events, such as moving two blocks at the same time, are easily represented as Happens formulae with identical time arguments, so long as the events are independent, that is so long as their effects are not cumulative (like putting two weights on a pair of scales at the same time) or cancelling (like trying to lift an object and pressing down on it at the same time). In Section 8, a version of the calculus is described which can cope with cumulative and cancelling concurrent events. Note that it is easy to write formulae which represent events whose exact order of occurrence is not known, using disjunctions of Happens formulae, or using Happens formulae with existentially quantified time arguments.

Suppose blocks A, B and C are initially clear and at locations X1, X2 and X3 respectively. Then we have

Initially(On(A,X1))	Initially(Clear(A))
Initially(On(B,X2))	Initially(Clear(B))
Initially(On(C,X3))	Initially(Clear(C))

If the conjunction of all the above Initiates and Terminates formulae and uniqueness-of-names axioms is denoted by D, and the conjunction of the Happens and Initially formulae by H, then in all models of $\mathrm{CIRC}_{ec}[\mathrm{EC} \wedge \mathrm{D} \wedge \mathrm{H}]$, we have, for example, HoldsAt(On(A,B),12). I will not attempt to prove this here, but the results of the next section will provide a basis for proving which properties hold at what times for any domain and history.

Intuitively, though, how has the frame problem been solved? Note that the only predicates needed to capture the Blocks World domain and to represent a history of events are Initiates, Terminates, HoldsIn, Happens and Initially. The circumscription policy for overcoming the frame problem minimises only these domain and history predicates. The results of temporal projection, on the other hand, are expressed in terms of the predicate HoldsAt, which doesn’t appear in domain and history formulae. So, temporal projection is independent of minimisation. This has been achieved by using HoldsIn in the representation of the domain, a predicate indexed on states. It would have been tempting to use HoldsAt instead, obviating the need for states altogether. But then the extensions of Initiates and Terminates would vary according to the outcome of temporal projection. The strong result of the next section would not then be applicable, and the Hanks-McDermott problem would arise.

4. Some Properties of the Calculus

As Lifschitz points out [15], we would like to be able to demonstrate that an approach to the frame problem yields correct results, not just with a single example, but with a significant class of examples. General results of this sort have been produced for the situation calculus by Lifschitz [15] and Lin and Shoham [18], but neither of these papers addresses continuous change, concurrent events, or events with non-deterministic effects. Lin and Shoham have extended their work to deal with concurrent events [19], but the general result they prove is built on a criterion of epistemological completeness which apparently excludes the possibility of events with non-deterministic effects.

In this section, I present a theorem which says that any collection of domain and history formulae of a certain form can be minimised independently from the axioms of the event calculus. The demands on the form of domain and history formulae are very liberal. Concurrent events are allowed, and in later sections I show that domains involving non-deterministic events and continuous change can also be represented in the required form. The theorem is very general, and applies not only to the calculus above, but also to any calculus having the right form. I will write $\underline{x}$ to denote a tuple of variables, and x_i to denote the i^{th} variable in such a tuple.

Definition 1. A formula is *chronological in argument* k if it has the form $\forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$, where q is a predicate whose k^{th} argument is a time point and $\phi(\underline{x})$ is a formula in which $\underline{x}$ is free, and all occurrences of q in $\phi(\underline{x})$ are in conjunctions of the form $q(\underline{z}) \wedge z_k < x_k$.

For example, Axiom (E3) is chronological in argument 1. Under certain conditions, it is easy to work out the consequences of circumscribing the conjunction of a formula with a chronological formula.

Theorem 1. Consider only models in which the time points are interpreted by the naturals, and in which < is interpreted accordingly. Let P* and Q* be sets of predicates such that Q* includes q. Let $\psi = \forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$ be a formula which is chronological in some argument. Let χ be a formula which doesn't mention the predicate q. Then $CIRC[\chi \wedge \psi ; P^*; Q^*] \vDash CIRC[\chi ; P^*; Q^*]$.

The proof of Theorem 1 is given in Appendix A. Since both Axioms (E3) and (E4) are chronological in one of their arguments, Theorem 1 ensures that any domain and history can be minimised independently from the event calculus axioms, so long as the domain and history formulae don't mention HoldsAt or State. After this minimisation, the event calculus axioms can be used classically to derive which properties hold at what times. Appendix B contains another theorem, which facilitates the construction of temporal projection algorithms.

It is important to see that meeting the conditions for applying Theorem 1 is not in itself sufficient to solve the frame problem. It is still necessary to get the minimisation right before (E3) and (E4) are added. Axioms (S1) to (S3) and (E1) and (E2) play a crucial role in this respect. Theorem 1 simply supplies certain conditions under which projection is guaranteed not to interfere with minimisation. This sort of interference is what gives rise to the Hanks-McDermott problem.

In order to minimise domains and histories, two other properties of circumscription will be useful. Theorems 2 and 3 are due to Lifschitz. They are reproduced here without proof, but proofs can be found in Lifschitz [16]. Let λ be any formula and $\delta(\underline{x})$ be any formula in which $\underline{x}$ is free.

Theorem 2. $CIRC[\lambda \wedge \forall \underline{x}\ \ p(\underline{x}) \leftarrow \delta(\underline{x}); p]$ is equivalent to $\lambda \wedge \forall \underline{x}\ \ p(\underline{x}) \leftrightarrow \delta(\underline{x})$ if λ and $\delta(\underline{x})$ are formulae containing no occurrences of the predicate p.

Theorem 3. If all occurrences of the predicates $p_1, p_2, ..., p_n$ in a formula λ are positive, then $CIRC[\lambda ; P^*]$, where $P^* = \{p_1, p_2, ..., p_n\}$, is equivalent to

$$CIRC[\lambda ; p_1] \wedge CIRC[\lambda ; p_2] \wedge . . . \wedge\ CIRC[\lambda ; p_n]$$

5. The Yale Shooting Problem

With the results above, it's easy to show that the event calculus can solve the Yale shooting problem (Hanks & McDermott [8]). The domain comprises three types of event — Load, Sneeze and Shoot — and two properties — Alive and Loaded. These events and properties are represented by the following formulae, whose conjunction along with the requisite uniqueness-of-names axioms will be denoted by D1. Note that there are no axioms for the Sneeze event, since it doesn't affect any property.

$$\text{Initiates(Load,Loaded,s)} \qquad \text{(D1.1)}$$

$$\text{Terminates(Shoot,Alive,s)} \leftarrow \text{HoldsIn(Loaded,s)} \qquad \text{(D1.2)}$$

The Yale shooting scenario can be represented by the following history formulae, whose conjunction will be denoted by H1, describing three events — Load then Sneeze then Shoot. The Sneeze event substitutes for the Wait action of the original Yale shooting scenario. Using the event calculus, waiting is most naturally represented simply as a pause between events, rather than as an event which has no effect. The Sneeze event here serves the same purpose as waiting in the original formulation: it provides an opportunity for the minimisation to go wrong.

$$\text{Initially(Alive)} \quad \text{(H1.1)} \qquad \text{Happens(Load,10)} \quad \text{(H1.2)}$$

$$\text{Happens(Sneeze,15)} \quad \text{(H1.3)} \qquad \text{Happens(Shoot,20)} \quad \text{(H1.4)}$$

Proposition 1. $\text{CIRC}_{ec}[\text{EC} \wedge \text{D1} \wedge \text{H1}] \vDash \neg\, \text{HoldsAt(Alive,25)}$.

Proof. Let χ be the conjunction of D1 with H1 and EC without (E3) and (E4). $\text{CIRC}_{ec}[\chi]$ is defined as the conjunction of,

$$\text{CIRC}[\chi \text{ ; Initially, Happens, Initiates, Terminates ; State, HoldsAt}]$$

with,

$$\text{CIRC}[\chi \text{ ; AbState ; Initially, Happens, Initiates, Terminates, State, HoldsAt}].$$

We will consider each conjunct in turn.

Take the first conjunct. Since all occurrences of Happens, Initiates and Terminates in χ are positive,

$$\text{CIRC}[\chi \text{ ; Initially, Happens, Initiates, Terminates}]$$

is equivalent to,

$$\text{CIRC}[\chi \text{ ; Initially}] \wedge \text{CIRC}[\chi \text{ ; Happens}] \wedge \\ \text{CIRC}[\chi \text{ ; Initiates}] \wedge \text{CIRC}[\chi \text{ ; Terminates}]$$

from Theorem 3. Then, by applying Theorem 2 to each conjunct in this formula, it can be seen that the completions of Happens, Initiates and Terminates are true in all of its models. In particular, we have the following.

$$\text{Terminates(e,p,s)} \leftrightarrow \\ \text{e=Shoot} \wedge \text{p=Alive} \wedge \text{HoldsIn(Loaded,s)} \qquad (1)$$

$$\text{Happens}(e,t) \leftrightarrow \qquad (2)$$
$$[e=\text{Load} \wedge t=10] \vee [e=\text{Sneeze} \wedge t=15] \vee$$
$$[e=\text{Shoot} \wedge t=20]$$

Since there are no occurrences of State or HoldsAt in χ, (1) and (2) are also true in all models of $\text{CIRC}_{ec}[\chi]$, where these predicates are allowed to vary.

Now let's look at the second conjunct of $\text{CIRC}_{ec}[\chi]$. Without any domain constraints in D1, the only abnormal combinations of truth elements are those which include both $\langle p,\text{False}\rangle$ and $\langle p,\text{True}\rangle$ for some p. So, in all models of CIRC[χ ; AbState ; Happens, Initiates, Terminates] we have,

$$\text{AbState}(s) \leftrightarrow \exists p[\langle p,\text{False}\rangle \in s \wedge \langle p,\text{True}\rangle \in s] \qquad (3)$$

Since there are no occurrences of State or HoldsAt in χ, allowing these predicates to vary does not affect the outcome of the circumscription, so (3) is also true in all models of $\text{CIRC}_{ec}[\chi]$.

Now, since (E3) and (E4) are chronological, by applying Theorem 1, first to add (E3) and then to add (E4), we can show that (1), (2) and (3) are also true in all models of $\text{CIRC}_{ec}[\text{EC} \wedge \text{D1} \wedge \text{H1}]$.

The combination of (3) with Axioms (S2), (S3) and (E1) ensures that every model includes a state in which Alive and Loaded hold. Given this, along with (1), (2) and (3), it's easy to show classically from D1 and H1 that in all models of the circumscription we have,

$$\exists s[\text{State}(20,s) \wedge \text{HoldsIn}(\text{Alive},s) \wedge \text{HoldsIn}(\text{Loaded},s)]$$

and therefore,

$$\exists s[\text{State}(25,s) \wedge \neg\, \text{HoldsIn}(\text{Alive},s)]$$

and $\neg$ HoldsAt(Alive,25). □

6. Ramifications

Domain constraints can be expressed as HoldsIn formulae. For example, the property Dead can be defined in the following way.

$$\text{HoldsIn}(\text{Dead},s) \leftrightarrow \neg\, \text{HoldsIn}(\text{Alive},s) \qquad (\text{D1.3})$$

Let D1' denote the conjunction of (D1.1) to (D1.3) with the requisite uniqueness-of-names axioms.

Proposition 2. $\text{CIRC}_{ec}[\text{EC} \wedge \text{D1}' \wedge \text{H1}] \vDash \text{HoldsAt}(\text{Alive},20) \wedge \text{HoldsAt}(\text{Dead},25)$.

Proof. The addition of (D1.3) does not substantially affect the proof for Proposition 1. Instead of (3) we have,

$$\begin{aligned}&\text{AbState}(f,s) \leftrightarrow \\ &\quad \exists p[\langle p,\text{False}\rangle \in s \wedge \langle p,\text{True}\rangle \in s] \vee \\ &\quad [\langle \text{Alive},\text{True}\rangle \in s \wedge \langle \text{Dead},\text{True}\rangle \in s] \vee \\ &\quad [\langle \text{Alive},\text{False}\rangle \in s \wedge \langle \text{Dead},\text{False}\rangle \in s]\end{aligned}$$

Every model still includes a state in which Alive and Loaded hold. But we also get,

$$\exists s[\text{State}(25,s) \wedge \text{HoldsIn}(\text{Dead},s)]$$

and therefore HoldsAt(Dead,25). □

Here we will only look at simple examples. But in general, a domain constraint could be any formula involving just the HoldsIn predicate in which the only situation term is a universally quantified variable. In addition, a domain constraint could involve any predicate apart from HoldsAt, Initiates, Terminates, Initially, and State. For example, to formalise a problem involving a chess board in the next section, I will introduce two predicates Black and White to be used in domain constraints.

Note that (D1.3) doesn't ensure that an event which initiates Dead also terminates Alive. A property that holds because it was initiated by an event, can only be terminated directly by an event, and not by an event which terminates a property on which it depends. What this amounts to is that a property which is initiated by an event must be considered as primitive until it is terminated.

This principle can be illustrated by introducing a new property Walking, and a further constraint that if Walking holds then Alive must hold (Baker [2]).[13] The obvious way to try to represent this is with a HoldsIn formula.

$$\text{HoldsIn}(\text{Walking},s) \rightarrow \text{HoldsIn}(\text{Alive},s) \qquad \text{(D1.4)}$$

However, this formula only yields intuitive conclusions under certain circumstances. With the addition of (D1.4), we can show $\neg$ HoldsAt(Walking,25). But suppose we add a new event type Walk and the following additional domain and history axioms.

$$\text{Initiates}(\text{Walk},\text{Walking},s) \qquad \text{(D1.5)}$$

$$\text{Happens}(\text{Walk},5) \qquad \text{(H1.5)}$$

Let D1" be the conjunction of (D1.1), (D1.2), (D1.4) and (D1.5) with the requisite uniqueness-of-names axioms. Let H1' be H1 $\wedge$ (H1.5). Do we still have $\neg$ HoldsAt(Walking,25) in all models of $\text{CIRC}_{ec}[\text{EC} \wedge \text{D1}'' \wedge \text{H1}']$? From Axiom (E3), we get,

$$\text{State}(25,s) \rightarrow [\langle \text{Alive},\text{False}\rangle \in s \wedge \langle \text{Walking},\text{True}\rangle \in s]$$

Therefore, from (D1.4), (E1) and (E2), we have,

$$\text{State}(25,s) \rightarrow \text{AbState}(s)$$

So we can no longer deduce anything interesting about time 25. A better way to represent a domain constraint like (D1.4) is to use Initiates and Terminates. This preserves the principle that a property which is initiated directly by an event must also be terminated directly by an event.

$$\text{Terminates(e,Walking,s)} \leftarrow \text{Terminates(e,Alive,s)} \qquad \text{(D1.6)}$$

Both (D1.4) and (D1.6) can be present in the same domain theory. Axiom (D1.4) would be used to deduce that Alive holds when Walking holds, given an event that initiates Walking but none that initiates Alive, and Axiom (D1.6) would be used to deduce that Walking does not hold after an event that terminates Alive. Let D1''' be the conjunction of (D1.1), (D1.2), (D1.5) and (D1.6) with the requisite uniqueness-of-names axioms.

Proposition 3. $\text{CIRC}_{ec}[\lambda] \vDash \neg$ HoldsAt(Walking,25), where λ is EC $\wedge$ D1''' $\wedge$ H1'.

Proof. Theorem 1 is applied in the usual way. $\text{CIRC}_{ec}[\lambda]$ yields (3) as in the proof of Proposition 1, but instead of (1) and (2), $\text{CIRC}_{ec}[\lambda]$ now gives,

$$\begin{aligned}&\text{Terminates(e,p,s)} \leftrightarrow && (4)\\ &\quad [\text{e=Shoot} \wedge \text{p=Alive} \wedge \text{HoldsIn(Loaded,s)}] \vee &&\\ &\quad [\text{e=Shoot} \wedge \text{p=Walking} \wedge \text{HoldsIn(Loaded,s)}] &&\end{aligned}$$

$$\begin{aligned}&\text{Happens(e,t)} \leftrightarrow && (5)\\ &\quad [\text{e=Walk} \wedge \text{t=5}] \vee [\text{e=Load} \wedge \text{t=10}] \vee &&\\ &\quad [\text{e=Sneeze} \wedge \text{t=15}] \vee [\text{e=Shoot} \wedge \text{t=20}] &&\end{aligned}$$

From (3), (4) and (5), it's easy to show classically that $\neg$ HoldsAt(Walking,25). □

7. Non-Deterministic Effects

The separation of temporal projection and minimisation permitted by Theorem 1 allows us to represent events whose effects are non-deterministic, knowing that this will not precipitate problems with temporal projection. Examples of events with non-deterministic effects are those which initiate a disjunction of properties or a property which is existentially quantified. Consider the following problem, due to Ray Reiter.[14] The action of moving an object onto a chess board either initiates the property that the object is on black, or the property that it is on white, or both properties at once (if it straddles two squares).

Can the effects of such an action be represented in the event calculus, whilst preserving the conditions for applying Theorem 1? There is a danger that minimising the effects of moving an object onto the board will exclude the possibility of moving it to a position where it is on both black and white, since such an action would initiate two properties when it could have initiated only one. The circumscription should allow at least one model in which such an action initiates both properties. Also, it is important to exclude models in which the object, once moved onto the board, flickers between black and white. If it is moved onto a particular location on the board, then it should stay there until it is moved again.

Events with non-deterministic effects, such as the one in this problem, can often be represented as initiating an intermediate property, which has non-deterministic ramifications. This is the approach taken in the following solution. The variable c is introduced for locations. A location is either black, white or mixed. There is one event type Move. The property Loc(c) represents that the object is at location c. There are two other properties: OnBlack and OnWhite.

$$\forall \text{s}\ \exists \text{c}\ [\text{Initiates(Move,Loc(c),s)}] \qquad \text{(D2.1)}$$

$$\forall \text{c}\ [\text{Black(c)} \vee \text{White(c)} \vee \text{Mixed(c)}] \qquad \text{(D2.2)}$$

$$\begin{aligned}&\text{Black}(c) \leftrightarrow && \text{(D2.3)}\\ &\quad \forall s\ [\text{HoldsIn}(\text{Loc}(c),s) \rightarrow \\ &\qquad \text{HoldsIn}(\text{OnBlack},s) \wedge \neg\, \text{HoldsIn}(\text{OnWhite},s)]\end{aligned}$$

$$\begin{aligned}&\text{White}(c) \leftrightarrow && \text{(D2.4)}\\ &\quad \forall s\ [\text{HoldsIn}(\text{Loc}(c),s) \rightarrow \\ &\qquad \text{HoldsIn}(\text{OnWhite},s) \wedge \neg\, \text{HoldsIn}(\text{OnBlack},s)]\end{aligned}$$

$$\begin{aligned}&\text{Mixed}(c) \leftrightarrow && \text{(D2.5)}\\ &\quad \forall s\ [\text{HoldsIn}(\text{Loc}(c),s) \rightarrow \\ &\qquad \text{HoldsIn}(\text{OnWhite},s) \wedge \text{HoldsIn}(\text{OnBlack},s)]\end{aligned}$$

It is easy to see that minimising Initiates does not affect HoldsIn, since HoldsIn is held fixed. Nor is HoldsIn affected by temporal projection (Lemma 9 in Appendix A guarantees that if there is a model of (D2.1) to (D2.5) in which there is a state s such that HoldsIn(OnBlack,s) $\wedge$ HoldsIn(OnWhite,s), then there will be such a model of EC conjoined with (D2.1) to (D2.5)). So, without further domain constraints, it will not be possible to show that $\neg$ HoldsAt(OnBlack,t) $\vee$ $\neg$ HoldsAt(OnWhite,t) for any time t after a Move event, unless there is another event between the Move event and t which terminates Loc(c). In other words, there will always be a model in which the object is both on black and on white after it is put on the chess board, which is the result we require.

8. Concurrent Events

A simple modification to the event calculus axioms, which I will describe in this section, permits the representation of concurrent events whose effects are cumulative or cancelling. Two or more events are *cumulative* if their simultaneous occurrence has effects that none of them has on its own. One event *cancels* the effect of another if their simultaneous occurrence prevents the second event from having an effect which it does have if it occurs on its own. It should be noted that, as with the suggestions in the previous section, none of the amendments I propose affects the applicability of Theorem 1, as Axioms (E3) and (E4) will remain chronological.

With the present axioms, it is possible to represent the simultaneous occurrence of two events, but not the fact that their effects are in any way dependent on each other. Let’s see this with an example. Suppose we want to formalise the following. If we push a supermarket trolley then it will move forwards. If we pull on it it will go backwards. But if we push on it at the same time as pulling on it, then it will spin around. The first two facts are easily represented by the following event calculus formulae.

Initiates(Push,Forwards,s)

Terminates(Push,Backwards,s)

Initiates(Pull,Backwards,s)

Terminates(Pull,Forwards,s)

Suppose we push the trolley at time 5 and then pull it at time 10.

Happens(Push,5)

Happens(Pull,10)

If the conjunction of the above Initiates and Terminates formulae with the requisite uniqueness-of-names axioms is denoted by D3, and the conjunction of the Happens formulae by H3, then in all models of $CIRC_{ec}[EC \wedge D3 \wedge H3]$ we have, for example, HoldsAt(Forwards,7) and HoldsAt(Backwards,12). Now consider what happens if we try to represent the additional information that we pull the trolley at time 5 as well as pushing it, and that we push the trolley at time 10 as well as pulling it.

Happens(Pull,5)

Happens(Push,10)

If we denote the conjunction of H3 with these formulae by H3', then in all models of $CIRC_{ec}[EC \wedge D3 \wedge H3']$, we still have HoldsAt(Forwards,7) and HoldsAt(Backwards,12). How can we represent the fact that simultaneously pushing and pulling on the trolley makes it spin around instead of moving either forwards or backwards? Two steps are required. First, we need to be able to write Initiates and Terminates formulae which describe the effect of two or more events occurring together. Then we need to be able to express the fact that one event can cancel the effect of another occurring at the same time. To achieve the first aim, an extra axiom is introduced, along with the infix function + which maps a pair of event types onto a third event type. A compound event of type e1 + e2 happens if events of type e1 and e2 happen concurrently.

$$\text{Happens}(e1 + e2,t) \leftarrow \text{Happens}(e1,t) \wedge \text{Happens}(e2,t) \wedge e1 \neq e2 \qquad \text{(E5)}$$

Note that this axiom will accumulate any number of concurrent events into a single event type.[15] Now, to represent the cumulative effect of two concurrent events, it is only necessary to write the appropriate Initiates and Terminates formulae in the usual manner. For the supermarket trolley example, the following extra Initiates and Terminates formulae will suffice.

Initiates(Push + Pull,Spinning,s)

Terminates(Push + Pull,Forwards,s)

Terminates(Push + Pull,Backwards,s)

For completeness there should be two further formulae representing the fact that Push and Pull each terminate Spinning. Now we will get the desired cumulative effect of both events, but we will still retain the unwanted previous conclusions about their individual effects. To overcome this problem, following Gelfond *et al.* [6] and Lin and Shoham [19], I will introduce a new predicate Cancels. The formula Cancels(e1,e2) represents that if an event of type e1 occurs then it cancels the effects of any event of type e2 occurring at the same time. Now we have to modify the definitions of Initiated and Terminated to take account of Cancels.

$$\begin{aligned}&\text{Initiated}(p,t2) \equiv_{def}\\ &\quad \exists e,t1,s[\text{Happens}(e,t1) \wedge t1 < t2 \wedge \text{State}(t1,s) \wedge\\ &\quad \text{Initiates}(e,p,s) \wedge \neg\, \text{Clipped}(t1,p,t2) \wedge \neg\, \text{Cancelled}(e,t1)]\end{aligned}$$

Terminated(p,t2) $\equiv_{def}$
$\exists$e,t1,s[Happens(e,t1) $\wedge$ t1 < t2 $\wedge$ State(t1,s) $\wedge$
Terminates(e,p,s) $\wedge$ $\neg$ Declipped(t1,p,t2) $\wedge$
$\neg$ Cancelled(e,t1)]

Cancelled(e1,t) $\equiv_{def}$ $\exists$e2[Happens(e2,t) $\wedge$ Cancels(e2,e1)]

Similar modifications are required for Clipped and Declipped.

Clipped(t1,p,t3) $\equiv_{def}$
$\exists$e,t2,s[Happens(e,t2) $\wedge$ t1 < t2 $\wedge$ t2 < t3 $\wedge$
State(t2,s) $\wedge$ Terminates(e,p,s) $\wedge$ $\neg$ Cancelled(e,t2)]

Declipped(t1,p,t3) $\equiv_{def}$
$\exists$e,t2,s[Happens(e,t2) $\wedge$ t1 < t2 $\wedge$ t2 < t3 $\wedge$
State(t2,s) $\wedge$ Initiates(e,p,s) $\wedge$ $\neg$ Cancelled(e,t2)]

Finally, Cancels must be minimised at the same priority as Happens, Initiates and Terminates, representing the assumption that events don't cancel each other's effects unless they are known to. It is interesting to note that there is no need for an axiom to ensure that a compound event inherits the effects of its component events, as required in the approach of Lin and Shoham [19]. Instead, a history of events is described entirely in terms of the separate occurrences of individual events, even if they are concurrent. Cancellation formulae may now be included in the description of a domain. To complete the supermarket trolley example, we have to add the following.

Cancels(Push,Pull)

Cancels(Pull,Push)

If we denote the conjunction of the modified event calculus axioms by EC' and the conjunction of D3 with the extra domain formulae by D3', then in all models of the circumscription of EC' $\wedge$ D3' $\wedge$ H3' according to the above policy, we have $\neg$ HoldsAt(Forwards,7) and $\neg$ HoldsAt(Backwards,12) as well as HoldsAt(Spinning,7) and HoldsAt(Spinning,12).

9. Continuous Change

This section extends the event calculus of Section 2 to deal with continuous as well as discrete change, as in Shanahan [35]. This is achieved through the introduction of two new predicates. Variables for elapsed time are introduced (a, a1, a2, etc.). The formula Trajectory(q,s,p,a) represents that, if the discrete property q is initiated in state s, then after a period of time a the continuous property p holds. For example, property q could be that a ball is moving at a certain velocity and property p could be that the ball has travelled a certain distance from its starting point.

The Trajectory predicate facilitates the representation of continuous change, such as the height of a falling object or the level of liquid in a filling vessel, but doesn't supply any means of representing that events occur when certain continuous properties hold. For example, when the level of liquid in a vessel reaches the vessel's rim, then an overflow event occurs. A second new predicate is introduced for this purpose. The formula Triggers(s,e) represents that an event of type e happens in the state s.

There are now two ways in which a property can be caused directly by an event. I will write Initiated_d(p,t) to represent that the discrete property p holds at time t and was initiated by an event, and $\text{Initiated}_c(p,t)$ to represent that the continuously varying property p holds at time t and was initiated by an event. The same property can be continuous at some times and discrete at others. The location of a ball, for instance, can be considered continuous while the ball is moving but discrete while it's stationary. We have the following replacement definitions for Initiated.

$$
\begin{aligned}
&\text{Initiated}_d(p,t2) \equiv_{def}\\
&\quad \exists t1,e,s[t1 < t2 \wedge \text{State}(t1,s) \wedge [\text{Happens}(e,t1) \vee \text{Triggers}(s,e)] \wedge\\
&\quad \text{Initiates}(e,p,s) \wedge \neg\, \text{Clipped}(t1,p,t2)]
\end{aligned}
$$

$$
\begin{aligned}
&\text{Initiated}_c(p,t2) \equiv_{def}\\
&\quad \exists t1,e,s,q[t1 < t2 \wedge \text{State}(t1,s) \wedge\\
&\quad [\text{Happens}(e,t1) \vee \text{Triggers}(s,e)] \wedge\\
&\quad \text{Initiates}(e,q,s) \wedge \neg\, \text{Clipped}(t1,q,t2) \wedge\\
&\quad a = t2 - t1 \wedge \text{Trajectory}(q,s,p,a)]
\end{aligned}
$$

$$
\text{Initiated}(p,t) \equiv_{def} \text{Initiated}_d(p,t) \vee \text{Initiated}_c(p,t)
$$

The rest of the definitions of Section 2 are modified to cope with Triggers, but are otherwise the same as before.

$$
\begin{aligned}
&\text{Clipped}(t1,p,t3) \equiv_{def}\\
&\quad \exists e,t2,s[t1 < t2 \wedge t2 < t3 \wedge \text{State}(t2,s) \wedge \text{Terminates}(e,p,s) \wedge\\
&\quad [\text{Happens}(e,t1) \vee \text{Triggers}(s,e)]]
\end{aligned}
$$

$$
\begin{aligned}
&\text{Declipped}(t1,p,t3) \equiv_{def}\\
&\quad \exists e,t2,s[t1 < t2 < t3 \wedge \text{State}(t2,s) \wedge \text{Initiates}(e,p,s) \wedge\\
&\quad [\text{Happens}(e,t1) \vee \text{Triggers}(s,e)]]
\end{aligned}
$$

$$
\begin{aligned}
&\text{Terminated}(p,t2) \equiv_{def}\\
&\quad \exists e,t1,s[t1 < t2 \wedge \text{State}(t1,s) \wedge [\text{Happens}(e,t1) \vee \text{Triggers}(s,e)] \wedge\\
&\quad \text{Terminates}(e,p,s) \wedge \neg\, \text{Declipped}(t1,p,t2)]
\end{aligned}
$$

Apart from these modifications, the axioms of EC are retained from Section 2. The circumscription policy to overcome the frame problem is to minimise AbState with high priority, then Trajectory, Triggers, Initially, Happens, Initiates and Terminates, allowing HoldsAt and State to vary. Given a formula λ, I will write $\text{CIRC}_{cec}[\lambda]$ to denote its circumscription according to this policy.

The statement of Theorem 1 and its proof in Appendix A assume that time points are interpreted by the natural numbers. To accommodate genuinely continuous change, time points need to be interpreted by the reals. A variant of Theorem 1 can be proved when time points are interpreted by the reals, but the proof is more complicated than that for the naturals. Also, a certain condition must hold. Intuitively, this condition states that it must be possible to map the real time line onto a well-founded structure (known as a marker set) in such a way that the recursive formula $\forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$ is also well-founded with respect to that structure. In the case of the event calculus axioms, this structure is the set of time points at which events occur. This condition clearly holds, for example, if there is a finite number of events.

Shortly I will illustrate the use of the axioms presented above with an example, but first I will present the extended version of Theorem 1.

Definition 2. A *marker set* is a subset S of $\mathbb{R}$ such that, for all T1 in $\mathbb{R}$, the set of T2 in S such that T2<T1 is finite.

From the definition, a marker set can be finite or infinite, but must be countable. Furthermore, the definition ensures that we can speak of the n^{th} element of a marker set and that this will be less than the $n+1^{th}$ element.

Definition 3. A formula ψ is *real-chronological in argument* k with respect to a formula χ and a marker set S if

a) It has the form $\forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$, where q is a predicate whose k^{th} argument is a time point and $\phi(\underline{x})$ is a formula in which $\underline{x}$ is free, and

b) All occurrences of q in $\phi(\underline{x})$ are in conjunctions of the form $q(\underline{z}) \wedge z_k < x_k \wedge \theta$, where $\chi \wedge \psi \vDash \neg\, \theta$ if $z_k \notin S$.

Under the right conditions, Axiom (E3) will be real-chronological in argument 1 with respect to a conjunction of domain and history formulae and a marker set corresponding to the set of time points at which events (including triggered events) occur.

Theorem 4. Consider only models in which the time points are interpreted by the reals, and in which < is interpreted accordingly. Let P* and Q* be sets of predicates such that Q* includes q. Let $\psi = \forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$ be a formula which is real-chronological in some argument with respect to a formula χ which doesn't mention the predicate q, and a marker set S. Then $\text{CIRC}[\chi \wedge \psi ; P^*; Q^*] \vDash \text{CIRC}[\chi ; P^*; Q^*]$.

The proof of Theorem 4 is given in Appendix C. Now let χ be the conjunction of a domain formula and a history formula with EC, but without (E3) and (E4). Let S be the set of time points at which events occur according to (E3) $\wedge\ \chi$. It can be seen that (E3) is real-chronological in argument 1 with respect to χ and S, so long as S is a marker set. So Theorem 4 can be used in exactly the same way as Theorem 1 to add (E3) and (E4) to a circumscribed χ without affecting the minimisation.

Under what circumstances does S constitute a marker set? Here are some examples. It is clear that S is a marker set if there is a finite number of events. It can easily be shown that it is also a marker set if there exists a smallest non-zero interval size between any two non-concurrent events. On the other hand, S is clearly not a marker set if an event occurs at every time point in the infinite series 0, 1/2, 1/4, 1/8, 1/16, and so on. Examples of this kind actually do arise, as in the idealised description of a bouncing ball (see Davis [5]), and the proof might be generalised to cope with these cases using transfinite induction.

To illustrate this approach to the representation of continuous change, I will formalise Sandewall's "ball and shaft" example (Sandewall [29]). A ball is moving horizontally along a surface towards a vertical shaft. When it reaches the shaft, it starts to fall, bouncing back and forth between the walls of the shaft, until it reaches the bottom where it comes to rest.

The "ball and shaft" can be represented as follows. New sorts are introduced for velocities, distances and heights, and I will consider only interpretations in which these sorts are interpreted by the reals.[16] There are four types of event: Propel(v), Drop, Bounce and Stop. A Propel(v) event sets the ball in motion with velocity v. A Drop event, which occurs when the ball is no longer supported, starts the ball falling. A Bounce event, which occurs when the ball hits a vertical surface, reverses the ball's direction of motion. A Stop event occurs when the ball comes to rest. The property Moving(v) represents that the ball is moving horizontally with velocity v, Falling represents that the ball is in freefall, Distance(d) represents that the ball is distance d from its starting point which is assumed to be 0, and Height(h) represents that the ball is at height h. The horizontal surface is assumed to be at height 0. The near wall of the shaft is at distance A, the far wall is at distance B and the bottom is at height C. I will consider only interpretations in which the arithmetic functions have their usual meanings.

Initiates(Propel(v),Moving(v),s) (D4.1)

Trajectory(Moving(v),s,Distance(d),a) ← d = v*a (D4.2)

Triggers(s,Drop) ← (D4.3)
HoldsIn(Distance(d),s) ∧ d = A ∧
HoldsIn(Moving(v),s) ∧ v > 0

Initiates(Drop,Falling,s) (D4.4)

Trajectory(Falling,s,Height(h),a) ← $h = -4.9*a^2$ (D4.5)

Triggers(s,Bounce) ← (D4.6)
HoldsIn(Distance(d),s) ∧ HoldsIn(Moving(v),s) ∧
[d=B ∧ v > 0] ∨ [d=A ∧ v < 0]

Initiates(Bounce,Moving(v2),s) ← (D4.7)
HoldsIn(Moving(v1),s) ∧ v1 = –v2

Terminates(Bounce,Moving(v),s) (D4.8)

Triggers(s,Stop) ← (D4.9)
HoldsIn(Height(h),s) ∧ h = C ∧ HoldsIn(Falling,s)

Terminates(Stop,Moving(v),s) (D4.10)

Terminates(Stop,Falling,s) (D4.11)

Eleven domain axioms are required, as in Sandewall's formulation [29]. Their conjunction along with the requisite uniqueness-of-names axioms will be denoted by D4. Now consider the following event.

Happens(Propel(5),0) (H4.1)

Proposition 4. $\mathrm{CIRC_{cec}}$[EC ∧ D4 ∧ (H4.1)] ⊨ ∃t[Happens(Stop,t) ∧ t > 0]. In other words, the ball eventually reaches the bottom of the shaft.

Proof. Let χ be the conjunction of EC without (E3) and (E4) with D4 ∧ (H4.1). The proof is similar to the proof of Proposition 1. Applying Theorem 3 then Theorem 2 to χ yields the completions of the predicates minimised by $\mathrm{CIRC_{cec}}$. Then, since (E3) and (E4) are still chronological with the new definitions, Theorem 4 can be applied, first to add (E3) to χ then to add (E4), to show that these completions are true in all models of $\mathrm{CIRC_{cec}}$[EC ∧ D4 ∧ (H4.1)]. There are no domain constraints, so $\mathrm{CIRC_{cec}}$ yields (3) as in the proof of Proposition 1. It is straightforward to show classically from these completions and the event calculus axioms that the proposition is true. □

The same example is formalised by Sandewall [29] using a form of chronological minimisation extended to cope with continuous change, in which discontinuities are postponed until as late as possible. Sandewall's formulation, although it has the same number of axioms, is more concise. It uses temporal modalities and doesn't introduce events where there are discontinuities. It is arguable whether this succinctness is achieved at the expense of expressive power. Indeed, when Sandewall does try to combine reasoning about action with reasoning about continuous change, a much more complex formalisation results [30]. One drawback to chronological minimisation, as noted by Kautz [10], is that it does not cope well with explanation problems. However, Sandewall has recently attempted to address this problem [31], [32]. A more thorough comparison of Sandewall's approach with that presented here would be useful.

Concluding Remarks

Arguably, the most pleasing solutions to default reasoning problems are modular. First, we represent in classical, monotonic logic what we know about the world — the propositions of which we are certain and from which we can draw definite conclusions. Then, we formalise the apparatus for jumping to reasonable, but defeasible, conclusions whose validity we cannot absolutely guarantee. The solution presented in this paper does not seem to conform to this ideal, since states were introduced into the ontology solely to facilitate the formalisation of default persistence by allowing domain and history formulae to be safely circumscribed separately from the axioms of the event calculus. It is tempting to conclude that there is a flaw in the solution, since it has forced the ontology of the formalism. On the other hand, it could be argued that the exercise of formalising default persistence has simply brought to light the desirability, independently of the need for default persistence, of an ontology which includes states.

This paper does not address the issue of explanation, that is reasoning from effects to causes. A number of temporal reasoning problems have an explanation component, such as the bloodless variation of the Yale shooting scenario (in which the victim is alive after the shot has been fired), and Kautz's stolen car problem [10]. In terms of the circumscriptive event calculus, a HoldsAt fact might be given, demanding an explanation in terms of Happens formulae. This conforms to a representational principle underlying the event calculus — that properties hold because events initiate them, whilst events themselves are "first causes".

There are two approaches to explanation with the event calculus: deductive and abductive. In the deductive approach, HoldsAt facts requiring explanation are conjoined with the domain, history and event calculus formulae, and explanations are expected to be among the logical consequences. This approach, using other formalisms, is common in the literature (Morgenstern & Stein [26], Lifschitz & Rabinov [17], Baker [1], [2]). In the abductive approach, which is less common (Shanahan [34], [35]), explanations are Happens formulae which, when conjoined with the domain and history formulae and event calculus axioms, yield the facts to be explained as logical consequences.

The abductive approach seems to fit better the above-mentioned representational principle. Furthermore, no extra representational apparatus is required to solve explanation problems using abduction. Only the reasoning mode changes. The deductive approach to explanation, on the other hand, violates the conditions for applying Theorems 1 and 4. If extra HoldsAt formulae are conjoined to Axiom (E4), the results of minimisation become unpredictable.

An interesting topic for further research is the relationship between the calculus of events presented here and the situation calculus of McCarthy and Hayes [23]. The two formalisms can be compared with respect to their ontologies, the set of basic relations they represent, and their treatment of persistence. I will briefly discuss each of these in turn. The ontologies of the two formalisms are similar — both include properties (or fluents) and event types (or actions). The event calculus also includes states — which strongly resemble situations — and time points, which featured in McCarthy and

Hayes’s 1969 paper, although they have rarely been employed in subsequent work using the situation calculus (but see Miller & Shanahan, [25]).

One of the main motivations of this paper was to extend approaches to the frame problem to deal with continuous change, a subject which has been almost entirely neglected in the situation calculus literature. It is not clear, at first glance, that the ontology of the situation calculus is adequate for representing continuous change, since it is centred on instantaneous snapshots of the world which are organised into a tree structure via the Result function. Gelfond *et al*. [6] argue that continuous change could be represented in the situation calculus through the introduction of infinitely divisible actions. Unfortunately, they do not explore this suggestion very far. It is particularly difficult to see how triggered events, which are needed to represent all but the simplest examples of continuous change, could be captured with an ontology that lacks real-valued narrative time.

How do the basic relations represented in the event calculus compare with those in the situation calculus? The event calculus relates properties to time points, whilst the situation calculus relates them to situations. Since situations are hypothetical whilst time points are actual, this leads to an emphasis in the event calculus on an actual narrative of events, whilst the situation calculus concentrates on hypothetical sequences of events. However, the event calculus of this paper includes states, which are hypothetical like situations, so there is no fundamental reason why it could not match the situation calculus in representing hypothetical sequences of events. Similarly, the situation calculus can incorporate predicates which distinguish actual from hypothetical situations (Pinto & Reiter [27]), or which map situations onto a narrative time line (Miller & Shanahan [25]). So, with respect to their ontologies and the basic relations they represent, it seems to be possible to extend both formalisms until they merge into each other.

Perhaps the most significant difference between the two formalisms is in the treatment of persistence. Frame axioms in the situation calculus literature usually relate the properties which hold in a situation to those which hold in the preceding situation, whilst the event calculus persistence axiom (E3) relates the properties which hold at a time point to earlier events. This has two important consequences. First, situation calculus frame axioms are usually bidirectional — persistence works backwards as well as forwards. Persistence in the event calculus works forwards only.[17] Second, properties only persist in the event calculus if they are initiated by events. Other properties are not “caught” by the persistence axiom. In the situation calculus literature, however, the frame axiom usually applies to all fluents, although in Lifschitz [14], [15] its application is restricted to a subset of the fluents known as the frame fluents.

Hopefully further insight into the relationship between the two formalisms can be gained in the future. Ultimately, what we would like to develop is a deeper understanding of the space of possible formalisms for representing change. Such an understanding would map out the possible ontologies, sets of basic relations, and approaches to persistence, and would highlight the implications of each choice. In the light of such an understanding, the apparent boundaries between particular formalisms would disappear altogether.

Acknowledgements

Thanks to Tony Kakas, Vladimir Lifschitz, Rob Miller, Ray Reiter and Pierre-Yves Schobbens. Special thanks to Mark Ryan for invaluable mathematical assistance. Thanks also to several anonymous referees for their helpful comments on this work. The author is supported by an SERC Advanced Research Fellowship.

References

1. A.B.Baker, A Simple Solution to the Yale Shooting Problem, in: Proceedings 1989 Knowledge Representation Conference, Toronto, Canada (1989) 11-20.

2. A.B.Baker, Nonmonotonic Reasoning in the Framework of the Situation Calculus, Artificial Intelligence 49 (1991) 5-23.

3. A.B.Baker and M.L.Ginsberg, Temporal Projection and Explanation, in: Proceedings IJCAI 89, Detroit, Michigan, USA (1989) 906-911.

4. J.M.Crawford and D.W.Etherington, Formalizing Reasoning about Change: A Qualitative Reasoning Approach, in: Proceedings AAAI 92, San Jose, California, USA (1992) 577-583.

5. E.Davis, Infinite Loops in Finite Time: Some Observations, in: Proceedings 1992 Knowledge Representation Conference, Boston, Massachusetts, USA (1992) 47-58.

6. M.Gelfond, V.Lifschitz and A.Rabinov, What Are the Limitations of the Situation Calculus? in: R.Boyer, ed., Essays for Bledsoe (Kluwer Academic, 1991) 167-179.

7. A.R.Haas, The Case for Domain-Specific Frame Axioms, in: Proceedings of the 1987 Workshop on the Frame Problem, Lawrence, Kansas, USA (1987) 343-348.

8. S.Hanks and D.McDermott, Nonmonotonic Logic and Temporal Projection, Artificial Intelligence 33 (1987) 379-412.

9. B.A.Haugh, Simple Causal Minimizations for Temporal Persistence and Projection, in: Proceedings AAAI 87, Seattle, Washington, USA (1987) 218-223.

10. H.Kautz, The Logic of Persistence, in: Proceedings AAAI 86, Philadelphia, Pennsylvania, USA (1986) 401-405.

11. R.A.Kowalski and M.J.Sergot, A Logic-Based Calculus of Events, New Generation Computing 4 (1986) 67-95.

12. V.Lifschitz, Pointwise Circumscription: Preliminary Report, in: Proceedings AAAI 86, Philadelphia, Pennsylvania, USA (1986) 406-410.

13. V.Lifschitz, Formal Theories of Action, in: Proceedings of the 1987 Workshop on the Frame Problem, Lawrence, Kansas, USA (1987) 35-57.

14. V.Lifschitz, Frames in the Space of Situations, Artificial Intelligence 46 (1990) 365-376.

15. V.Lifschitz, Toward a Metatheory of Action, in: Proceedings 1991 Knowledge Representation Conference, Boston, Massachisetts, USA (1991) 376-386.

16. V.Lifschitz, Circumscription, in: D.M.Gabbay, C.J.Hogger and J.A.Robinson, eds., The Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning (Oxford University Press, 1994) 297-352.

17. V.Lifschitz and A.Rabinov, Miracles in Formal Theories of Action, Artificial Intelligence 38 (1989) 225-237.

18. F.Lin and Y.Shoham, Provably Correct Theories of Action, in: Proceedings AAAI 91, Anaheim, California, USA (1991) 349-354.

19. F.Lin and Y.Shoham, Concurrent Actions in the Situation Calculus, in: Proceedings AAAI 92, San Jose, California, USA (1992) 590-595.

20. R.P.Loui, Response to Hanks and McDermott: Temporal Evolution of Beliefs and Beliefs about Temporal Evolution, Cognitive Science 11 (1987) 283-297.

21. J.McCarthy, Circumscription — A Form of Non-Monotonic Reasoning, Artificial Intelligence 13 (1980) 27-39.

22. J.McCarthy, Applications of Circumscription to Formalizing Common Sense Knowledge, Artificial Intelligence 26 (1986) 89-116.

23. J.McCarthy and P.J.Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in: D.Michie and B.Meltzer, eds., Machine Intelligence 4 (Edinburgh University Press, 1969) 463-502.

24. D.McDermott, We've Been Framed: Or Why AI Is Innocent of the Frame Problem, in: Z.W.Pylyshyn, ed., The Robots Dilemma: The Frame Problem in Artificial Intelligence (Ablex, 1987) 113-122.

25. R.S.Miller and M.P.Shanahan, Narratives in the Situation Calculus, to appear in: The Journal of Logic and Computation (1994).

26. L.Morgenstern and L.A.Stein, Why Things Go Wrong: A Formal Theory of Causal Reasoning, in: Proceedings AAAI 88, St.Paul, Minnesota, USA (1988) 518-523.

27. J.Pinto and R.Reiter, Temporal Reasoning in Logic Programming: A Case for the Situation Calculus, in: Proceedings of the Tenth International Conference on Logic Programming, Budapest, Hungary (1993) 203-221.

28. R.Reiter, The Frame Problem in the Situation Calculus: A Simple Solution (Sometimes) and a Completeness Result for Goal Regression, in: V.Lifschitz, ed., Artificial Intelligence and Mathematical Theory of Computation: Papers in Honour of John McCarthy (Academic Press, 1991) 359-380.

29. E.Sandewall, Combining Logic and Differential Equations for Describing Real-World Systems, in: Proceedings 1989 Knowledge Representation Conference, Toronto, Canada (1989) 412-420.

30. E.Sandewall, Filter Preferential Entailment for the Logic of Action in Almost Continuous Worlds, in: Proceedings IJCAI 89, Detroit, Michigan, USA (1989) 894-899.

31. E.Sandewall, Features and Fluents, Technical Report LiTH-IDA-R-92-30, Department of Computer and Information Science, Linköping University, Sweden (1992).

32. E.Sandewall, The Range of Applicability of Nonmonotonic Logics for the Inertia Problem, in: Proceedings IJCAI 93, Chambéry, France (1993) 738-743.

33. L.K.Schubert, Monotonic Solution of the Frame Problem in the Situation Calculus, in: H.Kyburg, R.Loui and G.Carlson, eds., Knowledge Representation and Defeasible Reasoning (Kluwer, 1989) 23-67.

34. M.P.Shanahan, Prediction Is Deduction but Explanation Is Abduction, in: Proceedings IJCAI 89, Detroit, Michigan, USA (1989) 1055-1060.

35. M.P.Shanahan, Representing Continuous Change in the Event Calculus, in: Proceedings ECAI 90, Stockholm, Sweden (1990) 598-603.

36. M.P.Shanahan, Explanation in the Situation Calculus, in: Proceedings IJCAI 93, Chambéry, France (1993) 160-165.

37. Y.Shoham, Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence (MIT Press, 1988).

Appendix A: Proof of Theorem 1

First we have some preliminary definitions.

Definition 4. A J-*indexed set* S is a set, every member of which is associated with exactly one element of the set J (its *index*). I will write S_j to denote the set of all s ∈ S such that j is the index of s.

Consider a many-sorted language $\mathcal{L}$ of first-order predicate calculus with a set $\mathcal{S}$ of sorts, an $\mathcal{S}^* \times \mathcal{S}$ -indexed set $\mathcal{F}$ of function symbols and an $\mathcal{S}^*$-indexed set $\mathcal{P}$ of predicate symbols.[18]

Definition 5. A *pre-interpretation* of $\mathcal{L}$ is a pair $\langle \mathcal{D},\mathrm{Fn} \rangle$, where $\mathcal{D}$ is an $\mathcal{S}$ -indexed set of objects, and for every n-ary function $f \in \mathcal{F}_{s_1,\ldots,s_n,s}$, Fn(f) is a mapping from $\mathcal{D}_{s_1} \times \ldots \times \mathcal{D}_{s_n}$ to $\mathcal{D}_s$.

Definition 6. A *partial interpretation* of $\mathcal{L}$ is a pair $\langle\langle \mathcal{D},\mathrm{Fn} \rangle,\mathrm{Pr}\rangle$, where $\langle \mathcal{D},\mathrm{Fn} \rangle$ is a pre-interpretation of $\mathcal{L}$, and for every n-ary predicate $P \in \mathcal{P}_{s_1,\ldots,s_n}$, Pr(P) is a partial mapping from $\mathcal{D}_{s_1} \times \ldots \times \mathcal{D}_{s_n}$ to TRUE or FALSE.

Definition 7. An *interpretation* of $\mathcal{L}$ is a partial interpretation $\langle \mathrm{Ip},\mathrm{Pr} \rangle$ of $\mathcal{L}$ such that for every predicate $P \in \mathcal{P}$, Pr(P) is a total mapping.

Definition 8. A *model* of a formula ψ of language $\mathcal{L}$ is an interpretation of $\mathcal{L}$ in which ψ is true, where truth in an interpretation is defined in terms of Pr and Fn in the standard way for predicate calculus.

Definition 9. If $\langle \mathrm{Ip},\mathrm{Pr}_1 \rangle$ is a partial interpretation of $\mathcal{L}$ and $\langle \mathrm{Ip},\mathrm{Pr}_2 \rangle$ is an interpretation of $\mathcal{L}$, then $\langle \mathrm{Ip},\mathrm{Pr}_2 \rangle$ *fills out* $\langle \mathrm{Ip},\mathrm{Pr}_1 \rangle$ for $\mathcal{L}$ if for all $P \in \mathcal{P}_{s_1,\ldots,s_n}$ and for all $X_i \in \mathcal{D}_{s_i}$, $\mathrm{Pr}_1(P)(X_1,\ldots,X_n) = V$ implies $\mathrm{Pr}_2(P)(X_1,\ldots,X_n) = V$.

Definition 10. A formula of $\mathcal{L}$ has the *truth value* TRUE in a partial interpretation I of $\mathcal{L}$ if it is true in every interpretation that fills out I for $\mathcal{L}$, and has the *truth value* FALSE in I if it is false in every interpretation that fills out I for $\mathcal{L}$. Otherwise, its truth value is not defined in I.

Definition 11. The n-ary predicate $P \in \mathcal{P}_{s_1,\ldots,s_n}$ is *uninterpreted* in a partial interpretation I of $\mathcal{L}$ if for all $X_i \in \mathcal{D}_{s_i}$, the truth value of $P(X_1,\ldots,X_n)$ is not defined in

I. Conversely, the predicate P is *fully interpreted* in I if for all $X_i \in \mathcal{D}_{s_i}$, the truth value of $P(X_1,...,X_n)$ is defined in I.

Now let the set $\mathcal{S}$ of sorts include the sorts $s_1...s_t...s_m$, where s_t is the sort of time points. Let the set $\mathcal{P}$ of predicates include $< \in \mathcal{P}_{s_t,s_t}$ and $q \in \mathcal{P}_{s_1,...,s_t,...,s_m}$. Consider only interpretations of $\mathcal{L}$ in which the time points are interpreted by the natural numbers, that is $\mathcal{D}_{s_t} = \mathbb{N}$, and in which < has its usual meaning. Let $Ip = \langle \mathcal{D},Fn \rangle$ be a pre-interpretation of $\mathcal{L}$. Let $\langle Ip,Pr \rangle$ be a partial interpretation of $\mathcal{L}$ in which q is uninterpreted, but in which every other predicate in $\mathcal{P}$ is fully interpreted. Finally, let $\psi = \forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$ be a formula of $\mathcal{L}$ which is chronological in argument t.

Definition 12. A q-*map* is a partial mapping from $\mathcal{D}_{s_1} \times ... \times \mathcal{D}_{s_m}$ to TRUE or FALSE.

A q-map can be used to extend an interpretation in which q is uninterpreted to one in which it is fully interpreted.

Definition 13. A q-map M is *total below* $T \in \mathbb{N}$ if for all $X_i \in \mathcal{D}_{s_i}$ where $X_t < T$, $M(X_1,...,X_m)$ is defined.

The proof of the theorem relies on several lemmas, which I will now prove.

Definition 14. Let M be a q-map and $P \in \mathcal{P}$ be an n-ary predicate. The function $\oplus$ is defined as follows.

$$(Pr \oplus M)(P) = \begin{cases} M \text{ if P is q} \\ Pr(P) \text{ otherwise} \end{cases}$$

Lemma 1. For any $T \in \mathbb{N}$, and any q-map M which is total below T, for all $X_i \in \mathcal{D}_{s_i}$ where $X_t \leq T$, the truth value of $\phi(X_1,...,X_m)$ is defined in $\langle Ip,Pr \oplus M \rangle$.

Proof. Since ψ is chronological, the only occurrences of q in $\phi(X_1,...,X_m)$ are in conjunctions of the form $q(Z_1,...,Z_m) \wedge Z_t < X_t$. If M is total below T and $X_t \leq T$, then the truth value of $q(Z_1,...,Z_m) \wedge Z_t < X_t$ is defined in $\langle Ip,Pr \oplus M \rangle$. Consequently, the truth value of $\phi(X_1,...,X_m)$ is defined in $\langle Ip,Pr \oplus M \rangle$, because $Pr \oplus M$ is a total mapping for all predicates occurring in ϕ apart from q. □

Definition 15. The q-map M_α is defined for any $\alpha \in \mathbb{N}$ as follows.

$$M_\alpha(X_1,...,X_m) = \begin{cases} V \text{ if } \alpha > 0 \text{ and } \phi(X_1,...,X_m) \text{ has truth value V in } \langle Ip,Pr \oplus M_{\alpha-1} \rangle \\ \text{undefined otherwise} \end{cases}$$

Note that M_0 is undefined everywhere. The mapping M_α will facilitate the definition of a q-map that extends a partial interpretation in which q is uninterpreted to a model of the chronological formula ψ. First I will prove some properties of M_α which will be required later.

Lemma 2. For any $\alpha \in \mathbb{N}$, M_α is total below α.

Proof. The proof is by induction. Clearly the proposition is true for the base case where α is 0. For the inductive case, consider any β and suppose that M_β is total below β. Then from Lemma 1, for all $X_i \in \mathcal{D}_{s_i}$ where $X_t \leq \beta$, the truth value of $\phi(X_1,...,X_m)$ is defined in $\langle Ip,Pr \oplus M_\beta \rangle$. So from Definition 13, $M_{\beta+1}$ is total below β+1. □

Lemma 3. For any $\alpha \in \mathbb{N}$ and all $X_i \in \mathcal{D}_{s_i}$, if $M_\alpha(X_1,\ldots,X_m) = V$ then $M_{\alpha+1}(X_1,\ldots,X_m) = V$.

Proof. The proof is by induction. Clearly the proposition is true for the base case where α is 0. For the inductive case, consider any β, and suppose that for all $X_i \in \mathcal{D}_{s_i}$, if $M_\beta(X_1,\ldots,X_m) = V$ then $M_{\beta+1}(X_1,\ldots,X_m) = V$. From this hypothesis, for all $X_i \in \mathcal{D}_{s_i}$, if the truth value of $\phi(X_1,\ldots,X_m)$ is defined in $\langle Ip,Pr \oplus M_\beta\rangle$, then it has the same truth value in $\langle Ip,Pr \oplus M_{\beta+1}\rangle$. Therefore, from Definition 15, if $M_{\beta+1}(X_1,\ldots,X_m) = V$ then $M_{\beta+2}(X_1,\ldots,X_m) = V$. □

Note that Lemma 3 doesn't imply that M_α is the same as $M_{\alpha+1}$, because M_α is only a partial interpretation.

Lemma 4. For any $\alpha \in \mathbb{N}$, and all $X_i \in \mathcal{D}_{s_i}$ where $X_t < \alpha$, if $q(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\alpha\rangle$ then $\phi(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\alpha\rangle$.

Proof. Clearly, if $q(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\alpha\rangle$ then $M_\alpha(X_1,\ldots,X_m) = V$, because q is uninterpreted in $\langle Ip,Pr\rangle$. But if $M_\alpha(X_1,\ldots,X_m) = V$ then, from Definition 15, $\phi(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_{\alpha-1}\rangle$. Therefore, from Lemma 3, $\phi(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\alpha\rangle$. □

Definition 16. The q-map M_ω is defined as follows.

$$M_\omega(X_1,\ldots,X_m) = V \text{ if there is some } \alpha \in \mathbb{N} \text{ such that } M_\alpha(X_1,\ldots,X_m) = V$$

Lemma 3 ensures that M_ω is a well-defined function. The mapping M_ω extends a partial interpretation in which q is uninterpreted to a model of the chronological formula ψ, as I will now show.

Lemma 5. M_ω is a total mapping from $\mathcal{D}_{s_1}\times\ldots\times\mathcal{D}_{s_m}$ to TRUE or FALSE.

Proof. Consider any $T \in \mathbb{N}$. From Lemma 2, M_{T+1} is total below T+1. That is, for all $X_i \in \mathcal{D}_{s_i}$ where $X_t \leq T$, $M_{T+1}(X_1,\ldots,X_m)$ is defined. From Definition 16, for all $X_i \in \mathcal{D}_{s_i}$, if $M_{T+1}(X_1,\ldots,X_m)$ is defined then $M_\omega(X_1,\ldots,X_m)$ is defined. Therefore M_ω is a total mapping from $\mathcal{D}_{s_1}\times\ldots\times\mathcal{D}_{s_m}$ to TRUE or FALSE. □

Lemma 6. For all $X_i \in \mathcal{D}_{s_i}$, if $q(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\omega\rangle$ then $\phi(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\omega\rangle$.

Proof. Since q is uninterpreted in $\langle Ip,Pr\rangle$, if $q(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\omega\rangle$ then $M_\omega(X_1,\ldots,X_m) = V$. From Definition 16, this implies that there is some α such that $M_\alpha(X_1,\ldots,X_m) = V$. Then from Definition 14 and Lemma 4, $\phi(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\alpha\rangle$. From Definitions 14 and 16, this implies that $\phi(X_1,\ldots,X_m)$ has truth value V in $\langle Ip,Pr \oplus M_\omega\rangle$. □

Lemma 7. $\langle Ip,Pr \oplus M_\omega\rangle$ is an interpretation of $\mathcal{L}$.

Proof. From Lemma 5, M_ω is a total mapping for q. Therefore, since Pr(P) is a total mapping for every predicate $P \in \mathcal{P}$ except q, $(Pr \oplus M_\omega)(P)$ is a total mapping for every predicate $P \in \mathcal{P}$. □

Lemma 8. $\langle Ip,Pr \oplus M_\omega\rangle$ is a model of $\psi = \forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$.

Proof. The proposition follows directly from Lemmas 6 and 7. □

Now, I can show that circumscribing the conjunction of a formula with a chronological formula is easy, under the right conditions. Let $\mathcal{L}$- be the same language as $\mathcal{L}$, but without the predicate q. Let χ be a formula of $\mathcal{L}$-. Note that since $\langle Ip,Pr \rangle$ is a partial interpretation of $\mathcal{L}$, it is also a partial interpretation of $\mathcal{L}$-, and recall that q is uninterpreted in $\langle Ip,Pr \rangle$.

Lemma 9. If $\langle Ip,Pr \rangle$ is a model of χ, then $\langle Ip,Pr \oplus M_\omega \rangle$ is a model of $\chi \wedge \psi$.

Proof. From Lemma 8, $\langle Ip,Pr \oplus M_\omega \rangle$ is a model of ψ. If $\langle Ip,Pr \rangle$ is a model of χ then, since χ doesn't mention q, $\langle Ip,Pr \oplus M_\omega \rangle$ is also a model of χ. □

Let P* and Q* be subsets of $\mathcal{P}$ such that Q* includes q.

Lemma 10. All models of CIRC[$\chi \wedge \psi$; P*; Q*] have the form $\langle Ip,Pr' \oplus M \rangle$ where $\langle Ip,Pr' \rangle$ is a model of CIRC[χ ; P*; Q*], and M is a q-map.

Proof. Any model $\langle Ip,Pr'' \rangle$ of CIRC[$\chi \wedge \psi$; P*; Q*] can be written in the form $\langle Ip,Pr' \oplus M \rangle$, where M is the q-map Pr''(q), q is uninterpreted in $\langle Ip,Pr' \rangle$, and Pr'(P) is Pr''(P) for every predicate P $\in \mathcal{P}$ except q. Then it remains to show that $\langle Ip,Pr' \rangle$ is a model of CIRC[χ ; P*; Q*].

Suppose $\langle Ip,Pr' \rangle$ is not a model of CIRC[χ ; P*; Q*]. Clearly $\langle Ip,Pr' \rangle$ is a model of χ, since $\langle Ip,Pr' \oplus M \rangle$ is a model of χ and M interprets only q which isn't mentioned in χ. So there must be some model $\langle Ip,Pr''' \rangle$ of χ which is smaller in P*, with Q* allowed to vary, than $\langle Ip,Pr' \rangle$. Then, from Lemma 9, $\langle Ip,Pr''' \oplus M_\omega \rangle$ is a model of $\chi \wedge \psi$. With q allowed to vary, $\langle Ip,Pr''' \oplus M_\omega \rangle$ must be smaller in P* than $\langle Ip,Pr' \oplus M \rangle$, because $(Pr''' \oplus M_\omega)(P)$ = Pr'''(P) for every predicate P $\in \mathcal{P}$ except q. Therefore $\langle Ip,Pr' \oplus M \rangle$ is not a model of CIRC[$\chi \wedge \psi$; P*; Q*], which is a contradiction. □

Lemma 11. Every formula which is true in all models of CIRC[χ ; P*; Q*] is also true in all models of CIRC[$\chi \wedge \psi$; P*; Q*].

Proof. Consider any model of CIRC[$\chi \wedge \psi$; P*; Q*]. From Lemma 10, it has the form $\langle Ip,Pr' \oplus M \rangle$ where $\langle Ip,Pr' \rangle$ is a model of CIRC[χ ; P*; Q*], and M is a q-map. The theorem then follows from the fact that any formula which is true in $\langle Ip,Pr' \rangle$ is also true in $\langle Ip,Pr' \oplus M \rangle$. □

Theorem 1 follows from directly from Lemma 11.

Appendix B: Temporal Projection Algorithms

Theorem 5, presented here, facilitates the construction of algorithms for temporal projection. Recall Definitions 4 to 14 and the corresponding assumptions from Appendix A.

Definition 17. The q-map C_α is defined for any $\alpha \in \mathbb{N}$ as follows.

$$C_\alpha(X_1,\ldots,X_m) = \begin{cases} V \text{ if } \alpha > 0\text{, and for all models } \langle Ip,Pr' \rangle \text{ of CIRC}[\chi\text{ ; P*; Q*}], \\ \quad \phi(X_1,\ldots,X_m) \text{ has truth value V in } \langle Ip,Pr' \oplus C_{\alpha-1} \rangle \\ \text{undefined otherwise} \end{cases}$$

Lemma 12. For any $\alpha \in \mathbb{N}$ and for all $X_i \in \mathcal{D}_{s_i}$, if $C_\alpha(X_1,\ldots,X_m)$ = V then $q(X_1,\ldots,X_m)$ has truth value V in all models of CIRC[$\chi \wedge \psi$; P*; Q*].

Proof. The proof is by induction. The proposition is trivially true for the base case where $\alpha = 0$. Consider any $\beta > 0$. Suppose $C_\beta(X_1,...,X_m) = V$ implies $q(X_1,...,X_m)$ has truth value V in all models of CIRC[$\chi \wedge \psi$; P*; Q*].

From Definition 17, $C_{\beta+1}(X_1,...,X_m) = V$ implies $\phi(X_1,...,X_m)$ has truth value V in every partial interpretation of the form $\langle Ip,Pr' \oplus C_\beta\rangle$, where $\langle Ip,Pr'\rangle$ is a model of CIRC[χ ; P*; Q*], which by definition means it has truth value V in every interpretation that fills out such a partial interpretation. But every model of CIRC[$\chi \wedge \psi$; P*; Q*] fills out such a partial interpretation (see below), so $\phi(X_1,...,X_m)$ has truth value V in all models of CIRC[$\chi \wedge \psi$; P*; Q*]. If $\phi(X_1,...,X_m)$ has truth value V in all models of CIRC[$\chi \wedge \psi$; P*; Q*] then so does $q(X_1,...,X_m)$.

To see that every model of CIRC[$\chi \wedge \psi$; P*; Q*] fills out a partial interpretation of the form $\langle Ip,Pr' \oplus C_\beta\rangle$, where $\langle Ip,Pr'\rangle$ is a model of CIRC[χ ; P*; Q*], recall Lemma 10. From Lemma 10, every model of CIRC[$\chi \wedge \psi$; P*; Q*] has the form $\langle Ip,Pr' \oplus M\rangle$ where $\langle Ip,Pr'\rangle$ is a model of CIRC[χ ; P*; Q*] and M is a q-map. From the induction hypothesis, $C_\beta(X_1,...,X_m) = V$ implies $M(X_1,...,X_m) = V$. So $\langle Ip,Pr' \oplus M\rangle$ fills out $\langle Ip,Pr' \oplus C_\beta\rangle$. □

Lemma 13. For any $\alpha \in \mathbb{N}$, C_α is total below α.

Proof. The proof is the same as for Lemma 2. □

Theorem 5. For any $\alpha \in \mathbb{N}$ and all $X_i \in \mathcal{D}_{s_i}$ where $X_t < \alpha$, CIRC[$\chi \wedge \psi$; P*; Q*] $\models$ $q(X_1,...,X_m)$ if and only if $C_\alpha(X_1,...,X_m) = TRUE$.

Proof. The lemma follows directly from Lemmas 12 and 13. □

Theorem 5 facilitates the construction of algorithms for deciding whether $q(X_1,...,X_m)$ is true (or false) in all models of CIRC[$\chi \wedge \psi$; P*; Q*], given any $X_i \in \mathcal{D}_{s_i}$. Let α be X_t. An algorithm built according to the following schema will compute the set S_P of all tuples $\langle Y_1,...,Y_m\rangle$ such that $C_\alpha(Y_1,...,Y_m) = TRUE$, and the set S_N of all tuples $\langle Y_1,...,Y_m\rangle$ such that $C_\alpha(Y_1,...,Y_m) = FALSE$. From Theorem 5, to check whether $q(X_1,...,X_m)$ is true (or false) in all models of CIRC[$\chi \wedge \psi$; P*; Q*], it is simply necessary to check for membership of S_P (or S_N).

$S_P := \{\}$
$S_N := \{\}$
FOR $Y_t := 0$ TO α
    $S_P := S_P \cup \{\langle Y_1,...,Y_m\rangle \mid TRUE(Y_1,...,Y_m)\}$
    $S_N := S_N \cup \{\langle Y_1,...,Y_m\rangle \mid FALSE(Y_1,...,Y_m)\}$

$TRUE(Y_1,...,Y_m)$ is shorthand for "$\phi(Y_1,...,Y_m)$ is true in all models of CIRC[χ ; P*; Q*]", and $FALSE(Y_1,...,Y_m)$ is shorthand for "$\phi(Y_1,...,Y_m)$ is false in all models of CIRC[χ ; P*; Q*]". Another algorithm is required for computing the sets $\{\langle Y_1,...,Y_m\rangle \mid TRUE(Y_1,...,Y_m)\}$ and $\{\langle Y_1,...,Y_m\rangle \mid FALSE(Y_1,...,Y_m)\}$ at each iteration. The details of this second algorithm depend on χ and ϕ, but it can exploit the fact that all occurrences of q in $\phi(Y_1,...,Y_m)$ are in conjunctions of the form $q(Z_1,...,Z_m) \wedge Z_t < Y_t$. On any given iteration of the algorithm, a conjunction of that form is true in all models of CIRC[χ ; P*; Q*] if and only if $\langle Z_1,...,Z_m\rangle \in S_P$, and is false if and only if $\langle Z_1,...,Z_m\rangle \in S_N$. In the case of (E3), since the definitions of Initiated and Terminated are clausal, resolution theorem proving techniques could be used to compute the required sets, and negation-as-failure could be used to minimise Horn clause fragments of the domain and history formulae.

There is nothing surprising about this algorithm schema, of course. It simply works forwards in time in the way we might expect. The purpose of Theorem 5 is to endorse the use of the obvious algorithm. Effectively, Theorem 5 allows us to forget about ψ computationally, in the same way that Theorems 1 and 4 allow us to forget about it from the point of view of minimisation.

Appendix C: Proof of Theorem 4

Consider the same language $\mathcal{L}$ as in Appendix A. Consider only interpretations of $\mathcal{L}$ in which time points are interpreted by the reals, that is $\mathcal{D}_{S_t} = \mathbb{R}$, and in which < has its usual meaning for real numbers. Let $\psi = \forall \underline{x}\ q(\underline{x}) \leftrightarrow \phi(\underline{x})$ be a formula of $\mathcal{L}$ which is real-chronological in argument t with respect to a formula χ of $\mathcal{L}$ and a marker set S. Let $Ip = \langle \mathcal{D}, Fn \rangle$ be a pre-interpretation of $\mathcal{L}$. Let $\langle Ip, Pr \rangle$ be a partial interpretation of $\mathcal{L}$ in which q is uninterpreted, but in which every other predicate in $\mathcal{P}$ is fully interpreted. Recall the definitions of a q-map, of M_α, and of M_ω from Appendix A.

Definition 18. A q-map M is *total up to* $T \in \mathbb{R}$ if for all $X_i \in \mathcal{D}_{S_i}$ where $X_t \leq T$, $M(X_1,...,X_m)$ is defined.

Lemma 14. For any $\alpha \in \mathbb{N}$ where $\alpha > 0$, if S has α or more elements then M_α is total up to any $T \in \mathbb{R}$ such that $T \leq \alpha^{th}$ element of S.

Proof. The proof is by induction. Clearly the proposition is true for the base case where α is 1. For the inductive case, consider any β, and suppose that M_β is total up to any $T \in \mathbb{R}$ such that $T \leq \beta^{th}$ element of S. Then, from the definition of M_α, we have to show that for all $X_i \in \mathcal{D}_{S_i}$ where $X_t \leq \beta{+}1^{th}$ element of S, the truth value of $\phi(X_1,...,X_m)$ is defined in $\langle Ip, Pr \oplus M_\beta \rangle$.

By hypothesis, $\phi(X_1,...,X_m)$ is defined for all $X_i \in \mathcal{D}_{S_i}$ where $X_t \leq \beta^{th}$ element of S. So consider any $X_t \leq \beta{+}1^{th}$ element of S such that $X_t > \beta^{th}$ element of S. Now, since ψ is real-chronological, it can be seen that the only occurrences of q in $\phi(X_1,...,X_m)$ are in conjunctions of the form $q(\underline{Z}) \wedge Z_t < X_t \wedge \theta$ which are false. To see this, consider Z_t. Clearly the conjunction is false if $Z_t \geq X_t$. On the other hand, if $Z_t < X_t$ then, since X_t is strictly between the β^{th} and $\beta{+}1^{th}$ elements of S, Z_t is not in S, and θ is therefore false from the definition of real-chronological. □

Lemma 15. For any $\alpha \in \mathbb{N}$, if S has fewer than α elements then M_α is total up to any $T \in \mathbb{R}$.

Proof. From Lemma 14, M_α is total up to the last element of S. So, from the definition of M_α, we have to show that for all $X_i \in \mathcal{D}_{S_i}$ where $X_t >$ the last element of S, the truth value of $\phi(X_1,...,X_m)$ is defined in $\langle Ip, Pr \oplus M_{\alpha-1} \rangle$. From the definition of real-chronological, the only occurrences of q in $\phi(X_1,...,X_m)$ are in conjunctions of the form $q(\underline{z}) \wedge \theta$, where θ is false if $X_t >$ the last element of S. Therefore the truth value of $\phi(X_1,...,X_m)$ is always defined if $X_t >$ the last element of S. So $\phi(X_1,...,X_m)$ is defined for any $X_i \in \mathcal{D}_{S_i}$, and M_α is total up to any $T \in \mathbb{R}$. □

The proofs of Lemmas 3 and 4 are unchanged from Appendix A.

Lemma 16. M_ω is a total mapping from $\mathcal{D}_{S_1} \times ... \times \mathcal{D}_{S_m}$ to TRUE or FALSE.

Proof. It is easy to show that for all $X_i \in \mathcal{D}_{S_i}$ there exists some α such that $M_\alpha(X_1,...,X_m)$ is defined. To see this, consider X_t. If there is no $T \in$ S before X_t, then $M_1(X_1,...,X_m)$ is defined, from Lemma 15. Otherwise, if there is no $T \in$ S after X_t, then from Lemma 14, $M_\alpha(X_1,...,X_m)$ is defined, where α is the number of elements in

S. Finally, if there is some $T \in S$ after or equal to X_t and some $T' \in S$ before X_t, then from Lemma 14, $M_\alpha(X_1,...,X_m)$ is defined, where the smallest $T \in S$ after or equal to X_t is the α^{th} element of S. The fact that S is a marker set ensures that these three cases are exhaustive. From the definition of M_ω, if $M_\alpha(X_1,...,X_m)$ is defined then $M_\omega(X_1,...,X_m)$ is defined. Therefore M_ω is a total mapping from $\mathcal{D}_{s_1} \times ... \times \mathcal{D}_{s_m}$ to TRUE or FALSE. □

The rest of the proof is the same as the proof of Theorem 1 in Appendix A, yielding Theorem 4 in the same way, but under the assumption that ψ is real-chronological in some argument with respect to the formula χ and some marker set.

[1] A particularly clear description of the frame problem is to be found in McDermott [24].

[2] Hanks and McDermott's celebrated example has become known as the Yale shooting scenario. Some authors, such as Loui [20], questioned the assumptions underlying Hanks and McDermott's analysis of this example. The task I set myself in the present paper is to address the frame problem. I will use the Yale shooting scenario simply as an example of where the frame problem arises. I am not interested in representing the Yale shooting scenario *per se*.

[3] I will not distinguish actions from events in this paper.

[4] A number of authors, notably Haas [7], Schubert [33], and Reiter [28], have proposed monotonic solutions to the frame problem for the situation calculus. Arguably, these solutions do not offer "elaboration-tolerance" (John McCarthy's term), because the acquisition of new knowledge about the domain necessitates the complete reconstruction of the domain theory. The present paper supplies an elaboration-tolerant solution, in which new knowledge is automatically absorbed into the existing theory.

[5] In what follows, I will use the term "event calculus" to refer to the formalism presented here, not to Kowalski and Sergot's formalism.

[6] This separation is also the basis of the more recent approach of Crawford and Etherington [4].

[7] Throughout the paper, all variables are universally quantified unless otherwise indicated.

[8] In fact, this is not quite true. The improper use of domain constraints can lead to time points for which no corresponding state exists. Thus, State is a predicate rather than a function.

[9] Because Initiated, Terminated, Clipped and Declipped don't have the status of predicates, we don't have to worry about them when designing circumscription policies.

[10] I will make several assumptions like this in the paper. The legitimacy of these manoeuvres depends on our ability, in principle, to fill the resulting gap between model theory and proof theory to our satisfaction. For example, a second-order axiomatisation of the "<" predicate could be provided. From the point of view of computation, the existence of numerical algorithms and the possibility of procedural attachment is a source of comfort.

[11] Blocks and locations should really have different sorts, but this would complicate the example.

[12] The axioms of the event calculus do not depend on this assumption, and in Section 9 on continuous change, I will consider the case in which time points are interpreted by the reals.

[13] This variant of the Yale shooting scenario is sometimes called the "walking turkey shoot."

[14] Personal communication.

[15] Axioms for the associativity and commutativity of the "+" operator are unnecessary.

[16] The set of fluents is now as large as the reals. However, note that Axioms (S1) to (S3) only insist on the existence of states with countably many fluents. In practice, only states with finitely many fluents will usually arise.

[17] This was not true in the original event calculus of Kowalski and Sergot [11].

[18] $\mathcal{S}^*$ is the set of all finite tuples of elements of $\mathcal{S}$.