1 Introduction

The idea of *action description languages* was first introduced by Gelfond and Lifschitz [7, 8], with their proposal for the Language $\mathcal{A}$. The intention was that such languages would have a specialised syntax and semantics, and would be at a high enough level of abstraction so as to be easily readable, understandable and intuitive. They could then be used as specifications for theories of action and change written in more general purpose or computation oriented formalisms. Since then, many variants and extensions of $\mathcal{A}$ have been proposed. The majority of these, like $\mathcal{A}$ itself, are based on an ontology inherited from the Situation Calculus.

The Language $\mathcal{E}$ [10] was introduced partly to illustrate that Gelfond and Lifschitz's methodology can be applied using ontologies other than that of the Situation Calculus. In particular, it showed that for *narrative reasoning* (that is, reasoning about actions which actually occur at various times, and reasoning about the properties that hold or do not hold at different times as a consequence of such occurrences), it is appropriate to use an ontology where the flow of time is independent of the action occurrences embedded within it. $\mathcal{E}$'s ontology of actions, time points and fluents facilitates a particularly simple and intuitive syntax and semantics for describing the effects of such occurrences. Our approach has been inspired by Kowalski and Sergot's Event Calculus [13] and its extensions (see [27] for a comprehensive overview).

A key characteristic of the $\mathcal{E}$'s semantics is that it is *modular*, and thus easily extendible. To illustrate what is meant by this more precisely, two extensions of the "core" semantics are briefly described in an appendix of [10], to deal with *qualifications* and *ramifications* respectively. The purpose of the present paper is to expand on the theme of ramifications, and in particular to show how the translation method described in [10], from basic $\mathcal{E}$ domain descriptions to Event Calculus style logic programs, may be extended to deal with domains containing ramification statements.

For generality, the semantics of $\mathcal{E}$ described both in this paper and in [10] is in terms of an arbitrary structure of time. All definitions and results apply whether time is linear, branching, continuous or discrete. Indeed, as shown in [10], as a special case time may be defined as a forward branching structure of situations, each characterised by a unique sequence of actions, as in various versions of the situation calculus. In this case, as shown in [10], the semantics of $\mathcal{E}$ is equivalent to that of $\mathcal{A}$ (and thus, by the results of Kartha [11], equivalent to the core Situation Calculi of Reiter [23], Pednault [22] and Baker [1]). Hence both the representational and the computational mechanisms reported here to deal with ramifications are, like those in [10], readily adaptable for use in the context of the Language $\mathcal{A}$ and the Situation Calculus.

The paper is organised as follows. In Section 2, a review is given of the syntax and semantics of the basic Language $\mathcal{E}$ (without ramification statements). In Section 3, we describe $\mathcal{E}$'s extension to deal with ramifications, illustrating its utility using two benchmark problems from the literature. In Section 4, a translation method into logic programs is given, followed by a proof of soundness with respect to the standard SLDNF (i.e. Prolog) proof procedure. As in [10], this proof holds even when complete information about some initial state of affairs is unavailable. Finally, in Section 5 we discuss the advantages and limitations of our approach, and compare it with other recent proposals for dealing with ramifications within logic based approaches to reasoning about the effects of actions.

2 A Review of the Basic Language $\mathcal{E}$

2.1 Syntax

Like the Language $\mathcal{A}$, the Language $\mathcal{E}$ is really a collection of languages. The particular vocabulary of each language depends on the domain being represented, but always includes a set of *fluent constants*, a set of *action constants*, and a partially ordered set of *time-points*. A *fluent literal* may either be a fluent constant or its negation, as shown in the following definitions.

Definition 1 [Domain Language] A *domain language* is a tuple $\langle \Pi, \preceq, \Delta, \Phi \rangle$, where $\preceq$ is a partial (possibly total) ordering defined over the non-empty set Π of *time points*, Δ is a non-empty set of *action constants*, and Φ is a non-empty set of *fluent constants*. □

Definition 2 [Fluent literal] A *fluent literal* of $\mathcal{E}$ is an expression either of the form F or of the form $\neg F$, where $F \in \Phi$. □

As in [10], for the remainder of this paper we assume a particular domain language $\mathcal{E} = \langle \Pi, \preceq, \Delta, \Phi \rangle$. We will often write $T_1 \prec T_2$ to mean $T_1 \preceq T_2$ and $T_1 \neq T_2$. Three types of statements are used to describe domains without ramifications; *h-propositions* ("h" for "happens"), *t-propositions* ("t" for "time point") and *c-propositions* ("c" for "causes"). Their intended meanings are clear from their definitions:

Definition 3 [h-proposition] An *h-proposition* in $\mathcal{E}$ is an expression of the form

$$A \textbf{ happens-at } T$$

where $A \in \Delta$ and $T \in \Pi$. □

Definition 4 [t-proposition] A *t-proposition* in $\mathcal{E}$ is an expression of the form

$$L \textbf{ holds-at } T$$

where L is a fluent literal of $\mathcal{E}$ and $T \in \Pi$. □

Definition 5 [c-proposition] A *c-proposition* in $\mathcal{E}$ is an expression either of the form

$$A \textbf{ initiates } F \textbf{ when } C$$

or of the form

$$A \textbf{ terminates } F \textbf{ when } C$$

where $F \in \Phi$, $A \in \Delta$, and C is a set of fluent literals of $\mathcal{E}$. □

"A **initiates** F **when** C" should be regarded as meaning "C is a minimally sufficient set of conditions for an occurrence of A to have an initiating effect on F". C-propositions of the form "A **initiates** F **when** $\emptyset$" and "A **terminates** F **when** $\emptyset$" can be written more simply as "A **initiates** F" and "A **terminates** F" respectively. For domains without indirect effects, a *domain description* in $\mathcal{E}$ is a triple $\langle \gamma, \eta, \tau \rangle$, where γ is a set of c-propositions, η is a set of h-propositions and τ is a set of t-propositions.

As a simple example domain, consider taking a photograph. Suppose that we need to refer to the actions of loading the camera with film, looking

at the view and taking a picture, so that $\Delta = \{Load, Look, Take\}$, and refer to the properties of the camera having film in it and a picture having been taken, so that $\Phi = \{Loaded, Picture\}$. For Π and $\preceq$ we will pick the real numbers with the usual ordering relation. Given this vocabulary, an example of a t-proposition is

$\neg Picture$ **holds-at** 1 (tsp1)

Examples of h-propositions are

$Load$ **happens-at** 2 (tsp2)
$Look$ **happens-at** 5 (tsp3)
$Take$ **happens-at** 8 (tsp4)

and examples of c-propositions are

$Load$ **initiates** $Loaded$ (tsp5)
$Take$ **initiates** $Picture$ **when** $\{Loaded\}$ (tsp6)

2.2 Semantics

The semantics of $\mathcal{E}$ is based on simple definitions of interpretations and models. Since we are primarily interested in inferences about what holds at particular time-points in Π, it is sufficient to define an interpretation as a mapping of fluent/time-point pairs to $true$ or $false$ (i.e. a "*holds*" relation). An interpretation *satisfies* a fluent literal or set of fluent literals at a particular time-point if it assigns the relevant truth values to each of the corresponding fluent constants:

Definition 6 [Interpretation] An *interpretation* of $\mathcal{E}$ is a mapping

$$H : \Phi \times \Pi \mapsto \{true, false\}$$

□

Definition 7 [Point satisfaction] Given a set of fluent literals C of $\mathcal{E}$ and a time point $T \in \Pi$, an interpretation H *satisfies* C *at* T iff for each fluent constant $F \in C$, $H(F, T) = true$, and for each fluent constant F' such that $\neg F' \in C$, $H(F', T) = false$. □

The definition of a model in $\mathcal{E}$ is parametric on the definitions of an *initiation point* and a *termination point*. For the basic $\mathcal{E}$ (without ramification statements), these definitions are particularly straightforward. Initiation and termination points are simply time-points where a c-proposition and an h-proposition combine to describe a direct effect on a particular fluent:

Definition 8 [Initiation/termination point (without ramifications)] Let H be an interpretation of $\mathcal{E}$, let $D = \langle \gamma, \eta, \tau \rangle$ be a domain description, let $F \in \Phi$ and let $T \in \Pi$. T is an *initiation-point* (respectively *termination-point*) *for* F *in* H *relative to* D iff there is an $A \in \Delta$ such that (i) there is both an h-proposition in η of the form "A **happens-at** T" and a c-proposition in γ of the form "A **initiates** F **when** C" (respectively "A **terminates** F **when** C") and (ii) H satisfies C at T. □

Having defined these preliminary notions, it is straightforward to describe which interpretations are models of a particular domain description.

Three basic properties need to be satisfied; (1) fluents change their truth values only via occurrences of initiating or terminating actions, (2) initiating a fluent establishes its truth value as $true$, and (3) terminating a fluent establishes its truth value as $false$. In addition, (4) every model must match with each of the t-propositions in the domain description:

Definition 9 [Model (without ramifications)] Given a domain description $D = \langle \gamma, \eta, \tau \rangle$ in $\mathcal{E}$, an interpretation H of $\mathcal{E}$ is a *model* of D iff, for every $F \in \Phi$ and $T, T', T_1, T_3 \in \Pi$ such that $T_1 \prec T_3$, the following properties hold:

1. If there is no initiation-point or termination-point T_2 for F in H relative to D such that $T_1 \preceq T_2 \prec T_3$, then $H(F, T_1) = H(F, T_3)$.

2. If T_1 is an initiation-point for F in H relative to D, and there is no termination-point T_2 for F in H relative to D such that $T_1 \prec T_2 \prec T_3$, then $H(F, T_3) = true$.

3. If T_1 is a termination-point for F in H relative to D, and there is no initiation-point T_2 for F in H relative to D such that $T_1 \prec T_2 \prec T_3$, then $H(F, T_3) = false$.

4. For all t-propositions in τ of the form "F **holds-at** T", $H(F, T) = true$, and for all t-propositions of the form "$\neg F$ **holds-at** T'", $H(F, T') = false$.

□

Definition 10 [Consistency] A domain description is *consistent* iff it has a model. □

Definition 11 [Entailment] A domain description D *entails* the t-proposition "F **holds-at** T", written "$D \models F$ **holds-at** T", iff for every model H of D, $H(F, T) = true$. D entails the t-proposition "$\neg F$ **holds-at** T" iff for every model H of D, $H(F, T) = false$. □

It is easy to see from the above definitions that, for example, the collection of t-, h- and c-propositions (tsp1)–(tsp6) entails the t-proposition

$$Picture \textbf{ holds-at } 10$$

Hence the Language $\mathcal{E}$ solves this (non-violent) version of the Yale Shooting Problem.

The framework described above allows for various modes of reasoning within various structures of time. As another example, we can construct a set Π_Δ of time-points out of the set $\Delta = \{Load, Look, Take\}$ of action constants used above. Π_Δ includes all sequences of these constants, e.g. $[Load, Look, Take]$ and $[Load, Load]$. We denote the empty sequence by S_0. It also includes auxiliary "start" points, $Start([Load, Look, Take])$, $Start([Load, Load])$, etc, for each such sequence other than S_0. The partial ordering $\prec_\Delta$ on this set is the obvious one. Given any sequence $[\alpha_1, \ldots, \alpha_m]$, where each $\alpha_i \in \Delta$, it is such that

$$S_0 \prec_\Delta Start([\alpha_1, \ldots, \alpha_{m-1}]) \prec_\Delta [\alpha_1, \ldots, \alpha_{m-1}] \prec_\Delta Start([\alpha_1, \ldots, \alpha_m])$$

Clearly, this gives a branching structure of time analogous to that used in many versions of the Situation Calculus. Now consider the domain description which includes the set η_Δ of all h-propositions of the form

$$\alpha_n \textbf{ happens-at } Start([\alpha_1, \ldots, \alpha_n])$$

and, in addition, the following c-propositions and t-propositions

$$Take \textbf{ initiates } Picture \textbf{ when } \{Loaded\}$$
$$\neg Picture \textbf{ holds-at } S_0$$
$$Picture \textbf{ holds-at } [Look, Take]$$

Again, it is easy to see that this domain entails the t-proposition

$$Loaded \textbf{ holds-at } S_0$$

Thus the semantics of $\mathcal{E}$ facilitates reasoning both forwards and backwards within various (branching or linear) structures of time. A general result in [10] shows that for domain descriptions with time structures such as $\langle \Pi_\Delta, \preceq_\Delta \rangle$ above and the set of h-propositions η_Δ, $\mathcal{E}$'s semantics is equivalent to that of the Language $\mathcal{A}$, so that $\mathcal{A}$ can be regarded as a special case of $\mathcal{E}$.[1] Note that proponents of the Situation Calculus generally regard reasoning with this type of branching structure of time as "hypothetical" reasoning about (possible future) actions (as opposed to reasoning about the "actual" time-line). Hence, viewed in this light, the h-propositions in η_Δ refer to hypothetical or possible future action occurrences. (This is in contrast to the very different use of the syntactically similar "occurrence facts" in [3], which refer to "actual" action occurrences.)

More generally, [10] also discusses how $\mathcal{E}$ can be used to formulate an intuitive notion of *planning*, and how various notions of *explanation* can be defined to deal with inconsistency within $\mathcal{E}$ domain descriptions. It also describes how Event Calculus style logic programs can be used to compute consequences (t-propositions) of $\mathcal{E}$ domain descriptions of the basic type described above, even in the absence of complete information about which fluents hold at some initial time-point.

3 Ramifications in $\mathcal{E}$

The *ramification problem* arises in domains whose description most naturally includes permanent constraints or relationships between fluents. In formalisms which allow for such statements, the effects of actions may sometimes be propagated via groups of these constraints. The problem is to adequately describe these propagations of effects, whilst retaining a solution to the frame problem – that is, the problem of succinctly expressing that most actions leave most fluents unchanged.

As briefly described in the appendix of [10], $\mathcal{E}$ domain descriptions can include ramification statements, called *r-propositions*, to describe constraints between fluents:

Definition 12 [r-proposition] An *r-proposition* in $\mathcal{E}$ is an expression of the form

$$L \textbf{ whenever } C$$

where L is a fluent literal and C is a set of fluent literals of $\mathcal{E}$. □

For the remainder of the paper, we will assume that domain descriptions are of the form $\langle \gamma, \eta, \tau, \rho \rangle$, where γ, η and τ are as before, and ρ is a set of r-propositions. The intended meaning of "L **whenever** C" is as follows: "at

[1] A slightly different notation is used in [10] to refer to the time-points within Π_Δ.

every time-point that C holds, L holds, and hence every action occurrence that brings about C also brings about L".

As a simple illustration of the use of r-propositions, consider the canonical "stuffy room" example. Imagine a room with a window and a vent. We can open and close both the window and the vent, with the obvious effects:

$$\begin{array}{ll} \mathit{CloseWindow} \textbf{ initiates } \mathit{WindowClosed} & \text{(sr1)} \\ \mathit{CloseVent} \textbf{ initiates } \mathit{VentClosed} & \text{(sr2)} \\ \mathit{OpenWindow} \textbf{ terminates } \mathit{WindowClosed} & \text{(sr3)} \\ \mathit{OpenVent} \textbf{ terminates } \mathit{VentClosed} & \text{(sr4)} \end{array}$$

Now suppose that we also want to assert that the room becomes stuffy whenever both the vent and the window are closed. We could of course attempt to describe this fact using extra c-propositions:

$$\begin{array}{c} \mathit{CloseWindow} \textbf{ initiates } \mathit{Stuffy} \textbf{ when } \{\mathit{VentClosed}\} \\ \mathit{CloseVent} \textbf{ initiates } \mathit{Stuffy} \textbf{ when } \{\mathit{WindowClosed}\} \\ \mathit{OpenWindow} \textbf{ terminates } \mathit{Stuffy} \\ \mathit{OpenVent} \textbf{ terminates } \mathit{Stuffy} \end{array}$$

However, there are problems with this approach. First, given the semantics of the previous section, these c-propositions fail to capture the "static" aspect of the intended constraint. In a domain description with no h-propositions or t-propositions at all, it would be possible to construct a model where, for example, *WindowClosed* and *VentClosed* were true at all time-points, but *Stuffy* was false. Second, this representation would fail to deal with simultaneous action occurrences. If both the vent and the window were initially open, closing both of them simultaneously would fail to make the room stuffy. If the vent was initially closed and the window initially open, simultaneously opening the vent and closing the window would cause an inconsistency.

The alternative is to use three r-propositions:

$$\begin{array}{ll} \mathit{Stuffy} \textbf{ whenever } \{\mathit{VentClosed}, \mathit{WindowClosed}\} & \text{(sr5)} \\ \neg\mathit{Stuffy} \textbf{ whenever } \{\neg\mathit{VentClosed}\} & \text{(sr6)} \\ \neg\mathit{Stuffy} \textbf{ whenever } \{\neg\mathit{WindowClosed}\} & \text{(sr7)} \end{array}$$

and to extend the semantics of $\mathcal{E}$ appropriately. As we shall see, the extension described below deals with the two problems described above, as well as with other problems raised by different examples. It is also easy to see that this approach has the additional advantage of being more succinct – imagine a slightly larger example where the room has a small window, a large window, a vent and a door, or where a number of different actions have the same effect of closing the vent.

In fact, we need make only very few modifications to the semantics of $\mathcal{E}$ in order to deal with r-propositions. The first modification is to extend condition (4) in Definition 9 of a model in the obvious way to cover the static aspect of r-propositions. Conditions (1), (2) and (3) stay the same (since they are parametric on the definitions of an initiation- and a termination-point, which are extended below), so that the new definition is:

Definition 13 [Model] Given a domain description $D = \langle \gamma, \eta, \tau, \rho \rangle$ in $\mathcal{E}$, an interpretation H of $\mathcal{E}$ is a *model* of D iff, for every $F \in \Phi$ and $T, T', T_1, T_3 \in \Pi$ such that $T_1 \prec T_3$, the following properties hold:

1. If there is no initiation-point or termination-point T_2 for F in H relative to D such that $T_1 \preceq T_2 \prec T_3$, then $H(F, T_1) = H(F, T_3)$.

2. If T_1 is an initiation-point for F in H relative to D, and there is no termination-point T_2 for F in H relative to D such that $T_1 \prec T_2 \prec T_3$, then $H(F, T_3) = true$.

3. If T_1 is a termination-point for F in H relative to D, and there is no initiation-point T_2 for F in H relative to D such that $T_1 \prec T_2 \prec T_3$, then $H(F, T_3) = false$.

4. For all t-propositions in τ of the form "F **holds-at** T", $H(F, T) = true$, for all t-propositions of the form "$\neg F$ **holds-at** T'", $H(F, T') = false$, and for all r-propositions in ρ of the form "L **whenever** C", if H satisfies C at T then H satisfies $\{L\}$ at T.

□

Using this new definition of a model, and given an arbitrary $T \in \Pi$, we can, for example, now infer the t-proposition

$$\neg WindowClosed \textbf{ holds-at } T$$

from the two t-propositions

$$\neg Stuffy \textbf{ holds-at } T$$
$$VentClosed \textbf{ holds-at } T$$

by considering the "contrapositive", static effect of the r-proposition (sr5).

It remains only to extend the definitions of an initiation point and a termination point to accommodate the possibility of fluents being indirectly initiated and terminated via r-propositions. Again the modification is straightforward. Intuitively, in order to find time-points at which the fluent literal L is established via the r-proposition "L **whenever** C", we need to look for time-points at which one or more of the conditions in C become established, and at which the remaining conditions are already and continue to be satisfied (up to some time-point beyond the point in question). This intuition translates directly into the following fixed-point definition[2]:

Definition 14 [Initiation/termination point] Let H be an interpretation of $\mathcal{E} = \langle \Pi, \preceq, \Delta, \Phi \rangle$, and $D = \langle \gamma, \eta, \tau, \rho \rangle$ be a domain description. Let $\mathcal{W}$ be the set $2^{\Phi \times \Pi} \times 2^{\Phi \times \Pi}$ and let the operator $\mathcal{F} : \mathcal{W} \mapsto \mathcal{W}$ be defined as follows. For each, $\langle \mathcal{I}n, \mathcal{T}e \rangle \in \mathcal{W}$ denote $\mathcal{F}(\langle \mathcal{I}n, \mathcal{T}e \rangle)$ by $\langle \mathcal{I}n', \mathcal{T}e' \rangle$. Then for any $F \in \Phi$ and $T \in \Pi$, (F, T) is in $\mathcal{I}n'$ (respectively in $\mathcal{T}e'$) iff one of the following two conditions holds.

1. There is an $A \in \Delta$ such that (i) there is both an h-proposition in η of the form "A **happens-at** T" and a c-proposition in γ of the form "A **initiates** F **when** C" (respectively "A **terminates** F **when** C") and (ii) H satisfies C at T.

2. There is an r-proposition in ρ of the form "F **whenever** C" (respectively "$\neg F$ **whenever** C") and a partition $\{C_1, C_2\}$ of C such that (i) C_1 is non-empty, for each fluent constant $F' \in C_1$, $(F', T) \in \mathcal{I}n$, and for each fluent literal $\neg F' \in C_1$, $(F', T) \in \mathcal{T}e$, and (ii) there is some $T_2 \in \Pi$, $T \prec T_2$, such that for all T_1, $T \preceq T_1 \preceq T_2$, H satisfies C_2 at T_1.

[2] Note that Michael Thielscher's online ETAI discussion question, at http://www.ida.liu.se/ext/etai/received/actions/003/aip.html#001, refers to a slightly earlier version of this paper with a more informal wording of this definition, in which the least fixed point construction is not described explicitly.

Let $\langle \mathcal{I}n^f, \mathcal{T}e^f \rangle$ be the least fixed point of the (monotonic) operator $\mathcal{F}$ starting from the empty tuple $\langle \emptyset, \emptyset \rangle$. T is an *initiation-point* (respectively *termination-point*) *for* F *in* H *relative to* D iff $(F,T) \in \mathcal{I}n^f$ (respectively $(F,T) \in \mathcal{T}e^f$). □

It is useful to note that any initiation or termination point at some time T relative to D, defined in this way, must refer to at least one known h-proposition at T in the domain D. Therefore a domain which contains a finite number of h-propositions will necessarily contain only a finite number of initiation and termination points.

We can see how the new definition of initiation and termination points works by extending the stuffy room example with some t- and h-propositions. For the sake of variety, suppose that the set Π of time-points is the set of natural numbers with the usual ordering relation, and that in our domain description $D_{sr} = \langle \gamma_{sr}, \eta_{sr}, \tau_{sr}, \rho_{sr} \rangle$, γ_{sr} consists of c-propositions (sr1)–(sr4), ρ_{sr} consists of r-propositions (sr5)–(sr7), η_{sr} contains the h-propositions

$$\mathit{CloseWindow} \textbf{ happens-at } 2 \qquad \text{(sr8)}$$
$$\mathit{CloseVent} \textbf{ happens-at } 2 \qquad \text{(sr9)}$$

and τ_{sr} contains the t-propositions

$$\neg \mathit{WindowClosed} \textbf{ holds-at } 0 \qquad \text{(sr10)}$$
$$\neg \mathit{VentClosed} \textbf{ holds-at } 0 \qquad \text{(sr11)}$$

Intuitively, we can see that at any given time-point after 2, say 4, the room should be stuffy. To show that the semantics matches with this intuition, we need to show that

$$D_{sr} \models \mathit{Stuffy} \textbf{ holds-at } 4$$

and that D_{sr} is consistent. It is sufficient to observe that the following interpretation H_{sr} is the only model of D_{sr}. For all time-points (naturals) $\pi \in \Pi$ and $\zeta \in \{\mathit{Stuffy}, \mathit{WindowClosed}, \mathit{VentClosed}\}$,

$$H_{sr}(\zeta, \pi) = \mathit{false} \quad \text{if } \pi \leq 2$$
$$H_{sr}(\zeta, \pi) = \mathit{true} \quad \text{if } \pi > 2$$

According to Definition 14, the only possible initiation point or termination point for any of the fluents in H_{sr} relative to D_{sr} is 2, since this is the only point to which an h-proposition refers. The combination of (sr1) with (sr8) and the combination of (sr2) with (sr9) ensure that (in any model) 2 is an initiation point for $\mathit{WindowClosed}$ and $\mathit{VentClosed}$, by condition (1) of Definition 14. Hence, since $\{\{\mathit{VentClosed}, \mathit{WindowClosed}\}, \emptyset\}$ is a partition of $\{\mathit{VentClosed}, \mathit{WindowClosed}\}$, (sr5) ensures that (in any model) 2 is an initiation point for Stuffy by condition (2) of Definition 14. It is now straightforward to check that H_{sr} satisfies all four conditions of Definition 13, and easy to see that there can be no other model of this domain. (Variations of this stuffy room example also behave correctly when represented in $\mathcal{E}$. In particular, if we replace (sr9) with "$\mathit{CloseVent}$ **happens-at** 3" our semantics does not give rise to the type of anomalous model which is problematic for some other approaches, e.g. [1], in which a change at 3 from $\neg\mathit{Stuffy}$ to Stuffy is avoided by incorporating an intuitively unjustified change from $\mathit{WindowClosed}$ to $\neg\mathit{WindowClosed}$.)

Our second example is taken from [28], and is problematic for what the author terms as "categorization-based" approaches to ramification such as [14] and [25] (see [28] for a detailed discussion of why this is so). An

electric circuit is arranged as in Figure 1, so that, when activated, the relay instantaneously pulls switch $Switch2$ open (i.e. disconnects it). We will use the fluent constants $Switch1$, $Switch2$ and $Switch3$ to represent that the respective switches are closed (i.e. making a connection), and the symbols $Relay$ and $Light$ respectively to indicate that these components are activated.

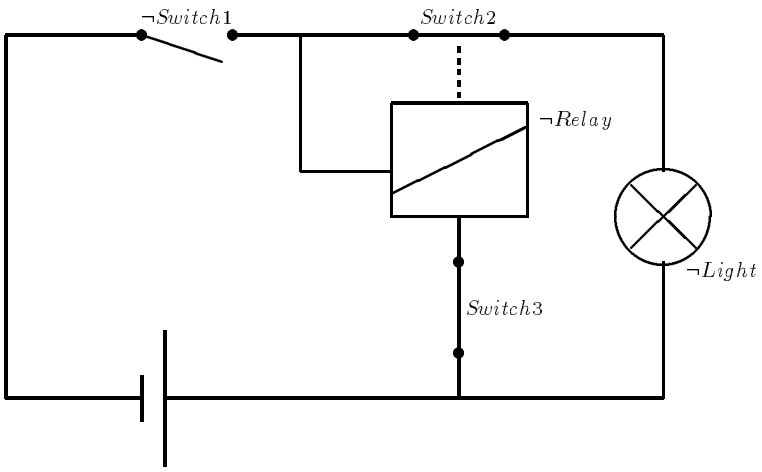

Figure 1

The representation of and reasoning about this domain is unproblematical for $\mathcal{E}$. The various permanent constraints on the state of the circuit can be represented with 7 r-propositions:

$Light$ **whenever** $\{Switch1, Switch2\}$ (c1)
$\neg Light$ **whenever** $\{\neg Switch1\}$ (c2)
$\neg Light$ **whenever** $\{\neg Switch2\}$ (c3)
$Relay$ **whenever** $\{Switch1, Switch3\}$ (c4)
$\neg Relay$ **whenever** $\{\neg Switch1\}$ (c5)
$\neg Relay$ **whenever** $\{\neg Switch3\}$ (c6)
$\neg Switch2$ **whenever** $\{Relay\}$ (c7)

The direct effects of opening (disconnecting) and closing (connecting) each of the three switches are captured in 6 c-propositions:

$CloseSwitch1$ **initiates** $Switch1$ (c8)
$OpenSwitch1$ **terminates** $Switch1$ (c9)
$CloseSwitch2$ **initiates** $Switch2$ **when** $\{\neg Relay\}$ (c10)
$OpenSwitch2$ **terminates** $Switch2$ (c11)
$CloseSwitch3$ **initiates** $Switch3$ (c12)
$OpenSwitch3$ **terminates** $Switch3$ (c13)

Taking the set Π of time-points to be the non-negative reals, we can describe the initial configuration of the circuit, pictured in Figure 1, as follows:

$\neg Switch1$ **holds-at** 0 (c14)
$Switch2$ **holds-at** 0 (c15)
$Switch3$ **holds-at** 0 (c16)
$\neg Light$ **holds-at** 0 (c17)
$\neg Relay$ **holds-at** 0 (c18)

(In fact, given condition (4) of Definition 13, only (c15) and (c16) are needed to adequately describe this initial configuration.)

If we add an occurrence of the action $CloseSwitch1$, say at 0.2, to this domain description:

$$CloseSwitch1 \textbf{ happens-at } 0.2 \qquad \text{(c19)}$$

it should follow from $\mathcal{E}$'s semantics that, at points after 0.2, the relay is activated and $Switch2$ is disconnected, so that the light remains off. So, for example, taking D_c to consist of propositions (c1)–(c19), it should be the case that

$$D_c \models \neg Light \textbf{ holds-at } 0.4$$

and that D_c is consistent. As in the previous example, it is sufficient to show that there is a single model H_c for D_c, where H_c is defined as follows. For all $\zeta \in \{Switch1, Switch2, Switch3, Relay, Light\}$ and time-points (non-negative reals) $\pi \in \Pi$,

$$\begin{array}{ll} H_c(\zeta, \pi) = true & \text{if } \pi \leq 0.2 \text{ and } \zeta \in \{Switch2, Switch3\} \\ H_c(\zeta, \pi) = true & \text{if } \pi > 0.2 \text{ and } \zeta \in \{Switch1, Switch3, Relay\} \\ H_c(\zeta, \pi) = false & \text{otherwise} \end{array}$$

Clearly, the only possible initiation point or termination point for any of the fluents in any interpretation relative to D_c is 0.2. The combination of (c8) with (c19) ensures that (in any model) 0.2 is an initiation point for $Switch1$, by condition (1) of Definition 14. Hence, by condition (2) of Definition 14, 0.2 is an initiation point for $Relay$ in H_c, because $\{\{Switch1\}, \{Switch3\}\}$ is a partition of the set $\{Switch1, Switch3\}$ which appears on the righthand side of (c4), and any time-point $T_2 > 0.2$ is such that, for all T_1 in the closed interval $[0.2, T_2]$, $H_c(Switch3, T_1) = true$. Therefore, again by condition (2) of Definition 14, 0.2 is a termination point for $Switch2$ in H_c, since $\{\{Relay\}, \emptyset\}$ is a partition for the set $\{Relay\}$ which appears on the righthand side of (c7). Notice that 0.2 is not a termination point for $Switch3$ in H_c, because $Switch3$ does not appear in the left-hand side of any r-proposition in D_c, and there is no h-proposition mentioning the action $OpenSwitch3$ in D_c. Neither is 0.2 an initiation point for $Light$ in H_c – although $\{\{Switch1\}, \{Switch2\}\}$ is a partition of the set $\{Switch1, Switch2\}$ which appears on the righthand side of (c1), condition (2) of Definition 14 is not applicable because there is no $T_2 > 0.2$ such that $H_c(Switch2, T_2) = true$.

From this argument it is clear that in every model of D_c, 0.2 must be an initiation point for $Switch1$ and $Relay$ and a termination point for $Switch2$, and that $Switch3$ and $Light$ cannot have any initiation or termination points. Hence H_c is the only possible model of D_c.

Note that variations of D_c involving simultaneous action occurrences also give the intuitively expected results. For example, consider the domain description D_c^+, which is the same as D_c except that it also includes the h-proposition

$$OpenSwitch3 \textbf{ happens-at } 0.2 \qquad \text{(c20)}$$

It is straightforward to check that the only model of D_c^+ is H_c^+, where for all $\zeta \in \{Switch1, Switch2, Switch3, Relay, Light\}$ and time-points (non-negative reals) $\pi \in \Pi$,

$H_c^+(\zeta, \pi) = true$ if $\pi \leq 0.2$ and $\zeta \in \{Switch2, Switch3\}$
$H_c^+(\zeta, \pi) = true$ if $\pi > 0.2$ and $\zeta \in \{Switch1, Switch2, Light\}$
$H_c^+(\zeta, \pi) = false$ otherwise

Thus, for example,

$$D_c^+ \models Light \textbf{ holds-at } 0.4$$

4 Logic Programs for $\mathcal{E}$ Domains with Ramifications

In this section we present a generalisation of the translation given in [10] (from $\mathcal{E}$ domains to logic programs) which can also be applied to domains containing r-propositions. We then show that the derivation of *HoldsAt* literals (translations of t-propositions) from these programs via SLDNF is sound with respect to $\mathcal{E}$'s notion of semantic entailment.

Before presenting the translation itself, we give some preliminary propositions and definitions which will be useful in proving soundness. The first proposition follows from a straightforward observation that in Definition 13 of a model we could replace condition (1) with two more specific conditions relating to the truth values $true$ and $false$:

Proposition 1 Let H be a model of the domain description D and let $T_1 \preceq T_3$. Then for every $F \in \Phi$

1. If $H(F, T_1) = true$ and there is no termination-point T_2 for F in H relative to D such that $T_1 \preceq T_2 \prec T_3$, then $H(F, T_3) = true$, and
2. If $H(F, T_1) = false$ and there is no initiation-point T_2 for F in H relative to D such that $T_1 \preceq T_2 \prec T_3$, then $H(F, T_3) = false$.

Proof: For the first part of the proposition, there are two cases to consider. (i) There is an initiation-point T', $T_1 \preceq T' \prec T_3$, for F in H relative to D. In this case this half of the proposition follows from condition (2) of Definition 13. (ii) There is no initiation-point T', $T_1 \preceq T' \prec T_3$, for F in H relative to D. In this case this part of the proposition follows from condition (1) of Definition 13. The proof of the second part of the proposition is analogous.

□

Aside from the generalisation to deal with r-propositions, we take the opportunity here to define a translation which, unlike that in [10], can deal with a limited class of domain descriptions containing an infinite number of h-propositions. This allows us, for example, to deal with domains containing a situation-calculus-like structure of time, such as the second example of Section 2.2, which contains an infinite number of h-propositions of the form

$$\alpha_n \textbf{ happens-at } Start([\alpha_1, \ldots, \alpha_n]).$$

However, our domain descriptions must be *occurrence-sparse*, i.e. cannot contain finite intervals of time in which an infinite number of actions occur. The formal definition of occurrence-sparsity is as in [10]:

Definition 15 [Occurrence Sparsity] Let $D = \langle \gamma, \eta, \tau, \rho \rangle$ be a domain description written in a language $\mathcal{E} = \langle \Pi, \preceq, \Delta, \Phi \rangle$. D and η are *occurrence-sparse* iff for any two points $T_1, T_2 \in \Pi$ there are only a finite number of h-propositions in η of the form "A **happens-at** T" such that $T_1 \preceq T \prec T_2$.
□

To avoid technical complications, we only consider domains with the following property; for all time-points at which some action occurs, there exists some future time-point. We will call such domains *end-free*. We could modify (and thereby complicate) the definitions to deal with occurrences at the "end(s) of time", but there seems little point, especially since no such points exist in most common representations of (both branching and linear) time, and since in any case we would not want to make predictions beyond such points.

Definition 16 [End-free domain description] A domain description D is *end-free* iff for every h-proposition "A **happens-at** T" in D, there is a $T' \in \Pi$ such that $T \prec T'$. □

Finally, translatable domain descriptions must contain only a finite number of t-, c- and r-propositions:

Definition 17 [T-c-r-finite domain] The domain description $\langle \gamma, \eta, \tau, \rho \rangle$ is *t-c-r-finite* iff γ, τ and ρ are all finite, and for each c-proposition in γ either of the form "A **initiates** F **when** C" or of the form "A **terminates** F **when** C", and for each r-proposition in ρ of the form "L **whenever** C", C is also finite. □

Our translation will also utilise an observation about Definition 14, the least fixed point definition of an initiation and a termination point. When the interpretation to which it refers is a model, and in the case that the domain description is occurrence-sparse and end-free, this definition can be re-cast as the following proposition. Proposition 2 below mirrors the statement of Definition 14, except that H is now a model, and condition (2)(ii) is re-stated in terms of the absence of initiation and termination points.

Proposition 2 Let H be a model of the occurrence-sparse, end-free domain description $D = \langle \gamma, \eta, \tau, \rho \rangle$, let $F \in \Phi$ and let $T \in \Pi$. Then T is an initiation-point (respectively termination-point) for F in H relative to D iff one of the following two conditions holds.

1. There is an $A \in \Delta$ such that (i) there is both an h-proposition in η of the form "A **happens-at** T" and a c-proposition in γ of the form "A **initiates** F **when** C" (respectively "A **terminates** F **when** C") and (ii) H satisfies C at T.

2. There is an r-proposition in ρ of the form "F **whenever** C" (respectively "$\neg F$ **whenever** C") and a partition $\{C_1, C_2\}$ of C such that (i) C_1 is non-empty, for each fluent constant $F' \in C_1$, T is an initiation point for F', and for each fluent literal $\neg F' \in C_1$, T is a termination point for F', and (ii) H satisfies C_2 at T, and for each fluent constant $F' \in C_2$, T is not a termination point for F', and for each fluent literal $\neg F' \in C_2$, T is not an initiation point for F'.

Proof: It is sufficient to show the equivalence of condition 2(ii) of the proposition statement and condition 2(ii) of Definition 14, under the conditions of the proposition that H is a model and D is occurrence-sparse and end-free, and in the particular case where $\langle \mathcal{I}n, \mathcal{T}e \rangle = \langle \mathcal{I}n^f, \mathcal{T}e^f \rangle$ (the least fixed point of the operator $\mathcal{F}$).
"Only if" part: Let T_2 be as defined in condition 2(ii) of Definition 14. Since D is occurrence-sparse there are a finite number (possibly zero) of h-propositions in D which refer to a time-point T^*, where $T \prec T^* \prec T_2$. Hence

there is a (minimal) time-point T'' such that there are no h-propositions in D which refer to a time-point T_1 such that $T \prec T_1 \prec T'' \preceq T_2$. Condition 2(ii) of the proposition, that we need to show, must hold since otherwise, by conditions (2) and (3) of Definition 13 of a model, H would not satisfy C_2 at T''.
"If" part: By Definition 16 of an end-free domain description, there exists a time-point greater than T. Let T' be such a time-point, so that $T \prec T'$. As above, since D is occurrence-sparse, there are a finite number (possibly zero) of h-propositions in D which refer to a time-point T^*, where $T \prec T^* \prec T'$. Hence there is a (minimal) time-point T'' such that there are no h-propositions in D which refer to a time-point T_1 such that $T \prec T_1 \prec T'' \preceq T'$. Hence, by condition (1) of Definition 13, condition 2(ii) of Definition 14 is satisfied with $T_2 = T''$. □

We are now in a position to describe the form of our logic program translations. In [10], this translation assumed a background program module able to correctly compute the order relation $\preceq$. Here we assume that, in addition, this background program is also able to correctly compute or generate the h-propositions of the domain in question. Clearly, where the domain contains only a finite number of h-propositions, this extra requirement can be satisfied simply by listing the (translations of) the h-propositions as individual facts. The background program module, or *basis program*, is defined as follows:

Definition 18 [Basis program] Given the domain description $D = \langle \gamma, \eta, \tau, \rho \rangle$ written in the language $\mathcal{E} = \langle \Pi, \preceq, \Delta, \Phi \rangle$, the program $B(\mathcal{E}, D)$ is a *basis program* for $\mathcal{E}$ and D iff

- for all $T, T' \in \Pi$, $B(\mathcal{E}, D)$ succeeds on the query $T \preceq T'$ if and only if $T \preceq T'$, and finitely fails otherwise.
- for all $T, T' \in \Pi$, $B(\mathcal{E}, D)$ succeeds on the query $T \prec T'$ if and only if $T \prec T'$, and finitely fails otherwise.
- for all $T \in \Pi$ and $A \in \Delta$, $B(\mathcal{E}, D)$ succeeds on the query *Happens-At*(A, T) if and only if the h-proposition "A **happens-at** T" is in D, and finitely fails otherwise.
- the set of answer substitutions for the variable pair $\langle a, t \rangle$ in the (unground) query *HappensAt*(a, t) is $\{\langle A, T \rangle \mid$ "A **happens-at** T" is in $D\}$.
- None of the following predicate symbols appear in $B(\mathcal{E}, D)$: *HoldsAt*, *Given*, *ClippedBetween*, *Whenever*, *PossiblyInitiates*, *Initiates*, *PossiblyTerminates*, *Terminates*, *Starts*, *PossiblyStarts*, *SetStarts*, *PossiblySetStarts*, *ContinuesToHold*, *PossiblyContinuesToHold*, *DisjunctiveForm*, *Partition*, *Resolve*, *NothingHoldsIn*, *ConverseList*.

□

An example basis program, for the situation-calculus-like domain in Section 2.2, is given in Appendix A.

As in [10], the programs described here are sound even when domain descriptions are 'incomplete', i.e. even when, for some fluent constant F and time T, neither "F **holds-at** T" nor "$\neg F$ **holds-at** T" is entailed by D. The

implicit completion often associated with logic programs is avoided by representing t-propositions containing negative fluent literals as *positive* program literals; "$\neg F$ **holds-at** T" is represented by the literal $HoldsAt(Neg(F), T)$. "F **holds-at** T" is represented by the literal $HoldsAt(Pos(F), T)$. The negative literal $not\ HoldsAt(F, T)$ is simply interpreted as 'the t-proposition "F **holds-at** T" is not provable'. (Similar techniques are also used in [20], [7], [6] and [2].) It is convenient to define translation operators λ and $\overline{\lambda}$ which convert $\mathcal{E}$ fluent literals to program terms such as $Pos(F)$ and $Neg(F)$:

Definition 19 [lp-term and lp-complement] Given a fluent literal L of $\mathcal{E} = \langle \Pi, \preceq, \Delta, \Phi \rangle$, the *lp-term* of L, written $\lambda(L)$, is defined to be

- $Pos(F)$ if $L = F$ for some $F \in \Phi$
- $Neg(F)$ if $L = \neg F$ for some $F \in \Phi$

and the *lp-complement* of L, written $\overline{\lambda(L)}$, is defined to be

- $Neg(F)$ if $L = F$ for some $F \in \Phi$
- $Pos(F)$ if $L = \neg F$ for some $F \in \Phi$

□

The full logic program translation of a domain D is defined as an extension of some basis program $B(\mathcal{E}, D)$ for D:

Definition 20 $[LP[B(\mathcal{E}, D)]]$ Given an occurrence sparse, end-free, t-c-r-finite domain description $D = \langle \gamma, \eta, \tau, \rho \rangle$ written in the language $\mathcal{E} = \langle \Pi, \preceq, \Delta, \Phi \rangle$, and a basis program $B(\mathcal{E}, D)$, the logic program $LP[B(\mathcal{E}, D)]$ is defined as the program $B(\mathcal{E}, D)$ augmented with the following domain-specific clauses

- For each t-proposition "L **holds-at** T" in τ, the clause

  $Given(\lambda(L), T).$

- For each r-proposition "L **whenever** $\{L_1, \ldots, L_n\}$" in ρ, the clause

  $Whenever(\lambda(L), [\lambda(L_1), \ldots, \lambda(L_n)]).$

- For each c-proposition "A **initiates** F **when** $\{L_1, \ldots, L_n\}$" in γ, the clause

  $Initiates(A, F, t) \leftarrow$
  $\quad HoldsAt(\lambda(L_1), t), \ldots, HoldsAt(\lambda(L_n), t).$

  and the clause

  $PossiblyInitiates(A, F, t) \leftarrow$
  $\quad not\ HoldsAt(\overline{\lambda(L_1)}, t), \ldots, not\ HoldsAt(\overline{\lambda(L_n)}, t).$

- For each c-proposition "A **terminates** F **when** $\{L_1, \ldots, L_n\}$" in γ, the clause

$$Terminates(A, F, t) \leftarrow$$
$$HoldsAt(\lambda(L_1), t), \ldots, HoldsAt(\lambda(L_n), t).$$

and the clause

$$PossiblyTerminates(A, F, t) \leftarrow$$
$$not\ HoldsAt(\overline{\lambda(L_1)}, t), \ldots, not\ HoldsAt(\overline{\lambda(L_n)}, t).$$

and the following general clauses

$$HoldsAt(l, t) \leftarrow Given(l, t). \tag{LP1}$$

$$HoldsAt(Pos(f), t_3) \leftarrow \tag{LP2a}$$
$$Given(Pos(f), t_1),\ t_1 \prec t_3,$$
$$not\ ClippedBetween(t_1, Pos(f), t_3).$$

$$HoldsAt(Pos(f), t_1) \leftarrow \tag{LP2b}$$
$$Given(Pos(f), t_3),\ t_1 \prec t_3,$$
$$not\ ClippedBetween(t_1, Neg(f), t_3).$$

$$HoldsAt(Neg(f), t_3) \leftarrow \tag{LP2c}$$
$$Given(Neg(f), t_1),\ t_1 \prec t_3,$$
$$not\ ClippedBetween(t_1, Neg(f), t_3).$$

$$HoldsAt(Neg(f), t_1) \leftarrow \tag{LP2d}$$
$$Given(Neg(f), t_3),\ t_1 \prec t_3,$$
$$not\ ClippedBetween(t_1, Pos(f), t_3).$$

$$HoldsAt(l, t_3) \leftarrow \tag{LP3}$$
$$Starts(l, t_1),\ t_1 \prec t_3,\ not\ ClippedBetween(t_1, l, t_3).$$

$$HoldsAt(l, t) \leftarrow \tag{LP4}$$
$$Whenever(l_1, c),\ Resolve(l_1, c, l, t).$$

$$Resolve(l_1, c, l, t) \leftarrow \tag{LP5}$$
$$DisjunctiveForm(l_1, c, d),$$
$$Partition(d, [l], c_1),\ NothingHoldsIn(c_1, t).$$

$$NothingHoldsIn([\,], t). \tag{LP6}$$

$$NothingHoldsIn([Pos(f)|c], t) \leftarrow \tag{LP7}$$
$$HoldsAt(Neg(f), t),\ NothingHoldsIn(c, t).$$

$$NothingHoldsIn([Neg(f)|c], t) \leftarrow \tag{LP8}$$
$$HoldsAt(Pos(f), t),\ NothingHoldsIn(c, t).$$

$$ClippedBetween(t_1, Pos(f), t_3) \leftarrow \tag{LP9}$$
$$t_2 \prec t_3,\ t_1 \preceq t_2,\ PossiblyStarts(Neg(f), t_2).$$

$$ClippedBetween(t_1, Neg(f), t_3) \leftarrow \tag{LP10}$$
$$t_2 \prec t_3,\ t_1 \preceq t_2,\ PossiblyStarts(Pos(f), t_2).$$

$$Starts(Pos(f), t) \leftarrow \quad \text{(LP11)}$$
$$\quad HappensAt(a, t),\ Initiates(a, f, t).$$

$$Starts(Neg(f), t) \leftarrow \quad \text{(LP12)}$$
$$\quad HappensAt(a, t),\ Terminates(a, f, t).$$

$$Starts(l, t) \leftarrow \quad \text{(LP13)}$$
$$\quad Whenever(l, c),\ Partition(c, [l_1|c_1], c_2),$$
$$\quad SetStarts([l_1|c_1], t),\ ContinuesToHold(c_2, t).$$

$$SetStarts([\,], t). \quad \text{(LP14)}$$

$$SetStarts([l_1|c_1], t) \leftarrow \quad \text{(LP15)}$$
$$\quad Starts(l_1, t),\ SetStarts(c_1, t).$$

$$ContinuesToHold([\,], t). \quad \text{(LP16)}$$

$$ContinuesToHold([Pos(f)|c], t) \leftarrow \quad \text{(LP17)}$$
$$\quad HoldsAt(Pos(f), t),\ not\ PossiblyStarts(Neg(f), t),$$
$$\quad ContinuesToHold(c, t).$$

$$ContinuesToHold([Neg(f)|c], t) \leftarrow \quad \text{(LP18)}$$
$$\quad HoldsAt(Neg(f), t),\ not\ PossiblyStarts(Pos(f), t),$$
$$\quad ContinuesToHold(c, t).$$

$$PossiblyStarts(Pos(f), t) \leftarrow \quad \text{(LP19)}$$
$$\quad HappensAt(a, t),\ PossiblyInitiates(a, f, t).$$

$$PossiblyStarts(Neg(f), t) \leftarrow \quad \text{(LP20)}$$
$$\quad HappensAt(a, t),\ PossiblyTerminates(a, f, t).$$

$$PossiblyStarts(l, t) \leftarrow \quad \text{(LP21)}$$
$$\quad Whenever(l, c),\ Partition(c, [l_1|c_1], c_2),$$
$$\quad PossiblySetStarts([l_1|c_1], t),$$
$$\quad PossiblyContinuesToHold(c_2, t).$$

$$PossiblySetStarts([\,], t). \quad \text{(LP22)}$$

$$PossiblySetStarts([l_1|c_1], t) \leftarrow \quad \text{(LP23)}$$
$$\quad PossiblyStarts(l_1, t),\ PossiblySetStarts(c_1, t).$$

$$PossiblyContinuesToHold([\,], t). \quad \text{(LP24)}$$

$$PossiblyContinuesToHold([Pos(f)|c], t) \leftarrow \quad \text{(LP25)}$$
$$\quad not\ HoldsAt(Neg(f), t),\ not\ Starts(Neg(f), t),$$
$$\quad PossiblyContinuesToHold(c, t).$$

$$PossiblyContinuesToHold([Neg(f)|c], t) \leftarrow \quad \text{(LP26)}$$
$$\quad not\ HoldsAt(Pos(f), t),\ not\ Starts(Pos(f), t),$$
$$\quad PossiblyContinuesToHold(c, t).$$

and the following auxiliary definitions (or equivalent):

$Partition([\,],[\,],[\,])$.

$Partition([h|r],[h|r_1],r_2) \leftarrow Partition(r,r_1,r_2)$.

$Partition([h|r],r_1,[h|r_2]) \leftarrow Partition(r,r_1,r_2)$.

$DisjunctiveForm(l,c,[l|c_1]) \leftarrow ConverseList(c,c_1)$.

$ConverseList([\,],[\,])$.

$ConverseList([Pos(f)|c_1],[Neg(f)|c_2]) \leftarrow ConverseList(c_1,c_2)$.

$ConverseList([Neg(f)|c_1],[Pos(f)|c_2]) \leftarrow ConverseList(c_1,c_2)$.

□

(Note that *Resolve* defined in clause (LP5) above is just a simple and naive implementation of a propositional, resolution-based theorem prover for positive or negative literals. $Resolve(l_1,c,l,t)$ means that it is possible to show that l holds by resolution starting from the clause corresponding to $Whenever(l_1,c)$ applied at the instant t. The program first transforms the "implication" of the r-proposition into disjunctive normal form, using the predicate *DisjunctiveForm*, then *Partition* picks out the literal l of interest, and finally shows through the predicate *NothingHoldsIn* that the rest of the disjunction is false, by showing that for each of its literals its negation holds.)

For domain descriptions without r-propositions and with only a finite number of h-propositions, the translation above essentially reduces to the translation given in [10]. However, even in this special case there are some improvements. In particular, the replacement of [10]'s clauses (LP1a) and (LP1b) with clauses (LP2a)–(LP2d) sometimes allows for extra derivations of t-propositions when domains are incompletely specified[3]. Intuitively, it is easy to see that clauses (LP2a)–(LP2d) are justified by Proposition 1. (LP2a) reflects statement (1) of the proposition, (LP2b) reflects the contrapositive of statement (2), (LP2c) reflects statement (2) and (LP2d) reflects the contrapositive of statement (1).

The soundness property for the translation to logic programs given in Definition 20 above is stated as follows.

Proposition 3 Let D be an occurrence sparse, end-free, t-c-r-finite domain description and let $B(\mathcal{E},D)$ be a basis program for $\mathcal{E}$ and D. Then for any fluent literal L of $\mathcal{E}$ and any $T \in \Pi$, if

$$LP[B(\mathcal{E},D)] \vdash_{SLDNF} HoldsAt(\lambda(L),T)$$

then

$$D \models L \textbf{ holds-at } T$$

□

[3] Thanks to Pascal Walheim for suggesting this improvement. He pointed out that, for example, the program in [10] could not derive (the translation of) "F_1 **holds-at** 3" from (the translation of) the domain description {"F_1 **holds-at** 1", "A **happens-at** 2", "A **initiates** F_1 **when** $\{F_2\}$"}.

The proof of this proposition given below is similar to that of the soundness result in [10] (for the translation of the basic $\mathcal{E}$ language into logic programs). It uses induction on the 'length' $length(\alpha)$ of SLDNF derivations α starting from goals of the form $HoldsAt(\lambda(L), T)$ or $Starts(\lambda(L), T)$, where $length(\alpha)$ is defined in Definition 21 below in terms of successful calls to *HappensAt* and *Whenever*. It is defined so that each SLDNF sub-derivation of a *HoldsAt* or *Starts* sub-goal within α has 'length' less than the top-level derivation α.

Definition 21 [$length(\alpha)$] Let α be a successful SLDNF derivation of the goal $HoldsAt(\lambda(L), T)$ or $Starts(\lambda(L), T)$ in $LP[B(\mathcal{E}, D)]$. $length(\alpha)$ is defined inductively as follows:

$$length(\alpha) = S + \sum_{\beta \in B_\alpha} size(\beta)$$

where

- S is the number of successful calls to *HappensAt* and *Whenever* at the top level of α, i.e. not called within a (negation-as-failure) finitely-failed subsidiary derivation of α.
- B_α is the set of all finitely-failed subsidiary derivations from negative calls (to *ClippedBetween* or *PossiblyStarts*) appearing at the top level of α.
- $size(\beta) = 0$ if there is no successful call to *HappensAt* or *Whenever* in any branch of β.
- $size(\beta) = 1 + \max(\{length(\alpha') \mid \alpha' \in A_\beta\})$ if there is a successful call to *HappensAt* or *Whenever* inside β, where A_β is the set of all successful SLDNF derivations of a *HoldsAt* or *Starts* goal, called as a negated sub-goal in one of the branches of β[4].

□

Proof of Proposition 3
Let α be a successful SLDNF derivation of a goal of the form $HoldsAt(\lambda(L), T)$ or $Starts(\lambda(L), T)$ in $LP[B(\mathcal{E}, D)]$. We will use induction on $length(\alpha)$ to show that, for any given model H of D,

- if α is a derivation of $HoldsAt(\lambda(L), T)$ then when $L = F$ for some $F \in \Phi$ then $H(F, T) = true$, and when $L = \neg F'$ for some $F' \in \Phi$ then $H(F', T) = false$.
- if α is a derivation of $Starts(\lambda(L), T)$ then when $L = F$ for some $F \in \Phi$ then T is an initiation-point for F relative to H, and when $L = \neg F'$ for some $F' \in \Phi$ then T is a termination-point for F relative to H.

In the sequel of this proof we will use the term "sound derivation" of a *HoldsAt* or *Starts* goal to mean that it satisfies the appropriate property above.

It is clear that the program $LP[B(\mathcal{E}, D)]$ is entirely symmetrical as regards treatment of literals of the form $Pos(f)$ and $Neg(f)$. This reflects

[4] We assume the convention that $\max(\emptyset) = 0$.

the symmetric nature of $\mathcal{E}$'s semantics. Therefore it is sufficient to consider only the case where $L = F$ for some $F \in \Phi$.

BASE CASE ($length(\alpha) = 0$):

Case 1: α is a derivation of $Starts(Pos(F), T)$.
Clearly, the form of (LP11) and (LP13) requires that a successful derivation of *Starts* has length at least 1. Hence in this case the base case is trivially satisfied.

Case 2: α is a derivation of $HoldsAt(Pos(F), T)$.
The goal can succeed only starting from clauses (LP1), (LP2a) or (LP2b), as success on (LP3) or (LP4) would require $length(\alpha) \geq 1$. We consider each of these possibilities in turn:

(i) Success on (LP1): Trivially, $H(F, T) = true$ by condition (4) of Definition 13 of a model.

(ii) Success on (LP2a): Clearly, the success of $Given(Pos(F), T_1)$ for some $T_1 \prec T$ means that "F **holds-at** T_1" $\in D$ and so $H(F, T_1) = true$ by condition (4) of Definition 13 of a model. Also, since the subsidiary derivation β for $ClippedBetween$ fails with $size(\beta) = 0$, the unground sub-calls to $HappensAt(a, t_2)$ and $Whenever(Neg(F), c)$ (inside the sub-call to $PossiblyStarts(Neg(F), t_2)$) fail. So there are no h-propositions in D, and hence no initiation or termination-points for F relative to H. Hence $H(F, T) = true$ by condition (1) of Definition 13 of a model.

(iii) Success on (LP2b): The proof is exactly analogous to possibility (ii).

INDUCTIVE STEP ($length(\alpha) = n$):

Suppose all successful SLDNF derivations α' of any *HoldsAt* or *Starts* goals which are less than length n are sound in the sense described at the beginning of this proof. To show that it then follows that α is also sound, it is useful to prove some intermediate results as follows.

Intermediate Result 3.1. Suppose for some set of literals C (represented as a list) and some time-point T, $PossiblyContinuesToHold(C, T)$ finitely fails, and that all of its subsidiary derivations of *HoldsAt* and *Starts* sub-goals are sound. Then either H (the model of D fixed in this proof) does not satisfy C at T, or for some fluent constant $F \in C$, T is a termination-point for F, or for some fluent literal $\neg F \in C$, T is an initiation-point for F. (In other words, the condition 2(ii) in Proposition 2 is not satisfied for this specific set C of fluent literals.)

Proof of Intermediate Result 3.1. Clearly, the form of clauses (LP24)–(LP26) ensures that the finitely failed call to *PossiblyContinuesToHold* must contain a successful call to *HoldsAt* or *Starts*, which are sound by assumption. *(End of proof of Intermediate Result 3.1.)*

Intermediate Result 3.2. Suppose for some fluent literal L and some time-point T, $PossiblyStarts(\lambda(L), T)$ finitely fails and that all of its subsidiary derivations of *HoldsAt* and *Starts* sub-goals are sound. Then if $\lambda(L) = Pos(F)$ for some F then T is not an initiation-point for F, and if $\lambda(L) = Neg(F)$ for some F then T is not a termination-point for F.

Proof of Intermediate Result 3.2. As remarked in the beginning of the main proof it is sufficient to consider only the case when $\lambda(L) = Pos(F)$ for some F. The proof uses a simple induction on the maximum number k of successful sub-calls to *HappensAt* and *Whenever* within any branch of the finitely failed derivation $PossiblyStarts(Pos(F), T)$. Clearly for the base case of $k = 0$ the result holds.

For $k > 0$ (the inductive step) we argue as follows. The failure of *PossiblyStarts* on clause (LP19) means that either $HappensAt(a, T)$ fails or $PossiblyInitiates(A, F, T)$ fails for all A such that $HappensAt(A, T)$ succeeds. Clearly for the former case there can be no initiation-point for F at T. In the latter case, either there are no matching c-propositions in the domain description D, or for any matching c-proposition "A **initiates** F **when** C" there exists an $L_i \in C$ such that $HoldsAt(\overline{\lambda(L_i)}, T)$ succeeds. Therefore, by the assumption of soundness for this call, $H(L_i, T) = false$, and hence condition (1) of Proposition 2 cannot be satisfied.

Failure of *PossiblyStarts* on clause (LP21) means that either there is no r-proposition of the form "F **whenever** C" and therefore condition (2) of Proposition 2 cannot be satisfied, or for any such r-proposition and any partition $\{C_1, C_2\}$ of C such that C_1 is non-empty, either $PossiblySetStarts(C_1, T)$ fails or $PossiblyContinuesToHold(C_2, T)$ fails. In the former case it is easy to see that $PossiblyStarts(L_i, T)$ fails for some $L_i \in C_1$, with maximum number of successful sub-calls to *HappensAt* and *Whenever* strictly less than k. Therefore by the (local) inductive hypothesis T is not an initiation-point for L_i if L_i is a positive literal and not a termination point for F_i if $L_i = \neg F_i$ for some F_i. Hence condition (2) of Proposition 2 cannot be satisfied. In the latter case (where $PossiblyContinuesToHold(C_2, T)$ fails) Intermediate Result 3.1 ensures that condition (2) of Proposition 2 cannot be satisfied. Hence in either case this condition is not satisfied. Therefore the required result holds by Proposition 2, as we have shown that neither of its two conditions are satisfied. *(End of proof of Intermediate Result 3.2.)*

Intermediate Result 3.3. Suppose for some set of literals C (represented as a list) and some time-point T, $ContinuesToHold(C, T)$ succeeds, and that all of its subsidiary derivations of *HoldsAt* and *Starts* sub-goals are sound. Then H satisfies C at T, and for each fluent constant $F \in C$, T is not a termination-point for F, and for each fluent literal $\neg F \in C$, T is not an initiation-point for F. (In other words, the condition 2(ii) in Proposition 2 is satisfied for this specific set C of fluent literals.)

Proof of Intermediate Result 3.3. The result is clearly true if C is empty. If C is not empty, we can easily see from the form of clauses (LP17) and (LP18) that for each literal $\lambda(L) \in C$, it is the case that both $HoldsAt(\lambda(L), T)$ succeeds and $PossiblyStarts(\overline{\lambda(L)}, T)$ fails. The result therefore follows directly from the assumption of soundness for subsidiary calls to *HoldsAt* and Intermediate Result 3.2. *(End of proof of Intermediate Result 3.3.)*

We now return to the main body of the inductive step, and consider the successful derivation α of length n.

Case 1: α is a derivation of $Starts(Pos(F), T)$.
There are two possibilities to consider:

(i) Success on (LP11): Clearly, since this involves a successful call to *HappensAt*, all subsidiary calls to *HoldsAt* within subsidiary calls to

domain-specific *Initiates* clauses are sound, and hence T is an initiation-point for F in H by condition (1) of Proposition 2.

(ii) Success on (LP13): Since this involves a successful call to *Whenever*, the subsequent successful derivations of *Starts* within *SetStarts* are of length strictly less than n and hence by the inductive hypothesis are sound. Similarly, the successful call to *ContinuesToHold* satisfies the condition of Intermediate Result 3.3 and hence the conclusion of this intermediate result holds. Hence condition (2) of Proposition 2 is satisfied so that T is an initiation-point for F in H.

Case 2: α is a derivation of $HoldsAt(Pos(F), T)$.
We first observe that the finitely failed call to $ClippedBetween$ required in clauses (LP2a), (LP2b) and (LP3) is given by a finitely failed sub-call to *PossiblyStarts*. From the forms of clauses (LP19)–(LP21) it is clear that any successful sub-derivation of *HoldsAt* or *Starts* within this sub-call to *PossiblyStarts* have length strictly less than n and hence by the inductive hypothesis are sound, and therefore Intermediate Result 3.2 applies. Hence, from the form of the clauses (LP9) and (LP10), the success of a $not\, ClippedBetween(T_1, \lambda(L), T_3)$ implies that when $L = F'$ for some fluent F' there is no termination-point for F' between T_1 and T_3, and when $L = \neg F'$ for some fluent F' there is no initiation-point for F' between T_1 and T_3.

There are four possibilities to consider:

(i) Success on (LP2a): In this case the success of $(Given(Pos(F), T_1)$, $T_1 \prec T)$ for some T_1 means that $H(F, T_1) = true$. Furthermore, by the observation above, the success of $not\, ClippedBetween(T_1, Pos(F), T)$ means that there is no termination-point for F between T_1 and T. Hence $H(F, T) = true$ by case (1) of Proposition 1.

(ii) Success on (LP2b): In this case the proof is exactly analogous to that of possibility (i), using (the contrapositive of) case (2) of Proposition 1 instead of case (1).

(iii) Success on (LP3): This success requires a successful derivation of $Starts(Pos(F), T_1)$ for some T_1 less than T, whose length m is not greater than n. Hence by the inductive hypothesis $(m < n)$ and the result of Case 1 above $(m = n)$, this derivation of $Starts(Pos(F), T_1)$ is sound, i.e. T_1 is an initiation point for F. Since $not\, ClippedBetween(T_1, Pos(F), T_3)$ also succeeds, by the observation above there is no termination-point for F between T_1 and T_3. Therefore by condition (2) of Definition 13 of a model, $H(F, T) = true$.

(iv) Success on (LP4): The program definition of $Resolve$ given in (LP5) implements a resolution procedure. Success on (LP4) reduces to a conclusion of this resolution procedure from a domain-dependent clause corresponding to an r-proposition "F **whenever** C". For a fixed T, this r-proposition can be treated as a classical implication in the model H by condition (4) of Definition 13 of a model. These resolution steps use subsidiary derivations of *HoldsAt* sub-goals which, due to the successful call to *Whenever* in (LP4), have length strictly less than n and therefore are sound. The resolution is therefore sound and thus the derived conclusion $HoldsAt(Pos(F), T)$ is sound, i.e. $H(F, T) = true$.

(End of proof of Proposition 3.) □

The program above has been used as the basis for Prolog implementations of various Language $\mathcal{E}$ domains with ramifications. (For some example

program listings, see http://www.ucl.ac.uk/~uczcrsm/abstract15.html.) In some of these implementations it was useful to introduce a loop checking mechanism for avoiding loops inside the *Resolve* predicate. It is possible that further extensions can be introduced to make this type of Prolog implementation more complete. This may involve further loop checking or extra rules to cover cases of reasoning not handled by this program, such as reasoning backwards to necessary preconditions of actions.

5 Discussion and Related Work

In this paper we have shown how the Language $\mathcal{E}$, extended with ramification statements or r-propositions, provides a simple, intuitive framework for representing and reasoning about direct and indirect effects of action occurrences embedded in an independent time structure. Time may be linear or branching, continuous or discrete. We have also given a general method of translating domain descriptions with r-propositions into ordinary (Prolog-like) Event Calculus style logic programs, and shown that this translation is sound.

In our formalisation, r-propositions play a double role. On the one hand, the proposition "L **whenever** C" acts as a static domain constraint – Definition 13 allows us to reason that whenever everything in C holds, L holds, and whenever L doesn't hold, something in C doesn't hold. On the other hand, it generates new causal information as far as the dynamic behaviour of the domain is concerned, through indirect initiation and termination. Roughly, when C is initiated, L is also initiated. This dymanic aspect is unidirectional; it is not the case that when L is terminated, C is terminated.

The issue of ramifications has recently received widespread attention in the context of reasoning about actions and change. As has been noted by authors such as Sandewall [26] and Thielscher [28], most recent approaches to ramification can be classed as either *categorization-based* or *causal-based*. (Not all approaches readily fit into this classification, however. For example the recent approach of McIlraith [18, 19] has both a causal and categorization flavour.) Categorization-based approaches overcome problems with more naïve applications of the principle of minimal change by (temporarily or permanently) categorizing fluents into two or more classes. The principle of minimisation of change can then be applied more stringently to some of these classes than to others. Variations on this theme can be found, for example, in the approaches of Lifschitz [14] (in his use of *frame* and *non-frame* fluents), in Sandewall [24] [26] and Gustafsson & Doherty [9] (in their use of *occlusion*), and in Kartha & Lifschitz [12] and Shanahan [27] (in their use of a *releases* mechanism or predicate). Such approaches deal easily with domains such as the stuffy room example described in Section 3 (for example, *Stuffy* would simply be a non-frame fluent in [14], or released in [12] and [27]). However, as pointed out in [28], these approaches cope less easily with domains such as Thielscher's circuit example (also described here, in Section 3).

The approach to ramification described here bears a stronger resemblance to causal-based approaches, such as Thielscher's [28], McCain & Turner's [16, 17], Lin's [15] and Turner's [29]. "backwards" through r-propositions, and this type of unidirectionality of indirect effects is a key characteristic of causal-based techniques. However, there are difficulties in forming detailed comparisons between $\mathcal{E}$ and some of these approaches. This is because, whereas $\mathcal{E}$ is a narrative-based formalism in which the notion of

an action occurrence within some explicit time structure is central, these formalisms are state-based and often do not incorporate an explicit notion of time or of an occurrence. The emphasis is instead on the type of "hypothetical reasoning" about sequences of actions supported by the Situation Calculus. Although $\mathcal{E}$ supports this type of reasoning (as illustrated in Section 2), it can also deal with more arbitrary, independent time structures. (In [17], McCain & Turner do introduce explicit time, but the time structure must be linear and discrete.) In our semantics, change is characterised by a series or collection of initiation and termination points, which are time-points at which particular actions occur. In other words, change is "caused" by explicitly represented, temporally embedded events. Explicit reference to the flow of time has made it easy to generalise the semantics so that it covers ramifications, i.e. indirect cause and effect. The unidirectionality of r-propositions is initimately related to the unidirectionality of this independent flow of time. Note that the idea of initiation and termination points is already implicit in the original Event Calculus [13]. Our contribution has been to make these notions explicit, and hence amenable to generalisation, in the semantics.

For Thielscher's approach [28] at least, the extension to deal with explicit time and narrative information seems likely to be a source of difficulty. Thielscher computes ramifications by applying an *action description* to an initial state, and then generating a sequence of temporary states using *causal relationships* until a *successor state* is reached which satisfies all the *domain constraints* of the system. If explicit time is added to this schema, the temporal status of these temporary states is questionable. If the temporary states are somehow assigned to successive time-points, the domain constraints may be violated at these time-points, and extra mechanisms will have to be introduced to deal with interruptions to the propagation of effects due to other action occurrences. If on the other hand these temporary states are "outside time", so that the overall change from initial state to successor state happens instantaneously or in a single time step, it seems that Thielscher's framework will sometimes give rise to anomalous models. For example, taking Thielscher's light-detector example (from [28], where a light detector is permanently activated whenever a light bulb emits light) we would, as explained in [28], be able to generate a model where the detector becomes activated even though there is no time-point at which the light bulb is on.

In their Language $\mathcal{ER}$ [30] [31] [32], Van Belleghem, Denecker & Dupré introduce some interesting extensions and variations to the ideas presented in the present paper and in [10]. Roughly speaking (and among numerous other innovations), they extend the expressivity given by r-propositions in $\mathcal{E}$ as follows. When re-expressing the r-proposition "F **whenever** C" as a *derived effect rule* in $\mathcal{ER}$, the syntax allows the user to stipulate that, when constructing a model, only some partitions of the set C may be applied within a definition analogous to condition (2) of Definition 14. This extra expressivity is motivated by an example involving the activation of an alarm system. Domain descriptions written in $\mathcal{ER}$ can also include simple *state constraints*, further adding to its expressivity. However, $\mathcal{ER}$ is arguably less general than $\mathcal{E}$ in the sense that it uses only a particular linear (real number) model of time. Aside from its specific contributions in terms of representational techniques, the development of $\mathcal{ER}$ illustrates that, as for action description languages such as $\mathcal{A}$ and $\mathcal{AR}$, "event description languages" (e.g. $\mathcal{E}$ and $\mathcal{ER}$) can form part of a spectrum, in which each individual language incorporates trade-offs between simplicity, succinctness

and range of applicability. (There are also some interesting technical differences between $\mathcal{E}$'s and $\mathcal{ER}$'s semantics, which are discussed in [30] and deserve further investigation. These regard certain combinations of ramification statements which give rise to non-determinism in $\mathcal{E}$ but inconsistency in $\mathcal{ER}$. See http://www.ucl.ac.uk/~uczcrsm/abstract15.html for one such example.)

The techniques discussed in [10] for *explanation* and *planning* in terms of $\mathcal{E}$ are equally applicable where domains include ramifications described by r-propositions. In particular, suppose that we wish to express narrative information as a sequence of observations of fluent values at different time-points (i.e. as a set of t-propositions), rather than as a sequence of action occurrences (i.e. a set of h-propositions). In these circumstances we may use the notion of *h-entailment*, defined in [10] as a generalisation of $\mathcal{E}$'s ordinary entailment, to compute other consequences of our domain theory. To illustrate, we take (and slightly modify) a simple example of Sandewall, which he introduces in [26] to illustrate a limitation of current causal approaches to ramification. Sandewall asks us to consider the case of a lamp LA which is turned on independently by two different switches. Sandewall describes this with the classical constraint $(LA \leftrightarrow (SW_1 \vee SW_2))$. In $\mathcal{E}$ we would use three r-propositions:

$\neg LA$ **whenever** $\{\neg SW_1, \neg SW_2\}$ (la1)
LA **whenever** $\{SW_1\}$ (la2)
LA **whenever** $\{SW_2\}$ (la3)

Closing either of the switches activates that switch, and (taking time as the naturals) initially the lamp is off (note that the semantics of $\mathcal{E}$ then implies that neither switch is initially closed):

$Close SW_1$ **initiates** SW_1 (la4)
$Close SW_2$ **initiates** SW_2 (la5)
$\neg LA$ **holds-at** 0 (la6)

Sandewall's point is essentially that at some later time-point, say 5, we might wish to simply state that LA becomes true, without explicitly stating that one of the switches has been closed. In terms of $\mathcal{E}$, we simply want to state:

LA **holds-at** 5 (la7)

We should then be able to infer that at 5 at least one of the switches is closed. Using the definitions of h-models, h-consistency and h-entailment in [10], it is straightforward to verify that the domain description consisting of (la1)-(la7) is h-consistent, and that in all h-models of this domain, either SW_1 is true or SW_2 is true at all times greater than or equal to 5, as intended by the example.[5] Note that for all other examples given in this paper, h-entailment coincides exactly with ordinary entailment.

There is considerable scope for further work on the themes of this paper. First, it would be instructive to investigate more formally and systematically the correspondence between $\mathcal{E}$ and other formalisms which deal with ramifications. Much of the discussion here has centered around a few benchmark

[5] A trivial extension to the syntax of $\mathcal{E}$ and to condition (4) in Definition 13 of a model would also allow us to write and assign meaning to t-propositions with disjunctions, such as "$(SW_1 \vee SW_2)$ **holds-at** 5". These kinds of extensions are discussed more generally in [31].

problems. Sandewall [24] and others have argued the benefits of a more systematic methodology for assessing and comparing different formalisms for reasoning about actions. As regards ramifications, Sandewall has already undertaken an extensive analysis of various categorization-based formalisms [26], comparing them both pairwise and against his own "semantics-level" template. It would certainly be beneficial to apply this same methodology to assess the range of applicability of $\mathcal{E}$. (In Sandewall's terminology, $\mathcal{E}$ might best be described as a "semantics-level" formalism, since by itself it offers no proof theory. Like Sandewall, our primary intention has been to give a natural and intuitive definition of a model-theoretic entailment relation, against which we have measured the correctness of a class of logic programs.)

We would also like to investigate the translation of $\mathcal{E}$ domains into other general-purpose formalisms, such as classical logic (Costello [4] has also conducted preliminary investigations in this respect), abductive logic programming, and argumentation frameworks. In this paper the translation was into logic programs with negation-as-failure (NAF). In analysing properties of this translation we have used only a basic (and limited) understanding of NAF, in terms of the the SLDNF proof procedure. Any of the more sophisticated semantics of logic programs subsumes this very procedural understanding. Our translation and analysis therefore provides and verifies a baseline implementation (which can be run more-or-less directly under a standard Prolog interpreter, see http://www.ucl.ac.uk/~uczcrsm/abstract15.html for examples) but could almost certainly be improved in several ways. For example, the incompleteness of the translation could probably be narrowed down if more general forms of NAF were to be used. Related to this is the issue of the ease with which it is possible to prove the correctness of a given translation. The proof of Proposition 3 is perhaps less than ideal, as it appeals to the specific structure of the particular program rather than utilizing more general results in logic programming. A translation into other more general or high-level frameworks, such as an argumentation framework, can perhaps make this task more straightforward. In fact, Dimopoulos & Kakas [5] have already developed a translation of the Language $\mathcal{A}$ into an argumentation framework, and shown this to be complete for consistent domains. We are currently investigating analogous translations of the Language $\mathcal{E}$ with similar properties. Turner's work [29] in providing a sound and complete translation of a causal extension of $\mathcal{A}$ into logic programs is also relevant here.

There are many ways in which the expressiveness of $\mathcal{E}$ might be extended (as has already been partly illustrated in [30], [31] and [32]). The syntax of $\mathcal{E}$ already allows for *concurrent* action occurrences, and it seems likely that it could be extended to enable the description of interactions between, and combinations of effects of, concurrent actions. Since in $\mathcal{E}$ the flow of time may be continuous, it is a good candidate formalism for dealing with *continuous change*, perhaps by adapting some of the ideas in [21]. Related to this is the issue of *triggered events* (or *natural actions*). Preliminary investigations have already shown that, since $\mathcal{E}$ is narrative-based, it is relatively straightforward to extend its syntax and semantics to include and assign meaning to statements such as "C **triggers** A" – i.e. "at time points when the set of fluent literals C holds, action A occurs". Finally, we believe that narrative based formalisms such as $\mathcal{E}$ can offer interesting perspectives on issues relating to the *qualification problem*. For example, they enable a distinction to be made between circumstances under which it is impossible for an action to occur, and circumstances under which is it impossible to

predict the effects of an (attempt to perform an) action. Some preliminary ideas for qualification statements in $\mathcal{E}$ are included in the appendix of [10], but much further work needs to be done.

Prolog listings for some of the example domains in this paper are available on the World Wide Web at http://www.ucl.ac.uk/~uczcrsm/abstract15.html.

Acknowledgements

We would particularly like to thank Kristof Van Belleghem, Tom Costello, Michael Gelfond, Michael Thielscher, Hudson Turner and Pascal Walheim for helpful comments relating to various aspects of this paper. This research was partially funded by the U.K. Engineering and Physical Sciences Research Council (EPSRC) and the University of Cyprus.

Appendix A: An Example Basis Program

The following is an example basis program (see Definition 18) for the situation-calculus-type domain described at the end of Section 2.2. (It assumes an appropriate definition of $Reverse$, where $Reverse(l_1, l_2)$ is true if the list l_2 is the reverse of the list l_1.)

$Action(Load).$

$Action(Look).$

$Action(Take).$

$ActionList([]).$

$ActionList([a, l]) \leftarrow ActionList(l),\ Action(a).$

$HappensAt(a, Start(l)) \leftarrow Reverse(l, [a|l_1]),\ ActionList(l_1),\ Action(a).$

$t_1 \prec_\Delta t_2 \leftarrow OneLessThan(t_1, t_2).$

$t_1 \prec_\Delta t_3 \leftarrow OneLessThan(t_2, t_3),\ t_1 \prec_\Delta t_2.$

$t \preceq_\Delta t.$

$t_1 \preceq_\Delta t_2 \leftarrow t_1 \prec_\Delta t_2.$

$OneLessThan(S_0, Start([a])) \leftarrow Action(a).$

$OneLessThan(Start([a|l]), [a|l]) \leftarrow ActionList([a|l]).$

$OneLessThan([a|l], Start([a_1, a_2|l_1])) \leftarrow$
$\quad Reverse([a_1, a_2|l_1], [a_3|l_2]),\ Reverse(l_2, [a|l]).$

The Prolog code for this basis program is available on the World Wide Web at http://www.ucl.ac.uk/~uczcrsm/abstract15.html.

References

[1] Andrew Baker, *Nonmonotonic Reasoning in the Framework of the Situation Calculus*, Artificial Intelligence, Vol 49(5-23), 1991.

[2] Chitta Baral and Michael Gelfond, *Representing Concurrent Actions in Extended Logic Programming*, Proceedings IJCAI 93, Morgan Kaufmann, page 866, 1993.

[3] Chitta Baral, Michael Gelfond and Alessandro Provetti, *Representing Actions: Laws, Observations and Hypotheses*, The Journal of Logic Programming, Vol 31(1-3) (Special Issue on Reasoning about Action and Change), Elsevier Science, 1997.

[4] Tom Costello, *Putting $\mathcal{E}$ into Classical Logic*, ETAI interaction pages, http://www.ida.liu.se/ext/etai/received/actions/003/aip.html#003, 1997.

[5] Yannis Dimopoulos and Antonis Kakas, *Logic Programming without Negation as Failure*, Proceedings of the 1995 International Logic Programming Symposium (ILPS'95), pages 369-383, MIT Press, 1995.

[6] Phan Minh Dung, *Representing Actions in Logic Programming and its Applications in Database Updates*, Proceedings of the Tenth International Conference on Logic Programming, ed David S. Warren, MIT Press, pages 222-238, 1993.

[7] Michael Gelfond and Vladimir Lifschitz, *Representing Actions in Extended Logic Programming*, Proceedings of the Joint International Conference and Symposium on Logic Programming, ed. Krzysztof Apt, MIT Press, page 560, 1992.

[8] Michael Gelfond and Vladimir Lifschitz, *Representing Action and Change by Logic Programs*, Journal of Logic Programming, volume 17 numbers 2, 3 and 4, North-Holland, pages 301-322, 1993.

[9] Joakim Gustafsson and Patrick Doherty, *Embracing Occlusion in Specifying the Indirect Effects of Actions*, in Proceedings of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR'96), ed.s Aiello, Doyle and Shapiro, Morgan Kaufmann, 1996.

[10] Antonis Kakas and Rob Miller, *A Simple Declarative Language for Describing Narratives with Actions*, The Journal of Logic Programming, Vol 31(1-3) (Special Issue on Reasoning about Action and Change), pages 157-200, Elsevier Science, 1997.

[11] G. Neelakantan Kartha, *Soundness and Completeness Theorems for Three Formalizations of Action*, Proceedings IJCAI 93, page 724, 1993.

[12] G. Neelakantan Kartha and Vladimir Lifschitz, *Actions with Indirect Effects (Preliminary Report)*, in Proceedings of the Fourth International Conference on Principles of Knowledge Representation and Reasoning, pages 341-350, 1994.

[13] Robert Kowalski and Marek Sergot, *A Logic-Based Calculus of Events*, New Generation Computing, vol 4, page 267, 1986.

[14] Vladimir Lifschitz, *Frames in the Space of Situations*, Artificial Intelligence Journal, Volume 46, pages 365-376, 1990.

[15] Fangzhen Lin, *Embracing Causality in Specifying the Indirect Effects of Actions*, In C. S. Mellish, ed., Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'95), pages 1985-1991, 1995.

[16] Norman McCain and Hudson Turner, *A Causal Theory of Ramifications and Qualifications*, In C. S. Mellish, ed., Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'95), pages 1978-1983, 1995.

[17] Norman McCain and Hudson Turner, *Causal Theories of Action and Change*, In Proceedings of AAAI'97, 1997.

[18] Sheila McIlraith, *A Closed-Form Solution to the Ramification Problem (Sometimes)*, In working notes of the IJCAI'97 workshop on Nonmonotonic Reasoning, Action and Change, 1997.

[19] Sheila McIlraith, *Representing Actions and State Constraints in Model-Based Diagnosis*, In Proceedings of AAAI'97, 1997.

[20] Rob Miller, *Situation Calculus Specifications for Event Calculus Logic Programs*, in Proceedings of the Third International Conference on Logic Programming and Non-monotonic Reasoning, Lexington, KY, USA, Springer Verlag, 1995.

[21] Rob Miller and Murray Shanahan, *Reasoning about Discontinuities in the Event Calculus*, in Proceedings of the Fifth International Conference on Principles of Knowledge Representation and Reasoning (KR'96), ed.s L. C. Aiello, J. Doyle & S. C. Shapiro, Morgan Kaufmann, 1996.

[22] Edwin Pednault, *ADL: Exploring the Middle Ground between STRIPS and the Situation Calculus*, In Ronald Brachman, Hector Levesque and Ray Reiter, ed.s, Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning, pages 324-332, 1989.

[23] Ray Reiter, *The Frame Problem in the Situation Calculus: a Simple Solution (sometimes) and a Completeness Result for Goal Regression*, in Vladimir Lifschitz, ed., Artificial Intelligence and Mathematical Theory of Computation, pages 359-380, Academic Press, 1991.

[24] Erik Sandewall, *Features and Fluents: A Systematic Approach to the Representation of Knowledge about Dynamical Systems*, Oxford University Press, 1994.

[25] Erik Sandewall, *Reasoning about Actions and Change with Ramification*, in Computer Science Today, volume 1000 of LNCS, Springer, 1995.

[26] Erik Sandewall, *Assessments of Ramification Methods that Use Static domain Constraints*, in Proceedings of the 5th International Conference on Principles of Knowledge Representation and Reasoning (KR'96), ed.s Aiello, Doyle and Shapiro, Morgan Kaufmann, 1996.

[27] Murray Shanahan, *Solving the Frame Problem: A Mathematical Investigation of the Commonsense Law of Inertia*, MIT Press, 1997.

[28] Michael Thielscher, *Ramification and Causality*, Artificial Intelligence, 1997.

[29] Hudson Turner, *Representing Actions in Logic Programs and Default Theories: a Situation Calculus Approach*, The Journal of Logic Programming, Vol 31(1-3) (Special Issue on Reasoning about Action and Change), Elsevier Science, 1997.

[30] Kristof Van Belleghem, *Open Logic Programming as a Knowledge Representation Language for Time-dependent Problem Domains*, Ph.D. thesis, K. U. Leuven, to be submitted, 1997.

[31] Kristof Van Belleghem, Marc Deneker and Daniele Theseider Dupré, *Ramifications and Dependencies in an Event-based Language*, submitted for publication, 1997.

[32] Kristof Van Belleghem, Marc Deneker and Daniele Theseider Dupré, *Ramifications in an Event-based Language*, in Proceedings of the 9th Dutch Artificial Intelligence Conference (NAIC 97), 1997.