TLP **10** (4–6): 675–690, 2010. © Cambridge University Press 2010
doi:10.1017/S1471068410000359

Logic programming for finding models in the logics of knowledge and its applications: A case study

C. BARAL and G. GELFOND
Department of Computer Science, Arizona State University, AZ, USA
E. PONTELLI, T. SON
Department of Computer Science, New Mexico State University, NM, USA

submitted 8 February 2010; revised 10 April 2010; accepted 17 May 2010

Abstract

The logics of knowledge are modal logics that have been shown to be effective in representing and reasoning about knowledge in multi-agent domains. Relatively few computational frameworks for dealing with computation of models and useful transformations in logics of knowledge (e.g., to support multi-agent planning with knowledge actions and degrees of visibility) have been proposed. This paper explores the use of logic programming (LP) to encode interesting forms of logics of knowledge and compute Kripke models. The LP modeling is expanded with useful operators on Kripke structures, to support multi-agent planning in the presence of both world-altering and knowledge actions. This results in the first ever implementation of a planner for this type of complex multi-agent domains.

KEYWORDS: planning, multi-agents, modal logics

1 Introduction

Modeling real-world planning scenarios that involve multiple agents has been an active area of research over the years. A considerable source of complexity derives from those scenarios where agents need to reason about knowledge and capabilities of other agents in order to accomplish their tasks. For example, a gambling agent needs to reason about what other players may know about the game in order to make the next move. Reasoning about knowledge and capabilities in multi-agent domains differs significantly from the same problem in single-agent domains. The complexity arises from two sources: *(1)* the representation of a planning domain needs to model not only the state of the world, but also the *knowledge/beliefs* of the agents; *(2)* the actions performed by an agent may lead to changes in knowledge and beliefs of different agents, e.g., action like announcements, cheating, lying, etc.

Various *modal logics* have been developed for reasoning about knowledge in multi-agent systems (see, e.g., Fagin *et al.* 1995; Halpern 1995; van Ditmarsch *et al.* 2007). The semantics of several of these logics is provided in terms of *Kripke structures*—where, intuitively, each Kripke structure captures the knowledge/beliefs

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

of all the agents. Naturally, when reasoning about knowledge of multiple agents in a *dynamic* environment, it is necessary to devise operations for updating Kripke structures after the occurrence of an action, as suggested in Baltag and Moss (2004), van Benthem *et al.* (2006), and van Ditmarsch *et al.* (2007). Two important questions, that have been less frequently considered, are: *(1)* how can one determine the initial model or the initial Kripke structure of a theory encoding the knowledge of the agents, and *(2)* what udpate operations on Kripke structures are necessary to lay the foundations of a high-level multi-agent action language.

The dominant presence of Kripke structures in different formalizations of multi-agent domains indicates that any system employing modal logic for reasoning and planning in multi-agent domains would need to implement some operations for the manipulation of Kripke structures. The lessons learned in the research in single-agent domains suggest that a high-level action language for multi-agent domains could be very useful. The development of such high-level language does not only help in modeling but also provides a new opportunity for the development of planning systems operating on top of this language. The complexity of various reasoning problems in modal logics (e.g., satisfaction, validity, etc.) and the difficulty in updating a Kripke structure after the occurrence of an event provide a computational challenge for logic programming. It is interesting to observe that, although there have been a few implementations of temporal logics in ASP (e.g., Heljanko and Niemelä 2003; Son *et al.* 2006), to the best of our knowledge, there has been no attempt in using logic programming for computing models of modal logics except in our recent work (Baral *et al.* 2010b). Our initial experience reveals that ASP-based implementation encounter severe difficulties, mostly associated to the grounding requirements of ASP. This also drives us to explore alternatives.

In this paper, we investigate the use of *Prolog* in the development of a computational framework for reasoning and planning in multi-agent domains. The advantage of Prolog lies in that it allows the step-by-step examination of parts of a Kripke structure without the need of constructing the complete structure. We focus on two initial aspects: computing the initial model of a theory of knowledge and manipulating Kripke structures. To this end, we will make use of a high-level action language for the specification of various types of actions in multi-agent domains.

2 Background: The logics of knowledge

In this paper we follow the notation established in Fagin *et al.* (1995). The modal language $\mathcal{L}\mathcal{K}_{\mathcal{A}}$ builds on a signature that contains a collection of propositions $\mathcal{F}$ (often referred to as *fluents*), the traditional propositional connectives, and a finite set of modal operators $\mathbf{K}_i$ for each i in a set $\mathcal{A}$. We will occasionally refer to the elements in the set $\mathcal{A}$ as *agents*, and the pair $\langle \mathcal{A}, \mathcal{F} \rangle$ as a *multi-agent domain.*

$\mathcal{L}\mathcal{K}_{\mathcal{A}}$ formulae are defined as follows. A *fluent formulae* is a propositional formula built using fluents and the standard Boolean operators $\vee, \wedge, \neg$, etc. A *modal formulae* is *(i)* a fluent formula, or *(ii)* a formula of the form $\mathbf{K}_i\psi$ where ψ is a modal formula, or *(iii)* a formula of the form $\psi \vee \phi$, $\psi \wedge \phi$, or $\neg\psi$, where ψ and ϕ are modal formulae. In addition, given a formula ψ and a non-empty set $\alpha \subseteq \mathcal{A}$:

- $\mathbf{E}_\alpha\psi$ denotes the set of formulae $\{\mathbf{K}_i\psi \mid i \in \alpha\}$.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

- $\mathbf{C}_\alpha\psi$ denotes the set of formulae of the form $\mathbf{E}^k_\alpha\psi$, where $k \geqslant 1$ and $\mathbf{E}^{k+1}_\alpha\psi = \mathbf{E}^k_\alpha\mathbf{E}_\alpha\psi$ (and $\mathbf{E}^1_\alpha\psi = \mathbf{E}_\alpha\psi$).

An $\mathscr{LK}_{\mathscr{A}}$ theory is a set of $\mathscr{LK}_{\mathscr{A}}$ formulae.

The semantics of $\mathscr{LK}_{\mathscr{A}}$ theories is given by *Kripke structures*. A *Kripke structure* is a tuple $(S, \pi, \{\mathscr{K}_i\}_{i\in\mathscr{A}})$, where S is a set of state symbols, π is a function that associates an interpretation of $\mathscr{F}$ to each state in S, and $\mathscr{K}_i \subseteq S \times S$ for $i \in \mathscr{A}$.

Different logic systems for reasoning with a $\mathscr{LK}_{\mathscr{A}}$ theory have been introduced; these differ in the additional axioms that the models are required to satisfy, e.g.,

- **P**: all instances of axioms of propositional logic
- **K**: $(\mathbf{K}_i\varphi \wedge \mathbf{K}_i(\varphi \Rightarrow \phi)) \Rightarrow \mathbf{K}_i\phi$
- **T**: $\mathbf{K}_i\varphi \Rightarrow \varphi$
- **4**: $\mathbf{K}_i\varphi \Rightarrow \mathbf{K}_i\mathbf{K}_i\varphi$
- **5**: $\neg\mathbf{K}_i\varphi \Rightarrow \mathbf{K}_i\neg\mathbf{K}_i\varphi$
- **D**: $\neg\mathbf{K}_i false$

For example, the following modal logic systems are frequently used (Halpern 1997): **S5** satisfies all of the above axioms with the exception of (**D**), **KD45** includes the four axioms **K**, **4**, **5**, and **D**, **S4** includes the axioms **K**, **T**, and **4**, and **T** includes the two axioms **K** and **T**.

Given a Kripke structure $M = (S, \pi, \{\mathscr{K}_i\}_{i\in\mathscr{A}})$ and a state $s \in S$, we refer to the pair (M, s) as a *pointed Kripke structure*—and s is referred to as the *real state*. The satisfaction relation between $\mathscr{LK}_{\mathscr{A}}$-formulae and a pointed Kripke structure (M, s) is defined as follows: (*i*) $(M, s) \models \varphi$ if φ is a fluent formula and $\pi(s) \models \varphi$; (*ii*) $(M, s) \models \mathbf{K}_i\varphi$ if $(M, s') \models \varphi$ for every s' such that $(s, s') \in \mathscr{K}_i$; and (*iii*) $(M, s) \models \neg\varphi$ iff $(M, s) \not\models \varphi$.

We will often view a Kripke structure M as a directed labeled graph, with S as its set of nodes, and whose set of arcs contains (s, i, t) iff $(s, t) \in \mathscr{K}_i$. (s, i, t) is referred to as an *arc*, from the state s to the state t. We identify special types of Kripke structures depending on the properties of $\mathscr{K}_i$:

- M is *r* if, for each $i \in \mathscr{A}$, $\mathscr{K}_i$ is reflexive relation;
- M is *rt* if, for each $i \in \mathscr{A}$, $\mathscr{K}_i$ is reflexive and transitive;
- M is *rst* if, for each $i \in \mathscr{A}$, $\mathscr{K}_i$ is reflexive, symmetric, and transitive;
- M is *elt* if, for each $i \in \mathscr{A}$, $\mathscr{K}_i$ is transitive, Euclidean (i.e., for all s_1, s_2, s_3, if $(s_1, s_2) \in \mathscr{K}_i$ and $(s_1, s_3) \in \mathscr{K}_i$ then $(s_2, s_3) \in \mathscr{K}_i$), and serial (i.e., for each s there exists s' such that $(s, s') \in \mathscr{K}_i$).

We use $M[S]$, $M[\pi]$, and $M[i]$, to denote the components S, π, and $\mathscr{K}_i$ of M.

3 A simple action language for multi-agent domains

In (Baral *et al.* 2010a), we proposed an action language $\mathbf{m}\mathscr{A}$ for multi-agent domains that considers various types of actions, such as world-altering actions, announcement actions, and sensing actions. In this language, a multi-agent theory is specified by two components: a $\mathscr{LK}_{\mathscr{A}}$ theory and an action description. The former describes the initial state of the world and knowledge of the agents. The latter describes the actions and their effects. We will next briefly describe the syntax of $\mathbf{m}\mathscr{A}$.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

- The *initial state* description is specified by axioms of the form:
  - *Modal logic system:* includes one statement of the form *system*(n), where $n \in \{\mathbf{T}, \mathbf{S4}, \mathbf{S5}, \mathbf{KD45}, \mathbf{none}\}$.
  - *Theory:* includes statements of the form *init*(φ), where φ is a $\mathscr{LK}_{\mathscr{A}}$ formula.
- The *action description* consists of axioms of the following forms:

$$a \textbf{ executable_if } \delta \tag{1}$$

$$a \textbf{ causes } \varphi \textbf{ if } \psi \textbf{ performed_by } \alpha \tag{2}$$

$$a \textbf{ announces } \phi \textbf{ performed_by } \alpha \textbf{ observed_by } \beta \tag{3}$$

$$a \textbf{ determines } f \textbf{ performed_by } \alpha \textbf{ observed_by } \beta \tag{4}$$

where a is an action, α, β are sets of agents from $\mathscr{A}$, δ is an arbitrary $\mathscr{LK}_{\mathscr{A}}$ formula, φ and ψ are conjunction of fluent literals (a fluent literal is either a fluent $f \in \mathscr{F}$ or its negation $\neg f$), ϕ is a restricted formula (a fluent formula, a formula $\mathbf{K}_i\varphi$, or a formula $\neg(\mathbf{K}_i\varphi \vee \mathbf{K}_i(\neg\varphi))$, where φ is a fluent formula), and f is a fluent.

An axiom of type (1) states the executability condition of the action a, (2) describes a world-altering action, (3) an announcement action, and (4) a sensing action. For the sake of simplicity, we assume that the conditions ψ in the axioms (2) for the same action a are mutually exclusive. An announcement will be referred to as a *public announcement* if $\alpha = \mathscr{A}$ and $\beta = \emptyset$, otherwise it will be referred to as a *private announcement*. When it is a private announcement, we restrict ϕ to be a fluent literal. By $(\mathscr{A}, \mathscr{F}, D, I)$ we denote a multi-agent theory over $\langle \mathscr{A}, \mathscr{F} \rangle$ with the initial state I and the action description D.

Observe that all axioms describing actions include a part indicating the agents participating in the action, i.e., who executes the action (**performed_by**) and who is aware of the action occurrence (**observed_by**). This is necessary, and dealing with this separation is one of the most difficult issues in reasoning about knowledge in multi-agent environments (Baltag and Moss 2004).

Given a multi-agent theory $(\mathscr{A}, \mathscr{F}, D, I)$, we are interested in queries of the form

$$\varphi \textbf{ after } [a_1; \ldots; a_n] \tag{5}$$

which asks whether φ, a $\mathscr{LK}_{\mathscr{A}}$ formula, holds after the execution of the action sequence $[a_1, \ldots, a_n]$ from the initial state.

Example 1

A, B, and C are in a room. On the table in the middle of the room is a box with a lock which contains a coin. No one knows whether the head or the tail is up and it is common among them that no one knows whether head or tail is up. Initially, A and C are looking at the box and B is not. One can peek at the coin to know whether its head or tail is up. To do so, one needs a key to the lock and looks at the box. Among the three, only A has the key for the lock. An agent can make another agent not to look at the box by distracting him/her.

```
init(looking_at_box(a)).  init(looking_at_box(c)). init(~looking_at_box(b)).
init(~k(a,tail)).         init(~k(b,tail)).        init(~k(c,tail)).
init(~k(a,~tail)).        init(~k(b,~tail)).       init(~k(c,~tail)).
init(has_key(a)).         init(~has_key(b)).       init(~has_key(c)).
```

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

```
peek(X,Y)     executable if looking_at_box(X), looking_at_box(Y), has_key(X).
distract(X,Y) executable if true.
peek(X,Y) determines tail performed_by  X observed_by Y.
distract(X,Y) causes ~looking_at_box(Y) if true performed by X.
```

4 Computing the initial state

Computing an initial state of a multi-agent theory $(\mathscr{A}, \mathscr{F}, D, I)$ means to compute a pointed Kripke structure (M, s) satisfying I where (M, s) satisfies I if

- (M, s) is a pointed Kripke structure w.r.t. the multi-agent domain $\langle \mathscr{A}, \mathscr{F} \rangle$;
- For each $init(\varphi) \in I$ we have that $(M, s) \models \varphi$;
- If $system(\mathbf{T}) \in I$ (resp. $system(\mathbf{S4})$, $system(\mathbf{S5})$, $system(\mathbf{KD45})$), then M is r (resp. rt, rst, elt).

4.1 Computing initial state using answer set programming

In Baral *et al.* (2010b), we proposed an answer set programming (ASP) (Marek and Truszczyński 1999; Niemelä 1999) implementation for computing the initial state of multi-agent theories. The implementation follows the basic idea of ASP by converting the initial state specification I to an ASP program $\Pi_I(m)$, with m being the number of states of the structure, whose answer sets are pointed Kripke structures satisfying I. The language of $\Pi_I(m)$ includes:

- A set of facts $fluent(f)$ ($agent(a)$) for each $f \in \mathscr{F}$ (resp. $a \in \mathscr{A}$);
- A set of constants $s_1, \ldots, s_n$, representing the names of the possible states;
- Atoms of the form $state(S)$, denoting the fact that S is a state in the Kripke structure being built;
- Atoms of the form $h(\varphi, S)$, indicating that the formula φ holds in the state S in the Kripke structure. These atoms represent the interpretation π associated with each state and the set of formulae entailed by the pointed Kripke structure;
- Atoms of the form $r(A, S_1, S_2)$, representing the accessibility relations of the Kripke structure, i.e., $(S_1, S_2) \in \mathscr{K}_A$ in the Kripke structure;
- Atoms of the form $t(S_1, S_2)$, used to represent the existence of a path from S_1 to S_2 in the Kripke structure present;
- Atoms of the form $real(S)$, to denote the real state of the world.

We assume formulae to be built from fluents, propositional connectives, and modal operators (in keeping with our previous definition). In particular, the fact that a formula of the form $\mathbf{K}_A\varphi$ holds in the current Kripke structure with respect to a state S is encoded by atoms of the form $k(A, S, \varphi)$.

The creation of the model requires selecting which states and which connections should be included in the Kripke structure and is expressed by choice rules:

$$
\begin{array}{lll}
1\{state(s_1), \ldots, state(s_n)\}m & \leftarrow & \\
0\{h(F, S)\}1 & \leftarrow & fluent(F), state(S) \\
0\{r(A, S_1, S_2) : state(S_1) : state(S_2)\}1 & \leftarrow & agent(A) \\
1\{real(S) : state(S)\}1 & \leftarrow &
\end{array}
$$

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

These rules will generate a pointed Kripke structure (M, s) whose states belong to $\{s_1, \ldots, s_n\}$ with s indicated by $real(s)$. The program is completed with rules determining the truth value of a formula in a given pointed structure. For instance, rules for checking whether an agent A knows a literal L are[1]:

$$\begin{array}{lcl} n_k(A,S,F) & \leftarrow & r(A,S,S_1), h(\neg F, S_1) \\ n_k(A,S,\neg F) & \leftarrow & r(A,S,S_1), h(F,S_1) \\ k(A,S,L) & \leftarrow & \texttt{not}\ \ n_k(A,S,L) \end{array}$$

In addition to the above rules, $\Pi_I(m)$ also contain rules encoding the additional properties imposed by a specific modal logic system:

$$\begin{array}{rcl} r(A,S,S) & \leftarrow & reflexivity, state(S), agent(A) \\ r(A,S_1,S_2) & \leftarrow & symmetry, r(A,S_2,S_1) \\ r(A,S_1,S_3) & \leftarrow & transitivity, r(A,S_1,S_2), r(A,S_2,S_3) \\ 1\{r(A,S,T) : state(T)\} & \leftarrow & serial, state(S) \\ r(A,S_2,S_3) & \leftarrow & euclidean, r(A,S_1,S_2), r(A,S_1,S_3) \end{array}$$

We can activate the necessary rules according to the modal logic system; e.g., if $system(\mathbf{S4}) \in I$, then we add the facts $reflexive$ and $transitive$ to $\Pi_I(m)$.

Finally, in order to ensure that a pointed structure satisfying I is generated, we add to $\Pi_I(m)$ the constraint: $\leftarrow init(\varphi), real(S), not\ h(\varphi, S)$.

The program $\Pi_I(m)$ can be used not only to generate models of I but also to retrospectively identify properties of I given a history of action occurrences. In Baral *et al.* (2010b), we used $\Pi_I(m)$ to solve the muddy children problem.

For example, one of the Kripke structures generated for the initial specification of Example 1 includes: *(1)* two states ($S = \{s_1, s_2\}$), *(2)* the interpretations $\pi(s_1)$ and $\pi(s_2)$ make `looking_at_box` true for a, c and false, `has_key` true for a and false for b, c, and they differ on the truth value of `tail` (e.g., `tail` is true in $\pi(s_1)$ and false in $\pi(s_2)$), and *(3)* $(s_i, s_j) \in \mathcal{K}_x$ for each $x \in \{a, b, c\}$ and $i, j \in \{1, 2\}$.

4.2 Computing Initial State Using Prolog

The elegance of ASP encoding has unfortunately to deal with the complexity of grounding imposed by modern ASP systems—and this will motivate our contribution of using Prolog for this task. Indeed, it is possible to find even simple theories of logics of knowledge that are beyond the capabilities of ASP.

Example 2 (*Sum and Product*)
An agent chooses two numbers $1 < x < y$ such that $x + y \leqslant 100$. The sum $x + y$ is communicated to agent s while the product $x * y$ is communicated to agent p. The following conversation takes place between the two agents:

- p states that it does not know the numbers x, y
- s indicates that it already knew this fact
- p states that now it knows the two numbers x, y
- s states that now it knows the two numbers x, y as well.

[1] Baral *et al.* (2010b) contain rules for checking whether (M, s) entails other types of formulae.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

Let us consider the problem of generating an initial Kripke structure for the above story. The two agents are operating on states that can be represented by a pair (x, y) satisfying the given conditions. It is reasonable to assume that initially, the agents will need to consider *all* possible states. This is encoded in ASP by the rule:

$$state(X, Y) \leftarrow 1 < X, X < Y, X + Y < 101$$

It is easy to see that the number of states is 2352. To generate the initial Kripke structure using ASP, we would have to use the rule

$$0\{r(A, state(X, Y), state(X_1, Y_1), 0) : state(X, Y) : state(X_1, Y_1)\}1 \leftarrow agent(A)$$

which will produce 2352^2 ground rules for each agent during the grounding. Computing a model for a program consisting of only this rule and the rule defining the states is already impossible[2]. A reasonable (but ad-hoc) way is to use the rules

$$r(s, state(X, Y), state(X_1, Y_1), 0) \leftarrow X + Y = X_1 + Y_1, state(X, Y), state(X_1, Y_1)$$
$$r(p, state(X, Y), state(X_1, Y_1), 0) \leftarrow X * Y = X_1 * Y_1, state(X, Y), state(X_1, Y_1)$$

which reduces the number of ground rules to less than 2352×110 (2352×99 for the sum, 2352×10 for the product) in total. Intuitively, these rules indicate that there is a link labeled s between $state(x, y)$ and $state(x', y')$ iff $x + y = x' + y'$, and there is a link labeled p between $state(x, y)$ and $state(x', y')$ iff $x \times y = x' \times y'$. Using these rules, Clingo/Smodels[3] is able to find a model (i.e., a possible Kripke structure) within a few seconds.

In order to verify that the generated Kripke structure satisfies the initial statements of the story, we will need to introduce fluents of the form $sum(S)$ and $product(S)$ to denote the sum and product of the two numbers of a state $S = (X, Y)$. We would also need to have fluents of the form $x(S)$ and $y(S)$ to represent that x and y are the two components of the state S.

The rules describing the truth values of the fluents are:

$$h(x(X), state(X, Y)) \leftarrow \qquad h(sum(S), state(X, Y)) \leftarrow S = X + Y$$
$$h(y(Y), state(X, Y)) \leftarrow \qquad h(product(P), state(X, Y)) \leftarrow P = X * Y$$

With these fluents, we can define a formula stating that an agent knows the value of the two numbers by $s_knows \equiv \bigvee_{state(x,y)} K_s(x(X) \wedge y(Y))$ and $p_knows \equiv \bigvee_{state(x,y)} K_p(x(X) \wedge y(Y))$.

The initial pointed Kripke structure is one that satisfies the formula $\neg p_knows$ and $k_s(\neg p_knows)$. This is what p and s state in their first announcement, respectively. Thus, to construct the initial pointed Kripke structure, we will need to add the rules defining the predicate $k(Agent, State, Formula)$ and other rules relating to this predicate. We will also need to add the constraints

$$\leftarrow real(X, Y), \texttt{not}\ k(p, state(X, Y), p_knows)$$
$$\leftarrow real(X, Y), \texttt{not}\ k(s, state(X, Y), \neg k(p, state(X, Y), p_knows))$$

The first constraint corresponds to p's first statement and the second to s's first statement. Adding these rules to the program, the ASP program does not find a

[2] We got the "Out of memory" message in Clingo. Lparse did not finish.
[3] `potassco.sourceforge.net`, `www.tcs.hut.fi/Software/smodels`

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

model within two hours. The main culprit is the number of ground rules that need to be generated before the answer set can be computed. For instance, the number of rules defining *n_k*, (as described in the previous section) is roughly 2352^2 (quadratic to the number of the states).

Many of these complications can be avoided by changing the model of computation from a bottom-up model (as used by ASP) to a top-down one (as used by Prolog). The Prolog encoding builds the similar clauses as described for ASP, to verify validity of a formula in a Kripke structure. The advantage is that, by operating top-down, the components of the structure are computed when requested (instead of being precomputed a priori during grounding).

We will now present a Prolog encoding for computing the initial pointed Kripke structure for a multi-agent theory. Each Kripke structure can be encoded using terms of the form `kripke(N,E)` where `N` is a list of nodes and `E` a list of edges. Each node is a term `node(Name, Interpretation)`, where `Name` is a name for the node and the interpretation is encoded as a list of terms `value(Fluent,true/false)`. A pointed Kripke structure is encoded as a term `sit(kripke(N,E), Node)` where `Node` is an element from the list `N`. The clauses to express the validity of a formula are very similar to the one presented earlier, e.g.,[4]

```
hold(F,sit(kripke(_Nodes,_Edges),N)) :-
  fluent(F),!, N = node(_Name,Int), member(value(F,true), Int).
hold(neg(F), Sit) :- \+hold(F,Sit), !.
hold(k(A,F), sit(kripke(N,E),node(X,_))) :-
  findall(M,member(edge(X,M,A),E),Nodes), iterated_hold(Nodes,F,kripke(N,E)).
iterated_hold([],_,_).
iterated_hold([A|B],F, kripke(N,E)) :-
  member(node(A,I),N), hold(F, sit(kripke(N,E),node(A,I))),
  iterated_hold(B,F,kripke(N,E)).
```

In the sum-and-product example, the initial Kripke structure can be implicitly defined using implicit rules to construct the edges of the graph:

```
edge(node(N1,Int1),node(N2,Int2),s) :-
   member(value(x(X),true),Int1), member(value(y(Y),true),Int1),
   member(value(x(X1),true),Int2), member(value(y(Y1),true), Int2), X+Y =:= X1+Y1.
edge(node(N1,Int1),node(N2,Int2),p) :-
   member(value(x(X),true),Int1), member(value(y(Y),true),Int1),
   member(value(x(X1),true),Int2), member(value(y(Y1),true), Int2), X*Y =:= X1*Y1.
```

The Prolog implementation of the announcements as constraints on the Kripke structure allows us to quickly converge on identifying $x = 4$ and $y = 13$ as the real state of the pointed Kripke structure in a matter of seconds (10.5 seconds on a MacBook Pro, 2.53GHz core duo).

5 The Semantics of m$\mathscr{A}$

We will now describe the semantics of **m**$\mathscr{A}$, based on the construction of a transition function. For the sake of simplicity, in this manuscript we will restrict our discussion

[4] The explicit representation of interpretations can be easily replaced with implicit encodings whenever specific domain knowledge is available.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

to domains where the initial state of the domain is described by one pointed Kripke structure and the actions are deterministic. Thus, the transition function is a map from a pointed Kripke structure and an action to another pointed Kripke structure[5].

5.1 Basic Kripke Structure Transformations

We start by introducing some basic transformations of Kripke structures necessary to model the evolution in models caused by the execution of actions. We also show how, following the representation of Kripke structures discussed in Section 2, it is possible to naturally encode these operators in Prolog.

Given a Kripke structure M, a set of states $U \subseteq M[S]$, and a set of arcs X in M, $M \overset{s}{\ominus} U$ is the Kripke structure M' defined by (*i*) $M'[S] = M[S] \setminus U$; (*ii*) $M'[\pi](s)(f) = M[\pi](s)(f)$ for every $s \in M'[S]$ and $f \in \mathscr{F}$; and (*iii*) $M'[i] = M[i] \setminus \{(t,v) \mid (t,v) \in M[i], \{t,v\} \cap U \neq \emptyset\}$ for every agent $i \in \mathscr{A}$. Intuitively, $M \overset{s}{\ominus} U$ is the Kripke structure obtained by removing from M all the states in U.

The $\overset{s}{\ominus}$ operator is encoded by the following Prolog rules:

```
node_minus(kripke(N,E), S, kripke(N1,E1)) :-
        delete(N,S,N1), remove_node_edges(S,E,E1).
remove_node_edges([],E,E).
remove_node_edges([A|B], E, NewEs) :-
        remove_one_node_edges(A,E, Es1), remove_node_edges(B, Es1, NewEs).
remove_one_node_edges(_,[],[]).
remove_one_node_edges(node(Node,Int), [edge(N1,N2,_)|Rest], Result) :-
        (N1=Node ; N2 = Node), !,
        remove_one_node_edges(node(Node,Int), Rest, Result).
remove_one_node_edges(N,[E|Rest],[E|Result]) :-
        remove_one_node_edges(N,Rest,Result).
```

$M \overset{a}{\ominus} X$ is a Kripke structure, M', defined by (*i*) $M'[S] = M[S]$; (*ii*) $M'[\pi] = M[\pi]$; and (*iii*) $M'[i] = M[i] \setminus \{(u,v) \mid (u,i,v) \in X\}$ for $i \in \mathscr{A}$. $M \overset{a}{\ominus} X$ is the Kripke structure obtained by removing from M all the arcs in X. The encoding of $\overset{a}{\ominus}$ is:

```
edge_minus(kripke(N,E), Es, kripke(N,E1)) :- delete(E,Es,E1).
```

Given two Kripke structures M_1 and M_2, we say that M_1 is *c-equivalent*[6] to M_2 ($M_1 \overset{c}{\sim} M_2$) if there exists a bijective function $c : M_1[S] \to M_2[S]$ such that: *(i)* for every $u \in S$ and $f \in \mathscr{F}$, $M_1[\pi](u)(f) = true$ iff $M_2[\pi](c(u))(f) = true$; and *(ii)* for every $i \in \mathscr{A}$ and $u, v \in M_1[S]$, $(u,v) \in M[i]$ iff $(c(u), c(v)) \in M_2[i]$.

M_1 and M_2 are *compatible* if for every $s \in M_1[S] \cap M_2[S]$ and every $f \in \mathscr{F}$, $M_1[\pi](s)(f) = M_2[\pi](s)(f)$.

For two compatible Kripke structures M_1 and M_2, we define $M_1 \overset{\kappa}{\cup} M_2$ to be the Kripke structure M', where (*i*) $M'[S] = M_1[S] \cup M_2[S]$; (*ii*) $M'[\pi] = M_1[\pi] \circ M_2[\pi]$;[7] and (*iii*) $M'[i] = M_1[i] \cup M_2[i]$. The encoding in Prolog of $M_1 \overset{\kappa}{\cup} M_2$ is

[5] It is easy to generalize this to maps where the possible configurations are described by sets of pointed Kripke structures and the actions are non-deterministic.

[6] This form of equivalence is similar to the notion of bisimulation in Baltag and Moss (2004).

[7] More precisely, $M'[\pi](s)(f) = M_1[\pi](s)(f)$ if $s \in M_1[S]$ and $M'[\pi](s)(f) = M_2[\pi](s)(f)$ if $s \in M_2[S] \setminus M_1[S]$.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

```
union_kripke(kripke(N1,E1), kripke(N2,E2), kripke(N3,E3)) :-
    append(N1,N2,N4), remove_dups(N4,N3), append(E1,E2,E4), remove_dups(E4,E3).
```

For a pair of Kripke structures M_1 and M_2 with $M_1[S] \cap M_2[S] = \emptyset$, $\alpha \subseteq \mathscr{A}$, and a one-to-one function $\lambda : M_2[S] \rightarrow M_1[S]$, we define $M_1 \uplus^{\lambda}_{\alpha} M_2$ to be the Kripke structure M', where (*i*) $M'[S] = M_1[S] \cup M_2[S]$; (*ii*) $M'[\pi] = M_1[\pi] \circ M_2[\pi]$; (*iii*) $M'[i] = M_1[i] \cup M_2[i]$ for each $i \in \alpha$, and $M'[i] = M_1[i] \cup M_2[i] \cup \{(u,v) \mid u \in M_2[S], v \in M_1[S], (\lambda(u),v) \in M_1[i]\}$ for each agent $i \in \mathscr{A} \setminus \alpha$.

Intuitively, the operators $\overset{\kappa}{\cup}$ and $\uplus^{\lambda}_{\alpha}$ allow us to combine different Kripke structures representing knowledge of different groups of agents, thereby creating a structure representing the knowledge of all the agents. The encoding of $\uplus^{\lambda}_{\alpha}$ in Prolog is:

```
knowledge_union(kripke(N1,E1), kripke(N2,E2), Alpha, Map, kripke(N3,E3)) :-
    check_k_union_properties(N1,N2,Map), union_list(N1,N2,N3),
    generate_k_union_edges(E1, N2, Alpha,Map, NewOnes),
    union_list(E1,E2,E4), union_list(E4,NewOnes,E3).
generate_k_union_edges([], _, _, _, []).
generate_k_union_edges([edge(_,_,Label)|Rest],Nodes,Alpha,Map,NRest) :-
    member(Label,Alpha), !, generate_k_union_edges(Rest,Nodes,Alpha,Map,NRest).
generate_k_union_edges([edge(Start,End,Label)|Rest], Nodes, Alpha,
                                                    Map, [NEdge|NRest]) :-
    member([N,Start], Map), NEdge = edge(N,End,Label),
    generate_k_union_edges(Rest,Nodes,Alpha,Map,NRest).
```

Several types of actions require the creation of "copies" of a Kripke structure, typically to encode the knowledge of the agents that are unaware of actions being executed. Let (M,s) be a pointed Kripke structure, and α be a set of agents:

- A pointed Kripke structure $(M',c(s))$ is a *replica* of (M,s) if $M' \overset{c}{\sim} M$ and $M'[S] \cap M[S] = \emptyset$;
- $(M,s)|_{\alpha} = (M \overset{a}{\ominus} X, s)$ where $X = \bigcup_{i \in \alpha}\{(u,i,v) \mid (u,v) \in M[i]\}$;

A replica of a pointed Kripke structure (M,s) refers to a copy of (M,s) with respect to a bijection c. The pointed Kripke structure $(M,s)|_{\alpha}$, referred to as *the restriction of* (M,s) *on* α, encodes the knowledge of the agents in the set $\mathscr{A} \setminus \alpha$. In Prolog:

```
replica(kripke(N,E), kripke(N1,E1), Map) :-
    replica_nodes(N,N1,Map), replica_edges(E,Map,E1).
replica_nodes([],[],[]).
replica_nodes([node(Name,Int)|Rest], [node(NewName,Int)|NewRest],
                    [[NewName,Name]|RestMap]) :-
    new_state_name(NewName), replica_nodes(Rest,NewRest,RestMap).
replica_edges([],_,[]).
replica_edges([edge(Start,End,Agent)|Rest], Map,
                    [edge(NewStart,NewEnd, Agent)|NewRest]) :-
    member([NewStart,Start], Map), member([NewEnd,End],Map),
    replica_edges(Rest,Map,NewRest).
remove_agent_edges(kripke(N,E), Alpha, K) :-
    collect_agent_edges(E,Alpha,Edges), edge_minus(kripke(N,E), Edges, K).
collect_agent_edges([],_,[]).
collect_agent_edges([edge(Start,End,Agent)|Rest], Alpha,
                            [edge(Start,End,Agent)|NewRest]) :-
    member(Agent,Alpha),!, collect_agent_edges(Rest,Alpha,NewRest).
```

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

```
collect_agent_edges([_|Rest], Alpha, NewRest) :-
    collect_agent_edges(Rest,Alpha,NewRest).
```

5.2 Action transition function

The following definitions are used to determine, given a pointed Kripke structure (M,s), what are the pointed Kripke structures obtained from the execution of an action. We will denote with $succ(a,(M,s))$ the pointed Kripke structure resulting from the execution of the action a in (M,s). Observe that in the following definitions we view a set of literals as the conjunction of its elements. We will define $succ(a,(M,s))$ for each type of action. Let a be an action. By $pre(a)$ we denote the formula δ, the condition in the law of the from (1) whose action is a. We begin with the public announcement action.

Public Announcement: Consider a pointed Kripke structure (M,s) and an action instance a occurring in an announcement law (3) such that $(M,s) \models \phi$. The successor pointed Kripke structure after the execution of a in (M,s) is defined as follows. If $(M,s) \not\models pre(a)$, then $succ(a,(M,s))$ is undefined, denoted by $succ(a,(M,s)) = \perp$. The first case deals with the assumption that the action should be executable in (M,s)—otherwise the resulting successor pointed Kripke structure is undefined.

If $(M,s) \models \phi$ then we have different cases depending on the structure of ϕ. If ϕ is a fluent formula then $succ(a,(M,s)) = (M \overset{s}{\ominus} U, s)$ for $U = \{s' \mid s' \in M[S], (M,s') \not\models \phi\}$. This can be captured by the following rules:

```
succ(sit(kripke(N,E),S), public(Phi), sit(kripke(N1,E1),S)) :-
    fluent_formula(Phi), !,
    get_nodes_not_satisfy_formula(Phi, kripke(N,E), Nodes),
    node_minus(kripke(N,E), Nodes, kripke(N1,E1)).
get_nodes_not_satisfy_formula(Phi, kripke(N,E), Nodes) :-
    iterate_not_satisfy(N,Phi,kripke(N,E), Nodes).
iterate_not_satisfy([],_,_,[]).
iterate_not_satisfy([A|B],F,K,Rest) :-
    hold(F,sit(K,A)), !, iterate_not_satisfy(B,F,K,Rest).
iterate_not_satisfy([A|B], F, K, [A|Rest]) :- iterate_not_satisfy(B,F,K,Rest).
```

If $\phi = \mathbf{K}_i\psi$, then $succ(a,(M,s)) = (M \overset{a}{\ominus} X, s)$ where $X = \{(u,i,v) \mid (u,v) \in M[i], (M,v) \not\models \psi\}$. This is captured by the following rules:

```
succ(sit(kripke(N,E),S), public(Phi), sit(kripke(N1,E1),S)) :-
    Phi = k(A,Psi), fluent_formula(Psi), !,
    get_edges_not_satisfy(E, A, Psi, kripke(N,E), Edges),
    edge_minus(kripke(N,E), Edges, kripke(N1,E1)).
get_edges_not_satisfy([],_,_,_, []).
get_edges_not_satisfy([Edge|Rest], Agent, Psi, kripke(N,E),[Edge|Result]) :-
    Edge=edge(Start,End,Agent), member(node(End,Int), N),
    \+ hold(Psi, sit(kripke(N,E), node(End,Int))), !,
    get_edges_not_satisfy(Rest,Agent,Psi,kripke(N,E), Result).
get_edges_not_satisfy([_|Rest],Agent,Psi,K,Result) :-
    get_edges_not_satisfy(Rest,Agent,Psi,K,Result).
```

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

If $\phi = \neg(\mathbf{K}_i\psi \vee \mathbf{K}_i\neg\psi)$, then $succ(a,(M,s)) = (M \overset{s}{\ominus} U, s)$ where $U = \{s' \mid s' \in M[S], (M,s') \models \mathbf{K}_i\psi \vee \mathbf{K}_i\neg\psi\}$. We omit the Prolog code due to lack of space.

If the announcement happens in a pointed Kripke structure (M,s) where $(M,s) \models \varphi$, then we need to ensure that φ is common knowledge (strictly speaking, this is the belief of the agents) among the agents of the system. This accounts for our removal of all nodes s' such that $(M,s') \not\models \varphi$. Similarly, if a formula $\mathbf{K}_i\psi$ is announced, we only remove arcs labeled by i from M whose existence invalidates the formula $\mathbf{C}_{\mathscr{A}}\mathbf{K}_i\psi$. On the other hand, when a formula $\neg(\mathbf{K}_i\psi \vee \mathbf{K}_i\neg\psi)$ is announced, we need to remove states in M whose existence invalidates the formula $\mathbf{C}_{\mathscr{A}}\neg(\mathbf{K}_i\psi \vee \mathbf{K}_i\neg\psi)$.

Continuing with Example 1, the execution of the action `peek(a,c)` in the initial Kripke structure described earlier transforms the Kripke structure as illustrated in Figure 1, where the double circle identifies the real state of the world, H denotes a state where `tail` is false and T denotes a state where `tail` is true.

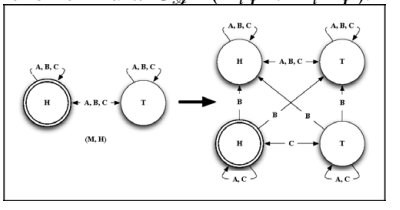


Fig. 1: Execution of `peek(a,c)`

Private Announcement: Consider a pointed Kripke structure (M,s) and a private announcement action a with $\beta \neq \mathscr{A}$, where the effect is the literal ℓ and $(M,s) \models \ell$. Let us also assume that f is the fluent used in the literal ℓ. The pointed Kripke structure (M',s') is the successor model after the execution of a in (M,s) if:

- $(M,s) \models pre(a)$ then $(M',s') = (M,s) \uplus^{c}_{\beta\cup\gamma} (M^r|_{\mathscr{A}\setminus(\beta\cup\gamma)} \overset{a}{\ominus} X, c(s))$ where $(M^r, c(s))$ is a replica of (M,s) and $X = \left\{(u,i,v) \middle| \begin{array}{l} i \in \beta, (u,v) \in M[i], \\ M[\pi](u)(f) \neq M[\pi](v)(f) \end{array}\right\}$;
- $(M,s) \not\models pre(a)$ and $(M',s') = \perp$.

The general successful case can be captured as:

```
succ(sit(kripke(N,E),S), private(Phi,Alpha,Beta), sit(kripke(N1,E1),St1)) :-
    get_literal_fluent(Phi,F), replica(kripke(N,E), kripke(N2,E2), Map),
    non_agents(Alpha,Beta,RemAgs),
    remove_agent_edges(kripke(N2,E2), RemAgs, kripke(N3,E3)),
    catch_discriminating_edges(E3,kripke(N3,E3),Alpha,F, X),
    edge_minus(kripke(N3,E3), X, kripke(N4,E4)),
    append(Alpha,Beta,AlphaBeta),
    knowledge_union(kripke(N,E), kripke(N4,E4), AlphaBeta, Map,kripke(N1,E1)),
    S = node(NameS,_), member([S1, NameS],Map), member(node(S1,Int1),N1),
    St1 = node(S1,Int1).
```

The predicate `non_agents` collects all agents not in `Alpha` and `Beta`. The predicate `catch_discriminating_edges` identifies edges for which the agents in `Alpha` see distinct truth values for the considered literal:

```
catch_discriminating_edges([],_,_,_,[]).
catch_discriminating_edges([Ed|Rest], kripke(N,E), Alpha, F, [Ed|NewRest]) :-
    Ed = edge(Start,End,Agent), member(Agent,Alpha),
    member(node(Start,Int1), N), member(node(End,Int2), N),
    member(value(F,V1), Int1), member(value(F,V2), Int2), V1 \= V2, !,
    catch_discriminating_edges(Rest, kripke(N,E), Alpha, F, NewRest).
catch_discriminating_edges([_|Rest], K, Alpha, F, NewRest) :-
    catch_discriminating_edges(Rest,K,Alpha,F,NewRest).
```

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

Sensing: Let (M,s) be a pointed Kripke structure and a be a sense action. A model (M',s') is a successor model after the execution of a in (M,s) if:

$$(M',s') = (M,s) \uplus^{c}_{\alpha\cup\beta} (M^r|_{\mathscr{A}\setminus(\alpha\cup\beta)} \overset{a}{\ominus} X, c(s))$$

for $X = \{(u,i,v) \mid i \in \alpha, (u,v) \in M^r[i], M^r[\pi](u)(f) \neq M^r[\pi](v)(f)\}$ and some replica $(M^r,c(s))$ of (M,s). The Prolog code is similar to the private announcement:

```
succ(sit(kripke(N,E),S), sense(F,Alpha,Beta), sit(kripke(N1,E1),St1)) :-
    replica(kripke(N,E), kripke(N2,E2), Map), non_agents(Alpha,Beta,RemAgs),
    remove_agent_edges(kripke(N2,E2), RemAgs, kripke(N3,E3)),
    catch_discriminating_edges(E3,kripke(N3,E3),Alpha,F, X),
    edge_minus(kripke(N3,E3), X, kripke(N4,E4)), append(Alpha,Beta,AlphaBeta),
    knowledge_union(kripke(N,E), kripke(N4,E4), AlphaBeta, Map,kripke(N1,E1)),
    S = node(NameS,_), member([S1, NameS],Map),
    member(node(S1,Int1),N1), St1 = node(S1,Int1).
```

World-altering Action: the computation of the successor pointed Kripke structure for world-altering actions requires the modification of the interpretations associated to certain states. Given an interpretation π and a set of literals φ, let us denote with $[\varphi]\pi$ the interpretation obtained by performing the minimal amount of modifications to π to ensure that $[\varphi]\pi$ makes all the literals in φ true. Let $u \in M[S]$ and let us consider an axiom of type (2); the axiom is applicable in u if $(M,u) \models \psi$. We define

$$res(a,u) = \begin{cases} [\varphi]M[\pi](u) & \text{if } a \textbf{ causes } \varphi \textbf{ if } \psi \textbf{ performed_by } \alpha \text{ applicable in } u \\ M[\pi](u) & \text{otherwise} \end{cases}$$

Let a be a world-altering action. By $Res(a,M,\alpha)$ we denote the Kripke structure M' which is obtained from M as follows:

- $M'[S] = \{r(a,u) \mid M[\pi](u) \models pre(a)\}$ where, for each state $u \in M[S]$, $r(a,u)$ denotes a new and distinguished state symbol;
- $M'[\pi](r(a,u)) = res(a,u)$;
- $(r(a,u),r(a,v)) \in M'[i]$ if $(u,v) \in M[i]$ and $r(a,u),r(a,v) \in M'[s]$ for $i \in \alpha$.

Let (M,s) be a pointed Kripke structure and let a be a world-altering action. The successor pointed Kripke structure of a in (M,s) is defined as follows.

1. If $(M,s) \not\models pre(a)$ then $succ(a,(M,s)) = \bot$.
2. If $(M,s) \not\models pre(a)$ then $succ(a,(M,s))$ is the pointed Kripke structure (M',s'):
   - $M'[S] = Q[S] \cup M[S]$;
   - $M'[\pi](u) = M[\pi](u)$ for $u \in M[S]$ and $M'[\pi](u) = Q[\pi](u)$ for $u \in Q[S]$;
   - $M'[i] = M[i] \cup Q[i] \cup R(i)$ where $R(i) = \emptyset$ for $i \in \alpha$ and $R(i) = \{(r(a,u),v) \mid r(a,u) \in Q[S], v \in M[S], (u,v) \in M[i]\}$; and
   - $s' = r(a,s)$.

   where $Q = Res(a,M,\alpha)$, and a **causes** φ **if** ψ **performed_by** α is the axiom of a applicable in s.

This can be mapped to the following Prolog encoding:

```
succ(sit(kripke(N,E),S), altering(If,Then,Alpha), sit(kripke(N1,E1),St1)) :-
    replica(kripke(N,E), kripke(N2,E2), Map), non_agents(Alpha,[],Gamma),
    remove_agent_edges(kripke(N2,E2), Gamma, kripke(N3,E3)),
    update_interpretations(N3,kripke(N3,E3),If, Then, N4),
```

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

```
    knowledge_union(kripke(N,E), kripke(N4,E3), Alpha, Map, kripke(N1,E1)),
    S = node(NameS,_), member([S1, NameS],Map),
    member(node(S1,Int1),N1), St1 = node(S1,Int1).
update_interpretations([],_,_,_,[]).
update_interpretations([node(Name,Inter)|Rest], K, If, Then,
                              [node(Name,NewInter)|NewRest]) :-
    (hold(If,sit(K,node(Name,Inter))) ->
        change_interpretation(Then,Inter,NewInter); NewInter = Inter),
    update_interpretations(Rest,K,If,Then,NewRest).
change_interpretation(F,I1,I2):- literal(F),!, single_update(F,I1,I2).
change_interpretation(and(F1,F2), I1, I2) :-
    change_interpretation(F1,I1,I3), change_interpretation(F2,I3,I2).
single_update(L, I1, I2):-  get_literal_fluent(L,F),delete(I1, value(F,_),I3),
    (fluent(L) -> I2 = [value(F,true)|I3]; I2 = [value(F,false)|I3]).
```

5.3 Query Answering and Planning in Prolog

The previous encoding can be used to support hypothetical reasoning and planning in multi-agent theories. To answer queries of the form (5) we can use the standard rules for computing the successor states of a sequence of action and verify the desirable properties in the final successor state. This is encoded as

```
hold(Query,Seq):- initial(sit(kripke(Nodes,Edges),Node)),
   succ(sit(kripke(Nodes,Edges),Node), Seq, S), hold(Query, S).
succ(sit(kripke(Nodes,Edges),Node), [], sit(kripke(Nodes,Edges),Node)).
succ(sit(kripke(Nodes,Edges),Node), [A|Seq], succ(S, Seq)):-
   succ(sit(kripke(Nodes,Edges),Node), A, S).
```

where $init(S)$ denotes that S is the initial pointed Kripke structure and $hold(F, S)$ is the predicate that determines whether F holds w.r.t. S.

For planning, we need to determine a sequence of actions that will accomplish a certain state of the world. The planning perspective supported by the implementation below is that of an external observer, who has complete knowledge about the capabilities of the different agents. A standard depth-first planner can be expressed using the following definition of the predicate `depthplan`$(S, Plan, N)$—where S is the initial pointed Kripke structure, N is a bound on the maximal length of the plan, and $Plan$ is the actual sequence of actions.

```
depthplan(S, Plan,N):-  depthplan(S,[],Plan,N).
depthplan(S, PlanIn, PlanIn,_):-  terminated(S).
depthplan(S0, PlanIn, PlanOut,N):- length(PlanIn,M), M<N, choose_action(A,S0),
    build_action(A,Fmt), succ(S0,Fmt,S1),  depthplan(S1,[A|PlanIn],PlanOut,N).
choose_action(A,S) :- action(A,_), valid(A,S).
```

The predicate `build_action` perform a simple syntactic rearrangement of the action representation (not reported). The predicate `valid` is used to ensure that the action is executable in the given pointed Kripke structure:

```
valid(A,S):- action(A,_Type), executability(A,Condition), hold(Condition,S).
```

The test for termination ensures that the goal has been achieved:

```
terminated(S) :- goal(Phi), hold(Phi,S).
```

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

As another example, it is easy to rewrite the planner to perform a breadth-first exploration of the search space.

```
breadthplan(S,Plan,N):- breadthplan(S,Plan,0,N).
breadthplan(S,Plan,M,N):- M<N, plan(S,Plan,S1), length(Plan,M), terminated(S1).
breadthplan(S,Plan, M, N):- M < N, M1 is M+1, breadthplan(S,Plan,M1,N).
plan(S1,[],S1).
plan(S0,[A|B],S):- plan(S0,B,S1), build_action(A,Format),
    valid(A,S1), succ(S1,Format,S).
```

Other forms of reasoning, e.g., switching from the perspective of an external observer to the perspective of an individual agent, can be similarly encoded.

6 Conclusion and discussion

In this paper, we investigated an application of logic programming technology to the problem of manipulating Kripke structures representing models of theories from the logic of knowledge. We illustrated how these foundations can be used to provide a computational background for a novel action language, $\mathbf{m}\mathcal{A}$, to encode multi-agent planning domains where agents can perform both world-altering actions as well as actions affecting agents' knowledge.

To the best of our knowledge, the encoding presented in this paper is the first implementation using Prolog of an action language with such a set of features. The use of logic programming allows a very natural encoding of the semantics of $\mathbf{m}\mathcal{A}$, and it facilitates the development of meta-interpreters implementing different forms of reasoning (e.g., observer-based planning). Let us underline that generic frameworks for reasoning with modal logics have been proposed (e.g., Horrocks 1998; Farinas del Cerro *et al.* 2001), and these could be used as external solvers to address the specific tasks of recomputing models—we will explore this option in our future work. Nevertheless, we believe the updates of Kripke structures directly performed in LP and the ability of embedding such process in a LP language, which offers other features (e.g., constraint solving) makes LP a more suitable avenue for implementing languages like $\mathbf{m}\mathcal{A}$. Although the prototype has not been subjected to formal testing (we are working on identifying interesting domains from the literature), simple planning tasks (e.g., develop a plan that enables agents a, c to learn about the status of the coin while b maintains its ignorance) can be solved within a few seconds of computation. A formal experimental evaluation will be part of the future work.

References

Baltag, A. and Moss, L. 2004. Logics for epistemic programs. *Synthese*, *139*, 2, 165–224.

Baral, C., Gelfond, G., Pontelli, E. and Son, T. C. 2010a. Planning with knowledge and world-altering actions in multiple agent domains. Unpublished manuscript.

Baral, C., Gelfond, G., Son, T. C. and Pontelli, E. 2010b. Using answer set programming to model multi-agent scenarios involving agents' knowledge about other's knowledge. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, IFAAMAS.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press

Fagin, R., Halpern, J., Moses, Y. and Vardi, M. 1995. *Reasoning about Knowledge*. MIT Press.

Fariñas del Cerro, L., Fauthoux, D., Gasquet, O., Herzig, A., Longin, D. and Massacci, F. 2001. Lotrec : The generic tableau prover for modal and description logics. *International Joint Conference on Automated Reasoning*, Springer Verlag, 453–458.

Halpern, J. Y. 1995. Reasoning about knowledge: A survey. In *Handbook of Logic in Artificial Intelligence and Logic Programming,* Vol. 4. Oxford University Press, 1–34.

Halpern, J. Y. 1997. A theory of knowledge and ignorance for many agents. *Journal of Logic and Computation 7,* 1, 79–108.

Heljanko, K. and Niemelä, I. 2003. Bounded LTL model checking with stable models. *Theory and Practice of Logic Programming 3,* 4–5, 519–550.

Horrocks, I. 1998. The FaCT system. *International Conference on Automated Reasoning with Analytical Tableaux and Related Methods*, Springer, 307–312.

Marek, V. and Truszczyński, M. 1999. Stable models and an alternative logic programming paradigm. In *The Logic Programming Paradigm: A 25-year Perspective*, Springer, 375–398.

Niemelä, I. 1999. Logic programming with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence 25,* 3–4, 241–273.

Son, T. C., Baral, C., Tran, N. and McIlraith, S. 2006. Domain-dependent knowledge in answer set planning. *ACM Transactions on Computational Logic 7,* 4, 613–657.

van Benthem, J., van Eijck, J. and Kooi, B. P. 2006. Logics of communication and change. *Information and Computation 204,* 11, 1620–1662.

van Ditmarsch, H., van der Hoek, W. and Kooi, B. 2007. *Dynamic Epistemic Logic*, Springer Synthese Library.

https://doi.org/10.1017/S1471068410000359 Published online by Cambridge University Press