

Reasoning with Knowledge, Action and Time in Dynamic and Uncertain Domains

Theodore Patkos and Dimitris Plexousakis

Institute of Computer Science

Foundation for Research and Technology Hellas

{patkos, dp}@ics.forth.gr

Abstract

We propose a new framework for reasoning about knowledge, action and time for domains that include actions with non-deterministic and context-dependent effects. The axiomatization is based on the Event Calculus and combines the expressiveness of possible worlds semantics with the efficiency of approaches that dispense the use of the accessibility relation. The framework is proved logically sound and, when restricted to deterministic domains, is also logically complete. To prove correctness of the approach, we construct a knowledge theory based on a branching version of the Event Calculus and study their correlation.

1 Introduction

Agents operating in complex dynamic worlds often need to achieve control over partially known and uncertain environments. Managing information acquired at execution time through sensing has proven to be an important cognitive skill. The topic of reasoning about action and knowledge has resulted in a wave of research triggered by Moore's adaptation of the possible worlds model in formal action theories [1985] and has since been extended in a multitude of formalisms; within the Situation Calculus Scherl and Levesque [2003] have adapted Reiter's solution to the frame problem which has also been extended to concurrent actions [Scherl, 2003], Thielscher [2000] has utilized the Fluent Calculus to also solve the inferential frame problem for knowledge and provide an elaborate notion of ability, while Lobo *et al.* [2001] introduced the action language \mathcal{A}_k that in addition considered non-deterministic action effects that cause loss of knowledge.

While these rigorous frameworks have provided very expressive formal accounts for knowledge and change, a serious impediment has been identified, owed to their dependence on the accessibility relation used to define knowledge in possible worlds (the semantics of epistemic states in \mathcal{A}_k are similarly defined); determining whether n atomic formulae are known potentially requires 2^n distinguishable worlds to check truth in, leading to computationally less promising implementations. Aiming at tractability, contemporary progress in the field explores alternative characterizations of knowledge that are disengaged from the accessibility relation. An

increasing number of recent approaches focus on restricting expressiveness, in order to perform efficient reasoning while preserving logical completeness with respect to the standard possible worlds specifications.

In this paper, we propose a new knowledge framework for enabling reasoning about a broad range of common-sense phenomena, beyond those considered by existing approaches that do not manipulate possible worlds. Specifically, we construct a unified theory of action, knowledge and time within the Event Calculus, to reason non-monotonically about conditional and indirect knowledge effects, ramifications of knowledge, as well as loss of knowledge caused by non-deterministic events. Moreover, due to the explicit representation of time, the framework can exploit the interaction of knowledge and time, in order to express temporal knowledge and delayed knowledge effects, cumulative and canceling knowledge effects of concurrent events or knowledge about continuously changing world aspects. The proposed approach treats knowledge in a way such that disjunctive or existentially quantified formulae are still derivable, approximating the possible worlds-based theories, but, when reasonably restricted, it can also perform tractable reasoning.

We illustrate logical soundness and completeness in deterministic domains with respect to the possible worlds specifications by constructing a branching Event Calculus knowledge theory and proving equivalence in the way knowledge changes in each case. We also suggest a way to extend the result and achieve correctness in the general case.

The paper proceeds as follows. We first review related approaches that do not utilize the accessibility relation. Then, we describe our proposed knowledge theory and also a theory based on a branching time representation that uses possible worlds. We study their correlation in Section 5. We conclude with a discussion on weaknesses and future objectives.

2 Related Work

To alleviate the computational intractability of reasoning with the possible worlds semantics (as well as other problematic issues, such as the logical omniscience side-effect), alternative approaches for reasoning about action and knowledge, disengaged from the accessibility relation, have been proposed. These approaches adopt restrictions about the type of knowledge formulae or domain classes that can be supported.

Maybe the first stimulating approach towards an alternative formal account for reasoning about knowledge and action is due to Demolombe and Pozos-Parra [2000] who introduced two different *knowledge fluents* to explicitly represent the knowledge that an ordinary fluent is true or false. Working on the Situation Calculus, they treated knowledge change as changing each of these fluents individually, the same way ordinary fluent change is performed in the calculus, thus reducing reasoning complexity by linearly increasing the number of fluents. Nevertheless, the expressive power of the representation was limited to knowledge of literals.

Petrick and Levesque [2002] have proved the correspondence of this approach to the possible worlds-based Situation Calculus axiomatization for successor state axioms of a restricted form. Specifically, when the conditions under which a fluent changes its truth value contain no fluent (context-free theories) or fluents in a restricted disjunctive normal form (literal-based theories). Moreover, they have defined a combined action theory that extended knowledge fluents to also account for first-order formulae when disjunctive knowledge is tautology-free, but enforcing it to be broken apart into knowledge of the individual parts. Recently, a decomposition property of even more expressive classes of action theories has been suggested, based on the notion of a Cartesian situation, that simplifies certain complex types of disjunctive formulae into equivalent components that only mention fluent literals [Petrick, 2008]. The price to pay is a requirement of definite knowledge of some of the disjunction components.

Regression used by standard Situation Calculus is considered impractical for large number of actions and introduces strong assumptions, such as closed-world and domain closure, which is problematic when reasoning with incomplete knowledge. Recent approaches deploy different forms of progression. Liu and Levesque [2005] for instance, study a class of incomplete knowledge that can be represented in so called *proper KBs* and perform progression on them. The approach is efficient and sound for actions with local-effects (i.e., when the properties of fluents changed are contained in the action), still proper KBs do not permit some general forms of disjunctions. The solution may even be complete when queries are in a certain normal form and the theory is context-complete (i.e., when there is complete knowledge about the context of context-dependent actions). The latter restriction is raised in [Vassos and Levesque, 2007], but with an extra cost of explicitly listing the possible values of each fluent, essentially rendering the theory propositional.

In all aforementioned approaches knowledge loss due to uncertain effects or unknown preconditions is not investigated. Beyond Situation Calculus, Son and Baral [2001] extend the action language \mathcal{A} capturing an agent's mental state with a simpler structure, as compared to traditional Kripke models and provide semantics for sound progression approximations of action transitions that have a more manageable state space and a lower complexity. Knowledge update is also suggested in [Amir and Russell, 2003] that uses transition of belief states instead of possible worlds. Algorithms for complete or approximate recursive state estimation (or logical filtering) are described, including non-deterministic domains. Both approaches are restricted to the propositional case.

3 DECKT Axiomatization

Our account of action and knowledge is formulated within the Linear Discrete Event Calculus (LDEC) [Mueller, 2006], a discrete version of the classical logic Event Calculus [Miller and Shanahan, 2002]. This formalism applies the *principle of inertia*, which captures the property that things tend to persist over time unless affected by some event; when released from inertia, a fluent may have a fluctuating truth value. It also uses *circumscription* [Lifschitz, 1994] to solve the frame problem and support default reasoning. A set of predicates is defined to express which fluents hold when (*HoldsAt*), what events happen (*Happens*), what their effects are (*Initiates*, *Terminates*, *Releases*) and whether a fluent is subject to inertia or released from it (*ReleasedAt*). If a fluent is initiated or terminated it becomes inertial at the next time instant.

The proposed Discrete Event Calculus Knowledge Theory (DECKT) assumes agents acting in dynamic, non-deterministic environments, having accurate but potentially incomplete knowledge and able to perform both knowledge-producing actions and actions that cause loss of knowledge, as well as actions with context-dependent effects. In a familiar trend, we treat knowledge as a fluent, namely the *Knows* fluent, to express knowledge about fluents and fluent formulae. The *Knows* fluent is always released from inertia, as a means to enable the application of rules in an elaboration tolerant manner, i.e., without having to write additional rules for each new axiom added to the theory. As a consequence, its truth value is only dependent on the state constraints triggered at each time instant that refer to it, preventing it from fluctuating. Still, as situations where the effects of actions are subject to inertia are very common, knowledge about them should persist until some event affects them. To "simulate" inertial knowledge, we introduce the *KP* fluent (for "knows persistently"). In brief, direct action effects cause inertial knowledge, while ramifications of knowledge, owed to state constraints, do not affect the *KP* fluent at all.

This intuition is captured in DECKT's core set of axioms presented in the next section. First, we introduce a minimal set of axioms and abbreviations needed to deploy the knowledge theory and clarify its properties. In the rest, event variables are represented by e , fluent variables by f , first-order fluent formulae by ϕ and variables of the timepoint sort by t , with subscripts where necessary.¹

(D) $HoldsAt(Knows(\phi), t) \Rightarrow HoldsAt(\phi, t)$

(K) $HoldsAt(Knows(\phi_1 \Rightarrow \phi_2), t) \Rightarrow$
 $(HoldsAt(Knows(\phi_1), t) \Rightarrow HoldsAt(Knows(\phi_2), t))$

(P1) $HoldsAt(Knows(\phi_1 \wedge \phi_2), t) \Leftrightarrow$
 $HoldsAt(Knows(\phi_1), t) \wedge HoldsAt(Knows(\phi_2), t)$

(P2) $HoldsAt(Knows(\phi_1), t) \vee$
 $HoldsAt(Knows(\phi_2), t) \Rightarrow HoldsAt(Knows(\phi_1 \vee \phi_2), t)$

(P3) $HoldsAt(Knows(\forall \vec{x}\phi), t) \Leftrightarrow$
 $\forall \vec{x} HoldsAt(Knows(\phi), t)$

(P4) $\exists \vec{x} HoldsAt(Knows(\phi), t) \Rightarrow$
 $HoldsAt(Knows(\exists \vec{x}\phi), t)$

This set is flexible enough not to adopt certain undesirable properties of epistemic logics, such as the necessitation rule

¹Free variables are implicitly universally quantified. The term *fluent literal* denotes either a fluent $f(\vec{x})$ or its negation $\neg f(\vec{x})$.

(NR), which dictates that all valid formulae are known, leading to implementations that avoid aspects of the logical omniscience problem. In general, in DECKT, instead of (D) we can apply the consistency axiom, but for the purposes of the present study we concentrate on a theory of knowledge, rather than belief. Furthermore, according to (P1)-(P4) no simplification concerning disjunctive or existentially quantified formulae is applied. Finally, we also introduce the abbreviation **(Kw)** $HoldsAt(Kw(\phi), t) \equiv HoldsAt(Knows(\phi), t) \vee HoldsAt(Knows(\neg\phi), t)$ and, similarly, $HoldsAt(KPw(\phi), t)$ for the KP fluent.

3.1 Core DECKT Axioms

DECKT consists of the following axiom sets:

Knowledge and the law of inertia. Knowledge is released from inertia at all times.

(KT1) $ReleasedAt(Knows(\phi), t)$

Knowledge persistence. This axiom captures the correlation between the $Knows$ and the KP fluent, introduced as a state constraint to the theory. KP is always subject to inertia.

(KT2) $HoldsAt(KP(\phi), t) \Rightarrow HoldsAt(Knows(\phi), t)$

Knowledge minimization. In DECKT knowledge is derived either from direct actions effects (see axiom sets (KT3-6) below) that affect the KP fluent or indirectly from state constraints of the form $HoldsAt(\phi_1, t) \Rightarrow HoldsAt(\phi_2, t)$ that affect the $Knows$ fluent according to (K). As knowledge is always released from inertia, we need a way to prevent the $Knows$ fluent from fluctuating whenever knowledge cannot be inferred, i.e., whenever no domain state constraint is triggered to produce either $HoldsAt(Knows(\phi_2), t)$ or $HoldsAt(Knows(\neg\phi_2), t)$. To obtain this result, we apply a form of default reasoning, assuming that by default at any timepoint knowledge about a fluent formula does not hold. Axiom (KT7) below performs exactly such a minimization to the extension of the $Knows$ fluent (in a style similar to performing circumscription to a formula for the purpose of minimizing the extension of a predicate).

Let $\phi_2(\vec{f}_i)$, $\phi_1(\vec{f}'_j)$ denote arbitrary formulae whose only free variables are fluents \vec{f}_i, \vec{f}'_j and which do not mention epistemic fluents. Axiom (KT7) is structured as follows:

(KT7) $HoldsAt(Kw(\phi_2(\vec{f}_i)), t) \Leftrightarrow \exists f'_1, \dots, f'_n (HoldsAt(Knows(\phi_1(\vec{f}'_j)), t) \wedge HoldsAt(Knows(\phi_1(\vec{f}'_j) \Rightarrow (\neg)\phi_2(\vec{f}_i)), t)) \vee HoldsAt(Knows(KPw(\phi_2(\vec{f}_i))), t)$

where f'_1, \dots, f'_n ($0 \leq n \leq j$) are those fluents in $\phi_1(\vec{f}'_j)$ that do not appear in $\phi_2(\vec{f}_i)$. The intuition is that an agent knows a formula iff there exists some state constraint known to be triggered at that particular time instant (therefore (KT2) is also accounted for).

By grounding (KT7) on the set of available state constraints of a particular domain axiomatization, its instantiation can significantly simplify the whole complex and complexity of the axiom. This set is well defined, even though it may be modified online according to occurring events and context. Suppose, for instance, that the only available to an agent state constraints are $HoldsAt(f_1, t) \Rightarrow HoldsAt(f, t)$ and

$HoldsAt(f_2, t) \Rightarrow HoldsAt(f, t)$, then (KT7) instance for f will be formulated as follows:

$HoldsAt(Kw(f), t) \Leftrightarrow HoldsAt(Knows(f_1 \vee f_2), t) \vee HoldsAt(KPw(f), t)$.

Events with known preconditions. If an agent knows all preconditions of a deterministic action, then it also knows its effect. $KP(f)$ and $KP(\neg f)$ cancel one another to preserve consistency. Specifically whenever $\bigwedge^i [HoldsAt(f_i, t)] \Rightarrow Initiates(e, f, t)$, i.e., the domain theory includes a *positive effect* axiom according to which event e initiates fluent f when the conjunction of fluent preconditions \vec{f}_i holds, then:

(KT3.1) $Happens(e, t) \wedge \bigwedge^i [HoldsAt(Knows(f_i), t)] \Rightarrow Initiates(e, KP(f), t)$

(KT3.2) $Happens(e, t) \wedge \bigwedge^i [HoldsAt(Knows(f_i), t)] \Rightarrow Terminates(e, KP(\neg f), t)$

Similarly, for *negative effect* axioms of the form $\bigwedge^j [HoldsAt(f_j, t)] \Rightarrow Terminates(e, f, t)$ we have

(KT3.3) $Happens(e, t) \wedge \bigwedge^j [HoldsAt(Knows(f_j), t)] \Rightarrow Initiates(e, KP(\neg f), t)$

(KT3.4) $Happens(e, t) \wedge \bigwedge^j [HoldsAt(Knows(f_j), t)] \Rightarrow Terminates(e, KP(f), t)$

Knowledge-producing (sense) events.

(KT4) $Initiates(sense(f), KPw(f), t)$

From this axiom, $Kw(f)$ is also implied, due to (KT2), (Kw).

Events with uncertain effects. If an action with *deterministic* effects occurs, which (a) has at least one precondition whose truth value the agent does not know (hence, the agent does not know whether the effect axiom is triggered), (b) there is no precondition that the agent knows it does not hold (otherwise, the agent would have been certain that the effect axiom would not be triggered) and (c) the agent does not already know the potential new truth value of the effect fluent, then the agent loses its knowledge about the state of the effect.

(KT5.1) $\bigvee^i [\neg HoldsAt(Kw(f_i), t)] \wedge \neg HoldsAt(Knows(\bigvee^i \neg f_i), t) \wedge \neg HoldsAt(Knows(f), t) \wedge Happens(e, t) \Rightarrow Terminates(e, KPw(f), t)$

whenever $\bigwedge^i [HoldsAt(f_i, t)] \Rightarrow Initiates(e, f, t)$.

(KT5.2) $\bigvee^j [\neg HoldsAt(Kw(f_j), t)] \wedge \neg HoldsAt(Knows(\bigvee^j \neg f_j), t) \wedge \neg HoldsAt(Knows(\neg f), t) \wedge Happens(e, t) \Rightarrow Terminates(e, KPw(f), t)$

whenever $\bigwedge^j [HoldsAt(f_j, t)] \Rightarrow Terminates(e, f, t)$.

Moreover, for actions with *non-deterministic* effects, i.e., whenever $\bigwedge^k [HoldsAt(f_k, t)] \Rightarrow Releases(e, f, t)$, then axiom (KT5.3) below expresses that if none of the preconditions is known not to hold (either the effect's preconditions are unambiguously satisfied or the agent does not know if they are satisfied; in either case knowledge about the effect is lost), then the agent cannot infer the state of the effect.

(KT5.3) $\neg HoldsAt(Knows(\bigvee^k \neg f_k), t) \wedge Happens(e, t) \Rightarrow Terminates(e, KPw(f), t)$

Hidden causal dependencies. The set of axioms (KT5) captures direct fluent effects due to unknown preconditions, causing knowledge to be lost. Still, knowledge about the ef-

fect becomes contingent on the preconditions; sensing the latter may provide information about the truth value of the former (assuming no event happens in the meantime affecting them). Such a dependency, which we call *hidden causal dependency* (HCD), is inherently modeled in possible worlds by appropriately decreasing the number of accessible worlds, even when the agent’s epistemic state remains unaffected. In DECKT we treat HCDs by means of a new fluent, $Implies(\phi, f)$, which encodes a notion of epistemic causality in the sense that if future knowledge brings about ϕ is also brings about f :

$$(Im) \text{ HoldsAt}(Implies(\phi, f), t) \Rightarrow \text{HoldsAt}(Knows(\phi \Rightarrow f), t)$$

HCDs are created each time an action occurs whose effect preconditions are unknown to the agent. For instance, for positive effect axioms we have that:

$$(KT6.1) \bigvee^i [\neg \text{HoldsAt}(Kw(f_i), t)] \wedge \neg \text{HoldsAt}(Knows(\bigvee^i \neg f_i), t) \Rightarrow \text{Initiates}(e, Implies(\bigwedge^j f_j, f), t)$$

where \vec{f}_j are those \vec{f}_i that the agent does not know. In addition, if the agent knows that before the action the effect fluent did not hold, another HCD is created between the effect and the precondition (sensing that f is true means that all preconditions must have been true, as well):

$$(KT6.2) \neg \text{HoldsAt}(Knows(\bigvee^i \neg f_i), t) \wedge \bigvee^i [\neg \text{HoldsAt}(Kw(f_i), t)] \wedge \text{HoldsAt}(Knows(\neg f), t) \Rightarrow \bigwedge^j [\text{Initiates}(e, Implies(f, f_j), t)]$$

where \vec{f}_j are those \vec{f}_i that the agent does not know.

On the other hand, a HCD expires whenever any of the involved fluents is potentially affected.² The axiomatization for negative and release effect axioms is defined in a similar way, but, due to lack of space, we do not explicitly represent all axioms here, instead refer to their conjunction as (KT6).

As an example with HCDs consider the effect axioms

$$\text{HoldsAt}(f_p, t) \Rightarrow \text{Initiates}(e, f, t) \\ \text{HoldsAt}(f'_p, t) \Rightarrow \text{Terminates}(e, f', t)$$

If initially no fluent is known, after $Happens(e, T)$ axiom (KT6.1) results in $\text{HoldsAt}(Implies(f_p, f), T + 1)$ and $\text{HoldsAt}(Implies(f'_p, f'), T + 1)$. Although still no fluent is known, if the agent later senses either f_p or f'_p it can also infer knowledge about f or f' , respectively.

Similarly, if initially the agent knows that $\text{HoldsAt}(Knows(f_p \vee f'_p), T)$, at $T + 1$ it can now infer that $\text{HoldsAt}(Knows(f \vee \neg f'), T + 1)$, due to (KT6.1), (Im) and (K). Nevertheless, still none of f, f' is known!

4 BDECKT Axiomatization

To prove correctness of the DECKT treatment of knowledge in terms of standard possible worlds semantics, we need to show its correspondence to a theory that utilizes the latter approach. Yet, no such theory exists for the Event Calculus and, moreover, the classical Event Calculus is based on

²Yet, new HCDs may be created, for instance if the event affects some of the involved fluents non-deterministically. We substantiate in detail the complete HCD axiomatization for even broader domain classes that also involve natural actions in a further work.

a linear time representation, where parallel worlds cannot be imitated. Therefore, in this section we construct a knowledge theory that is based on the Branching Event Calculus [Mueller, 2007] (BDEC) and study their equivalence.

BDEC is a modified version of the Linear Discrete Event Calculus; the two formalisms are proved logically equivalent on the grounds of a bridging set of axioms (see Section 5). BDEC replaces the timepoint sort with the sort of situations and introduces the relation $S(s_1, s_2)$ to express that situation s_2 is a successor of s_1 (in the rest, we also interchange sorts for DECKT as well, to maintain consistency with Mueller’s work). BDEC’s axiomatization shares the same set of axioms with LDEC, with the addition of a second order induction axiom similar to the one for the Situation Calculus. The difference to LDEC is that a situation is allowed to have zero or more successors, therefore all event-related predicates are modified to also specify the resulting situation. Thus, the following axiom is also introduced to capture the intuition:

$$(BDEC12) \text{ Happens}(e, s_1, s_2) \Rightarrow S(s_1, s_2)$$

Our Branching Event Calculus Knowledge Theory (BDECKT) follows Moore’s [1985] formalization of possible world semantics in action theories, where the number of K -accessible worlds remains unchanged upon ordinary event occurrences and reduces as appropriate when sense actions occur. Similar to Scherl and Levesque’s approach for the Situation Calculus [2003], BDECKT generalizes BDEC in that there is no single initial situation in the tree of alternative situations, but rather a forest of trees each with its own initial situation. Note, also, that for reasons that will become clear in the next section, we restrict non-determinism to only allow for events with unknown preconditions; although release effect axioms are still modeled, we implicitly assume that whenever a known fluent becomes released it remains known, subject to some state constraint.

To axiomatize knowledge in BDECKT we introduce the predicate $K(s', s)$ denoting that world (or situation) s' is accessible from s . The theory proceeds as follows:

$$(BKT1) \text{ HoldsAt}(Knows_B(f), s) \equiv \forall s' K(s', s) \Rightarrow \text{HoldsAt}(f, s')$$

$$(BKT2) K(s'_1, s_1) \Rightarrow$$

$$\exists s_2, s'_2 (Happens_B(e, s_1, s_2) \Leftrightarrow Happens_B(e, s'_1, s'_2))$$

$$(BKT3) \text{ Initiates}_B(e, f, s_1, s_2) \Leftrightarrow \text{Initiates}_B(e, f, s'_1, s'_2)$$

$$(BKT4) \text{ Terminates}_B(e, f, s_1, s_2) \Leftrightarrow \text{Terminates}_B(e, f, s'_1, s'_2)$$

(BKT1) is the standard definition for knowledge. Axiom (BKT2) states that an event happens in all K -related states, while axioms (BKT3) and (BKT4) require for an event to have the same effect regardless of the world it occurs in.

Ordinary events. When an ordinary event occurs in a situation, then all successor situations of the K -related to it situations (and only them) are K -related to its successor.

$$(BKT5) \text{ Happens}_B(e, s_1, s_2) \Rightarrow (K(s'_2, s_2) \Leftrightarrow \exists s'_1 (S(s'_1, s'_2) \wedge K(s'_1, s_1)))$$

Knowledge-producing (sense) events. Sensing a fluent ensures that it will be known in the successor situation, i.e., it will have the same truth value in all possible worlds.

$$(BKT6) \text{ Happens}_B(\text{sense}(f), s_1, s_2) \Rightarrow (K(s'_2, s_2) \Leftrightarrow \exists s'_1 (S(s'_1, s'_2) \wedge K(s'_1, s_1) \wedge (\text{HoldsAt}(f, s_1) \Leftrightarrow \text{HoldsAt}(f, s'_1))))$$

5 DECKT - BDECKT Equivalence Result

In [2007] Mueller has established a set of mapping rules L between the Linear and the Branching Event Calculus and showed that these two versions can be logically equivalent. This set restricts BDEC to a linear past (i.e., only one branch at a time can have an equivalent linear axiomatization), which introduces an important limitation to our attempt to relate DECKT and BDECKT; non-determinism due to release effect axioms cannot be supported, as it requires new worlds to emerge from a single one. Still, even this restriction can be lifted, as we discuss in the next section.

The objective now is to show that knowledge evolves the same way in both DECKT and BDECKT, i.e., the set of known formulae is the same after a sequence of actions. Therefore, in contrast to the linear and branching versions of the classical Event Calculus, whose equivalence is based on a one-to-one mapping of all their axioms, our intention here is to only focus on those axioms of the two knowledge theories that manipulate knowledge change. In particular, we focus on proving logical equivalence between axiom sets (KT3,4,5,6) and (BKT5,6) and define a set M that serves as a bridge between DECKT and BDECKT for that purpose.

The set L , defined by Mueller, comprises the following axioms, where $S_L(s_1)$ denotes the successor situation of s_1 in the LDEC:

- (L1) $S(s_1, s_2) \Leftrightarrow S_L(s_1) = s_2$
- (L2) $Happens_B(e, s_1, s_2) \Leftrightarrow Happens(e, s_1) \wedge S_L(s_1) = s_2$
- (L3) $Initiates_B(e, f, s_1, s_2) \Leftrightarrow Initiates(e, f, s_1)$
- (L4) $Terminates_B(e, f, s_1, s_2) \Leftrightarrow Terminates(e, f, s_1)$
- (L5) $Releases_B(e, f, s_1, s_2) \Leftrightarrow Releases(e, f, s_1)$

Let M be the following set of axioms:

- (M1) $HoldAt(Knows(f), s) \Leftrightarrow HoldAt(Knows_B(f), s)$
- (M2) $S_L(s_1) = s_2 \Rightarrow (K(s'_2, s_2) \Rightarrow \exists s'_1 (S(s'_1, s'_2) \wedge K(s'_1, s_1)))$
- (NR) If $\vdash HoldAt(\phi, t)$, then $\vdash HoldAt(Knows(\phi), t)$

Axiom (M2) relates event occurrences in DECKT with BDECKT's accessibility relation. Specifically, it disallows a world to be K -related to worlds other than those whose precedents were K -related to its own precedent (apart from the initial state, of course, which has no precedent). In other words, it prohibits unexpected world appearances, as well as it prevents worlds to be accessibly related to others that belong to future or past situations. The necessitation rule (NR) is produced as a side-effect in BDECKT, due to the definition of knowledge, therefore we need to explicitly include it in the set M in order to accomplish equivalence.

For our equivalence result we will use some lemmas. The proofs proceed in a similar style, therefore we only sketch the most representative one.

Lemma 1. $DECKT \wedge BDEC \wedge LDEC \wedge L \wedge M \vdash (BKT5)$

Proof sketch: Let e an arbitrary event, f, f_i arbitrary fluents and s_1, s_2 arbitrary situations such that $Happens_B(e, s_1, s_2)$, where $(\bigwedge^i [HoldAt(f_i, t)] \Rightarrow Initiates_B(e, f, s_1, s_2))$ (the procedure for negative effect axioms is similar). Let us first consider the case where the preconditions are known,




Figure 1: First step in proving Lemma 1.

i.e., $HoldAt(Knows(\bigwedge^i f_i), s_1)$. From (L2) and (L3), axiom (KT3.1) is triggered which, due to (KT2) and (M1), results in $HoldAt(Knows_B(f), s_2)$. According to (BKT1) f is true in all situations K -related to s_2 . These situations, as instructed by (M2), must have a precedent that is K -related to s_1 . Therefore, e happens and initiates f in all of them, due to (BKT2) and (BKT3). What has been proved so far is that $K(s'_2, s_2) \Rightarrow \exists s'_4 (S(s'_4, s'_2) \wedge K(s'_4, s_1))$, which is schematically depicted in the graph of Figure 1, where nodes represent situations, solid arcs accessibility relations and dashed arcs successor situations due to event occurrences. To also prove that there cannot be situations K -related to s_1 whose successors are not K -related to s_2 recall that, due to inertia, any fluent not affected by e and known (unknown) in s_1 must still be known (unknown, respectively) in s_2 . Consequently, any information, i.e., fluent truth values and accessibility relations, must be "transferred" to s_2 , which is proved by contradiction, since if some successor situation is not accessible from s_2 inertia may be violated. This leads us to $K(s'_2, s_2) \Leftrightarrow \exists s'_1 (S(s'_1, s'_2) \wedge K(s'_1, s_1))$, as required. The proof proceeds in a similar style if initially $\bigvee^i [\neg HoldAt(Kw(f_i), s_1)] \wedge \neg HoldAt(Knows(\bigvee^i \neg f_i), s_1)$ by replacing axiom (KT3.1) above with (KT5.1), while it is trivially shown if initially $HoldAt(Knows(\bigvee^i \neg f_i), s_1)$. \square

Lemma 2. $DECKT \wedge BDEC \wedge LDEC \wedge L \wedge M \vdash (BKT6)$
Using Lemmas 1 and 2 we can conclude:

Theorem 1. (Completeness) *The conjunction of DECKT, L and M axioms produces all BDECKT epistemic derivations.*

Lemma 3. $BDECKT \wedge BDEC \wedge LDEC \wedge L \wedge M \vdash (KT3.1) \wedge (KT3.2) \wedge (KT3.3) \wedge (KT3.4)$

Lemma 4. $BDECKT \wedge BDEC \wedge LDEC \wedge L \wedge M \vdash (KT4)$

Lemma 5. $BDECKT \wedge BDEC \wedge LDEC \wedge L \wedge M \vdash (KT5.1) \wedge (KT5.2)$

Lemma 6. $BDECKT \wedge BDEC \wedge LDEC \wedge L \wedge M \vdash (KT6)$
Using Lemmas 3 to 6 we can conclude:

Theorem 2. (Soundness) *All epistemic derivations produced by the conjunction of DECKT, L and M axioms are also produced by BDECKT.*

Corollary 1 *After any ground sequence of actions with deterministic effects but with potentially unknown preconditions, a fluent formula ϕ is known whether it holds in DECKT if and only if it is known whether it holds in BDECKT, under the bridging set of axioms L and M.*

6 Discussion

The proposed approach combines the expressiveness of possible-worlds specifications with the efficiency of reasoning without the accessibility relation. Whereas with possible worlds an agent needs to perform an exponential, with the number of fluents, number of reasoning tasks after each action (one for each possible world), DECKT requires a single reasoning task using the same set of inference rules, but to a larger number of fluents, ought to the different epistemic fluents. Specifically, although knowledge about conjunctions is decomposable according to (P1), disjunctions are not, resulting in $2^n - 1$ epistemic fluents with n domain fluents. Nevertheless, the efficiency of reasoning does not deteriorate analogously, as the set of state constraints remains unaffected and most inferences for epistemic fluents become trivial. For instance, when the only state constraint available is $HoldsAt(f, t) \Rightarrow HoldsAt(f_1, t)$, determining the truth value of $HoldsAt(Knows(f_1 \vee \bigvee^i f_i), t)$ is the same for any arbitrary i , because all the agent needs to consider is whether $HoldsAt(Knows(f_1), t)$ and consequently $HoldsAt(Knows(f), t)$ is true; axiom (KT7) for each fluent formula is grounded to the available state constraints. Moreover, if DECKT is restricted to domains that satisfy the property that knowledge of disjunctions can be broken apart to knowledge of the individual fluents, as most alternative approaches require, then reasoning can become tractable. Nevertheless, compared to them, DECKT is also applicable to non-deterministic domains, non-context-complete theories and can also express time-dependent knowledge effects.

One restriction of the current equivalence result is that logical completeness for non-deterministic domains cannot be proved using the current BDECKT semantics. Extending correspondence between DECKT and BDECKT for actions with non-deterministic effects cannot be based on producing multiple successor situations from an initial one, as this would ruin equivalence between the underlying linear and branching formalisms. We plan to resolve this issue by introducing the notion of *parallel universes* in the definition of knowledge for BDECKT, where in each sub-universe only deterministic actions may occur. The idea is that whenever a non-deterministic action is ought to happen, then in one sub-universe this action deterministically initiates a fluent and in another it deterministically terminates it. Knowledge is understood as universal in all sub-universes; a fluent is considered known if it is locally known in every sub-universes. Another restriction is the inability to support functional fluents (eg. see [Vassos and Levesque, 2007]). They can be implicitly represented though, by means of relational fluents, but with a less elegant manipulation.

Beyond the aforementioned research issues, the current framework offers also the means to go beyond the strict tenets of possible worlds specifications. One example is the modeling of agents that remember and forget, which has been used to accommodate knowledge-producing actions about both inertial and continuously changing world properties in a uniform manner. Representing introspection is also a challenging issue. But more important, we also intend to investigate ways to adopt solutions to aspects of the logical omniscience

problem, inherent in the possible worlds specification. An appropriate modification of (KT7) axiom can play the role of an *awareness set*, where only knowledge about specified fluents can be derived.

References

- [Amir and Russell, 2003] Eyal Amir and Stuart J. Russell. Logical filtering. In *IJCAI'03*, pages 75–82, 2003.
- [Demolombe and Pozos-Parra, 2000] R. Demolombe and MP Pozos-Parra. A Simple and Tractable Extension of Situation Calculus to Epistemic Logic. *Twelfth International Symposium on Methodologies for Intelligent Systems (ISMIS-00)*, pages 515–524, 2000.
- [Lifschitz, 1994] V. Lifschitz. Circumscription. *Handbook of Logic in Artificial Intelligence and Logic Programming*, 3:297–352, 1994.
- [Liu and Levesque, 2005] Yongmei Liu and Hector J. Levesque. Tractable Reasoning with Incomplete First-Order Knowledge in Dynamic Systems with Context-Dependent Actions. In *IJCAI'05*, pages 522–527, 2005.
- [Lobo et al., 2001] Jorge Lobo, Gisela Mendez, and Stuart R. Taylor. Knowledge and the Action Description Language A. *Theory and Practice of Logic Programming*, 1(2):129–184, 2001.
- [Miller and Shanahan, 2002] Rob Miller and Murray Shanahan. Some Alternative Formulations of the Event Calculus. In *Computational Logic: Logic Programming and Beyond, Essays in Honour of Robert A. Kowalski, Part II*, pages 452–490, London, UK, 2002. Springer-Verlag.
- [Moore, 1985] R. C. Moore. A Formal Theory of Knowledge and Action. In *Formal Theories of the Commonsense World*, pages 319–358. J. Hobbs, R. Moore (Eds.), 1985.
- [Mueller, 2006] Erik Mueller. *Commonsense Reasoning*. Morgan Kaufmann, 1st edition, 2006.
- [Mueller, 2007] Erik Mueller. Discrete Event Calculus with Branching Time. In *8th International Symposium on Logical Formalizations of Commonsense Reasoning*, pages 126–131, 2007.
- [Petrick and Levesque, 2002] Ronald P. A. Petrick and Hector J. Levesque. Knowledge Equivalence in Combined Action Theories. In *KR'02*, pages 303–314, 2002.
- [Petrick, 2008] Ronald P. A. Petrick. Cartesian Situations and Knowledge Decomposition in the Situation Calculus. In *KR'08*, pages 629–639, 2008.
- [Scherl and Levesque, 2003] Richard B. Scherl and Hector J. Levesque. Knowledge, Action, and the Frame Problem. *Artificial Intelligence*, 144(1-2):1–39, 2003.
- [Scherl, 2003] R.B. Scherl. Reasoning about the Interaction of Knowledge, Time and Concurrent Actions in the Situation Calculus. In *IJCAI'03*, volume 18, pages 1091–1098, 2003.
- [Son and Baral, 2001] Tran Cao Son and Chitta Baral. Formalizing Sensing Actions—A Transition Function based Approach. *Artificial Intelligence*, 125(1-2):19–91, 2001.
- [Thielscher, 2000] Michael Thielscher. Representing the Knowledge of a Robot. In *KR'00*, pages 109–120, 2000.
- [Vassos and Levesque, 2007] Stavros Vassos and Hector Levesque. Progression of Situation Calculus Action Theories with Incomplete Information. In *IJCAI'07*, pages 2024–2029, 2007.