Linköping Electronic Articles in
Computer and Information Science
Vol. 1(1996): nr 3

Assessments of Ramification Methods that Use Static Domain Constraints

Erik Sandewall

Linköping University Electronic Press
Linköping, Sweden

http://www.ep.liu.se/ea/cis/1996/003/

Published on November 5, 1996 by
Linköping University Electronic Press
581 83 Linköping, Sweden

Linköping Electronic Articles in
Computer and Information Science
ISSN 1401-9841
Series editor: Erik Sandewall

Typeset by the author using LaTeX
Formatted using étendu *style*

Recommended citation:
<Author>. <Title>. Linköping Electronic Articles in Computer and Information Science, Vol. 1(1996): nr 3. http://www.ep.liu.se/ea/cis/1996/003/. *November 5, 1996.*

This URL will also contain a link to the author's home page.

Abstract

The paper analyzes and compares methods for ramification, or for update, which are based on minimization of change while respecting static domain constraints. The main results are:

- An underlying semantics for propagation oriented ramification, called *causal propagation semantics*
- *Comparative assessments* of a number of previously proposed entailment methods for ramification, in particular, methods based on minimization of change after partitioning of the state variables (fluents). This includes the methods previously proposed by del Val and Shoham, by Kartha and Lifschitz, and by our group.
- Assessment of two *ranges of sound applicability* for one of those entailment methods, MSCC, relative to causal propagation semantics.
- Identification of an essential feature for minimization-based approaches to ramification, namely *changeset-partitioning.*

This article is concurrently published in:
L. Aiello, J. Doyle, and S. Shapiro (eds):
Principles of Knowledge Representation and Reasoning.
Proceedings of the Fifth International Conference.
Morgan-Kaufmann, 1996.

A few minor technical mistakes in the proceedings version have been corrected in the present issue.

1 Topic and approach

1.1 The ramification problem

The *ramification problem* for reasoning about actions and change has been described as follows by Ginsberg and Smith, [6]:

> The difficulty is that it is unreasonable to explicitly record all of the consequences of an action, even the immediate ones. For example [...] For any given action there are essentially an infinite number of possible consequences that depend upon the details of the situation in which the action occurs.

More concretely, the difference between ramification and the simpler case of strict inertia is that strict inertia assumes action laws to specify *all* the effects of the action explicitly, whereas in ramification the action law only needs to specify *some of* the effects; the others are to be deduced. On the other hand, the logic cannot do magic; those other effects must also be declared somehow. One can think of two specific reasons why ramification is useful:

1. In order to avoid duplication of the same information on the consequent side of several action laws. Suppose, for a trivial example, that both $alive(x)$ and $dead(x)$ occur as state variables. Under strict inertia, every action that causes a change in $alive(x)$ must also specify the corresponding change in $dead(x)$.

2. In order to avoid configuration-dependent action laws. Suppose the scenario world consists of a number of interconnected objects, where a change in one of them may cause changes in neighboring objects, and different scenarios in the same world contain different configurations. If the same action law is going to be used in all scenarios, it must apply regardless of the choice of configuration, and then it is not possible to represent the secondary changes explicitly in the action law.

We shall refer to these as the *redundancy* reason and the *propagation* reason for ramification, respectively.

1.2 Current approaches

Most solutions which have been proposed for the ramification problem require action laws to specify the most significant effects of the action, and to rely on *domain constraints* for specifying additional changes that are due to the action. However, the imposition of domain constraints alone is not sufficient. Two courses have been taken. *Minimization based approaches* assume that the total set of changes in state variables that result from the action, consists of those changes that are explicitly stated in the action law, and a minimal set of other changes while satisfying the domain constraints. *Causation oriented approaches* require instead that some information is available about how changes in one state variable may "cause" changes in another state variable, and accept those models which satisfy action laws and domain constraints, where only correctly "caused" changes are present. Domain-constraint and influence information are sometimes integrated, sometimes separated.

The *update problem* is in principle the same problem: given a world state before an action, some information about the world state after the action, and certain general restrictions on permissible world states, which are the possible or preferred world states resulting from the action? The main difference is one of presentation: the ramification problem is presented in the context of reasoning about actions and change, where the actions themselves as well as their timing and sequencing is explicit in the logical language being used. For the update problem, by contrast, one only considers one update, and leaves aside the question of timing and sequencing of the multiple actions.

Within the general frameworks of minimization based and causation oriented approaches, there are a number of different specific proposals which have been advanced by different authors, and which differ in their details. So far, these methods have only been motivated on the basis of (1) intuitive and informal reasons why a particular method is the "natural" one, and (2) selected test examples where competing methods do not give intended results.

In earlier articles [17, 18] we have argued that this kind of evidence is not sufficient, and that proposed methods for nonmonotonic common-sense reasoning should be *assessed* with respect to an *underlying semantics* which formally defines the intended conclusions for a class of scenario descriptions. We have also performed such assessment of the range of applicability for about one dozen entailment methods dealing with strict inertia, that is, persistence without ramification. Yi, in [25], has generalized the underlying semantics to allowing concurrent actions, and has extended the assessment of range of applicability to one method for this case.

1.3 Varieties of assessment in the case of ramification

The present paper describes the first assessment results for current approaches to ramification. It turns out that the assessment methodology has to be revised in this case, for the following reasons. In the case of strict inertia, it was possible to obtain the range for completeness, that is, the range where the set of selected models *equals* the set of intended models. Ramification appears to be more difficult, and with the underlying semantics proposed below we have only been able to obtain certain results on the range of soundness, that is, the range where the set of selected models *contains* the set of intended models.

For many practical purposes, soundness is good enough. Methods which are sound but not complete involve of course a certain loss of information, but information is lost anyway in several other ways in any practical system. However, if we merely have soundness, then it is also of interest to know something about how well or how badly the set of selected models approximates the set of intended models.

For this reason, we propose to complement the *soundness assessment* with a *comparative assessment* where entailment methods are compared pairwise with respect to the extent of the model sets that they select. The present paper begins with an overview of minimization based entailment methods for ramification. Then follow the results on comparative assessment, then the proposed underlying semantics which is called *causal propagation semantics*, and finally some results on assessment of range of sound applicability wrt causal propagation semantics. This ordering is natural since the comparative assessment is technically easier to follow, and since it has been performed for a larger number of entailment methods.

The comparative assessment results are based on a more extensive departmental technical report [20] containing the full formal machinery and the full proofs. The formal aspects have previously been summarized in

[19], although the main topic of that article was different from the present one.

2 Pros and cons of current approaches

2.1 Plain minimization of change

The original method, plain minimization of change, is originally due to Winslett [24], at least in the context of the A.I. literature. The idea is simple: identify a set of *static domain constraints*, that is, restrictions on the possible states of the world. For a given action, the effects specified by the action law are obligatory; a minimal set of additional changes is selected so that the domain constraints are satisfied.

It was observed already by Winslett that plain minimization of change sometimes obtains unintended results. She suggested to solve this by imposing a priority ordering on the state variables, but this line of research has apparently not been pursued since. The following approaches have been developed instead.

2.2 Categorization of state variables

Categorization of state variables is a fairly simple modification of plain minimization of change: instead of establishing certain changes in the world, and then minimizing changes across all (other) state variables, one divides the state variables into two or more classes which are treated differently with respect to minimization.

One instance of this concept was the introduction of *occlusion* by Sandewall [15, 17], also represented by the causal explanation predicate of Persson and Staflin [13], the *persists* predicate in [14], the *release* operation of Kartha and Lifschitz [8] and the *persists* predicate of del Val and Shoham in [4]. In the simplest case, one separates the state variables into two classes A and B, where the action laws specifying the effects of an action only specify changes in class A, whereas domain constraints are allowed to refer to state variables in both classes. Then, change is minimized for state variables in class B, whereas those in class A are allowed to vary freely to the extent permitted by the action laws.

Another categorization method is to single out some state variables as *dependent* on others, and to stipulate that minimization of change does not apply to dependent state variables. The definitions of the dependencies are part of the domain constraints. This approach was first proposed by Lifschitz as a distinction between frame and no-frame fluents [10].

2.3 Causation oriented approaches

Plain minimization of change is symmetric: a constraint of the form $a \rightarrow b$, if imposed on a state where a is true and b is false, can be satisfied either by making a false, or by making b true. This does not always correspond to the intentions in the application at hand, since sometimes the influence applies only one way.

The distinction between dependent and independent state variables is often sufficient for dealing with this problem. However, Thielscher [23] has shown an example, involving relays, where the distinction between dependent and independent variables is not sufficient. In his paper, and in concurrent papers by other authors [12, 11], a "causal" approach is proposed.

The basic idea is to express *causal dependencies*, which say that change in one state variable necessitates or enables a change in another state variable. Similar methods have been proposed earlier by Cordier and Siegel [2, 3] and by Brewka and Hertzberg [1]. Recently, Gustafsson and Doherty [7] have showed how a causation oriented approach can be reduced to a categorization approach without change minimization.

The fundamental limitation of current causal approaches is that they assume the explicitly stated effects of the action to be the first ones in a causal chain, and that indirect effects are downstream in that chain. For a counterexample, consider the case of a lamp la which is turned on independently by two different switches: $la \leftrightarrow sw_1 \vee sw_2$. la, sw_1, and sw_2 are the only state variables, and there is no other domain constraint. The action of "turning on the lamp" is defined simply by stating that la becomes true. Presumably, ramification shall allow us to conclude that either sw_1 or sw_2 has become true as well, since those are the only ways of turning on the lamp. However, a causal approach will observe that the causal link is from the switches to the lamp, and not vice versa, so it will not deal correctly with this example.

I submit that this is not an odd example: is is often desirable to specify actions in terms of their ultimate and main effects, rather than in terms of the specifics of how to perform them. Reasoning from main effect to instrumental operations, and from there to other effects of the same operations is therefore an important part of the ramification problem.

2.4 Approaches addressed here

The present paper is restricted to an analysis of approaches involving minimization, either plain or under categorization of state variables. Our first assessment results for a causation oriented approach have been reported in [21], but they do not fit into the present conference paper. Anyway, the semantics that is proposed and used here is distinctly causation-oriented.

3 Underlying semantics

The purpose of the underlying semantics is to define the set of intended models, and thereby the set of intended conclusions, in a precise, simple, and intuitively convincing fashion. We shall begin with a naive variant of underlying semantics, and then refine it twice.

3.1 State-transition semantics

The following basic concepts are used. Let $\mathbb{R}$ be the set of possible states of the world, formed as the Cartesian product of the finite range sets of a finite number of *state variables*. Also, let $\mathbb{E}$ be the set of possible actions, and let the *main next-state function* $\mathbf{N}(E, r)$ be a function from $\mathbb{E} \times \mathbb{R}$ to non-empty subsets of $\mathbb{R}$. The function $\mathbf{N}$ is intended to indicate the set of possible result states if the action E is performed when the world is in state r. The assumption $\mathbf{N}(E, r) \neq \emptyset$ expresses that every action E can always be executed in every starting state r. (The symbol E is derived from "event"; we consider actions and events to be closely related).

The assumption of a finite number of state variables is motivated as long as one can assume a fixed and finite number of objects *within each scenario*, although different scenarios may involve different object sets. This means that each scenario is essentially propositional in character, even though a first-order syntax is used in the logical formulae.

Now, consider a chronicle (= scenario description) consisting of a specification of the object set, action laws, schedule, and observations, where the action laws specify a next-state function $\mathbf{N}$ using some appropriate notation, and where time is assumed to be taken as the non-negative integers. Like in previous publications, we understand the "schedule" to specify the actions and their (absolute or relative) time of occurrence, and the "observations" to specify the values of certain state variables (= fluents) at some points in time, before, during, or after the actions. An interpretation for such a chronicle shall be chosen as a history over time, that is, a mapping from timepoints to states. Assignments of values to constant symbols etc. are also needed, but are of marginal importance for the present purpose.

The set of intended models for a chronicle is defined using an *ego-world game*, as follows. First, the observations are removed from the chronicle, so that only the action laws and the schedule remain. The following process is performed using the observation-free chronicle: Start with time as zero and the world in an arbitrary state. At each point in time, perform the move of the "ego" or agent, and the move of the world. The move of the ego is dictated by the schedule, and consists of selecting the action that is indicated by the schedule for the current point in time, or no action if nothing is specified. The move of the world is as follows: if the state of the world at the current time t is r, and the ego has just chosen the action E, then select some r' which is a member of $\mathbf{N}(E, r)$, make it the state of the world at time $t + 1$, and increase the current time by 1. If no move is required, assign the state r to the next timepoint $t + 1$ as well.

The described procedure will non-deterministically generate a set of possible histories of the world. The given definition applies trivially if all actions are specified explicitly by the schedule. Less trivial cases arise if actions are specified in other ways, for example as a disjunction between actions ("either action A or B is performed"), or with an imprecisely specified timing. However, those are technical details, and the general idea remains the same.

Among the world histories that are generated by the ego-world game, one then selects a subset consisting of those histories where all the observations are satisfied. That subset is the set of intended models for the given chronicle. In this case, which is the simplest one of all, the set of intended models will of course be exactly the set of models for the chronicle, using a conventional concept of models. We proceed to less trivial cases, where the characteristic aspects of the frame problem enter the picture.

3.2 Partial-state transition semantics

The purpose of the partial-state transition semantics is to capture the basic assumption of inertia ("no change except as the explicit result of an action") while retaining the general flavor of the state-transition semantics. We use $\mathbb{R}$, $\mathbb{E}$, and $\mathbf{N}$, like before, but we also introduce an *occlusion function* $\mathbf{X}(E, r)$ which is a function from $\mathbb{E} \times \mathbb{R}$ to subsets of the set of state variables. The occlusion function is intended to specify, for a given action E and starting state r, the set of those state variables whose new value after the action is going to be explicitly stated in the action law.

In the particular case of strict inertia, then, it must be the case that any state variable which is not a member of $\mathbf{X}(E, r)$ has the same value in r and in all the members of $\mathbf{N}(E, r)$. (Thus, it would be correct to select $\mathbf{X}(E, r)$ as a the set of all state variables, but it is intended that smaller sets shall be chosen whenever possible).

The functions $\mathbf{N}$ and $\mathbf{X}$ together determine the action laws down to the level of semantic equivalence: if these two functions are given, then the action law for an action E shall be constructed so that for each $r \in \mathbb{R}$, it specifies what $\mathbf{N}(E, r)$ says about the possible new values of the state variables in $\mathbf{X}(E, r)$, and it shall not say anything about the new values of other state variables. (Notice, however, that the new values of state variables in $\mathbf{X}(E, r)$ may equally well *depend on* the values of state variables outside $\mathbf{X}(E, r)$ in the starting state r of the action).

Since we are only interested in the set of new states that are permitted by the action law, and not in its actual syntactic form, it is not necessary to introduce any assumptions about the choice of logical language or about the exact structure of action laws. Using a much simpler approach, we shall introduce a derived function $\mathbf{N}'$, the *stated next-state function* which is obtained from the combination of $\mathbf{N}$ and $\mathbf{X}$, and which is used for defining and analyzing the nonmonotonic entailment methods. This function $\mathbf{N}'$ is obtained in the following manner. For every result state $r' \in \mathbf{N}(E, r)$, identify the set of all states which agree with r' for all state variables in $\mathbf{X}(E, r)$, but which have arbitrary values elsewhere. Construct $\mathbf{N}'(E, r)$ as the union of all those sets. It follows immediately that $\mathbf{N}'(E, r) \supseteq \mathbf{N}(E, r)$.

One might represent $\mathbf{N}'(E, r)$ equivalently as the set of partial states which are restrictions of members of $\mathbf{N}(E, r)$ to the state variables in $\mathbf{X}(E, r)$. This is the reason for the name "partial-state transition semantics".

The function $\mathbf{N}'$ that is obtained in this way is a direct counterpart of the usual, logical formulation of an action law, although we have turned the situation around in one important respect when it comes to ramification later on. Usually, one formulates the action law as a partial description of the effects of the action, and then proceeds to asking "what other effects should one also be able to infer?" Here, instead, we start with the function $\mathbf{N}$ that defines *all* the effects, and then proceed to the function $\mathbf{N}'$ that corresponds to the description of the action.

Notice that $\mathbf{N}'$ is uniquely determined by $\mathbf{N}$ and $\mathbf{X}$. We do not write $\mathbf{X}$ as an index on $\mathbf{N}$ because we are going to need another index below, and double indices are cumbersome. Notice, also, that $\mathbf{N}(E, r) = \emptyset \leftrightarrow \mathbf{N}'(E, r) = \emptyset$. In other words, our assumption of the universal applicability of actions is preserved in the transformations from $\mathbf{N}$ to $\mathbf{N}'$.

4 Semantics-level vs. logic-level analysis

Generally speaking, our task is to analyze logics of actions and change, either by relating them to an underlying semantics, or by relating two such logics to each other. In principle, this should require one to go through the usual routine of defining the syntax of logic formulae, interpretations, etc. In the present case, however, it is possible to do the essential parts of the analysis on the semantic level, and to rely on earlier results for the relation to the conventional parts of the logic.

4.1 Ramification on the semantic level

On the semantic level, we view ramification or update as follows. Let $\mathbb{R}_c$ be a subset of $\mathbb{R}$ consisting of those states [called *admitted states*] that satisfy the static domain constraints, and assume from $\mathbf{N}$ that

$$r \in \mathbb{R}_c \rightarrow \mathbf{N}(E, r) \subseteq \mathbb{R}_c. \tag{1}$$

Given $\mathbf{N}$, $\mathbf{X}$, and $\mathbb{R}_c$, and having derived $\mathbf{N}'$, our task is to obtain from $\mathbf{N}'$ the best possible approximation to $\mathbf{N}$, and ideally to obtain back $\mathbf{N}$ itself.

The first step must be to make full use of $\mathbb{R}_c$, the set of legal world states. Clearly, we only require $\mathbf{N}$ and its approximation to agree when the second argument is a member of $\mathbb{R}_c$, and since we have assumed that $\mathbf{N}(E,r) \subseteq \mathbb{R}_c$ in such cases, we obtain at once

$$r \in \mathbb{R}_c \rightarrow \mathbf{N}(E,r) \subseteq \mathbf{N}'(E,r) \cap \mathbb{R}_c. \tag{2}$$

If it happens to be the case that equality is obtained here, then no minimization (or other additional measures) are required at all. The use of $\mathbf{N}'(E,r) \cap \mathbb{R}_c$ as an approximation of $\mathbf{N}(E,r)$ will be called *straight ramification*. Then, in logical terms, one can take the action law (characterizing the new values of the state variables occuring in $\mathbb{R}_c$, after the action E has been performed from the starting state r), add the domain constraints, and draw the conclusions using the well-understood methods for strict inertia.

One may think that straight ramification is too simple, and that it will not work in practice. However, it does have the advantage of soundness: it may not draw all intended conclusions, but at least it does not draw any unintended conclusions. In terms of the two rationale for ramification that were mentioned initially, it would seem intuitively that straight ramification should be appropriate for redundancy-motivated ramification, but probably not for propagation-motivated ramification.

4.2 Ramification on the logic level

When the ramification problem is posed on the level of the concrete logic, one encounters a host of difficulties: the choice of logical language, dealing with the various parts of the scenario description, the relationship between the next-state functions and the corresponding action laws, the problem of non-initial observations which lead to postdiction problems, etc. Fortunately, however, all of these problems have been adequately handled already in the context of strict inertia, and can be considered as routine matters. We refer to [18] for all the necessary details. Therefore, we can proceed entirely on the semantic level here.

5 Measure systems

We proceed now to the analysis of cases where straight ramification is not sufficient. For our analysis of the commonly used, minimization-based approaches, we introduce the following uniform framework. A *measure system* is a threetuple

$$\langle \mathbb{Q}, d, \prec \rangle, \tag{3}$$

where $\mathbb{Q}$ is a domain of *quantities*, d is a *difference function* $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{Q}$, $\prec$ is a strict partial order on $\mathbb{Q}$, and where $q \not\prec d(r,r)$ for all r and q. Notice that even $d(r,r') = d(r',r)$ is not required, which is why we call this a "difference" and not a "distance". We shall say that r' is *less remote* than r'' from r if $d(r,r') \prec d(r,r'')$.

If m is a measure system $\langle \mathbb{Q}, d, \prec \rangle$, then $\mathbf{N}'_m$ shall be the function such that

$$\mathbf{N}'_m(E,r) \subseteq \mathbf{N}'(E,r) \cap \mathbb{R}_c \tag{4}$$

and where $\mathbf{N}'_m(E,r)$ consists exactly of those members of $\mathbf{N}'(E,r) \cap \mathbb{R}_c$ which are *minimally remote from r according to the measure system m.*

Formally, for any member r' of $\mathbf{N}'_m(E,r)$, there must not be any other member r'' of $\mathbf{N}'(E,r) \cap \mathbb{R}_c$ which is less remote than r' from r, and any $r' \in \mathbf{N}'(E,r) \cap \mathbb{R}_c$ having that property must be a member of $\mathbf{N}'_m(E,r)$.

This concept of measure system makes it straightforward to define a number of current minimization methods. For example, the simple case where one partitions state variables into "occluded" and "unoccluded" ones, and minimizes the set of changes in the unoccluded state variables, is obtained by defining $\mathbb{Q}$ as the power set of the set of state variables, $d(r,r')$ as the set of unoccluded state variables where r and r' assign different values, and $\prec$ as the strict subset relation.

It is easily verified that the approximation $\mathbf{N}'_m$ for arbitrary m satisfies all the Katsuno-Mendelzon postulates [9] **U1** through **U8**, except **U2**. The appropriateness of **U2** has already been cast in doubt by del Val and Shoham in [4], where they observed that **U2** is not well chosen in the case of nondeterministic actions.

6 Comparative assessments

With the given definitions, a measure system m is *correct* for a main next-state function $\mathbf{N}$ and occlusion function $\mathbf{X}$ iff $\mathbf{N}(E,r) = \mathbf{N}'_m(E,r)$ for every action E and every $r \in \mathbb{R}_c$. The *range of applicability* of a measure system m is the set of main next-state functions $\mathbf{N}$ where m is correct.

If one can not find a correct measure system for a given application, then presumably one prefers a measure system that is *sound* in the sense that it does not allow any unwarranted conclusions. In our framework, the measure system m is sound for a main next-state function $\mathbf{N}$ and occlusion function $\mathbf{X}$ iff $\mathbf{N}(E,r) \subseteq \mathbf{N}'_m(E,r)$ for every action E and every $r \in \mathbb{R}_c$.

It is evident that no measure system is correct for all $\mathbf{N}$, and that straight ramification is the only measure system that is sound for all $\mathbf{N}$. (The latter follows by choosing $\mathbf{N}(E,r)$ so that it equals $\mathbf{N}'(E,r) \cap \mathbb{R}_c$, so that no minimization is desired).

We shall compare different minimization methods with respect to how restrictive they are: a minimization method based on measure system m is *more restrictive* than one based on measure system m' iff $\mathbf{N}'_m(E,r) \subseteq \mathbf{N}'_{m'}(E,r)$ for all E and for all $r \in \mathbb{R}_c$. Thus, a more restrictive method is one which obtains a smaller set of preferred models, and thereby a larger set of conclusions, and a smaller set of pairs $\langle \mathbf{N}, \mathbf{X} \rangle$ for which it is sound. Clearly, straight ramification is the unique least restrictive method.

The following methods are of particular interest.

MSC: Minimal secondary change. State variables are divided into two categories: *occluded* ones specified by $\mathbf{X}$, and *secondary* ones which are the rest. Action laws specify the possible new values for occluded state variables. For given E and r, one selects those members of $\mathbf{N}'(E,r) \cap \mathbb{R}_c$ which minimize the set of secondary state variables that change compared to r. This is the method proposed by del Val and Shoham in [4].

MSCC: Minimal secondary change with changeset-partitioning. Similar to MSC for deterministic actions, but differs for nondeterministic actions. In MSCC, one divides $\mathbf{N}'(E,r) \cap \mathbb{R}_c$ into partitions each consisting of those proposed new states which are equal with respect to occluded state variables. In each of the partitions, one selects the subset of those members which are minimally remote from r, using the same criterion as for MSC. Then one forms the union of those closest subsets. It follows at once that MSCC is less restrictive than MSC, although more restrictive than straight ramification.

The intuition behind this modification is as follows. Consider an action which is nondeterministic already from the point of view of the action law, that is, the action law specifies several alternative outcomes of the action. Then one would like to minimize changes for each of the outcomes separately. For example, the action of playing a game on a gambling machine may result in different states in the machine; let us say just *win* and *lose* for simplicity. Suppose that the action law only specifies this outcome, and there is a static domain constraint specifying that *win* is accompanied by *money-falling-out*. Then MSC will always prefer *lose* over *win* because it minimizes the set of side-effects. MSCC, by comparison, considers *win* and *lose* to be equally preferred. This method was first proposed by Sandewall in [16], page 77.

MRC: Minimal remanent change. State variables are divided into three disjoint categories: occluded, remanent, and dependent ones. The occluded state variables are those whose new values are specified in the action laws, like before. Changes in remanent state variables are minimized, whereas changes in dependent state variables are disregarded for the purpose of minimization. This is essentially the method $\mathcal{AR}_0$ proposed by Kartha and Lifschitz in [8], if one identifies "occluded" with "frame released", "remanent" with "frame unreleased", and "dependent" with "non-frame" state variables. It is also used in PMON(R) proposed by Doherty and Peppas in [5].

MRCC: Minimal remanent change with changeset-partitioning. MRCC uses the same threeway partitioning as in MRC, but minimizes changes separately for each outcome in occluded state variables, like in MSCC.

The intuition behind the three-way partitioning is as follows: even if one chooses to divide the state variables into two categories, one where change is minimized and one where it is not, there may still be two distinct reasons why a state variable is in the latter category. On one hand, there are those state variables where the outcome or set of possible outcomes is already specified in the action law, so that further minimization is not needed and often not desired (because it may remove some of the intended outcomes of nondeterministic actions). On the other hand, there are those state variables which are functionally dependent on some of the others, and where the dependency is specified by the static domain constraints. In our formulation of the problem, the first subgroup is indicated by $\mathbf{X}$, since $\mathbf{X}$ is used in two ways: it defines $\mathbf{N}'$ from $\mathbf{N}$, and it is part of the definition of minimization. The other subgroup can not be included in $\mathbf{X}$ – the formal account in [20] introduces the required notation.

One might think that since the dependent state variables are functionally determined by the others, it would not hurt to bring them along in the minimization. The problem is, as was shown already by Winslett in the 'Aunt Agatha' scenario, that there may be undesired tradeoffs so that an unintended model containing a change in a remanent state variable may be be accepted because it contains a smaller set of changes in dependent variables. In such scenarios it is necessary to exclude dependent state variables from minimization.

Finally, the list should also include the classical method:

MC: Minimization of change. No partitioning of the state variables. One selects the subset of $\mathbf{N}'(E, r) \cap \mathbb{R}_c$ consisting of those members minimizing the set of state variables where the value differs from r. This is the method originally proposed by Winslett [24].

It is straightforward to define these methods in terms of measure systems. The definitions can be obtained from [20].

Writing $m \leq m'$ when m is more restrictive than m', and **SR** for straight ramification, we have proved the following inclusion results:

Theorem 1

$$\mathbf{MC} \leq \mathbf{MSCC} \leq \mathbf{SR} \tag{5}$$

$$\mathbf{MSC} \leq \mathbf{MSCC} \leq \mathbf{SR} \tag{6}$$

$$\mathbf{MRC} \leq \mathbf{MRCC} \leq \mathbf{SR} \tag{7}$$

The proofs are found in our technical report [20]. There does not seem to be any similar result relating e.g. MSCC and MRCC.

These results suggest at once that unless there are specific reasons to the contrary, one ought to prefer methods with changeset-partitioning over methods without, that is, MSCC over MSC and MRCC over MRC. This is because one will then obtain a method with a larger range of sound applicability. In other words, although possibly one will fail to obtain some intended conclusions, at least there is a smaller possibility that the system actually infers non-intended conclusions.

The methods with changeset-partitioning also have another advantage, namely that they are monotone with respect to the set of occluded features, $\mathbf{X}$. The precise result is as follows.

Theorem 2 If MSCC is incorrect for a certain choice of $\mathbf{N}$ and $\mathbf{X}$, and $\mathbf{X}'(E,r) \subseteq \mathbf{X}(E,r)$ for every choice of E and r, then MSCC is incorrect for $\mathbf{N}$ and $\mathbf{X}'$ as well.

The proof is found in [20]. The same theorem holds for MRCC as long as the set of dependent state variables is held fixed.

Therefore, it is meaningful to ask for a *smallest* choice of action laws, in the spirit of the initial quotation from Ginsberg and Smith: a choice of action laws that minimizes the set of state variables where the laws specify the new values after the action has been performed, and still the inference mechanism obtains the intended conclusions. The corresponding methods without changeset-partitioning do not have that monotonicity property.

7 Causal propagation semantics

We now proceed to the underlying semantics which is the necessary prerequisite for the analysis of range of (sound) applicability. The purpose of the underlying semantics is to define the set of intended models, and thereby the set of intended conclusions, in a precise, simple, and intuitively convincing fashion. In the case of underlying semantics for ramification, we consider it particularly important to capture the propagation-oriented reason for ramification. We propose to use the following *causal propagation semantics*, which is obtained through a second modification of the basic state transition semantics described above.

The domain $\mathbb{R}$ of world states and the domain $\mathbb{E}$ of actions are defined like before. A binary non-reflexive *causal transition relation* C between states is introduced; if $C(r,r')$ then r' is said to be the *successor* of r. A state r is called *stable* iff it does not have any successor. The set $\mathbb{R}_c$ of *admitted states* is chosen as a subset of $\mathbb{R}$ all of whose members are stable with respect to C. (It may or may not be chosen to contain all the stable states).

Furthermore, we assume an *action invocation relation* $G(E, r, r')$ where $E \in \mathbb{E}$ is an action, r is the state where the action E is invoked, and r' is the new state where the instrumental part of the action has been executed. For every E and r there must be at least one r' such that $G(E, r, r')$, that is, every action is always invokable.

An *action system* is a tuple $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G\rangle$ whose five elements have been selected as just described. It is intended to be used as follows.

A *transition chain* for a state $r \in \mathbb{R}_c$ and an action E is a finite or infinite sequence of states

$$r_1, r_2, ...(r_k)$$

where $G(E, r, r_1)$, $C(r_i, r_{i+1})$ for every $i \geq 1$ for an infinite sequence and for every i where $1 \leq i < k$ for a finite sequence, and where r_k is a stable state again in the case of a finite sequence.

The last element r_k of a finite transition chain is called a *result state* of the action E in the sitution r.

We intend, of course, that it shall be possible to characterize the result states r_k in terms of E and r, but without referring explicitly to the details of the intermediate states. The following assumptions (or something similar) are needed in order to make this work as intended.

Definition 3 If three states r, r_i, and r_{i+1} are given, we say that the pair r_i, r_{i+1} **respects** r iff

$$r_i(f) \neq r_{i+1}(f) \rightarrow r_i(f) = r(f),$$

for any state variable f that is defined in $\mathbb{R}$. Then, an action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G\rangle$ is said to be **respectful** iff, for every $r \in \mathbb{R}_c$, every $E \in \mathbb{E}$, r is respected by every pair r_i, r_{i+1} in every transition chain, and the last element of any finite chain is a member of $\mathbb{R}_c$. $\Box$

This condition amounts to a "write-once" or "single-assignment" property: if the action E is performed from state r, the world may go through a sequence of states, but in each step from one state to the next, there cannot be changes in state variables which have already changed previously in the sequence, nor can there be any additional change in a state variable that changed in the invocation transition from r to r_1. This condition is also analogous to Lin's definition of a stratified system in [11].

As a consequence of these definitions, we obtain that if $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G\rangle$ is a respectful action system, if $r \in \mathbb{R}_c$, $E \in \mathbb{E}$, and $G(E, r, r')$ holds for some r', then all the transition chains that emerge from (E, r) are finite and cycle-free. The set of result states will be denoted $\mathbf{N}(E, r)$, using the same notation as in the state-transition semantics. Thus, the main next-state function $\mathbf{N}(E, r)$ which was previously taken as given, is here taken as derived from the elementary relations G and C in the action system. Also, $\mathbf{N}(E, r) \neq \emptyset$.

Respectful action systems are intended to capture the basic intuitions of actions with indirect effects which are due to causation, as follows. Suppose the world is in a stable state r, and an action E is invoked. The immediate effect of this is to set the world in a new state, which is not necessarily stable. If it is not, then one allows the world to go through the necessary sequence of state transitions, until it reaches a stable state. *That whole sequence* of state transitions is together viewed as the action, and the resulting admitted state is viewed as the result state of the action.

The assumption of sequentiality of actions within a chronicle is then taken to mean that the agent or "ego" abstains from invoking any new action before the previous action has reached a stable state. The exact interpretation of the state transitions may vary: as spontaneous transitions in the physical world, or as more or less logical conclusions about a world where all changes happen (or are modelled as happening) concurrently. One may even include the case where additional interventions by the agent (viewed as a controller) are necessary along the path from the starting state to the result state of the action.

For example, in the lamp scenario that was mentioned in section 2.3, if the state variables represent the first switch, the second switch, and the lamp, respectively, the invocation relation G will contain argument sequences such as $(E_1, \mathtt{FFF}, \mathtt{TFF})$ for turning on the first switch; the causal transition relation will contain the argument sequence $(\mathtt{TFF}, \mathtt{TFT})$ expressing that if the first switch is on, then the lamp goes on as well. The action E_2 which is instrumentally defined as turning on either of the two switches is true for both $G(E_2, \mathtt{FFF}, \mathtt{TFF})$ and $G(E_2, \mathtt{FFF}, \mathtt{FTF})$.

The idea with ramification, now, is that one does not wish to specify *all* the effects of an action in the action law; one wishes to restrict the action law to the most important effects, and other effects should be deducible if needed. In other words, the action law shall not specify all components of the result state, and not even all the components that change, but only some of them. Other changes are deduced using information about the relations C and G.

Causal propagation semantics is not applicable to all kinds of ramification. For example, it would be completely unnatural in the context of physical movement, where the displacement of one object may cause another object to move as well. We are currently working on other semantics for systems involving coupled continuous change.

Notice that it is not necessary to let the action law characterize the invocation relation G. Alternatively, the action law may characterize some of the later changes along the chain, namely, if the indirect changes are considered to be the most important ones in terms of describing the action. For example, the action law for E_2 in the second-previous paragraph may specify the change in the state variable indicating that "the lamp is lit", leaving the "upstream" change of one or the other of the switches to be inferred. In general, the non-explicit effects of the action may be a combination of changes that are causationally "downstream" as well as "upstream" of the explicit ones.

8 Soundness assessments for MSCC

After having related the various minimization methods to each other, we proceed now to a first analysis of how they relate to the underlying semantics.

The general problem is viewed as follows. Let the following be given:

- A respectful action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G \rangle$. Its corresponding main next-state function will be written $\mathbf{N}$, as before.
- An occlusion function, $\mathbf{X}$.
- A measure system $m = \langle \mathbb{Q}, d, \prec \rangle$.

Informally, it is still the case that $\mathbf{X}$ indicates what are the state variables whose new values are indicated in the action law. The functions $\mathbf{N}$ and $\mathbf{X}$ together define the stated next-state function $\mathbf{N}'$, as described above. The measure system is said to be *correct* for the action system and the occlusion function iff

$$\mathbf{N}'_m(E, r) = \mathbf{N}(E, r)$$

for every E and r. Likewise, the measure system is said to be *sound* for the action system and the occlusion function iff

$$\mathbf{N}'_m(E, r) \supseteq \mathbf{N}(E, r)$$

for every E and r. Notice how the entities in this equation have been derived: $\mathbf{N}$ from the action system alone, $\mathbf{N}'$ from $\mathbf{N}$ and $\mathbf{X}$, and $\mathbf{N}'_m$ from $\mathbf{N}'$ and m.

In particular, given one of the specific measure systems and entailment methods that were described above, for example MRCC, a particular partitioning of the state variables is said to be *correct* for the action system, the occlusion function, and the method at hand iff the corresponding measure system is correct for them. Soundness is defined similarly.

The corresponding *assessment problem* for an entailment method is defined as the following question: what are the requirements on the combination of a respectful action system, an occlusion function, and a partitioning whereby the given entailment method is correct (or sound)?

8.1 Occlusion of invocation- changed state variables

We consider separately the case where action laws describe invocations, and the case where they describe downstream effects.

Definition 4 A respectful action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G \rangle$ is said to have **fixed control** iff there is some subset $\mathbb{F}_c$ of the set $\mathbb{F}$ of state variables, such that

$$G(E, r, r') \rightarrow dom(r - r') \subseteq \mathbb{F}_c$$

and

$$C(r, r') \rightarrow dom(r - r') \cap \mathbb{F}_c = \emptyset$$

The set $\mathbb{F}_c$ is called the *control set.* $\Box$

Here and henceforth the following conventions are used: $-$ denotes set subtraction, mappings are viewed as sets of argument/value pairs, and $dom(r)$ maps a mapping (in this case, a partial state) to its domain.

Members of the control set will be called *controlled* state variables; the others will be called *obtained* state variables.

This definition means that invocation transitions can only influence state variables in the control set, and causal transitions can never influence state variables in the control set. Thielscher's relay example in [22] is an example of a system which does not have fixed control.

Definition 5 Given a state domain $\mathbb{R}$, a state r'' is said to be a **hybrid** between two states r and r' *with respect to* a set $F \subseteq \mathbb{F}$ of state variables iff

(1) for every state variable f,

$$r''(f) = r(f) \vee r''(f) = r'(f),$$

(2) if $f \in F$ then $r'(f) = r''(f)$, and

(3) r'' is different from both r and r'. $\Box$

Definition 6 A respectful action system is said to be **hybrid free** with respect to a set F of state variables iff no two members r and r' of $\mathbb{R}_c$ have a hybrid with respect to F which is also a member of $\mathbb{R}_c$. $\Box$

It follows that if an action system is hybrid free with respect to a set F, then it is hybrid free with respect to any superset of F. (Any hybrid with respect to the superset would be a hybrid with respect to F). We immediately obtain:

Theorem 7 MSCC is sound for the combination of a respectful action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G \rangle$ and an occlusion function $\mathbf{X}$ if the following conditions are satisfied:

1. The action system has fixed control with a control set $\mathbb{F}_c$
2. The action system is hybrid free with respect to $\mathbb{F}_c$
3. If $G(E, r, r')$ then $dom(r - r') \subseteq \mathbf{X}(E, r)$

Proof We are to prove that $\mathbf{N}'_{\mathrm{MSCC}}(E, r) \supseteq \mathbf{N}(E, r)$. Suppose the opposite is the case, that is, for some E and r there is some $r^* \in \mathbf{N}(E, r)$ which is not a member of $\mathbf{N}'_{\mathrm{MSCC}}(E, r)$. Since it is a member of $\mathbf{N}(E, r)$, r^* must be a member of $\mathbf{N}'(E, r)$ and of $\mathbb{R}_c$, so its non-membership of $\mathbf{N}'_{\mathrm{MSCC}}(E, r)$ must be due to the existence of some $r' \in \mathbf{N}'(E, r) \cap \mathbb{R}_c$ which is less remote than r^* from r.

By the definition of MSCC, r' must agree with r^* in all occluded state variables. By assumptions 1 and 3, r^* agrees with r in all non-occluded controlled state variables, so r' must do so as well. Therefore, r^* and r' agree in all state variables in $\mathbb{F}_c$. But r' differs from r in a subset of the state variables where r^* differs from r. Therefore, either r' is a hybrid between r and r^* with respect to $\mathbb{F}_c$, or $r' = r$. The former case is inconsistent with assumption 2 given that $r' \in \mathbb{R}_c$. The latter, bizarre case requires that r and r^* are equal in all controlled state variables, which means that $\mathbf{N}(E, r) = \{r\}$ and $r^* = r$, making it impossible for any r' to be less remote from r than r^* is. □

8.2 Occlusion of non-invoked state variables

We proceed now to the case where action laws characterize other changes of state variables than those involved in the invocation.

Definition 8 A respectful action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G \rangle$ is **unary controlled** iff $G(E, r, r')$ implies that $dom(r - r')$ is a singleton, for any E, r, and r'. □

Note that this property is realistic for low-level actions in engineered systems, where controlled state variables represent the actuators.

Definition 9 Consider the combination of a respectful action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G \rangle$ having fixed control, and an occlusion function $\mathbf{X}$. A state $r' \in \mathbb{R}_c$ is called a **shortcut** for (E, r) in this combination, where $r \in \mathbb{R}_c$, if and only if there is some $r'' \in \mathbf{N}(E, r)$ such that

(1) for every state variable f,

$$r'(f) = r(f) \vee r'(f) = r''(f),$$

(2) r' equals r in all controlled state variables, and
(3) r' equals r'' for all state variables in $\mathbf{X}(E, r)$. □

Example 10 The following is an example of a shortcut. Consider an electric circuit consisting of a number of switches, lamps, and relays, where the switches and only them are controlled. In the current state r of the circuit, changing a particular switch s will turn on a lamp l which is presently off;

there may also be other effects. Let r'' be the resulting state. An action E is defined so that $\mathbf{X}(E,r) = \{l\}$ and r'' is a member of $\mathbf{N}(E,r)$, that is, the action is defined in terms of the effect of turning on the lamp, and not in terms of changing the switch. However, there is also a sequence of other actions where different switches are changed in succession, and which ends up in a state r' where all switches are in the same position as in the present state r, the lamp l is on anyway, and only a subset of the changes from r to r'' occur in r'. □

It would seem that the property of having a shortcut is not a particularly natural one for an action system to have. However, there is an important special case, involving the combination of an action system and an occlusion function where an action possibly has a null stated effect, that is, there is some $r'' \in N(E,r)$ for which r and r'' are equal for all occluded state variables. In this case, r itself is a shortcut for (E,r).

We immediately obtain:

Theorem 11 **MSCC** is sound for the combination of a respectful action system $\langle \mathbb{R}, \mathbb{E}, C, \mathbb{R}_c, G \rangle$ and an occlusion function $\mathbf{X}$ if the following conditions are satisfied:

1. The action system has fixed control with a control set $\mathbb{F}_c$.
2. The action system is hybrid free with respect to $\mathbb{F}_c$.
3. $\mathbf{X}(E,r)$ is disjoint from $\mathbb{F}_c$ for every E and $r \in \mathbb{R}_c$.
4. The action system is unary controlled.
5. The combination of the action system and the occlusion function does not have any shortcut.

Proof We are to prove that $\mathbf{N}'_{\mathrm{MSCC}}(E,r) \supseteq \mathbf{N}(E,r)$. Suppose the opposite is the case, that is, for some E and r there is some $r^* \in \mathbf{N}(E,r)$ which is not a member of $\mathbf{N}'_{\mathrm{MSCC}}(E,r)$. Since it is a member of $\mathbf{N}(E,r)$, r^* must be a member of $\mathbf{N}'(E,r)$ and of $\mathbb{R}_c$, so its non-membership of $\mathbf{N}'_{\mathrm{MSCC}}(E,r)$ must be due to the existence of some $r' \in \mathbf{N}'(E,r) \cap \mathbb{R}_c$ which is less remote than r^* from r.

Consider the causation chain leading up to r^*, and its first element r_1. Because of the unary control property, r and r_1 will differ in exactly one state variable; this state variable is a controlled one because of assumption 1, and it is not occluded because of assumption 3. There are two possibilities with respect to r':

(1) r' equals r_1 in all controlled state variables. Then r' is a hybrid between r and r^* with respect to $\mathbb{F}_c$, violating assumption 2.

(2) r' equals r in all controlled state variables. Since r' is a member of $\mathbf{N}'(E,r)$, it must agree with some member r'' of $\mathbf{N}(E,r)$ in all occluded state variables. Therefore, r' is a shortcut for (E,r), which violates assumption 5. □

The results of these theorems do not apply in general for MSC because of interference between primary outcomes of the action. Generalization of these results to MRCC encounters technical complications in the generalization of the "hybrid" concept.

These assessments may be compared with the previous results on assessment of entailment methods with respect to the assumption of strict inertia. Those assessments [18] made use of the concept of minimal-change-compatible, which in the present context may be equivalently defined as follows: An action system is **minimal-change-compatible** iff it is not the case, for any r, E, r', and r'', that $r' \in \mathbf{N}(E,r)$, $r'' \in \mathbf{N}(E,r)$, and r''

is a hybrid between r and r' with respect to $\emptyset$. We showed in [18] that PCMF (prototypical minimization of change with filtering) is correct for all minimal-change-compatible chronicles with single-timestep actions. We observe that if an action system is hybrid free with respect to $\emptyset$, then it is both minimal-change-compatible and hybrid free with respect to an arbitrary set F.

8.3 Examples

The soundness property means that the entailment method MSCC does not miss any intended models, but it still allows for the possibility of including some unintended models. The following two examples illustrate how unintended models may be included in $\mathbf{N}'_{\mathrm{MSCC}}(E, r)$.

Example 12 For the first example, suppose $\mathbb{R}$ is selected as $\{\mathtt{T}, \mathtt{F}\} \times \{\mathtt{1}, \mathtt{2}\} \times \{\mathtt{1}, \mathtt{2}\}$, where the first component is controlled and always occluded, and the other two are secondary. Suppose further that the causal transition function C is true for the following pairs:

$$C(\mathtt{T11}, \mathtt{T12}), \quad C(\mathtt{T22}, \mathtt{T21}),$$

$$C(\mathtt{F12}, \mathtt{F11}), \quad C(\mathtt{F21}, \mathtt{F22})$$

and that $\mathtt{T12}, \mathtt{T21}, \mathtt{F11}, \mathtt{F22}$ are in $\mathbb{R}_c$. The invocation relation G is supposed to contain

$$G(E_1, \mathtt{F11}, \mathtt{T11}), \quad G(E_1, \mathtt{F22}, \mathtt{T22})$$

as well as some suitable transitions when the controlled state variable is initially T. It is clear that $\mathbf{N}(E_1, \mathtt{F11}) = \{\mathtt{T12}\}$ and $\mathbf{N}'_{\mathrm{MSCC}}(E_1, \mathtt{F11}) = \{\mathtt{T12}, \mathtt{T21}\}$, so soundness but not correctness is obtained. □

The problem here, of course, is that the entailment method MRCC is based on a notion of minimal change which does not use any of the causal transition information from C.

Example 13 The following example illustrates another possibility for the introduction of unintended models. Suppose now that $\mathbb{R}$ is selected as $\{\mathtt{T}, \mathtt{F}\} \times \{\mathtt{T}, \mathtt{F}\} \times \{\mathtt{1}, \mathtt{2}\}$, where the first two components are controlled and the third one is obtained. Suppose further that the transition function C satisfies

$$C(\mathtt{TF1}, \mathtt{TF2}),$$

and that $\mathtt{FF1}, \mathtt{FT1}, \mathtt{TF2}, \mathtt{TT1}$ are in $\mathbb{R}_c$. The invocation relation G is supposed to contain

$$G(E_1, \mathtt{FF1}, \mathtt{TF1}), \quad G(E_1, \mathtt{FT1}, \mathtt{TT1})$$

as well as some suitable invocation transitions for the other members of $\mathbb{R}_c$. Then $\mathbf{N}(E_1, \mathtt{FF1}) = \{\mathtt{TF2}\}$. In this case $\mathbf{X}(E_1, \mathtt{FF1})$ needs only contain the first state variable, so both the second and the third state variable can be selected as secondary. Using MSCC with that assumption one obtains $\mathbf{N}'_m(E_1, \mathtt{FF1}) = \{\mathtt{TF2}, \mathtt{TT1}\}$, so again soundness but not correctness is obtained. □

In this case, the unintended model entered because $\mathbf{X}(E, r)$ is only required to be a subset of the controlled state variables. This opens the door for admitting some model which satisfies the action laws with respect to occluded variables, but which also includes some change in a controlled, non-occluded variable (in this case, the second state variable), provided that it thereby obtains a smaller set of changes in remanent state variables.

The source of incompleteness that was illustrated here can be removed by requiring that all controlled state variables must be occluded, that is, their new value must be stated in the action law. This means effectively that the action law must not merely specify the new values for some state variables that change their value, but also specify explicitly and un-overridably that some state variables do not change their value in the current action. An additional possibility is to introduce a fourth category for those state variables which *must not* change as the result of an action.

9 Summary and conclusions

We published our main results on underlying semantics and assessment of range of applicability for strict inertia in 1993-94. It has been surprisingly difficult to find a natural extension to the case of ramification, and the causal propagation semantics proposed here is the first real advancement. At the same time, we observe that only soundness assessments have been obtained so far, and we think it is likely that soundness assessments will continue to predominate. Therefore, there will be a continuing need for comparative assessments, providing information on the relative size of the selected model sets for different entailment methods.

The main concrete results in the present article are therefore:

- The causal propagation semantics
- The comparative assessments of a number of entailment methods
- The soundness assessment of one of the entailment methods, MSCC
- The observation of the usefulness of changeset-partitioning in minimization-based approaches.

Acknowledgements

The first part of this research was performed while the author was a visiting researcher at LAAS-CNRS, Toulouse, France. We acknowledge valuable comments on that part of the work by Marie-Odile Cordier and Pavlos Peppas.

References

[1] Gerhard Brewka and Joachim Hertzberg. How to do things with worlds: On formalizing actions and plans. *Journal of Logic and Computation*, 1993.

[2] Marie-Odile Cordier and Pierre Siégel. A temporal revision model for reasoning about world change. In *International Conference on Knowledge Representation and Reasoning*, pages 732–739, 1992.

[3] Marie-Odile Cordier and Pierre Siégel. Prioritized transitions for updates. In C. Froidevaux and J. Kohlas, editors, *Proc. European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty*, pages 142–151. Springer Verlag, 1995.

[4] Alvaro del Val and Yoav Shoham. Deriving properties of belief update from theories of actions (II). In *International Joint Conference on Artificial Intelligence*, pages 732–737, 1993.

[5] Patrick Doherty and Pavlos Peppas. A comparison between two approaches to ramification: PMON(R) and $\mathcal{AR}_0$. In *Proc. Australian AI Conference*, 1995.

[6] Matthew L. Ginsberg and David E. Smith. Reasoning about actions I: A possible worlds approach. In Frank Brown, editor, *Proceedings of the 1987 Workshop on the Frame Problem in Artificial Intelligence*, pages 233–258. Morgan Kaufmann, Lawrence, Kansas, 1987.

[7] Joakim Gustafsson and Patrick Doherty. Embracing occlusion in specifying the indirect effects of actions. In *International Conference on Knowledge Representation and Reasoning*, 1996.

[8] G. Neelakantan Kartha and Vladimir Lifschitz. Actions with indirect effects (preliminary report). In *International Conference on Knowledge Representation and Reasoning*, pages 341–350, 1994.

[9] Hirofumi Katsuno and Alberto O. Mendelzon. A unified view of propositional knowledge base update. In *International Joint Conference on Artificial Intelligence*, pages 1413–1419, 1989.

[10] Vladimir Lifschitz. Frames in the space of situations. *Artificial Intelligence*, 46:365–376, 1990.

[11] Fangzhen Lin. Embracing causality in specifying the indirect effects of actions. In *International Joint Conference on Artificial Intelligence*, 1995.

[12] Norman McCain and Hudson Turner. A causal theory of ramifications and qualifications. In *International Joint Conference on Artificial Intelligence*, 1995.

[13] Tommy Persson and Lennart Staflin. A causation theory for a logic of continuous change. In *European Conference on Artificial Intelligence*, pages 497–502, 1990.

[14] Erik Sandewall. Non-monotonic entailment for reasoning about time and action, part I: Sequential actions. Technical report, IDA, 1988.

[15] Erik Sandewall. Filter preferential entailment for the logic of action in almost continuous worlds. In *International Joint Conference on Artificial Intelligence*, pages 894–899, 1989.

[16] Erik Sandewall. Features and fluents. Review version of 1991. Technical Report LiTH-IDA-R-91-29, IDA, 1991.

[17] Erik Sandewall. The range of applicability of nonmonotonic logics for the inertia problem. In *International Joint Conference on Artificial Intelligence*, 1993.

[18] Erik Sandewall. *Features and Fluents. The Representation of Knowledge about Dynamical Systems. Volume I.* Oxford University Press, 1994.

[19] Erik Sandewall. Reasoning about actions and change with ramification. In J. van Leeuwen, editor, *Computer Science Today*. Springer Verlag, Berlin. Lecture Notes in Computer Science, Anniversary Volume 1000, 1995.

[20] Erik Sandewall. Systematic comparison of approaches to ramification using restricted minimization of change. Technical Report LiTH-IDA-R-95-15, IDA, 1995.

[21] Erik Sandewall. Transition cascade semantics and first assessments results for ramification. In Oliviero Stock, editor, *Spatial and Temporal Reasoning*. Kluwer Publishing Company, 1997.

[22] Michael Thielscher. An analysis of systematic approaches to reasoning about actions and change. In P. Jorrand, editor, *International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA)*, Sofia, Bulgaria, 1994. World Scientific Publishing Co.

[23] Michael Thielscher. Computing ramifications by postprocessing. In *International Joint Conference on Artificial Intelligence*, 1995.

[24] Marianne Winslett. Reasoning about actions using a possible models approach. In *National Conference on Artificial Intelligence*, pages 89–93, 1988.

[25] Choong-Ho Yi. Towards the assessment of logics for concurrent actions. In Dov M. Gabbay and Hans Jürgen Olbach, editors, *Conference on Formal and Applied Practical Reasoning, FAPR 96*, pages 679–690. Springer Verlag, 1996.