The PMA and Relativizing Change for Action Update

Patrick Doherty
Dept. of Computer Science
Linköping University
S-58183 Linköping, Sweden
patdo@ida.liu.se

Witold Łukaszewicz
Ewa Madalińska-Bugaj
Institute of Informatics
Warsaw University
Warsaw, Poland

Abstract

Using intuitions from the temporal reasoning community, we provide a generalization of the PMA, called the modified PMA (MPMA), which permits the representation of disjunctive updates and the use of integrity constraints interpreted as causal constraints. In addition, we provide a number of syntactic characterizations of the MPMA, one of which is constructed by mapping an MPMA update of a knowledge base into a temporal narrative in a simple temporal logic (STL). The resulting representation theorem provides a basis for computing entailments of the MPMA and could serve as a basis for further generalization of the belief update approach for reasoning about action and change.

1 Introduction

Recently, much effort has been invested in applying the belief revision/update (BR/U) paradigm to the domain of reasoning about action and change. In a similar vein, much effort has been invested in applying temporal logics (TL) to reasoning about action and change. Unfortunately, there have been few attempts at trying to bridge the gap between the two paradigms by analyzing one approach in terms of the evaluation methodology of the other, or by applying solutions generated in one paradigm to similar open problems in the other paradigm (although, see [10, 20, 3, 24, 22]). We believe that there is much to be gained by a cross-fertilization of the two paradigms and in this paper, we will try to do just that.

We begin by noting that much of the recent research in the temporal reasoning paradigm has focused on two issues: 1) proper representation of non-deterministic actions [18, 21, 15], and 2) proper representation of the indirect effects of actions, where the use of causal or fluent dependency constraint approaches has received a great deal of attention [17, 19, 13, 23]. On the other hand, the BR/U approach applied to reasoning about action and change, does not appear to have reached any satisfactory consensus on the proper approach to disjunctive update and is continually plagued with an inability to represent casual constraints in terms of standard integrity constraints.

Beginning with the PMA, we will apply techniques from the TL paradigm resulting in a generalization of the PMA, called the modified PMA (MPMA), which permits the representation of disjunctive update together with integrity constraints interpreted as causal constraints. In addition, we provide a number of syntactic characterizations of MPMA with integrity constraints and show that a very simple temporal logic, STL, which is a fragment of a 1st-order temporal logic described in [4, 13] with origins from the Features and Fluents framework [21], can be used to compute entailments of the MPMA with integrity constraints. The representation theorem we provide opens up the possibility of additional generalizations of the BR/U paradigm applied to reasoning about action and change.

Another interesting result of this work is the discovery that several existing approaches to reasoning about action and change [15, 23, 13], suffer from a form of syntactic sensitivity which we call the *redundant atom anomaly*, where if one is not careful about the syntactic form of consequents to action or causal rules, certain arguably unintuitive direct and indirect effects of action invocation may result. In the paper, we point out the problem using MPMA, and provide a remedy which can also be applied to the temporal logics which suffer from similar problems.

The generalization to the PMA can succinctly be described as the *relativization* or weakening of the minimal change policy to a subset of atoms in the language rather than to all atoms. The technique is similar to the use of varied predicates in circumscription or the use of an *occlusion*, *release*, or *noninert* predicate in the temporal logic paradigm. The proposed solution to the use of integrity constraints is to view them as causal constraints representing fluent dependency information as described in Gustafsson & Doherty [13], where the technique was used to deal with the ramification problem.

The paper is structured as follows: In Section 2, we introduce the PMA and discuss some well known criticisms associated with its use as a formalism for reasoning about action and change. In Section 3, we modify the PMA, by relativizing or weakening the minimal change policy built into the definition of distance between interpretations. We discuss the redundant atom anomaly and provide a means of automatically preprocessing updates to remove redundancy and choose the atoms to be excluded from the minimal change policy. A syntactic characterization of MPMA in terms of *eliminants* is considered along with an analysis of the MPMA in terms of the Katsuno and Mendelzon [14] postulates. In Section 4, we introduce the simple temporal logic, STL, a fragment of a highly expressive logic used to reason about temporal narratives. In Section 5, we reformulate the MPMA in STL and provide a representation theorem. A decision procedure for MPMA in terms of STL is then provided. In Section 6, we extend the MPMA with integrity constraints represented as causal rules or dependency constraints.In Section 7, we provide a comparison with related work and in Section 8, conclude with a discussion.

Due to page limitations, proofs are excluded, but may be found in Doherty *et al.* [6], an extended version of this paper which also includes an additional syntactic characterization of the MPMA in terms of a Dijkstra semantics based approach together with a corresponding decision procedure, and additional comparisons with related work.

2 The Classical PMA

We present the classical PMA semantics, originally developed by Winslett [25, 26, 27].

2.1 The Language $\mathcal{L}_{pma}$

We start with a language $\mathcal{L}_{pma}$ of classical propositional logic based on a finite fixed set $ATM = \{p, q, r, \ldots\}$ of atoms. Formulas are built in the usual way using the connectives $\wedge, \vee, \neg, \rightarrow, \leftrightarrow$, $\top$ (truth) and $\bot$ (falsity). If α and β are formulas and p is an atom, then we write $\alpha[p \leftarrow \beta]$ to denote the formula which is obtained from α by replacing all occurrences of p by β. A *literal* is an atom or its negation.

2.2 Semantics

Interpretations are identified with maximal consistent sets of literals. For any formula α, we write $\mid \alpha \mid$ to denote the set of all *models* of α, *i.e.* interpretations satisfying α. A formula α is said to *correspond* to an interpretation u iff $\mid \alpha \mid = u$. To construct such a formula, it suffices to take the conjunction of all literals occurring in u. Similarly, a formula α is said to correspond to a finite set of interpretations $\{u_1, \ldots, u_n\}$ if $\mid \alpha \mid = \{u_1, \ldots, u_n\}$. To obtain such a formula, it suffices to take the disjuncion $\beta_1 \vee \cdots \vee \beta_n$, where, for each $1 \leq i \leq n$, β_i is the formula corresponding to u_i. For instance, the formula corresponding to the set $\{\{p, \neg q\}, \{p, q\}\}$ is $(p \wedge \neg q) \vee (p \wedge q)$ which is equivalent to p.

Definition 1 Let w, v be two interpretations. The *distance* between w and v, written $DIST(w, v)$, is a set of atoms that have different truth-values in w and u. ■

For instance, the distance between $w = \{p, q, r, \neg s\}$ and $v = \{q, s, \neg p, \neg r\}$ is $\{p, r, s\}$.

Definition 2 The *update of an interpretation w by a set of interpretations V*, written $w \star V$, is the set of those elements of V that are closest to w, *i.e.* whose distance to w is minimal. More formally:

$$w \star V = \{v \in V : \text{there is no } v' \in V \text{ such that } DIST(w, v') \subset DIST(w, v)\}. ■$$

For instance,
if $w = \{\neg p, \neg q\}$ and $V = \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\}$, then $w \star V = \{\{p, \neg q\}, \{q, \neg p\}\}$.

Definition 3 The *update of a set of interpretations U by a set of interpretations V*, written $U \star V$, is given by

$$U \star V = \bigcup_{w \in U} w \star V. ■$$

In the classical PMA, the update $KB \star \alpha$ of a knowledge base KB by a formula α is identified with the formula corresponding to the set of interpretations $\mid KB \mid \star \mid \alpha \mid$.

2.3 Examples

Example 1 Let $ATM = \{p, q\}$, $KB = p$ and $\alpha = p \vee q$. Then $|KB| \star |\alpha| =$

$$\{p, \neg q\} \star \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} \cup \{p, q\} \star \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} = \{p, \neg q\} \cup \{p, q\}.$$

Thus $KB \star \alpha$ is the formula $(p \wedge \neg q) \vee (p \wedge q)$ which is equivalent to p.

Example 2 Let $ATM = \{p, q\}$, $KB = \neg p \wedge \neg q$ and $\alpha = p \vee q$. Then $|KB| \star |\alpha| =$

$$\{\neg p, \neg q\} \star \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} = \{\{p, \neg q\}, \{q, \neg p\}\}.$$

Thus $KB \star \alpha$ is the formula $(p \wedge \neg q) \vee (\neg p \wedge q)$.

3 The Modified PMA

A number of researchers ([1], [11]) have observed that the classical PMA may lead to unintuitive results, particularly when disjunctive update is involved. To see this, reconsider Example 1. If all we know about the world is p and the effect of an action performed is $p \vee q$, then the description of the new state of the world should be $p \vee q$ rather than p. If p was accepted as the description of the new state, then this would imply that our knowledge about the current state makes an unpredictable action predictable. Example 2 also illustrates an undesirable phenomenon: updating a knowledge base by the inclusive *or* results in the exclusive *or*. What these two examples have in common is the fact that if an action effect follows from the database, the action will have no effect. The culprit in both cases is the fact that the PMA minimal change policy should be relaxed somewhat, but in a principled manner. This is the insight gained from the TL community and incorporated in current solutions to similar problems.

In order to incorporate a relaxed minimal change policy for the PMA, we will change it in two ways. Firstly, the distance between two interpretations will always be relativized to a chosen set of atoms. Intuitively, this set represents atoms that are allowed to vary their values during the action execution. Secondly, the update of an interpretation w by a set of interpretations U will be defined as those elements of U whose distance to w is the empty set. We call this new form of the PMA the *modified* PMA and we denote it by MPMA. The details follow.[1]

[1] The motivation behind MPMA will become clear when we provide its syntactic characterization in Section 3.2. A policy similar to MPMA is mentioned by Winslett, but rejected as not representing minimal change. MPMA also reflects similar minimization policies used in the TL paradigm as we will show.

Definition 4 Let w and v be two interpretations and suppose that P is a set of atoms. The *distance* between w and v wrt P is $DIST(w, v) - P$. ■

For instance, the distance between $\{p, q, r, \neg s\}$ and $\{q, s, \neg p, \neg r\}$ wrt $\{p\}$ is $\{r, s\}$.

Definition 5 The *update of an interpretation w by a set of interpretations V* wrt a set of atoms P, written $w \star^P V$, is the set of those elements of V whose distance to w wrt P is $\emptyset$. ■

Definition 6 The *update of a set of interpretations U by a set of interpretations V* wrt a set of atoms P, written $U \star^P V$, is given by

$$U \star^P V = \bigcup_{w \in U} w \star^P V. \blacksquare$$

In the modified PMA, the update $KB \star^P \alpha$ of a knowledge base KB by a formula α wrt a set of atoms P is identified with the formula corresponding to the set of interpretations $|KB| \star^P |\alpha|$.

Example 1 (continued) To properly deal with this example, the atoms p and q should be allowed to vary. $|KB| \star^{\{p,q\}} |\alpha| =$

$$\{p \neg q\} \star^{\{p,q\}} \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} \cup \{p, q\} \star^{\{p,q\}} \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} = \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} \cup \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\} = \{\{p, \neg q\}, \{q, \neg p\}, \{p, q\}\}.$$

Thus $KB \star^{\{p,q\}} \alpha$ is the formula $(p \wedge \neg q) \vee (q \wedge \neg p) \vee (p \wedge q)$ which is equivalent to $p \vee q$. ■

Example 2 (continued) As before, the variable atoms should be p and q. We leave it to the reader to check that $KB \star^{\{p,q\}} \alpha \equiv p \vee q$. ■.

It should be noted that the MPMA need not preserve consistency. For instance, if $ATM = \{q\}$, $KB = \neg q$, $\alpha = q$ and $P = \emptyset$, then $KB \star^P \alpha \equiv \perp$. However, as we shall see in the next section, the consistency preservation is guaranteed if variable atoms are properly chosen.

3.1 Determining Variable Atoms

In this section we discuss the fundamental question: how can one provide a mechanism to determine variable atoms which can be applied automatically and still guarantee intuitively reasonable conclusions?

Consider a knowledge base KB and an update formula α. Since α represents effects of a performed action, all atoms occurring in α are potential candidates for variable atoms. However, one should be cautious. The reason is that some atoms from α may be redundant. For example, if α is $p \vee (p \wedge q)$, then $\alpha \equiv p$ and so q should not be considered as a variable atom.

Example 3 Suppose that $ATM = \{p, q\}$, $KB = \neg p \wedge q$ and $\alpha = p$. Taking p as the only variable atom, we get $KB \star^{\{p\}} \alpha \equiv p \wedge q$.

Suppose now that we introduce a redundant atom to α, replacing α by its equivalent $\alpha' = (p \wedge q) \vee (p \wedge \neg q)$. Taking p and q as variables, we obtain $KB \star^{\{p,q\}} \alpha' \equiv p$. This result is clearly undesirable. ■

This is an example of what we call the *redundant atom anomaly.* It appears to be a flaw that several formalisms [15, 17, 13], based on the use of an *Occlusion* or *Release* predicate, succumb to. *Occlusion* and *Release* have the role of distinguishing variable fluents from non-variable fluents. In fact, it can be shown that other causal formalisms such as Thielscher [23] suffer from the same flaw when syntactic rules are not preprocessed to remove redundant atoms. Semantic approaches toward modeling action effect axioms such as Sandewall [21], where a *full trajectory normal form* (FTNF) is generated from a model-theoretic description, do not suffer from this problem because there are no redundant atoms in the normal form. On the other hand, the claim that the FTNF can be replaced with a logically equivalent formula does not appear to be correct for those logics dependent on the use of an *Occlude* predicate, where the choice of occluded fluents is dependent on syntax.

Definition 7 Let α be a formula. An atom p occurring in α is said to be *redundant* for α iff $\alpha[p \leftarrow \top] \equiv \alpha[p \leftarrow \bot]$.[2] ■

As follows from the above definition, an atom is redundant for a formula iff the logical value of the formula does not depend on the logical value of the atom.

Our choice of variable atoms is to identify them with the set of non-redundant atoms for the update formula. This poses the question of how these atoms are to be selected. Definition 7 is rather impractical from the computational point of view. Below, we present a generally more efficient method based on the notion of the *Blake canonical form* of a formula. Our discussion follows Brown [2].

We start with preliminary terminology.

A *term* is either $\top$, $\bot$, or a conjunction of literals in which no atom appears more than once. A formula is said to be in *disjunctive normal form* (*DNF*, for short) if it is a disjunction of different terms.[3] It is well-known that each formula can be constructively transformed into its equivalent in *DNF*. We say that a term t_1 *absorbs* a term t_2 if either t_1 is $\top$, or t_2 is $\bot$, or t_1 is a subterm of t_2. For instance, the term p absorbs the term $p \wedge q$. Let α be a formula in *DNF*. We write $ABS(\alpha)$ to denote the formula obtained from α by deleting all absorbed terms. Clearly, α and $ABS(\alpha)$ are equivalent.

Two terms are said to have an *opposition* if one of them contains an atom p and the other the atom $\neg p$. For instance, the terms $\neg q \wedge r$ and $q \wedge s$ have a single opposition, in the atom q.

Suppose that two terms, t_1 and t_2, have exactly one opposition. Then the *consensus* of t_1 and t_2, written $c(t_1, t_2)$, is the term obtained from the conjunction $t_1 \wedge t_2$ by deleting the opposed atoms as well as any repeated atoms. For example, $c(\neg q \wedge p, q \wedge r)$ is $p \wedge r$.

Let α be a formula. The *Blake canonical form of* α, written $BCF(\alpha)$, is the formula obtained from α by the following construction.

(1) Replace α by its disjunctive normal form. Denote the resulting formula by β.

(2) Repeat as long as possible:
If β contains a pair t_1 and t_2 of terms whose consensus exists and no term of β is a subformula of $c(t_1, t_2)$, then $\beta := \beta \vee c(t_1, t_2)$.

(3) Take $ABS(\beta)$. This is $BCF(\alpha)$.

The following results can be found in Brown [2].

Theorem 1
(1) Formulas α and $BCF(\alpha)$ are equivalent.
(2) All atoms occurring in $BCF(\alpha)$ are non-redundant for α. ■

[2] Recall that $\alpha[p \leftarrow \top]$ (resp. $\alpha[p \leftarrow \bot]$) is the formula obtained from α by replacing all occurrences of p by $\top$ (resp. $\bot$).

[3] In the logical literature *DNF* is often defined as a disjunction of terms where a term is understood as either $\top$, or $\bot$ or any conjunction of literals. Note, however, that we can always restrict ourselves to terms in which no atom appears more than once: each repeated occurence of an atom p can be removed from a term, whereas any term including p and $\neg p$ can be replaced by $\bot$.

Example 4 Let α be $(\neg p \wedge q) \vee (p \wedge r \wedge \neg s) \vee (p \wedge \neg s \wedge \neg q)$. Since α is in disjunctive normal form, $\beta = \alpha$. After performing step (2), the result is,

$$(\neg p \wedge q) \vee (p \wedge r \wedge \neg s) \vee (p \wedge \neg s \wedge \neg q) \vee (r \wedge q \wedge \neg s). \quad (1)$$

Since $ABS((1)) = (1)$, the formula (1) is the Blake canonical form of α. ■

As we noted in the previous section, the MPMA need not preserve consistency. However, if all non-redundant atoms of the update formula are considered as variable atoms, consistency preservation is assured. In the sequel, we write $atm(\alpha)$ to denote the set of all atoms occurring in a formula α.

Theorem 2 For any knowledge base $KB \not\equiv \bot$ and any formula $\alpha \not\equiv \bot$, $P = atm(BCF(\alpha))$ implies $KB *^P \alpha \not\equiv \bot$. ■

3.2 Syntactic Characterization of the MPMA

In this section, we present a syntactic characterization of the MPMA. We start with some terminology.

Let p be an atom and suppose that α is a formula. We write $\exists p.\alpha$ to denote the formula $\alpha[p \leftarrow \top] \vee \alpha[p \leftarrow \bot]$. If $P = \{p_1, \ldots, p_n\}$ is a set of atoms and α is a formula, then $\exists P.\alpha$ stands for $\exists p_1 \cdots \exists p_n.\alpha$.

A formula of the form $\exists P.\alpha$, where $P = \{p_1, \ldots, p_n\}$, is called an *eliminant of* $\{p_1, \ldots, p_n\}$ *in* α. Intuitively, such an eliminant can be viewed as a formula representing the same knowledge as α about all atoms from $ATM - P$ and providing no information about the atoms in P. The reader interested in a detailed theory of eliminants should consult Brown [2].

The following result, which easily follows from the definition of an eliminant, provides its semantical characterization.

Theorem 3 Let α be a formula and suppose that P is a set of atoms. Then $|\exists P.\alpha| =$

$\{u :$ exists $w \in |\alpha|$ such that $DIST(w, u) - P = \emptyset\}$. ■

We are now ready to provide a syntactic characterization of the MPMA.

Theorem 4 Let KB, α and $P = atm(BCF(\alpha))$ be a knowledge base, a formula and a set of atoms, respectively. Then $KB \star^P \alpha \equiv \alpha \wedge \exists P.KB$. ■

Theorem 4 clearly shows how the MPMA works. First, we select the atoms that may vary their values when the action corresponding to the update formula α is performed. Next, we weaken the knowledge base KB by eliminating all those variable atoms. Finally, we strengthen the knowledge base $\exists P.KB$ by combining it with the update formula.

3.3 Computing the MPMA

In view of Theorem 4, all we need to compute the MPMA is the ability to compute eliminants. The following method can be found in Brown [2].

Let α be a formula and suppose that $P = \{p_1, \ldots, p_k\}$ is a set of atoms. The eliminant $\exists P.\alpha$ is the formula obtained by the following construction.

(1) Replace α by its disjunctive normal form $t_1 \vee \cdots \vee t_n$, where $t_1, \ldots, t_n$ are terms. Denote the resulting formula by β.

(2) From each t_i $(1 \leq i \leq n)$, remove all occurrences (both positive and negative) of $p_1, \ldots, p_k$. If during this process all literals have been removed from some term t_i, stop and return $\top$ as $\exists P.\alpha$. Otherwise, return the resulting formula as $\exists P.\alpha$.

Example 5 Let $\alpha = (p \rightarrow q) \wedge (\neg q \vee r)$. We compute $\exists q.\alpha$. Converting α into its disjunctive normal form, we get $\beta = (\neg p \wedge \neg q) \vee (\neg p \wedge r) \vee (q \wedge r)$. Eliminating the atom q, we finally obtain $\neg p \vee (\neg p \wedge r) \vee r$ which is equivalent to $\neg p \vee r$.

Example 6 This is a classical example from Winslett [25]. We have two atoms, b and m, standing for "a book is on the floor" and "a magazine is on the floor", respectively. $KB = (b \wedge \neg m) \vee (\neg b \wedge m)$. The update formula α is m. Since α is already in the Blake canonical form, the only variable atom is m. We first compute $\exists m.KB$. KB is already in disjunctive normal form. Eliminating all occurrences of m in KB, we get $b \vee \neg b$ which is equivalent to $\top$. Hence, $KB *^m \alpha$ is $m \wedge \top$ which is equivalent to m. ■.

3.4 Properties of the MPMA

Katsuno & Mendelzon [14] proposed eight postualtes which, as they claimed, should be satisfied by each update operator. In this section, we analyse the MPMA in the context of these postulates.

Let KB be a knowledge base and suppose that α is an update formula. In what follows, we shall write $KB \star \alpha$ as an abbreviation for $KB \star^P \alpha$, where P is the set of all atoms occurring in the Blake canonical form of α.

Katzuno and Mendelzon set up the following postulates:

(1) $KB \star \alpha$ implies α.

(2) If KB implies α, then $KB \star \alpha$ is equivalent to KB.

(3) If both KB and α are satisfiable, then $KB \star \alpha$ is also satisfiable.

(4) If $KB_1 \equiv KB_2$ and $\alpha_1 \equiv \alpha_2$, then $KB_1 \star \alpha_1 \equiv KB_2 \star \alpha_2$.

(5) $(KB \star \alpha_1) \wedge \alpha_2$ implies $KB \star (\alpha_1 \wedge \alpha_2)$.

(6) If $KB \star \alpha_1$ implies α_2 and $KB \star \alpha_2$ implies α_1, then $KB \star \alpha_1 \equiv KB \star \alpha_2$.

(7) If KB is complete, *i.e.* has at most one model, then $(KB \star \alpha_1) \wedge (KB \star \alpha_2)$ implies $KB \star (\alpha_1 \vee \alpha_2)$.

(8) $(KB_1 \vee KB_2) \star \alpha \equiv (KB_1 \star \alpha) \vee (KB_2 \star \alpha)$.

We have already seen (Example 1 (continued)) that postulate (2) does not generally hold in MPMA. This is also the case as regards postulates (5), (6) and (7).

To see that (5) does not generally hold, take $ATM = \{m, b\}$, $KB = \neg m \wedge b$, $\alpha_1 = \neg b \vee m$ and $\alpha_2 = b \vee m$. $(KB \star \alpha_1) \wedge \alpha_2 = (\neg b \vee m) \wedge (b \vee m) \equiv m$. On the other hand, $KB \star (\alpha_1 \wedge \alpha_2) \equiv KB \star m \equiv b \wedge m$. Clearly, m does not imply $b \wedge m$. We claim that our result is intuitively plausible. Let us interpret b and m as "a book is on the table" and "a magazine is on the table", respectively. KB states that initially the magazine is not on the table and the book is on it. α_1 corresponds to the action "remove the book from the table or put the magazine on it (or perform both these subactions together). It is obvious that the updated knowledge base should be $\neg b \vee m$. If we take the conjunction of the new knowledge base and α_2, we will obtain m. On the other hand, $\alpha_1 \wedge \alpha_2$ is equivalent to m, which means that the action it corresponds to is "put the magazine on the table". Now, if we apply this action to our original knowledge base, the updated knowledge base should be $m \wedge b$.

It turns out that the following weaker form of postulate (5) holds for MPMA:

Theorem 5 If $atm(BCF(\alpha_1)) \subseteq atm(BCF(\alpha_1 \wedge \alpha_2))$, then $(KB \star \alpha_1) \wedge \alpha_2$ implies $KB \star (\alpha_1 \wedge \alpha_2)$. ■

Intuitively, Theorem 5 states that postulate (5) holds, provided that the action corresponding to $\alpha_1 \wedge \alpha_2$ may vary every atom that can be varied by the action corresponding to α_1.

To see that (6) need not hold, suppose that $ATM = \{b, m\}$, $KB = \neg b \wedge m$, $\alpha_1 = (b \vee m)$ and $\alpha_2 = \top$. $KB \star \alpha_1 \equiv b \vee m$ implies α_2. Similarly, $KB \star \alpha_2 \equiv \neg b \wedge m$ implies α_1. On the other hand, $KB \star \alpha_1$ and $KB \star \alpha_2$ are not equivalent. We leave it to the reader to check that the interpretation of the atoms b and m as "a book is on the table" and "a magazine is on the table", respectively, makes our result plausible.

The following weaker form of postulate (6) holds for the MPMA:

Theorem 6 Let $atm(BCF(\alpha_1)) = atm(BCF(\alpha_2))$. If $KB \star \alpha_1$ implies α_2 and $KB \star \alpha_2$ implies α_1, then $KB \star \alpha_1 \equiv KB \star \alpha_2$. ■

Intuitively, Theorem 6 states that postulate (6) holds, provided that the actions corresponding to α_1 and α_2 may vary exactly the same atoms.

To see that (7) need not to hold, suppose that $ATM = \{b, m, n\}$, $KB = m \wedge b \wedge n$, $\alpha_1 = (b \wedge m) \vee n$ and $\alpha_2 = (b \wedge \neg m) \vee n$. It is easily verified that $(KB \star \alpha_1) \wedge (KB \star \alpha_2) = \alpha_1 \wedge \alpha_2$ which is equivalent to n. On the other hand, $\alpha_1 \vee \alpha_2$ reduces to $b \vee n$, when the redundant atom m is eliminated. So, $KB \star (\alpha_1 \vee \alpha_2)$ is equivalent to $(b \vee n) \wedge m$ which is not a consequence of n. We leave it to the reader to check that our result is plausible if the atoms b, m and n are interpreted as "a book is on the table", "a magazine is on the table" and "a newspaper is on the table", respectively.

The following version of postulate (7) holds for the MPMA:

Theorem 7 If $atm(BCF(\alpha_1) \subseteq atm(BCF(\alpha_1 \vee \alpha_2))$ and $atm(BCF(\alpha_2) \subseteq atm(BCF(\alpha_1 \vee \alpha_2))$, then $(KB \star \alpha_1) \wedge (KB \star \alpha_2)$ implies $KB \star (\alpha_1 \vee \alpha_2)$.[4] ■

Intuitively, Theorem 7 says that postulate (7) holds, provided that the set of variable atoms of the action corresponding to α_1 and the set of variable atoms of the action corresponding to α_2 are both subsets of the set of variable atoms of the action corresponding to $\alpha_1 \vee \alpha_2$.

The remaining postulates of Katsuno & Mendelzon hold for the MPMA:

Theorem 8 The MPMA satisfies postulates (1), (3)-(4) and (8). ■

4 A Simple Temporal Logic

Since our primary concern in this section is to compare the knowledge-base update and temporal logic

[4] The assumption that KB is complete is not necessary here.

approaches to reasoning about action and change, we will only be concerned with action scenarios specified in terms of an initial state description S, a formula α describing an effect of an action executed in S and a resulting state description S'. This will correlate with an initial KB, an update formula α added to the KB, and the resulting KB. The above restrictions allow us to work with a very simple pseudo-temporal logic, referred to as STL. Actually, STL is the classical propositional logic which simulates a fragment of the 1st-order temporal logic in Doherty [4, 13], restricted to two time-points.

4.1 The Language $\mathcal{L}_{stl}$

Given the language $\mathcal{L}_{pma}$, based on a set of atoms ATM, we define the language $\mathcal{L}_{stl}$ as that of classical propositional logic over ATM_{stl}, where

$$ATM_{stl} = \{p^I : p \in ATM\} \cup \{p^R : p \in ATM\} \cup \{O^p : p \in ATM\}.$$

In other words, in the language $\mathcal{L}_{stl}$, each atom $p \in ATM$ is replaced by two copies p^I and p^R. In addition, for each atom $p \in ATM$, we introduce an auxiliary atom O^p. Intuitively, p^I and p^R represent the values of p in the initial state and in the resulting state, respectively. The atom O^p informally states that p is *occluded*, where occlusion plays the role of releasing an atom from keeping its value persistent from the initial state to the final state.

By an *I-formula* (resp. *R-formula*) we mean any formula constructable using the atoms of the form p^I (resp. p^R). For any formula $\alpha \in \mathcal{L}_{pma}$, we write α^I (resp. α^R) to denote the result of replacing each atom symbol p in α by p^I (resp. p^R).

The following straightforward result will be useful.

Proposition 1 For any $\alpha \in \mathcal{L}_{pma}$ and any $\beta \in \mathcal{L}_{pma}$, $\alpha \models \beta$ iff $\alpha^R \models \beta^R$, where $\models$ denotes the entailment relation of classical propositional logic. ■

Let α be an R-formula and suppose that $p_1^R, \ldots, p_n^R$ are all the atoms occurring in α. We write $Occlude(\alpha)$ as an abbreviation for $O^{p_1} \wedge \cdots \wedge O^{p_n}$.

4.2 Action Scenarios

Action scenarios, or narratives, are often used in the TL paradigm to represent sequences of action occurrences and observations together with a set of timing constraints which provide a partial or total order on the action occurrences. One then proposes an entailment policy, views the narrative as a theory and deduces facts about the narrative. Below, we provide a lightweight version of scenarios appropriate for our two state sequences.

Definition 1 (STL State Description) An *initial state description* is any I-formula. ■

Definition 2 (STL Action Description) An *action description* is the conjunction $\alpha \wedge Occlude(\alpha)$, where α is any R-formula. ■

Definition 3 (Action Scenario) An *action scenario* $\mathcal{A}$ is the conjunction $\alpha \wedge \beta$, where α is an initial state description and β is an action description. ■

4.3 The Minimization Policy

The minimization policy for STL is based on the PMON policy introduced by Sandewall [21] and investigated and extended in Doherty *et al.* [4, 5, 13].

A *nochange axiom*, written NCA, is the conjunction

$$\bigwedge_{p \in ATM} \neg O^p \rightarrow (p^I \leftrightarrow p^R).$$

Intuitively, the nochange axiom states that the value of each non-occluded atom persists from the initial state to the resulting state.

Given the language $\mathcal{L}_{stl}$, we write OCC to denote the set of all atoms of the form O^p.

We are now ready to state the minimization policy. Given an action scenario $\mathcal{A}$ we will use the following filtered preferential entailment policy:

$$NCA \wedge Circ[\mathcal{A}; OCC],$$

where $Circ[\Gamma; P]$ represents the standard second-order circumscription axiom which minimizes the predicates P relative to the finite formula Γ. In other words, we minimize the atoms of the form O^p relative to the action description conjoined with the initial state description, and then filter with the nochange axiom NCA. Since the predicates we minimize are actually 0-ary ones, $Circ[\mathcal{A}; OCC]$ is reducible to propositional logic. Consequently, any efficient decision procedure for classical propositional logic will suffice for computing preferential entailment.

Definition 4 (Preferential Entailment) Let $\mathcal{A}$ be an action scenario and suppose that $\alpha \in \mathcal{L}_{stl}$. We say that α is *preferentially entailed* from $\mathcal{A}$ in STL, written $\mathcal{A} |\!\approx_{stl} \alpha$, iff $NCA \wedge CIRC[\mathcal{A}; OCC] \models \alpha$, where $\models$ denotes the entailment relation of classical propositional logic. ■

5 Reformulating the MPMA in STL

In this section, we provide a reformulation of the MPMA in terms of STL.

Definition 5 Let $KB \in \mathcal{L}_{pma}$ and $\alpha \in \mathcal{L}_{pma}$ be a knowledge base and an update formula, respectively. An action scenario $\mathcal{A}$ *corresponding* to KB and α is given by

$$\mathcal{A} = KB^I \wedge \alpha^R \wedge Occlude(\alpha).$$

Example 7 Suppose that $KB = p$ and $\alpha = p \vee q$. The action scenario corresponding to KB and α is $p^I \wedge (p^R \vee q^R) \wedge O^p \wedge O^q$. ■

5.1 Proving Equivalence between MPMA and STL

Theorem 9 Assume that the language $\mathcal{L}_{pma}$ is based on a set of atoms ATM. Let $KB \in \mathcal{L}_{pma}$ and $\alpha \in \mathcal{L}_{pma}$ be a knowledge base and an update formula, respectively. Suppose that α is in Blake canonical form and let P be the set of all atoms from ATM occurring in α. If $\mathcal{A}$ is an action scenario corresponding to KB and α, then for any formula $\beta \in \mathcal{L}_{pma}$

$$KB \star^P \alpha \models \beta \quad \text{iff} \quad \mathcal{A} | \!\!\approx_{stl} \beta^R. \blacksquare$$

5.2 The Decision Procedure for MPMA

Theorem 9 provides the basis for a simple decision procedure for computing the MPMA using STL.[5]

Let $\mathcal{L}_{pma}$ and $\mathcal{L}_{stl}$ be the languages used to describe knowledge base update and action scenario queries, respectively. Given an MPMA update query

$$KB \star^P \alpha \models \beta,$$

where KB and α are formulas in $\mathcal{L}_{pma}$ and α is in Blake canonical form, we first translate the update $KB \star \alpha$ into a corresponding action scenario $\mathcal{A}$, and we translate the query β in $\mathcal{L}_{pma}$ into its correlate β' in $\mathcal{L}_{stl}$. We then solve the equivalent problem

$$NCA \wedge CIRC[\mathcal{A}; OCC] \models \beta'$$

using a decision procedure for classical propositional logic.

Example 8 Let $ATM = \{p, q\}$, $KB = p \wedge q$ and $\alpha = \neg q$. We use Theorem 9 to show that p is entailed by $KB \star^{\{q\}} \alpha$. An action scenario corresponding to KB and α is $\mathcal{A} = p^I \wedge q^I \wedge \neg q^R \wedge O^q$. It is easily checked that $CIRC[\mathcal{A}, OCC] \equiv p^I \wedge q^I \wedge \neg q^R \wedge O^q \wedge \neg O^p$ and so $NCA \wedge CIRC[\mathcal{A}, OCC] \equiv p^I \wedge q^I \wedge \neg q^R \wedge O^q \wedge \neg O^p \wedge (p^I \leftrightarrow p^R)$. Since $NCA \wedge CIRC[\mathcal{A}, OCC] \models p^R$, we conclude that $KB \star^{\{q\}} \alpha \models p$. ■

[5]For an on-line implementation of a restricted first-order version of STL, called TAL (Temporal Action Logic), see http://www.anton.ida.liu.se/vital/vital.html.

6 The MPMA and Integrity Constraints

In this section, we provide a means of extending the MPMA operator for integrity constraints. To deal with integrity constraints in the framework of the classical PMA, it has been proposed in Katsuno & Mendelzon [14] that one define such an extension by $KB \star (A \wedge IC)$, where $\star$ denotes the classical PMA operator and IC is the conjunction of the set of integrity constraints under consideration. Unfortunately, as has been observed by many researchers (see for instance Herzig [11], this solution is very problematic. To show this, we use a classical example in Herzig [11], originally due to Lifschitz [16].

Example 9 There are three atoms $sw1$, $sw2$ and l standing for "switch 1 is up", "switch 2 is up" and "light is on", respectively. The integrity constraint $IC = l \leftrightarrow (sw1 \leftrightarrow sw2)$. We would like to update the database $KB = l \wedge sw1 \wedge sw2$ with the update formula $\alpha = \neg sw1$. It is well known that describing the problem with the classical PMA as

$$KB \star (\alpha \wedge IC),$$

will not provide intuitive results. The new knowledge base has two models, w_1 and w_2, given by $\{\neg sw1, \neg l, sw2\}$ and $\{\neg sw1, l, \neg sw2\}$, respectively. The model w_2 is obviously an unintended one.

It is interesting to note that using this solution, but replacing the PMA with the MPMA, gives the same two models. ■

The problem which is well understood by now is that integrity constraints must include fluent dependency information in one form or another. In this paper, we shall employ *causal rules* as described in Gustafsson & Doherty [13]. These are expressions of the form

$$\alpha \gg \beta \tag{2}$$

where α and β are formulas, referred to as an *antecedent* and a *consequent* of the rule, respectively. A rule of the form (2) has the following intuitive interpretation:

(1) The formula $\alpha \to \beta$ holds in both the initial and the updated knowledge base.

(2) If $\neg\alpha$ held in the initial knowledge base and α holds in the updated knowledge base, then there is a cause for β to hold in the updated knowledge base.[6]

Recall that in the MPMA we minimize change, excluding from the minimization the atoms occurring in the update formula. In the MPMA with integrity constraints we, in addition, exclude from the minimization the atoms occurring in the consequents of *active* causal rules, where a rule is said to be active iff its antedecent changed its truth-value from *False* to *True* during the update.

Example 9 (continued) To encode causal information contained in IC, we introduce two rules: $(sw1 \leftrightarrow sw2) \gg l$ and $\neg(sw1 \leftrightarrow sw2) \gg \neg l$. Suppose we want to check whether $w_1 = \{\neg sw1, \neg l, sw2\}$ should be considered as a model of KB updated by $\neg sw1$ under the integrity constraints IC and the above causal rules. Denote by u the only model of $KB \wedge IC \wedge [(sw1 \leftrightarrow sw2) \rightarrow l] \wedge [\neg(sw1 \leftrightarrow sw2) \rightarrow \neg l]$, *i.e.* $u = \{l, sw1, sw2\}$.[7] Note that the antecedent of the rule $\neg(sw1 \leftrightarrow sw2) \gg \neg l$ has changed its value from $False$ to $True$ during the update. Therefore, we are allowed to ignore the atom l while minimizing change. The atom $sw1$ is also ignored since it occurs in the update formula. On the other hand, the atom $sw2$ has the same value in both u and w_1. Accordingly, w_1 is to be viewed as an intended model of the considered update. Consider now the model $w_2 = \{\neg sw1, l, \neg sw2\}$. It is readily checked that we are allowed to ignore $sw1$ and l while minimizing change. However, the models u and w_1 still differ in the atom $sw2$. Acordingly, w_2 should not be viewed as the intended model of the considered update. ■

We now formalize the above idea.

Definition 6 A causal rule $\alpha \gg \beta$ is said to be *active* wrt a pair of interpretations $\langle w, v\rangle$ iff the truth-values of α in w and u are $False$ and $True$, respectively. ■

Definition 7 Let w and u be two interpretations and suppose that α is a formula in Blake canonical form. Assume further that $CR = \{\alpha_i \gg \beta_i\}$ is a set of causal rules such that all β_i's are in Blake canonical form. Let P be a set of atoms given by $P =$

$$atm(\alpha) \cup \bigcup \{atm(\beta_i) : \alpha_i \gg \beta_i \text{ is active wrt } \langle w, u\rangle\}.$$

The *distance between* w *and* u *wrt* α *and* CR is $DIST(w,u) - P$. ■

Definition 8 Let α and CR be as in Definition 7. The *update of an interpretation* w *by a set of interpretations* V wrt α and CR, written $w \star_{\alpha,CR} V$, is the set of those elements from V whose distance wrt α and CR is $\emptyset$. ■

Definition 9 Let α and CR be as in Definition 7. Let KB be a knowledge base and suppose that IC is a formula representing integrity constraints. Let U be the set of all models of $KB \wedge IC \wedge T(CR)$, where $T(CR)$ denotes the formula $\bigwedge_{i=1}^{n}(\alpha_i \rightarrow \beta_i)$ and let V be the set of all models of $\alpha \wedge IC \wedge T(CR)$. The *update* of KB by α wrt IC and CR, written $KB *_{IC,CR} \alpha$, is given by

$$KB *_{IC,CR} \alpha = \bigcup_{w \in KB \wedge IC \wedge T(CR)} w \star_{\alpha,CR} V. \blacksquare$$

6.1 STL and Integrity Constraints

To represent the MPMA with integrity constraints in STL, we first slightly generalize the notion of an action scenario.

Let $CR = \{\alpha_i \gg \beta_i : i = 1, \ldots, n\}$ be a set of causal rules. We write $Trans(CR)$ to denote the formula

$$\bigwedge_{i=1}^{n}(\alpha_i^I \rightarrow \beta_i^I) \wedge \bigwedge_{i=1}^{n}(\alpha_i^R \rightarrow \beta_i^R) \wedge \bigwedge_{i=1}^{n}(\neg\alpha_i^I \wedge \alpha_i^R \rightarrow Occlude(\beta_i)).$$

Definition 10 An *action scenario under integrity constraints* IC *and a set of causal rules* CR is the conjunction $IC^I \wedge IC^R \wedge Trans(CR) \wedge \alpha \wedge \beta$, where α is an initial state description and β is an action description. ■

Definition 11 Let KB, α, IC and CR be a knowledge base, an update formula, a formula representing integrity constraints and a set of causal rules, respectively. An action scenario $\mathcal{A}$ *corresponding* to KB, α, IC and CR is given by

$$\mathcal{A} = KB^I \wedge IC^I \wedge \alpha^R \wedge Occlude(\alpha) \wedge IC^R \wedge Trans(CR).$$

■

[6] That is, if $\neg\alpha$ held in the initial knowledge base and α holds in the updated knowledge base, we are allowed to change the values of atoms occurring in β to guarantee that this formula holds in the updated knowledge base.

[7] If KB is a knowledge base, IC is a formula representing integrity constraints and the set of causal rules is $\cup_{i=1}^{n}\{\alpha_i \gg \beta_i\}$, then the original knowledge base should be considered not as KB, but rather as $KB \wedge IC \wedge \bigwedge_{i=1}^{n}(\alpha_i \rightarrow \beta_i)$.

Theorem 10 Let KB, α, IC and CR be a knowledge base, an update formula, a formula representing the set of integrity constraints and a set of causal rules, respectively. Suppose that α and all consequents of causal rules are in Blake canonical form. If $\mathcal{A}$ is an action scenario corresponding to KB, α, IC and CR, then for any formula $\beta \in \mathcal{L}_{pma}$

$$KB \star_{IC,CR} \alpha \models \beta \quad \text{iff} \quad \mathcal{A} | \!\approx_{stl} \beta^R. \blacksquare$$

Practically, the theorem shows that we have provided an alternative syntactic characterization and implementation of MPMA with integrity constraints in terms of STL. Since STL is a very restricted version of a richer temporal logic TAL [4, 13], it would be straightforward to provide additional generalizations to MPMA, based on intuitions from TAL.

7 Comparisons with Existing Work

In this section, we compare our approach with that of Herzig [11], whose work is most closely related to ours. Additional comparisons and observations are made with del Val and Shoham [24] and Sandewall [22] in Doherty *et al.* [6].

7.1 Herzig

Herzig [11] provides a sound and complete decision procedure for the PMA by constructing an equivalent presentation of the PMA in terms of conditional logic.[8] A syntactic characterization of an update $w \star U$ of an interpretation w by a set of interpretations U is provided in terms of a conditional operator $>$ where $A > C$ is read as a hypothetical update: "if the current database is updated with A then C follows". Herzig proves the following proposition:

Proposition 1 (Herzig [11]) *Let KB, A and C be classical. Then $KB \star A \models C$ iff $KB \models A > C$.* $\blacksquare$

Herzig then shows how the conditional $A > C$ can be rewritten to a classical formula using normal forming which together with the proposition above provides the decision procedure.

In Section 8 of his paper, Herzig sketches how one might deal with integrity constraints by introducing a new update operator which after careful analysis can be shown to be quite similar in spirit to the MPMA operator presented in this paper. Although similar in concept, Herzig's decision procedure is still based on translating update queries into conditionals and then translating via normal forming to a classical formula. In addition, instead of computing one PMA update, he is forced to compute 2^n classical PMA updates per query, where n is the number of atoms dependent on any of the atoms in the update formula.

To show the relation between the MPMA update operator $\star^P$ and Herzig's new update operator, written $\star_{IC,DEP}$, we will use Example 9 which is also used in Herzig [11], but with the atoms renamed.

In addition to the integrity constraint IC, one must also include fluent dependency information in one form or another. Herzig does this in the following manner: He first specifies a dependency function, DEP, from atoms to sets of atoms such that $p \in DEP(p)$, for each atom p. In addition, $DEP(p)$ may contain atoms other than p. Intuitively, $q \in DEP(p)$ means that updates concerning p may change the truth-value of q. For a formula A, $DEP(A)$ stands for $\bigcup_{[p \in atm(A)]} DEP(p)$.

In the example, $DEP(l) = \{l, sw1, sw2\}$, $DEP(sw1) = \{sw1, l\}$, and $DEP(sw2) = \{sw2, l\}$. The dependence function DEP encodes the change dependency inherent in a causal reading of IC and represents a fluent dependency graph.

An update operator $\star_{IC,DEP}$ under a set of integrity constraints IC and a dependence function DEP is then defined by

$$KB \star_{IC,DEP} A = \left(\bigvee_{B \in CTX(A)} (KB \star B) \right) \wedge IC \wedge A$$

where $\star$ is the classical PMA operator and $CTX(A)$ is $\{l_1 \wedge \cdots \wedge l_n : l_i = p_i \,\text{or}\, l_i = \neg p_i\}$ if $DEP(A) = \{p_1, \ldots, p_n\}$. So, for example, if $A = \neg sw1$ and the function DEP is as specified above where $DEP(sw1) = \{sw1, l\}$, then

$$CTX(sw1) = \{sw1 \wedge l, sw1 \wedge \neg l, \neg sw1 \wedge l, \neg sw1 \wedge \neg l\}.$$

This means that if the cardinality of $DEP(\alpha)$ is n then the cardinality of $CTX(\alpha)$ is 2^n.

The role of $CTX(A)$ is central to the approach and essentially does the job that the set P of atoms in the MPMA operator $\star^P$ does together with the eliminant $\exists P.KB$, or that the OCC atoms together with the NCA axiom plays in the translation to STL. Taking the disjunction of PMA updates of KB with each of the formulas in $CTX(sw1)$, has the effect of generating the same models as one would with the eliminant of KB where $P = \{sw1, l\}$. It can easily be shown that

$$KB \star_{IC,DEP} \neg sw1 \equiv KB \star_{IC,CR} \neg sw1,$$

[8] A relationship between the PMA and conditionals has also been studied in Fariñas del Cerro *et al.* [7, 8, 9] and Grahne [12].

where,

$$KB = l \wedge sw1 \wedge sw2,$$
$$IC = (sw1 \leftrightarrow sw2) \leftrightarrow l,$$

and

$$CR = \{[(sw1 \leftrightarrow sw2) \gg l], \\ [\neg(sw1 \leftrightarrow sw2) \gg \neg l]\},$$

or the translation into STL where,

$$\mathcal{A} = KB^I \wedge IC^I \wedge IC^R \wedge \neg sw1^R \wedge O^{sw1} \wedge \mathit{Trans}(CR).$$

The advantage of using the latter is that the decision procedure is much more efficient and that the dependency information included in DEP is implicit in the expansion of the abbreviation $\gg$.

Herzig's new update operator can be also used for the empty set of integrity constants by putting $IC = \top$. It turns out, that the operator $\star_{\top,DEP}$ is just the MPMA operator, with the set of variable atoms identified with the set of all the atoms occurring in the update formula.[9] More precisely:

Theorem 11 For each KB and A

$$KB \star_{\top,DEP} A \equiv KB \star^P (A)$$

where $P = atm(A)$. ■

8 Discussion

We have demonstrated the benefits of applying intuitions derived from research on action and change in the temporal reasoning community to the BR/U paradigm. The result is a generalization of the PMA to MPMA which handles disjunctive updates and integrity constraints interpreted as causal constraints. We have provided a number of syntactic characterizations of the MPMA and shown how an MPMA query can be mapped into a query of a temporal narrative represented in STL, a simple temporal logic for reasoning about action and change. Since STL is the propositional fragment of a highly expressive 1st-order temporal logic, it should be straightforward to generalize the results described here to the 1st-order case, where it has already been shown that the circumscription formula associated with the minimization policy in the full first-order version of STL is reducible to a first-order formula. An additional generalization would include the modeling of iterated belief update in MPMA in terms of narratives with more than two timepoints. Because the minimization policy associated with STL is similar to many other current approaches using temporal logics, we also believe that additional generalizations to the BR/U paradigm such as distinguishing between observations and action effects, or concurrent update can be added to MPMA or at least fully understood in this context.

[9] This means that Herzig's new update operator does not distinguish between redundant and non-redundant atoms occurring in the update formula. In consequence, it is syntax dependent and may lead to counterintuitive results (see Example 3).

Acknowledgements

Patrick Doherty and Witold Łukaszewicz have been partially supported by grants from the Wallenberg Foundation and the Swedish Research Council for Engineering Sciences (TFR). Ewa Madalińska-Bugaj is supported by grant KBN 8 T11C 035 111. Thanks to Jonas Kvarnström and Patrik Haslum for comments.

References

[1] G. Brewka, J. Hertzberg. How to Do Things with Worlds: on Formalizing Actions and Plans. *Journal of Logic and Computation.* Vol. 3, No. 5, 1993, 517-532.

[2] F. M. Brown. *Boolean Reasoning.* Kluwer Academic Publishers, 1990.

[3] M-O. Cordier, P. Siegel. A Temporal Revision Model for Reasoning about World Change. In *Int'l. Conference on Principles of Knowledge Representation and Reasoning*, 732-739, Morgan Kaufmann, 1992.

[4] P. Doherty. Reasoning about Action and Change Using Occlusion. In: *Proc. of the 11th European Conference on Artificial Intelligence*, 1994, 401-405.

[5] P. Doherty, W. Łukaszewicz. Circumscribing Features and Fluents. In: *Proc. of the 1st Int'l. Conference on Temporal Logic*, Lecture Notes on Artificial Intelligence, vol. **827**, Springer-Verlag, 1994, 82-100.

[6] P. Doherty, W. Łukaszewicz, E. Madalińska-Bugaj. The PMA and Relativizing Minimal Change for Action Update: Extended Report. *Linköping Electronic Articles in Computer and Information Science*, 1998. `http://www.ep.liu.se/ea/cis/`.

[7] L. Fariñas del Cerro, A. Herzig. An Automated Modal Logic for Elementary Changes. In: P. Smets, A. Mamdami, D. Dubois and H. Prade (eds.), *Non-Standard Logics for Automated Reasoning*, Academic Press, 1988, 63-79.

[8] L. Fariñas del Cerro, A. Herzig. A Conditional Logic for Updating in the Possible Model Approach. In: *Proc. of the 18th German Conference on Artificial Intelligence*, Springer-Verlag, LNAI, **861**, 1994.

[9] L. Fariñas del Cerro, A. Herzig. All Other Things Being Equal: On a Notion of Inertia in Conditional Logic. In: E. Ejerhed and S. Lindström (eds.), *Logic, Action and Cognition*, Kluwer Academic Press, to appear.

[10] N. Y. Foo, A. S. Rao. Minimal Change and Maximal Coherence: A Basis for Belief Revision and Reasoning about Actions. *In Proc. of the 11th Int'l. Joint Conference on Artificial Intelligence*, 1989.

[11] A. Herzig. The PMA Revisited. In: *Proc. of the 5th Int'l. Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1996, 40-50.

[12] G. Grahne. Updates and Counterfactuals. In *Proc. of the 2nd Int'l. Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1991, 269-276.

[13] J. Gustafsson, P. Doherty. Embracing Occlusion in Specifying the Indirect Effects of Actions. In: *Proc. of the 5th Int'l. Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1996, 87-98.

[14] H. Katsuno, A. O. Mendelzon. On the Difference between Updating a knowledge base and revising it. In: *Proc. of the 2nd Int'l. Conference on Principles of Knowledge Representation and Reasoning*, KR'91, Morgan Kaufmann, 1991, 387-395.

[15] G.N. Kartha, V. Lifschitz. Actions with Indirect Effects (Preliminary Report). In *Proc. of the 4th Int'l. Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1994.

[16] V. Lifschitz. Frames in the Space of Situations. *Artificial Intelligence J.*, **46**, 1990, 365-376.

[17] F. Lin. Embracing Causality in Specifying the Indirect Effects of Actions. *In Proc. of the 14th Int'l. Joint Conference on Artificial Intelligence*, 1995.

[18] F. Lin. Embracing Causality in Specifying the Indeterminate Effects of Actions. *In Proc. of the National Conference on Artificial Intelligence*, 1996.

[19] N. McCain, H. Turner. A Causal Theory of Ramifications and Qualifications. In: *Proc. of the 14th Int'l. Joint Conference on Artificial Intelligence*, Morgan Kaufmann, 1995, 1978-1984.

[20] P. Peppas, W. Wobcke. On the Use of Epistemic Entrenchment in Reasoning about Action. In: *Proc. of the 10th European Conference on Artificial Intelligence*, 1994, 403-407.

[21] E. Sandewall. *Features and Fluents. A Systematic Approach to the Representation of Knowledge about Dynamical Systems.* Volume 1. Oxford University Press, 1994.

[22] E. Sandewall. Assessments of Ramification Methods that use Static Domain Constraints. In *Proc. of the 5th Int'l. Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann, 1996.

[23] M. Thielscher. Ramification and causality. *Artificial Intelligence J.*, 1997.

[24] A. del Val, Y. Shoham. Deriving Properties of Belief Update from Theories of Actions (II). In *Proc. of the 13th Int'l. Joint Conference on Artificial Intelligence*, pages 732–737, 1993.

[25] M. Winslett. Reasoning about Action Using a Possible Model Approach. In: *Proc. AAAI-88*, St. Paul, MN, 1988, 89-93.

[26] M. Winslett. *Updating Logical Databases.* Cambridge Tracks in Theoretical Computer Science. Cambridge University Press, 1990.

[27] M. Winslett. Updating Logical Databases. In: D. Gabbay, A. Galton, J. A. Robinson (eds.). *Handbook of Logic in Artificial Intelligence and Logic Programming.* Vol. 4 (Epistemic and Temporal Reasoning), 1995, Oxford University Press, 133-174.