Representing Continuous Change in The Event Calculus

Murray Shanahan

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

June 1989

Revised January 1990

Abstract

The Event Calculus of Kowalski and Sergot only deals with discrete change. This paper introduces a simplified version of the Event Calculus and extends it to deal with continuous change, as in the height of a falling object or the level of liquid in a filling vessel. The idea of autotermination is introduced. A period of continuous change autoterminates if it brings about the event which terminates it. For example, when the increasing level of water in a sink reaches the overflow, it ceases to increase. The formulation is applied to a simple example with liquid filling a sink, and to a more complicated one with many tanks discharging liquid into another tank.

In Proceedings ECAI 90, pages 598-603

Introduction

The Situation Calculus of McCarthy and Hayes [1969] represents a changing world by a discrete and strictly ordered sequence of "snapshots", each representing the complete state of the world at a given instant. The ontology of the Situation Calculus makes it hard to represent partially ordered or simultaneous events, or continuous change. To overcome these shortcomings, several authors have developed richer calculi for representing change [McDermott, 1982, Hayes, 1985, Allen, 1984, Kowalski and Sergot, 1986].

The Event Calculus of Kowalski and Sergot [1986] is one such formalism. In the Event Calculus, events initiate periods during which properties hold. A property which has been initiated continues to hold by default until some event occurs which terminates it. Partially ordered and simultaneous events can be represented. However, in the Event Calculus described by Kowalski and Sergot, all change is discrete. It needs to be extended to be able to represent continuously changing quantities, such as the height of a falling object, the position of a rotating wheel or the level of liquid in a filling vessel. This paper presents such an extension.

I introduce a simplified version of the Event Calculus, and then present an extension which copes with continuous change. The idea of *autotermination* is introduced — if a period of continuous change carries on for long enough, it can cause the event which terminates it, as when a sink overflows or a falling object hits the ground. The extended Event Calculus is then applied to a number of examples of continuous change involving vessels filling and emptying.

1. The Event Calculus

In Kowalski and Sergot's Event Calculus [Kowalski and Sergot, 1986] and its variants [Kowalski, 1989], the ontological primitives are *events*, which initiate and terminate periods during which *properties* hold. The Horn clause subset of the Predicate Calculus is used, augmented with negation-as-failure. The Event Calculus used in this paper is a simplified version of that given by Kowalski and Sergot in their paper. Only two clauses are necessary, as follows.

holds-at(P,T2) if (1.1)
happens(E) and time(E,T1) and T1 < T2 and
initiates(E,P) and not clipped(T1,P,T2)

clipped(T1,P,T3) if (1.2)
happens(E') and time(E',T2) and terminates(E',P) and
T1 ≤ T2 and T2 < T3

The formula *holds-at(P,T)* represents that property *P* holds at time *T*. The formula *happens(E)* represents that the event *E* occurs and the formula *time(E,T)* represents that the time of event *E* is *T*. Times are ordered by the usual comparative operators. The formula *initiates(E,P)* represents that the event *E* initiates a period during which property *P* holds, and *terminates(E,P)* represents that the event *E* terminates any ongoing period during which property *P* holds. The *not* operator is interpreted as negation-as-failure. The use of negation-as-failure in Axiom (1.1) gives a form of default persistence. The formula *clipped(T1,P,T2)* represents that the property *P* ceases to hold at some time between *T1* and *T2*. Note that a property does not hold at the time of the event which initiates it, but does hold at the time of the event which terminates it.

The problem domain is (partly) captured by a set of *initiates* and *terminates* clauses. For example, the Blocks World is described by the following clauses. The term *on(X,Y)* names the property that block *X* is on top of block *Y* or at location *Y*, and the term *clear(X)* names the property that block or location *X* has nothing on top of it. The term *move(X,Y)* names the event or act type of moving block *X* onto block or location *Y*.

initiates(E,on(X,Y)) if act(E,move(X,Y)) (2.1)

$$initiates(E,clear(Z)) \text{ if } act(E,move(X,Y)) \text{ and } time(E,T) \text{ and } holds\text{-}at(on(X,Z),T) \text{ and } Z \neq Y \quad (2.2)$$

$$terminates(E,clear(Y)) \text{ if } act(E,move(X,Y)) \quad (2.3)$$

$$terminates(E,on(X,Z)) \text{ if } act(E,move(X,Y)) \text{ and } Z \neq Y \quad (2.4)$$

A particular course of events is represented by a set of *happens*, *act* and temporal ordering clauses. For example, to represent that block *a* was moved to location *x* and then to location *y*, we write

$$happens(e1) \quad (3.1)$$

$$act(e1,move(a,x)) \quad (3.2)$$

$$happens(e2) \quad (3.3)$$

$$act(e2,move(a,y) \quad (3.4)$$

$$e1 > e2 \quad (3.5)$$

We can use these axioms deductively to predict the locations of blocks at different times. We can also use the Blocks World Axioms (2.1) to (2.4) abductively to generate explanations of the locations of blocks in terms of possible events [Shanahan, 1989].

2. Events Which Are Caused

In the simplified form of the Event Calculus described above, like the original Kowalski and Sergot formalism, it is not clear how to represent that a certain type of event invariably follows a certain other type of event, or that a certain type of event occurs when some property holds. We will need to be able to do this to represent autoterminating periods of continuous change. In the full Predicate Calculus, the natural way to represent this sort of law is to use an existentially quantified conjunction as the consequent of an implication. For example, we might write something of the form

$$\forall E_1[\exists E_2[happens(E_2) \wedge P_1 ... P_n] \leftarrow [happens(E_1) \wedge Q_1 ... Q_m]] \quad (4.1)$$

Using only Horn clauses plus negation-as-failure, it is necessary to eliminate the existential quantifier by skolemisation. We get a clause of the form

$$happens(f(E_1,\underline{X})) \text{ if } happens(E_1) \text{ and } Q_1 \ldots Q_m \quad (4.2)$$

and *n* clauses of the form

$$P_i \text{ if } happens(E_1) \text{ and } Q_1 \ldots Q_m \quad (4.3)$$

The skolem function $f(E_1,\underline{X})$, where $\underline{X}$ is a tuple of all universally quantified variables in (4.1), names an event caused by E_1. In fact, the event can often be uniquely named without including all these variables in $\underline{X}$. The function name *f* identifies the type of the event.

Suppose, for example, that we wanted to represent a law that says whenever an alarm is set at time *t* it goes off at time *t+n*. The following axioms will suffice.

$$happens(go\text{-}off(E)) \text{ if } happens(E) \text{ and } act(E,set\text{-}alarm) \quad (4.5)$$

$$time(go\text{-}off(E),T2) \text{ if } time(E,T1) \text{ and } T2 = T1 + n \quad (4.6)$$

3. Continuous Change

The Event Calculus provides a representational framework for a variety of temporal reasoning problems. Problems involving discrete changes, such as Blocks World problems, are easily represented. But many temporal reasoning problems, especially in the domain of commonsense physics, demand an ability to represent continuous variation in some quantity, and this is not provided by the Event Calculus as it stands. For example, we may wish to represent the changing height of a falling object, or the rotation of a cog or wheel, or the motion of a billiard ball across a table.

A continuously varying quantity takes on values from some *quantity space*. The quantity space in question could be many-dimensioned, such as the position in physical space of a moving object, or it could be one-dimensional, such as the height of a falling object, the angular position of a wheel or the level of liquid in a vessel. A quantity space which is in reality a continuum can be represented as such, that is *quantitatively*, or alternatively can be represented *qualitatively* as a discrete, ordered set of values of interest.

The approach I will take is to axiomatise a particular example first — that of the rising level of water in a kitchen sink — and then to consider some axioms for continuous change in general. I will consider both the quantitative and qualitative representations of the quantity space. Finally I will axiomatise a more complicated example involving a number of tanks discharging into each other.

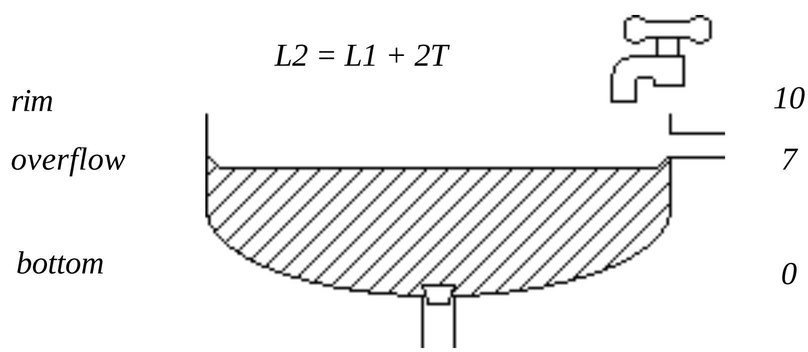

Figure 1: A Kitchen Sink

Consider the kitchen sink in Figure 1. When the tap is turned on, it starts to fill up. The level of water increases continuously until either the tap is turned off or an outlet is reached. The task is to represent this scenario in the Event Calculus. We might consider representing the level of water either quantitatively or qualitatively, depending on our purposes and the information at hand.

If we choose to represent it quantitatively, then the level of water is a real number $L2$, and it is given by the equation $L2 = L1 + R * T$, where $L1$ is the level of water before the tap is turned on, R is the rate of filling and T is the time elapsed since the tap was turned on. The bottom is at level 0, the overflow is at level 7, and the rim is at level 10. We will assume that the rate of filling R is 2 units of length per unit of time.

For various reasons, we might choose to represent the level qualitatively, perhaps because an exact equation of filling is not known or perhaps because we want to represent and reason at a higher level of abstraction than exact numerical values. We can represent the level as simply one of three values; the bottom, the overflow and the rim. This set of values represents all the points at which anything interesting can happen, such as the sink overflowing. Rather than having a quantitative equation characterising the increase in the level, it simply passes through each of these points in turn.

4. Representing the Kitchen Sink

I will consider the quantitative case first. To capture the continuous variation in the level of water once the tap has been turned on requires a new *holds-at* axiom, since Axiom (1.1) only caters for discretely changing properties. Let us suppose that *level(L)* names the property that the water is at level *L*. We can write

holds-at(level(L2),T2) if (5.1)
happens(E) and act(E,tap-on) and time(E,T1) and
*holds-at(level(L1),T1) and L2 = L1 + 2 * (T2 - T1)*

Using this axiom, we can infer the level of the sink at any time while it is filling. But the sink can stop filling either because the tap is turned off or because it overflows. To represent this, we can introduce another property *filling*, which represents that the level of water in the sink is increasing. This property is initiated when the tap is turned on and terminated when it is turned off or when the sink overflows. Whilst *filling* holds, the level of water is a function of the time that has elapsed since the tap was turned on. Replacing (5.1) we have

holds-at(level(L2),T2) if (6.1)
happens(E) and act(E,tap-on) and time(E,T1) and
*holds-at(level(L1),T1) and L2 = L1 + 2 * (T2 - T1) and*
not clipped(T1,filling,T2)

initiates(E,filling) if act(E,tap-on) (6.2)
terminates(E,filling) if act(E,tap-off) (6.3)
terminates(E,filling) if act(E,spill) (6.4)

Finally, we need to add axioms which describe how a spill event can occur. These *autotermination* axioms are very common with problems of continuous change. They capture the fact that, if left alone, the period of filling terminates itself. In general, it represents that if the continuously changing quantity reaches a threshold value, an event occurs which terminates the period of continuous change. Using the convention described in Section 2, for this example we write

happens(spills(E)) if (6.5)
happens(E) and time(E,T1) and initiates(E,filling) and
holds-at(level(L1),T1) and outlet(L2) and
*L2 = L1+ R * (T2 - T1) and T1 < T2 and*
not clipped(T1,filling,T2)

time(spills(E),T2) if (6.6)
happens(E) and time(E,T1) and initiates(E,filling) and
holds-at(level(L1),T1) and outlet(L2) and
*L2 = L1+ R * (T2 - T1) and T1 < T2 and*
not clipped(T1,filling,T2)

act(spills(E),spill) (6.7)

Recall that a property still holds at the time of the event which terminates it. Without this convention, autoterminating periods would be difficult to represent. We add to Axioms (6.1) to (6.7) a description of the particular problem at hand; there are outlets at levels *7* and *10*, and the tap is turned on at time *0*.

outlet(7) (7.1)

outlet(10) (7.2)
happens(e0) (7.3)
time(e0,0) (7.4)
act(e0,tap-on) (7.5)
holds-at(level(0),0) (7.6)

We can then prove, for example

holds-at(level(7),3.5) (8.1)

and therefore that

happens(spills(e0)) (8.2)
time(spills(e0),3.5) (8.3)

That is, the level reaches the overflow at time *3.5*. The level does not, of course, reach the rim because the filling autoterminates when the level reaches the overflow. So we cannot prove either of

holds-at(level(10),5) (9.1)
time(spills(e0),5) (9.2)

If the overflow were blocked and Axiom (7.1) were absent, then we would not be able to prove (8.3), but would be able to prove (9.1) and (9.2).

5. The General Case

From the example above, we can see how, in general, to write Event Calculus axioms which describe continuous change. Generalising from (6.1), we write

holds-at(P,T2) if (10.1)
happens(E) and time(E,T1) and
T1 < T2 and initiates(E,Q) and
not clipped(T1,Q,T2) and trajectory(Q,T1,P,T2)

This introduces the idea of the *trajectory* of a continuously changing quantity, which is a path plotted against time through the corresponding quantity space. For example, the level of liquid in the sink increases linearly from the time of the start of a period of filling. The formula *trajectory(Q,T1,P,T2)* represents that property *P* holds at time *T2* on the trajectory of the period of continuous change *Q* which starts at time *T1*. Effectively, this formulation isolates the function which describes the continuous change from the *holds-at* axiom. So, for the kitchen sink problem we simply write

trajectory(filling,T1,level(L2),T2) if (11.1)
*holds-at(level(L1),T1) and L2 = L1 + 2 * (T2 - T1)*

We can also generalise the autotermination axioms (6.4) to (6.7). The term *end(E,Q)* is introduced to denote the event which terminates the period of continuous change *Q* which was initiated by the event *E*. Instead of *spills(E)* we write *end(E,filling)*. The general axioms are

happens(end(E,Q)) if (12.1)
happens(E) and time(E,T1) and initiates(E,Q) and
autoterminates(P,Q) and trajectory(Q,T1,P,T2) and
not clipped(T1,Q,T2)

time(end(E,Q),T2) if (12.2)
happens(E) and time(E,T1) and initiates(E,Q) and

autoterminates(P,Q) and trajectory(Q,T1,P,T2) and
not clipped(T1,Q,T2)

terminates(end(E,Q),Q) (12.3)

The formula *autoterminates(P,Q)* represents that the period of continuous change *Q* is terminated when and if the property *P* holds. For the kitchen sink problem we write

autoterminates(level(L),filling) if outlet(L) (13.1)

To use the general axioms for the sink problem, Axiom (6.1) is replaced by (10.1) and (11.1), and Axioms (6.4) to (6.7) are replaced by (12.1) to (12.3) and (13.1). Then, the same consequences follow from the example represented by (7.1) to (7.6).

6. The Qualitative Version

Perhaps the exact, quantitative equation describing a quantity's trajectory is not known. For example, the rate of filling of the sink might be unknown, or the sink might fill up eratically rather than linearly. However, we still know that the liquid will reach every level in the sink up to the level at which it spills, and we can predict that the spilling will take place at the level of the overflow. This can be axiomatised by substituting a qualitative version of *trajectory* for the definition of (11.1).

A qualitative quantity space can be represented as an ordered set of *landmark values* [Kuipers, 1986]; those points in the corresponding quantitative space which are of interest for some reason. The continuously changing quantity in question can either be at one of the landmark values or between two adjacent ones. There are only three points of interest in the qualitative space of liquid levels in the kitchen sink; the bottom, the overflow and the rim. We write *follows(P1,P2,Q)* to express that *P1* and *P2* are adjacent landmark values in the trajectory of the period of continuous change *Q*. Also, we write *in(P,Q)* to express that the property *P* is in the trajectory of the period of continuous change *Q*. So, for the sink we have

in(level(bottom),filling) (14.1)
in(level(overflow),filling) (14.2)
in(level(rim),filling) (14.3)

follows(level(overflow),level(bottom),filling) (14.4)
follows(level(rim),level(overflow),filling) (14.5)

outlet(overflow) (14.6)
outlet(rim) (14.7)

Then, the following axioms define the general case of the qualitative version of the *trajectory* predicate.

trajectory(Q,T,P,time-of(P,Q,T)) if in(P,Q) (15.1)

trajectory(Q,T,between(P1,P2),T2) if (15.2)
time-of(P1,Q,T) < T2 < time-of(P2,Q,T) and
follows(P2,P1,Q)

time-of(P,Q,T) > T if in(P,Q) (15.3)

time-of(P2,Q,T) > time-of(P1,Q,T) if follows(P2,P1,Q) (15.4)

The term *time-of(P,Q,T)* represents the time at which the property *P* holds on the trajectory of the period of continuous change *Q* which starts at time *T*. The comparative operators now range over these terms, as well as the real numbers, keeping all their usual properties, such as transitivity. The property *between(P1,P2)* represents that the continuously changing quantity is between the adjacent landmark values *P1* and *P2*. To describe a scenario in which the tap is turned on at time *0*, we write as before

happens(e0) (16.1)
time(e0,0) (16.2)
act(e0,tap-on) (16.3)

We can then prove

happens(spill(e0)) (17.1)
time(spill(e0),time-of(level(overflow),filling,0)) (17.2)

but we cannot prove

holds-at(level(rim),T) (17.3)

7. The Many Tanks Problem

As a final exercise, I will present an axiomatisation of a slightly more complex problem, involving many tanks discharging liquid into another tank. Figure 2 shows an example of the two tank case. Each tank has an associated tap which can be turned on or off and which discharges liquid at a given rate. This problem is more elaborate than the kitchen sink problem because both taps can be discharging liquid at the same time.

Suppose that initially both taps are off. First, *tap1* is turned on, then *tap2* is turned on, then *tap2* is turned off, and then *tap1* is turned off. There are two ways of looking at this scenario. We might consider this as a *single* period of continuous change. After all, the bottom tank is filling up uninterruptedly from the time the first tap is turned on until it is turned off again. Then, to axiomatise it using the Event Calculus, we would have to write a suitable set of *trajectory* clauses which captures the variation in the rate of filling which occurs when the second tap is turned on and off.

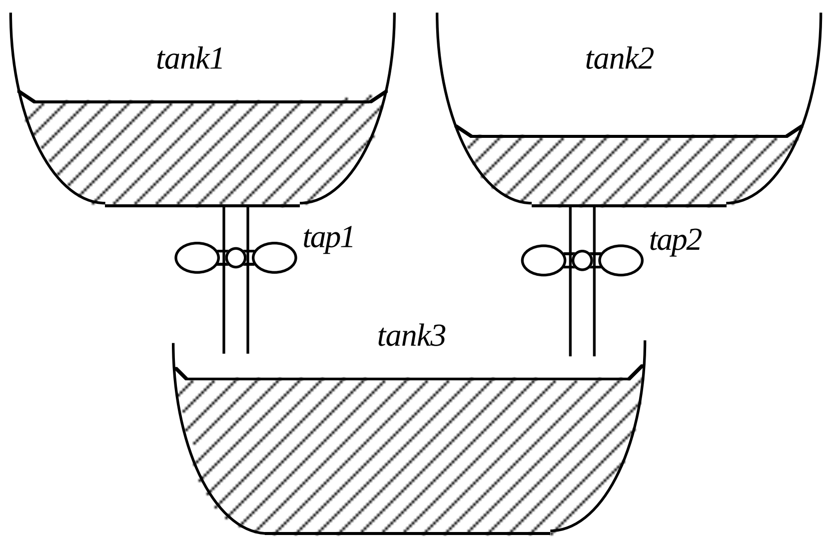

Figure 2: The Two Tank Problem

However, regarding this as a single period of continuous change would be a mistake for two reasons. First, it turns out to be very hard to write a suitable set of *trajectory* clauses. Second, it confuses discrete and continuous change. A better way is to consider this scenario as composed of four connected periods of continuous change. When the *tap1* is turned on, a period of continuous change is initiated. But when the *tap2* is turned on, a *discrete* change takes place — the rate of filling is (effectively) instantaneously increased by some amount. So the previous period of continuous change is terminated and a new one is initiated in which the sink is filling at a different rate.

8. Representing the Many Tanks Problem

The following set of axioms is for the *n* tanks problem. The terms *tap-on(N)* and *tap-off(N)* denote the acts of turning tap *N* on and off respectively. The term *filling(V,R)* represents the property that tank *V* is filling continuously at rate *R*. A tank cannot be filling at rate *0*. Rather, it is not filling at all. The formula *connects(N,V1,V2)* represents that tank *V1* discharges into tank *V2* via tap *N*, and the formula *rate(N,R)* represents that tap *N* discharges liquid at rate *R*. First we have axioms which describe how a tank fills.

initiates(E,filling(V2,R3)) if (18.1)
act(E,tap-on(N)) and rate(N,R1) and time(E,T) and
holds-at(filling(R2),T) and R3 = R1 + R2 and
connects(N,V1,V2)

terminates(E,filling(V2,R)) if (18.2)
act(E,tap-on(N)) and connects(N,V1,V2)

initiates(E,filling(V2,R3)) if (18.3)
act(E,tap-off(N)) and rate(N,R1) and time(E,T) and
holds-at(filling(R2),T) and R3 = R2 - R1 and R3 > 0
connects(N,V1,V2)

terminates(E,filling(V2,R)) if (18.4)
act(E,tap-off(N)) and connects(N,V1,V2)

initates(E,filling(V2,R2)) if (18.5)
act(E,tap-on(N)) and rate(N,R2) and time(E,T) and
not holds-at(filling(R1),T) and connects(N,V1,V2)

The term *emptying(V,R)* denotes the property that tank *V* is emptying at rate *R*. The next two axioms describe how a tank empties. If tank *V1* discharges into *V2* via tap *N*, then the act of turning tap *N* on initiates two periods; one in which *V2* is filling at some rate and another in which *V1* is emptying.

initiates(E,emptying(V1,R)) if (18.6)
act(E,tap-on(N)) and rate(N,R) and connects(N,V1,V2)

terminates(E,emptying(V1,R)) if (18.7)
act(E,tap-off(N)) and connects(N,V1,V2)

When a tank is neither emptying nor filling, its level remains constant. So, when the last tap is turned off, a period is initiated in which the level of the tank remains constant. Likewise, when the first tap is turned on, this terminates a period during which the level in the tank remains constant. In the kitchen sink example, this was ignored.

initiates(E,level(V2,L)) if (18.8)
act(E,tap-off(N)) and connects(N,V1,V2) and
time(E,T) and rate(N,R) and holds-at(filling(V2,R),T) and

holds-at(level(V2,L),T)

initiates(E,level(V1,L)) if (18.9)
act(E,tap-off(N)) and connects(N,V1,V2) and
time(E,T) and holds-at(level(V1,L),T)

terminates(E,level(V2,L)) if (18.10)
act(E,tap-on(N)) and connects(N,V1,V2) and
time(E,T) and not holds-at(filling(V2,R),T)

terminates(E,level(V1,L)) if (18.11)
act(E,tap-on(N)) and connects(N,V1,V2)

The trajectory of the filling of a tank is defined in a way similar to that for the kitchen sink. Also, the trajectory of the emptying of a tank needs to be defined. The term *level(V,L)* represents the property that the liquid in tank *V* is at level *L*.

trajectory(filling(V,R),T1,level(V,L2),T2) if (18.12)
*holds-at(level(V,L1),T1) and L2 = L1 + R * (T2 - T1)*

trajectory(emptying(V,R),T1,level(V,L2),T2) if (18.13)
*holds-at(level(V,L1),T1) and L2 = L1 - R * (T2 - T1)*

As with the kitchen sink problem, periods of filling and emptying can autoterminate. If a tank is filled to capacity, then its filling autoterminates. Similarly, if a tank drains its contents completely then its emptying autoterminates. The formula *rim(V,L)* represents that the rim of tank *V* is at level *L*.

autoterminates(level(V,L),filling(V,R)) if rim(V,L) (18.14)

autoterminates(level(V,0),emptying(V,R)) (18.15)

Given any sequence of *tap-on* and *tap-off* events, we can use this formulation to predict the levels of the tanks at different times. Unfortunately, the formulation does not work if two taps are turned on or off simultaneously. This is because the Event Calculus, as it stands, cannot easily be used to represent events which occur simultaneously and which have a cumulative effect on the same property, the rate of filling in this instance. This problem is separate from the problems of representing continuous change, and is beyond the scope of this paper.

Concluding Remarks

Some techniques have been presented for representing continuous change in the Event Calculus, to complement its exisitng capability for representing discrete change. A simplified version of Kowalski and Sergot's formulation was introduced and extended to cope with continuous change. A simple example involving a kitchen sink filling with liquid was axiomatised. The idea of autotermination was introduced, whereby a period of continuous change can bring about the event which terminates it, as when the level of liquid in the sink overflows. A more complex example involving many tanks discharging into another tank was also presented.

Since the axioms presented are all Horn clauses, they can be used with a Prolog interpreter to solve a variety of prediction problems. As they stand, they require some transformation to prevent looping, and for some problems a delaying mechanism for inequalities is required. A CLP(R) interpreter has this sort of mechanism built in, and can answer many queries which ordinary Prolog cannot.

Only a handfull of other authors have attended to the problem of using logic to represent continuous change. McDermott [1982, Section 4] outlines a framework for representing continuous change, based on a different ontology to that used in this paper but producing a similar effect. Hayes's work [1985] concentrates on continuous change. He presents a set of axioms which represent the behaviour of liquid under various conditions, but his approach is entirely different to that of this paper. Hayes's ontology is based on the idea of a *history*, which is a piece of spacetime. Histories lead to a different style of axiomatisation, in which the axioms constrain the ways histories may be connected together. More recently, Sandewall [1989] has described a framework which combines logic and differential equations to represent continuous change. A thorough comparison of these approaches is the subject of further work.

The field of qualitative reasoning is also concerned with representing continuous change. DeKleer and Brown's calculus of *confluences* [1985], for example, can represent a large variety of problems involving continuous change, and in some respects is a more powerful formalism than that offered in this paper. In particular, the omnidirectionality of a confluence (a qualitative differential equation) cannot easily be represented using the formalism described here. On the other hand, the formalisms used for qualitative reasoning are weak when it comes to representing discrete change. A combined formalism for temporal and qualitative reasoning would seem to be the way ahead [Shanahan, 1990].

Acknowledgements

Thanks to Bob Kowalski for making significant improvements to my earlier formulations of continuous change. Thanks also to Marek Sergot for valuable discussions on continuous change, especially in relation to the many tanks problem. Much of this work was done while the author was employed on Esprit project 2409 (Equator).

References

[Allen, 1984] J.F.Allen, Towards a General Theory of Action and Time, *Artificial Intelligence*, vol 23 (1984), p123.

[deKleer and Brown, 1985] J.deKleer and J.S.Brown, A Qualitative Physics Based on Confluences, in *Formal Theories of the Commonsense World*, ed Hobbs J.R. and Moore R.C., Ablex (1985).

[Hayes, 1985] P.J.Hayes, Ontology for Liquids, in *Formal Theories of the Commonsense World*, ed Hobbs J.R. and Moore R.C., Ablex (1985).

[Kowalski and Sergot, 1986] R.A.Kowalski and M.J.Sergot, A Logic-Based Calculus of Events, *New Generation Computing*, vol 4 (1986), p267.

[Kowalski, 1989] R.A.Kowalski, Database Updates in the Event Calculus, Imperial College Department of Computing Technical Report (1989).

[Kuipers, 1986] B.Kuipers, Qualitative Simulation, *Artificial Intelligence*, vol 29 (1986), p289.

[McCarthy and Hayes, 1969] J.McCarthy and P.J.Hayes, Some Philosophical Problems from the Standpoint of Artificial Intelligence, in *Machine Intelligence 4*, ed Michie D. and Meltzer B., Edinburgh University Press (1969).

[McDermott, 1982] D.McDermott, A Temporal Logic for Reasoning about Processes and Plans, *Cognitive Science*, vol 6 (1982), p101.

[Sandewall, 1989] E.Sandewall, Combining Logic and Differential Equations for Describing Real-World Systems, *Proceedings 1989 Knowledge Representation Conference*, p412.

[Shanahan, 1989] M.P.Shanahan, Prediction Is Deduction but Explanation Is Abduction, *Proceedings IJCAI 89*, p1055.

[Shanahan, 1990] M.P.Shanahan, A Calculus for Temporal and Qualitative Reasoning, in preparation.

Appendix: The Collected Axioms

Here are the collected general purpose axioms of the Event Calculus, extended to deal with qualitative and quantitative continuous change, and autotermination.

holds-at(P,T2) if (1.1)
happens(E) and time(E,T1) and T1 < T2 and
initiates(E,P) and not clipped(T1,P,T2)

clipped(T1,P,T3) if (1.2)
happens(E') and time(E',T2) and terminates(E',P) and
T1 ≤ T2 and T2 < T3

holds-at(P,T2) if (10.1)
happens(E) and time(E,T1) and
T1 < T2 and initiates(E,Q) and
not clipped(T1,Q,T2) and trajectory(Q,T1,P,T2)

happens(end(E,Q)) if (12.1)
happens(E) and time(E,T1) and initiates(E,Q) and
autoterminates(P,Q) and trajectory(Q,T1,P,T2) and
not clipped(T1,Q,T2)

time(end(E,Q),T2) if (12.2)
happens(E) and time(E,T1) and initiates(E,Q) and
autoterminates(P,Q) and trajectory(Q,T1,P,T2) and
not clipped(T1,Q,T2)

terminates(end(E,Q),Q) (12.3)

trajectory(Q,T,P,time-of(P,Q,T)) if in(P,Q) (15.1)

trajectory(Q,T,between(P1,P2),T2) if (15.2)
time-of(P1,Q,T) < T2 < time-of(P2,Q,T) and
follows(P2,P1,Q)

time-of(P,Q,T) > T if in(P,Q) (15.3)

time-of(P2,Q,T) > time-of(P1,Q,T) if follows(P2,P1,Q) (15.4)