

An Epistemic Logic for Labelled Transition Systems

Sophia Knight
Ecole Polytechnique
and INRIA Saclay

Radu Mardare
University of Aalborg

Prakash Panangaden
McGill University

November 26, 2011

Abstract

We define an epistemic logic for labelled transition systems by introducing equivalence relations for the agents on the states of the labelled transition system. The idea is that agents observe the dynamics of the system modulo their ability to distinguish states and in the process learn about the current state and past history of the execution. This is in the spirit of dynamic epistemic logic but is a direct combination of Hennessy-Milner logic and epistemic logic. We give an axiomatization for the logic and prove a completeness theorem with respect to the class of models obtained by unfolding labelled transition systems.

1 Introduction

Concurrency theory has been built upon the implicit assumption of omniscience of all the agents involved, but for many purposes – notably security applications – it is crucial to incorporate and reason about what agents “know” or do not know. Tracking the flow of information is the essence of analyses of security protocols. Equally crucial is the idea that different participants may have different views of the system and hence know different things. The purpose of this paper is to meld traditional concurrency concepts with epistemic concepts and define a logic with both dynamic and epistemic modalities.

Epistemic logic has been a major theme within distributed systems ever since the groundbreaking paper of Halpern and Moses [HM84], but has been strangely slow to influence concurrency theory. A few investigations have appeared but, as far as we know, there has not been a thorough integration of epistemic concepts with the traditional theory of labelled transition systems. Typically one sees a multimodal logic closely tied to the syntax of some particular process calculus with reasoning principles that are not proven complete in any sense [CDK09]. Such logics are interesting and useful but their close tie to a particular process formalism obscures the general principles. Another closely related strand

is, of course, dynamic epistemic logic [vDvdHK08] which, as the name makes manifest, is all about how knowledge evolves. However, the bulk of this work is about actions that communicate information, perhaps through messages or announcements, rather than about general transitions that could change basic facts about the state. A few papers indeed deal with so-called fact-changing actions but, as far as we know, the theory is still geared toward communication actions. Our goals are to develop the theory for a suitable general class of labelled transition systems and to formulate axioms that are provably complete with respect to this class of models. We provide more detailed comparisons with related work in a later section after the presentation of our framework.

The standard route to modelling epistemic concepts is to use Kripke models: these are sets of states equipped with indistinguishability (equivalence) relations [FHMV95]. We will equip the states with a labelled transition system structure as well and impose coherence conditions between the two kinds of relations. The resulting modal logic is a blend of Hennessy-Milner logic, epistemic logic and temporal modalities. The essential point is that one can reason about how knowledge changes as transitions occur. There are many variations that one could contemplate and the particular formalism that we have developed is geared toward representing the unfolding of a labelled transition system through time taking into account different agents' differing views of the labelled transition system.

The paper is organized as follows. In the next section we review background material on labelled transition systems and Hennessy-Milner logic. In Section 3 we define the class of transition systems that we work with; they are called *history labelled transition systems* and are unfoldings of the usual labelled transition systems. In Section 4 we define the logic and its semantics. In Section 5 we prove the weak completeness theorem. There is an easy argument, which we present in Section 5, that shows that a strong completeness theorem is not possible. The final sections discuss related work and conclusions.

2 Background

We assume familiarity with basic concepts like labelled transition systems (LTSs) and epistemic logic. For the benefit of readers who may not be familiar with these ideas we give a brief overview in this section. An excellent exposition of general modal logics is the text book by Blackburn et al. [BdRV01] called *Modal Logic*.

Definition 2.1. A *labelled transition system* is a triple $(S, \mathcal{A}, \rightarrow \subseteq S \times \mathcal{A} \times S)$, where S is a set of *states*, \mathcal{A} is a set of *actions* and \rightarrow is a *labelled transition relation*. We will write $s \xrightarrow{a} s'$ when $(s, a, s') \in \rightarrow$.

The idea is that S represents the possible states of a dynamical system. The system can perform certain actions and these cause a change in the state. The resulting state is not completely determined by the initial state and the action so that one has a transition relation rather than a function. Some actions may

not be possible in some states, if an action a is possible from state s we say that a is *enabled* in s .

There are various senses in which states may be deemed to be equivalent. A canonical one is called *bisimulation*. The idea of bisimulation is that if the actions possible from two states *and their successors* do not distinguish them, they should be deemed equivalent. Here is a formal definition.

Definition 2.2. We say that an equivalence relation R on the state space S of an LTS is a *bisimulation relation* if whenever sRt and $\xrightarrow{s} as'$ then there exists some t' such that $\xrightarrow{t} at'$ with $s'Rt'$. We say that s and t are *bisimilar* if there is some bisimulation relation relating them.

Since R is required to be an equivalence relation it follows that the analogous condition holds with the roles of s and t exchanged. The properties of bisimulation are discussed at length in the concurrency theory literature, see, for example [Mil89, Mil90] or in the modal logic literature, see, for example [Pop94].

There is a remarkable theorem due independently to van Benthem and to Hennessy and Milner that gives a modal characterization of bisimulation. The logic has come to be called Hennessy-Milner logic. The basic constructs are the boolean connectives and a modal operator written $\langle a \rangle$ or its dual $[a]$, where the a 's appearing in the formulas are actions associated with the LTS being studied. The definition of satisfaction for these formulas follows the standard inductive construction due to Tarski with only the modal operator requiring explicit explanation. This is given by $s \models \langle a \rangle \phi$ iff $\xrightarrow{s} as'$ and $s' \models \phi$. The fundamental theorem is the following.

Theorem 2.3. Assume that $(S, \mathcal{A}, \rightarrow)$ is a labelled transition system with the property that for a given s and a the set of s' such that $\xrightarrow{s} as'$ is finite¹. Then two states s and t are bisimilar iff they satisfy all the same formulas of Hennessy-Milner logic.

The basic setup for modelling epistemic logic is due to Kripke [Kri63]; see *Reasoning About Knowledge* by Fagin et al. [FHMV95]. Consider the set S of possible states² of some system. We have a finite set of *agents*, typically written $\mathcal{I} = \{i, j, \dots\}$. We define a modal operator – one for each agent – written K_i . The idea is that the formula $K_i \phi$ means that the agent i *knows* the fact ϕ . The axioms usually used are due to Hintikka [Hin62]:

- | | | |
|----|--|---------------------------|
| 0. | All propositional tautologies. | |
| 1. | $K_i \phi \Rightarrow \phi$ | Only truths can be known. |
| 2. | $K_i(\phi \Rightarrow \psi) \Rightarrow (K_i \phi \Rightarrow K_i \psi)$ | Deductive closure. |
| 3. | $K_i \phi \Rightarrow K_i K_i \phi$ | Positive introspection. |
| 4. | $\neg K_i \phi \Rightarrow K_i \neg K_i \phi$ | Negative introspection. |

¹Such systems are said to be *image finite*.

²They are often called *possible worlds* in the philosophical literature.

These are used together with the following rule of inference.

$$\frac{\phi}{K_i\phi}$$

The semantics for this logic is given in terms of *indistinguishability relations*. The idea is that a particular agent has only limited awareness of everything that might be true in a state. In particular, an agent might not be able to distinguish two states. We associate with each agent an equivalence relation that models its ability to distinguish two states.

Definition 2.4. A *Kripke structure* is a set S of states, a finite set \mathcal{I} of agents, a set \mathcal{P} of primitive propositions, for each state s a set $\pi(s) \subset \mathcal{P}$ and for each $i \in \mathcal{I}$ an equivalence relation \sim_i on S .

The meaning of an atomic proposition is built into the definition of the Kripke structure: $s \models p$ iff $p \in \pi(s)$; the meaning of the boolean connectives is standard. We define the meaning of the modal formula as follows: $s \models K_i\phi$ iff for every state s' such that $s \sim_i s'$, $s' \models \phi$. The fundamental completeness theorem is that a formula is provable from the Hintikka axioms iff it is true in all Kripke structures.


3 Histories

The main contribution of this paper is to study how an agent’s knowledge changes as transitions occur in a labelled transition system. The basic picture is that the agent has a limited view of the states of the labelled transition system and this is modelled by an equivalence relation on the states of the system just as in a Kripke structure. The agent does not choose the actions to perform but can see which action has happened and tries to deduce from this where it is. Our temporal-epistemic logic will be designed to handle this type of reasoning.

The semantics of the formulas will be given in terms of histories or runs, as with the semantics of Halpern and Moses [HM84, HM90], but we view the runs as coming from the executions of a labelled transition system (LTS). In fact, we will view the set of runs as forming a labelled transition system in its own right. This will give a “branching-time” logic rather than a linear-time logic. We will use the box and diamond modalities of Hennessy-Milner logic [HM85] rather than the “always” and “eventually” modalities of temporal logic. In this section, we motivate the need for this particular combination of modalities.

The basic set up for a purely epistemic (static) logic is a set of states with equivalence relations, one for each agent. If we wish to incorporate this into a given labelled transition system the natural step is to define equivalence relations on the states of the labelled transition system. If one does this naïvely one gets situations where one cannot say what an agent has learned from its history.

Example 3.1. Consider the following simple labelled transition system:



where the wiggly line refers to the indistinguishability equivalence relation of agent i and the proposition \mathbf{p} holds in the state s_3 and in no other state. The agent i in state s_0 cannot tell whether he is in s_0 or on s_1 . Similarly, in s_2 he cannot tell whether he is in s_2 or in s_3 . However, if the agent is in s_0 and then observes an a action then he “knows” he must have been in s_0 and further than he is in s_2 now. No purely state-based semantics can say this. It is only because the agent “remembers” how he got there that one can say anything. Thus, a purely state based semantics is not adequate for even the simplest statements about evolving knowledge. ♣

The basic paradigm that we have in mind is that the agent is observing a transition system: the agent can see the actions and can remember the actions but cannot *control* the actions nor see *which actions are available at a given state*. The extent to which an agent can “see” the state is what the indistinguishability relation spells out.

In order to give the semantics of the epistemic modalities we need to extend the equivalence relation from states to histories. We formalize, labelled transition systems, histories and this equivalence relation as follows.

Definition 3.2. A **labelled transition system** is a set of states, S , a finite set of actions \mathcal{A} , and, for every $a \in \mathcal{A}$, a binary relation, written \xrightarrow{a} , on the states. We write $s \xrightarrow{a} s'$ instead of $(s, s') \in \xrightarrow{a}$. In addition, there is a finite set of **agents**, denoted by letters like i, j, \dots . For each agent i there is an equivalence relation, written \sim_i defined on S .

The relation \xrightarrow{a} can be nondeterministic and does not have to be image-finite³. We also assume that all actions are visible, there are no hidden actions commonly denoted by τ .

Definition 3.3. A **history** is a finite alternating sequence of states and actions

$$s_0 a_1 s_1 a_2 s_2 \dots a_n s_n,$$

where, for each $l \in \{0, \dots, n-1\}$, $s_l \xrightarrow{a_{l+1}} s_{l+1}$.

Given a pair of histories, an agent can tell immediately that they are not the same if they do not have exactly the same action. In order to say this it will be convenient to define the notation $\text{act}(h)$ to mean the action sequence extracted from the history h , it has an evident inductive definition. Given a history h , we write $h[n]$ for the n^{th} state in h . Thus if $h = s_0 a_1 s_1 a_2 s_2 a_3 s_3$, $\text{act}(h) = a_1 a_2 a_3$

³“Image finite” means that for a given s and a the set $\{s' | s \xrightarrow{a} s'\}$ is finite.

and $h[0] = s_0$ while $h[2] = s_2$. We write $|h|$ for the length of the sequence of states in h .


Definition 3.4. We say that the histories h_1 and h_2 are **indistinguishable by agent i** , written $h_1 \sim_i h_2$, if: (i) $\text{act}(h_1) = \text{act}(h_2)$ and (ii) for all $0 \leq n \leq |h_1| (= |h_2|)$, $h_1[n] \sim_i h_2[n]$.

The use of the same notation for indistinguishability of states and histories should not occasion anxiety for the reader as the context will disambiguate which we mean; this usage is meant to emphasize the tight connection between the concepts.

It is useful to have both past and future modalities. We will define the syntax precisely in the next section, for the moment we note that $\langle - \rangle$ means one step in the past and $\langle + \rangle_a$ means *possibly* after an a -step into the future. Consider the labelled transition system we have used for our example above. Suppose we introduce the proposition $\text{@}s$ to mean “at the state s ” then we want to be able to say things like $s_0 a s_2 \models K_i \langle - \rangle \text{@}s_0$. Note that we cannot say $s_0 \models K_i \text{@}s_0$, so we need the past operator to express the idea that agent i learns where he was, or, in general, learns that a fact used to be true. Note that, for this example, $s_0 a s_2 \models \langle - \rangle K_i \text{@}s_0$ does not hold.

Note that every history has a beginning, every state has a finite number of predecessors: in short the prefix order on histories is well founded. This will cause most of the difficulties in the completeness proof.

Example 3.5. Why do we need the Hennessy-Milner like modalities indexed by actions? Consider the following simple labelled transition system:



which is like the previous example except for the addition of the extra state and transition and the fact that \mathbf{p} is true in s_2 instead of s_3 . We would like to be able to say $s_0 \models \langle + \rangle_a K_i \mathbf{p}$. Note that s_4 can be distinguished by i from any other state.

♣


The logic, though its semantics is given in terms of runs, is actually a branching time logic. It is applied to a very specific type of transition system that arises as the set of histories of general labelled transition system. The “states” are histories and the transitions are of the form

$$s_0 a_1 s_1 \dots a_n s_n \xrightarrow{a} s_0 a_1 s_1 \dots a_n s_n a s$$

whenever $s_n \xrightarrow{a} s$ is a transition of the underlying labelled transition system. The key features of these labelled transition systems of histories are: a well


foundedness property for the backward transitions, determinacy for the backward transitions and a few other properties⁴ In the course of the completeness proof we will spell out these properties and then proceed with the axiomatization and completeness theorem.

Example 3.6. Here is an example about why the identity of actions is important.




If this system starts out in s_0 and an a action occurs, then agent i will not know which state he is in, because s_1 and s_2 are equivalent for the agent. But if the system does a b action, then the agent knows he is in s_1 because he observes the b action and s_1 is the only state that a b action leads to. Similarly, if the system does a c action, then the agent knows that the system is in s_2 . ♣

Example 3.7. This example shows why we want to be able to combine epistemic modalities and (past or future) temporal modalities. Here p represents some proposition.



If the system starts out in s_0 or s_1 , then after an a action, the agent does not know whether p is true, but he does know that if p is true now, then it must have been true in the first state, and if p is false now, it must have been false in the first state. ♣

Example 3.8.



If this system starts out in s_0 or s_1 and then an a action occurs, then after the action, the agent does not know whether p is true, but he knows that if p is true now, then it was true in the start state. But he also knows that if p is not true now, then p may or may not have been true in the start state. ♣

⁴In fact, such transition systems arise naturally as unfoldings of general labelled transition systems.

3.1 History Systems

First we will explain how to translate any LTS with equivalence classes into an equivalent history LTS: an LTS with designated starting states, where the entire history of any run starting from a starting state is determined by its current state.

Definition 3.9. Given the LTS $(S_0, \mathcal{A}, \mathcal{I}, \xrightarrow{0}, \sim^0)$, where S_0 is the set of states, \mathcal{A} is the set of actions, \mathcal{I} the set of agents, $\xrightarrow{0} \subseteq S_0 \times \mathcal{A} \times S_0$ is the transition relation and $\sim^0 \subseteq S_0 \times \mathcal{I} \times S_0$ is the indistinguishability relation, inductively construct the *unfolding* $(S_1, \mathcal{A}, \mathcal{I}, \xrightarrow{+}, \xrightarrow{-}, \xrightarrow{+}^*, \xrightarrow{-}^*, \sim^1)$, where $\xrightarrow{+} \subseteq S_1 \times \mathcal{A} \times S_1$, $\xrightarrow{-} \subseteq S_1 \times \mathcal{A} \times S_1$, $\xrightarrow{+}^* \subseteq S_1 \times S_1$ and $\xrightarrow{-}^* \subseteq S_1 \times S_1$, as follows:

1. If $s \in S_0$ then $s \in S_1$.
2. If $s_0.a_1.s_1.a_2...s_n \in S_1$ and $s_n \xrightarrow{0} s$ then $s_0.a_1...s_n.a.s \in S_1$ and $s_0.a_1...s_n \xrightarrow{+} s_0.a_1...s_n.a.s$.
3. If $s_0.a_1...s_n, s_0.a_1...s_n.a.s \in S_1$ then $s_0.a_1...s_n.a.s \xrightarrow{-} s_0.a_1...s_n$.
4. If $s_0.a_1...s_n, s_0.a_1...s_n.a_{n+1}...a.s \in S_0$ then $s_0.a_1...s_n \xrightarrow{+}^* s_0.a_1...s_n.a_{n+1}...a.s$.
5. If $s_0.a_1...s_n, s_0.a_1...s_n.a_{n+1}...a.s \in S_0$ then $s_0.a_1...s_n.a_{n+1}...a.s \xrightarrow{-}^* s_0.a_1...s_n$.
6. If $s, t \in S_0$ and $s \sim_i^0 t$ then $s \sim_i^1 t$.
7. If $s, t \in S_1$ and $s \sim_i^1 t$ and $s \xrightarrow{+} s.a.s'$ and $t \xrightarrow{+} t.a.t'$ and $s' \sim_i^0 t'$ then $s.a.s' \sim_i^1 t.a.t'$.

Definition 3.10. An LTS with agent equivalence classes and with transition relations $\xrightarrow{+} \subseteq S_1 \times \mathcal{A} \times S_1$, $\xrightarrow{-} \subseteq S_1 \times \mathcal{A} \times S_1$, $\xrightarrow{+}^* \subseteq S_1 \times S_1$ and $\xrightarrow{-}^* \subseteq S_1 \times S_1$ is called a *history-LTS* if it satisfies the following properties:

1. Forward and backward transitions are converse: $s \xrightarrow{+} t$ iff $t \xrightarrow{-} s$.
2. There is only one way to reach each state: if $s \xrightarrow{+} t$ then for all states s' and all actions b , if $s' \xrightarrow{+} t$ then $s = s'$ and $a = b$.
3. If we let $\xrightarrow{+} = \bigcup_{a \in \mathcal{A}} \xrightarrow{+}^a$, then $\xrightarrow{+}^*$ is the transitive reflexive closure of $\xrightarrow{+}$.
4. If we let $\xrightarrow{-} = \bigcup_{a \in \mathcal{A}} \xrightarrow{-}^a$, then $\xrightarrow{-}^*$ is the transitive reflexive closure of $\xrightarrow{-}$.
5. There are no infinite backward paths: it is impossible to have an infinite chain $s_0 \xrightarrow{-} s_1 \xrightarrow{-} \dots \xrightarrow{-} s_n \xrightarrow{-} \dots$.
6. \sim_i is transitive, reflexive and symmetric for each agent i .

7. If $s_1 \sim_i t_1$ and there exists a state s_0 and an action a such that $s_0 \xrightarrow{+}_a s_1$ then there exists a state t_0 such that $t_0 \xrightarrow{+}_a t_1$ and $s_0 \sim_i t_0$.

These properties capture the idea that a history LTS is exactly what we get when we unfold the paths of an LTS with agent equivalence relations; a formal proof is straightforward. At each stage there is possible future branching but the past is determined in a particular history. Thus the past modalities are like LTL modalities but not the future modalities. The starred modalities give one the power of “always” and “eventually” operators in temporal logics. A history is assumed to have a starting point so it must be well founded.

4 The Logic and its Semantics

In this section we present the logic. It allows us to discuss what is true at a certain state, what was true in the past, what agents know at the current state, and what may or must be true in the future.

We assume a finite set of agents \mathcal{I} , a finite set of actions \mathcal{A} , and a countable set of propositions Q . In the following definition, $a \in \mathcal{A}$, $i \in \mathcal{I}$, and $q \in Q$.

Definition 4.1 (Syntax).

$$\phi := \top \mid q \mid \langle + \rangle_a \phi \mid \langle - \rangle_a \phi \mid \langle + \rangle_* \phi \mid \langle - \rangle_* \phi \mid K_i \phi \mid \neg \phi \mid \phi \wedge \phi$$

As usual, we assume the boolean constants $\perp = p \wedge \neg p$ and $\top = \neg \perp$ and the boolean operators \Rightarrow, \vee, \iff . In addition we define

$$\begin{aligned} [-]_a \phi &= \neg \langle - \rangle_a \neg \phi & [+]_a \phi &= \neg \langle + \rangle_a \neg \phi, \\ [-]^* \phi &= \neg \langle - \rangle_* \neg \phi, & [+]^* \phi &= \neg \langle + \rangle_* \neg \phi, \\ \langle - \rangle \phi &= \bigvee_{a \in \mathcal{A}} \langle - \rangle_a \phi, & \langle + \rangle \phi &= \bigvee_{a \in \mathcal{A}} \langle + \rangle_a \phi, \\ [-] \phi &= \neg \langle - \rangle \neg \phi, & [+] \phi &= \neg \langle + \rangle \neg \phi. \end{aligned}$$

In order to define the semantics we consider the (oriented) labeled graphs over \mathcal{A} . These capture sets of histories as we defined them in the previous section. The nodes of the graph are states and the transitions are labelled by actions in \mathcal{A} . A path through the graph is a history.

If $G = (S, \xrightarrow{a}_{a \in \mathcal{A}})$ is a labelled graph, we denote by \rightarrow the relation $\bigcup_{a \in \mathcal{A}} \xrightarrow{a}$ and

by $\xrightarrow{*}$ the reflexive-transitive closures of \rightarrow respectively.

Definition 4.2 (Labelled forest⁵). A *labelled forest over \mathcal{A}* is a labelled graph $G = (S, \xrightarrow{a}_{a \in \mathcal{A}})$ such that

1. for arbitrary $s, s', s'' \in S$, $s' \rightarrow s$ and $s'' \rightarrow s$ implies $s' = s''$;
2. there exists no infinite sequence $s_0, s_1, \dots, s_k, \dots \in S$ such that $s_{i+1} \rightarrow s_i$ for each $i \in \mathbb{N}$; i.e. it is well-founded to the past.

⁵We call it labelled forest because it is a set of labelled trees.

The *support* of a forest \mathcal{F} , denoted by $\text{supp}(\mathcal{F})$, is the set of its nodes. Given a labelled forest \mathcal{F} , we say that an equivalence relation $\approx \subseteq \text{supp}(\mathcal{F}) \times \text{supp}(\mathcal{F})$ *reflects the branching structure* if whenever $s \approx t$, the existence of a transition $s' \xrightarrow{a} s$ implies the existence of $t' \in \text{supp}(\mathcal{F})$ such that $t' \xrightarrow{a} t$ and $s' \approx t'$. Notice that this is a *backward* bisimulation property; it is a backward preservation property.

Definition 4.3 (Epistemic Frame). Given a set \mathcal{I} (of *agents*), an *epistemic frame* is a tuple $\mathcal{E} = (\mathcal{F}, (\approx_i)_{i \in \mathcal{I}})$, where \mathcal{F} is a labelled forest over \mathcal{A} and $(\approx_i)_{i \in \mathcal{I}}$ is an indexed set of equivalence relations on $\text{supp}(\mathcal{F})$ such that for each $i \in \mathcal{I}$, \approx_i preserves the branching structure.

We call the relation \approx_i the *indistinguishability* relation of agent $i \in \mathcal{I}$. Observe that an epistemic frame defines a unique history-LTS and a history-LTS is supported by a unique epistemic frame.

In the following definition we write s, t, r with or without subscripts for states, p and variants for propositions, ϕ, ψ for formulas and a for actions and i for agents.

Definition 4.4 (Semantics). The semantics is defined for an epistemic frame $\mathcal{E} = (\mathcal{F}, (\approx_i)_{i \in \mathcal{I}})$, a state $s \in \text{supp}(\mathcal{F})$ and an interpretation function $\text{Prop} : \text{supp}(\mathcal{F}) \rightarrow 2^{\mathcal{P}}$, as follows.

$$\begin{aligned}
s \models \top & \quad \text{for all } s. \\
s \models p & \quad \text{if } p \in \text{Prop}(s). \\
s \models \langle + \rangle_a \phi & \quad \text{if there exists a state } t \text{ such that } s \xrightarrow{a} t \text{ and } t \models \phi. \\
s \models \langle - \rangle_a \phi & \quad \text{if there exists a state } r \text{ such that } r \xrightarrow{a} s \text{ and } r \models \phi. \\
s \models \langle + \rangle_* \phi & \quad \text{if there exist } s_1, \dots, s_n \in S \text{ and } a_1, \dots, a_n \in \mathcal{A} \text{ such that} \\
& \quad s \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \xrightarrow{a_3} \dots \xrightarrow{a_{n-1}} s_{n-1} \xrightarrow{a_n} s_n \text{ and } s_n \models \phi. \\
s \models \langle - \rangle_* \phi & \quad \text{if there exist } s_0, \dots, s_{n-1} \in S \text{ and } a_1, \dots, a_n \in \mathcal{A} \text{ such that} \\
& \quad s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_{n-1}} s_{n-1} \xrightarrow{a_n} s \text{ and } s_0 \models \phi. \\
s \models K_i \phi & \quad \text{if for all } t \text{ such that } s \sim_i t, t \models \phi. \\
s \models \neg \phi & \quad \text{if it is not the case that } s \models \phi. \\
s \models \phi_1 \wedge \phi_2 & \quad \text{if } s \models \phi_1 \text{ and } s \models \phi_2.
\end{aligned}$$

Now we have defined our basic operators. For convenience, we also define other

operators as shorthand for certain combinations of these basic operators:

$$\begin{aligned}
\langle + \rangle \phi &:= \bigvee_{a \in \mathcal{A}} \langle + \rangle_a \phi \\
\langle - \rangle \phi &:= \bigvee_{a \in \mathcal{A}} \langle - \rangle_a \phi \\
[+]_a \phi &:= \neg \langle + \rangle_a \neg \phi \\
[-]_a \phi &:= \neg \langle - \rangle_a \neg \phi \\
[+] \phi &:= \bigwedge_{a \in \mathcal{A}} [+]_a \phi \\
[-] \phi &:= \bigwedge_{a \in \mathcal{A}} [-]_a \phi \\
[+]_* \phi &:= \neg \langle + \rangle_* \neg \phi \\
[-]_* \phi &:= \neg \langle - \rangle_* \neg \phi \\
L_i \phi &:= \neg K_i \neg \phi
\end{aligned}$$

Note that $[+] \phi = \neg \langle + \rangle \neg \phi$ and $[-] \phi = \neg \langle - \rangle \neg \phi$. The semantics of these derived operators are:

$$\begin{aligned}
s &\models \perp \text{ never.} \\
s &\models [+]_a \phi \text{ iff for any } t \in \text{supp}(\mathcal{F}) \text{ s.t. } s \xrightarrow{a} t, t \models \phi, \\
s &\models [-]_a \phi \text{ iff for any } t \in \text{supp}(\mathcal{F}) \text{ s.t. } t \xrightarrow{a} s, t \models \phi, \\
s &\models [+]_* \phi \text{ iff for any } t \in \text{supp}(\mathcal{F}) \text{ s.t. } s \xrightarrow{*} t, t \models \phi, \\
s &\models [-]_* \phi \text{ iff for any } t \in \text{supp}(\mathcal{F}) \text{ s.t. } t \xrightarrow{*} s, t \models \phi.
\end{aligned}$$

If we have an epistemic frame \mathcal{E} , a *valuation* is a map $\rho : \text{supp}(\mathcal{F}) \rightarrow 2^{\mathcal{P}}$ which provides an interpretation of the propositions in the states of \mathcal{E} . If a formula ϕ is true in a given epistemic frame \mathcal{E} and state s with a valuation ρ we write $\mathcal{E}, s, \rho \models \phi$ and we say that (\mathcal{E}, s, ρ) is a model of ϕ . In this case we say that ϕ is *satisfiable*. Given an arbitrary $\phi \in \mathcal{L}$, if for any epistemic frame $\mathcal{E} = (\mathcal{F}, (\approx_i)_{i \in \mathcal{I}})$, any state $s \in \text{supp}(\mathcal{F})$ and any valuation ρ , $\mathcal{E}, s, \rho \models \phi$ we say that ϕ is *valid* and write $\models \phi$. We also write $\mathcal{E}, s, \rho \models \Phi$, where Φ is a set of formulas if it models every formula in the set Φ . We write $\Gamma \models \phi$ if any model of Γ is a model of ϕ .


Example 4.5. Here is a more complicated example with multiple agents⁶ which we describe as an illustration of our logic.

The situation is as follows: There are three agents, one diamond, and a bag. The diamond can either be held by one of the agents or it can be in the box. Each agent can perform two actions: reach into the bag and take the diamond *if it is there*, and drop the diamond into the bag, or pretend to drop it. After dropping or pretending to drop the diamond, the agent shows the other agents that his hands are empty, so it is impossible to keep the diamond while pretending to drop it. On the other hand, if the agent does not have the diamond, he can

⁶This example was developed by Caitlin Phillips.

still pretend to drop it in the box. If the agent reaches into the box to take the diamond, he will take it if it is there, and will not take it if it is not there.

Here is the transition system:



The agents are i , j , and l . In state N , no one has the diamond, and in states I , J , and L , agents i , j , and l respectively have the diamond. Action P_i represents agent i picking up or pretending to pick up the diamond and action D_i represents agent i dropping or pretending to drop the diamond.

The equivalence classes are as follows:

$$\begin{aligned} N &\sim_i J \sim_i L \\ N &\sim_j I \sim_j L \\ N &\sim_l I \sim_l J. \end{aligned}$$

We use as propositions $@I$, $@J$, $@L$ and $@N$; each proposition is true only in the corresponding state and in each state only the corresponding proposition is true. For example, the only proposition true in state I is $@I$. We write Prop for this set of 4 propositions. Now we consider the formulas

$$\begin{aligned} \phi_1 &= \bigwedge_{X \in \text{Prop}} X \Rightarrow K_l X \\ \phi_2 &= \langle - \rangle_{P_l} @N \\ \phi_3 &= \bigvee_{X \in \text{Prop}} K_l X \end{aligned}$$

The first formula says that if any of the propositions are true then l knows it: in short l knows where the diamond is. Of course this formula is not universally true, it might or might not be true depending on the situation. The second formula is true for a history where the immediately preceding action is P_l (l picks up the diamond) and in the immediately preceding state nobody had the diamond (i.e. it was in the bag). In other words ϕ_2 describes the situation

where the diamond was in the bag and l has just picked it up. The formula ϕ_3 says whatever the state happens to be, l knows it. Here are two formulas that are true in every state of the unfolded labelled transition system (the history LTS):

$$\phi_2 \Rightarrow [+]^* \phi_1 \text{ and } \phi_3 \Rightarrow [+]^* \phi_1.$$

The first is true because l has picked up the diamond and can now track its movements precisely for all future moves since all actions are visible to him. The second statement is slightly more general, it says that once l knows where the diamond is he can track its future exactly.

Here is another example of reasoning within this system. We define ϕ_4 to be like ϕ_1 except that we have K_i instead of K_l and ϕ_5 is like ϕ_1 except that K_j replaces K_l . Now we can conclude that the following formula is true in every state

$$\langle - \rangle_{D_i} \langle - \rangle_{D_j} \langle - \rangle_{D_l} \Rightarrow [+]^* (\phi_1 \wedge \phi_2 \wedge \phi_3).$$

What we cannot say in this logic is that the location of the diamond is common knowledge. ♣

5 A Complete Axiomatization

We assume the axioms and rules of classical propositional logic. Because we have 5 independent modalities in our logic (K_i , $\langle + \rangle_a$, $\langle - \rangle_a$, $\langle + \rangle^*$ and $\langle - \rangle^*$) we expect to have, in addition, five classes of axioms (one for each modality) reflecting the behaviour of that modality in relation to Booleans. In addition, we will have a few other classes of axioms describing the relations between various modalities. For instance, $\langle + \rangle_a$ and $\langle - \rangle_a$ are in a certain duality supported by our intuition about time, so we expect to have some axioms relating these two. Similarly between $\langle + \rangle^*$ and $\langle - \rangle^*$. We also have some clear intuition about the relation between time transition and knowledge update that will be characterized by some axioms combining dynamic and epistemic operators.

The axioms of \mathcal{L} are presented in Table 1.

Many of the lemmas apply generically to $\langle \rangle$ or \square modalities and the proofs are essentially identical for the different variants. To streamline some proofs, we use the tuple of symbols (\diamond, \square) to represent an arbitrary tuple of type $(\langle - \rangle_a, [-]_a)$, $(\langle + \rangle_a, [+]_a)$, $(\langle - \rangle, [-])$, or $(\langle + \rangle, [+])$. Similarly, (\diamond^*, \square^*) represents $(\langle + \rangle^*, [+]^*)$ or $(\langle - \rangle^*, [-]^*)$. We also use (\diamond^x, \square^x) to represent an arbitrary tuple of type $(\langle - \rangle_a, [-]_a)$, $(\langle + \rangle_a, [+]_a)$, $(\langle - \rangle, [-])$, $(\langle + \rangle, [+])$, $(\langle + \rangle^*, [+]^*)$ or $(\langle - \rangle^*, [-]^*)$. With these notations, the axioms (A1),(A2), (B1), (B2), (C1), (C2) and (D1), (D2) can be regarded as instances of (X1), (X2). Similarly, (C3), (C4) and (D3), (D4) are instances of (X3), (X4).

- (X1): $\vdash \square^x \phi \wedge \square^x (\phi \Rightarrow \psi) \Rightarrow \square^x \psi$
- (X2): If $\vdash \phi$ then $\vdash \square^x \phi$
- (X3): $\vdash \square^* \phi \leftrightarrow (\phi \wedge \square \square^* \phi)$
- (X4): $\vdash \square^* (\phi \Rightarrow \square \phi) \Rightarrow (\phi \Rightarrow \square^* \phi)$

- (A1): $\vdash [+]_a \phi \wedge [+]_a (\phi \Rightarrow \psi) \Rightarrow [+]_a \psi$
(A2): If $\vdash \phi$ then $\vdash [+]_a \phi$
- (B1): $\vdash [-]_a \phi \wedge [-]_a (\phi \Rightarrow \psi) \Rightarrow [-]_a \psi$
(B2): If $\vdash \phi$ then $\vdash [-]_a \phi$
(B3): $\vdash \langle - \rangle_a \top \Rightarrow \bigwedge_{a \neq b} [-]_b \perp$
(B4): $\vdash \langle - \rangle_a \phi \Rightarrow [-]_a \phi$
- (AB1): $\vdash \phi \Rightarrow [+]_a \langle - \rangle_a \phi$
(AB2): $\vdash \phi \Rightarrow [-]_a \langle + \rangle_a \phi$
- (C1): $\vdash [+]^* \phi \wedge [+]^* (\phi \Rightarrow \psi) \Rightarrow [+]^* \psi$
(C2): If $\vdash \phi$ then $\vdash [+]^* \phi$
(C3): $\vdash [+]^* \phi \leftrightarrow (\phi \wedge [+] [+]^* \phi)$
(C4): $\vdash [+]^* (\phi \Rightarrow [+] \phi) \Rightarrow (\phi \Rightarrow [+]^* \phi)$
(D1): $\vdash [-]^* \phi \wedge [-]^* (\phi \Rightarrow \psi) \Rightarrow [-]^* \psi$
(D2): If $\vdash \phi$ then $\vdash [-]^* \phi$
(D3): $\vdash [-]^* \phi \leftrightarrow (\phi \wedge [-] [-]^* \phi)$
(D4): $\vdash [-]^* (\phi \Rightarrow [-] \phi) \Rightarrow (\phi \Rightarrow [-]^* \phi)$
- (BD1): $\vdash \langle - \rangle^* [-] \perp$
- (E1): $\vdash K_i \phi \wedge K_i (\phi \Rightarrow \psi) \Rightarrow K_i \psi$
(E2): If $\vdash \phi$ then $\vdash K_i \phi$
(E3): $\vdash K_i \phi \Rightarrow \phi$
(E4): $\vdash K_i \phi \Rightarrow K_i K_i \phi$
(E5): $\vdash \neg K_i \phi \Rightarrow K_i \neg K_i \phi$
- (BE1): $\vdash \langle - \rangle_a K_i \phi \Rightarrow K_i \langle - \rangle_a \phi$

Table 1: Hilbert-style axiomatization for \mathcal{L}

From (X1) and (X2) alone we can prove a lemma which can be instantiated to all the particular instances. This is a standard lemma of modal logic.

Lemma 5.1. 1. If $\vdash \phi \Rightarrow \psi$, then $\vdash \Box^x \phi \Rightarrow \Box^x \psi$ and $\vdash \Diamond^x \phi \Rightarrow \Diamond^x \psi$.

2. If $\vdash \phi \Rightarrow \psi$, then $\vdash K_i \phi \Rightarrow K_i \psi$.

3. $\vdash \langle - \rangle_a \phi \Rightarrow [-]_a \phi$ and $\vdash \langle - \rangle \phi \Rightarrow [-] \phi$.

Proof. 1. From (X2), $\vdash \phi \Rightarrow \psi$ implies $\vdash \Box^x (\phi \Rightarrow \psi)$. If we use this with $\vdash \Box^x (\phi \Rightarrow \psi) \Rightarrow (\Box^x \phi \Rightarrow \Box^x \psi)$, which is equivalent to (X1), we obtain $\vdash \Box^x \phi \Rightarrow \Box^x \psi$.

To prove the second implication, we start from $\vdash \neg \psi \Rightarrow \neg \phi$ and apply the first result which gives us $\vdash \Box^x \neg \phi \Rightarrow \Box^x \neg \psi$. Using De Morgan we derive $\vdash \Diamond^x \phi \Rightarrow \Diamond^x \psi$.

2. It is proved in the same way as 1; in fact K is a box-like modality.
3. From (B4) we have $\vdash \langle - \rangle_a \phi \Rightarrow \bigwedge_a [-]_a \phi$ which implies $\vdash \langle - \rangle_a \phi \Rightarrow [-]_a \phi$. The same axiom implies $\vdash \bigwedge_a (\langle - \rangle_a \phi \Rightarrow [-] \phi)$ which is equivalent to $\vdash \bigvee_a \langle - \rangle_a \phi \Rightarrow [-] \phi$ which implies $\vdash \langle - \rangle \phi \Rightarrow [-] \phi$. ■

As usual, we say that a formula $\phi \in \mathcal{L}$ is *provable*, denoted by $\vdash \phi$, if it can be proved from the axioms in Table1 and boolean rules. We say that ϕ is *consistent*, if $\neg \phi$ is not provable from the axioms.

Given $\Phi, \Psi \subseteq \mathcal{L}$, Φ proves Ψ if from the formulas of Φ and the axioms we can prove each $\psi \in \Psi$; we write $\Phi \vdash \Psi$. Let $[\Phi] = \{\psi \in \mathcal{L} \mid \Phi \vdash \psi\}$; this is the deductive closure of Φ . Φ is consistent if it is not the case that $\Phi \vdash \perp$.

For a sublanguage $L \subseteq \mathcal{L}$, we call Φ *L-maximally consistent* if Φ is consistent and no formula of L can be added to it without making it inconsistent. The following lemma follows directly from the definition of maximal consistency.

Lemma 5.2. If Γ is a consistent set of formulas then the following assertions are true.

1. if $\diamond^x \top \in [\Gamma]$ and $\diamond^x \phi \notin [\Gamma]$, then $\{\psi \in \mathcal{L} \mid \square^x \psi \in [\Gamma]\} \cup \{\neg \phi\}$ is consistent.
2. if $\square^x \phi \notin [\Gamma]$, then $\{\psi \in \mathcal{L} \mid \square^x \psi \in [\Gamma]\} \cup \{\neg \phi\}$ is consistent.

Proof. Let $\Lambda = \{\psi \in \mathcal{L} \mid \square^x \psi \in [\Gamma]\}$. Suppose that $\Lambda \cup \{\neg \phi\}$ is inconsistent. Then there is a finite set $\{f_1, \dots, f_n\} \subseteq \Lambda$ s.t. $\vdash f_1 \wedge \dots \wedge f_n \Rightarrow \phi$. Hence, $\vdash \square^x (f_1 \wedge \dots \wedge f_n) \Rightarrow \square^x \phi$ implying further $\vdash (\square^x f_1 \wedge \dots \wedge \square^x f_n) \Rightarrow \square^x \phi$. Hence, $\square^x \phi \in [\Gamma]$.

1. If $\diamond^x \top \in [\Gamma]$, from $\square^x \phi \in [\Gamma]$ we obtain $\diamond^x \phi \in [\Gamma]$ - contradiction.
2. $\square^x \phi \notin [\Gamma]$ is again contradictory. ■

A basic theorem that holds for the axiom system is the soundness property.

Theorem 5.3 (Soundness). The axiomatic system of \mathcal{L} is sound, i.e., for any $\phi \in \mathcal{L}$,

$$\vdash \phi \text{ implies } \models \phi.$$

The proof is a routine structural induction. It is sufficient to prove that each axiom is sound and that each rule preserves the soundness.

The more interesting result is the completeness of the axiom system. Moreover, we will show that for each consistent formula a *finite* model can be constructed.

Recall that there are two notions of completeness: *strong* completeness and *weak* completeness. Strong completeness says that

$$\Gamma \models \phi \iff \Gamma \vdash \phi.$$

An important easy consequence of strong completeness is the so-called *compactness* property. A logic is said to be compact if every inconsistent set of formulas has a *finite* inconsistent subset. Our logic is not compact. For example, the set of formulas

$$\{p, [+]p, [+] [+]p, [+]^3 p, \dots, \neg [+]^* p\}$$

is not consistent but any finite subset is consistent. Therefore we cannot hope to prove strong completeness. Instead we prove weak completeness

$$\models \phi \iff \vdash \phi.$$

Many of the basic completeness proofs in the literature are strong completeness proofs and are much easier than weak completeness proofs. The proof that we present shares many of the features of the weak completeness proof for PDL.

Before proceeding with these proofs we establish some notation that will be useful for future constructions.

We extend, canonically, all the logical operators from formulas to sets of formulas. Thus for arbitrary $\Phi, \Psi \subseteq \mathcal{L}$, $\Phi \wedge \Psi = \{\phi \wedge \psi \mid \phi \in \Phi, \psi \in \Psi\}$, $\langle + \rangle_a \Phi = \{\langle + \rangle_a \phi \mid \phi \in \Phi\}$, and so on for all the modal operators.

If $\Phi \subseteq \mathcal{L}$ is finite, we use Φ to also denote $\bigwedge_{\phi \in \Phi} \phi$; it should be clear from the context when Φ denotes a set of formulas and when it denotes the conjunction of its elements.

A key step in the proof is the construction of models by using maximally consistent sets as states. However, because we are trying to prove a weak completeness theorem we have to ensure that we are constructing finite sets of formulas. The liberal notion of maximal consistency used in strong completeness proofs is not available to us. If we wish to construct a model of a formula ϕ , we need to define a special family of formulas associated with ϕ from which we will construct maximal consistent subsets. Furthermore we need to ensure that the collection of formulas we construct is finite. We adapt a construction due to Fischer and Ladner [FL79] developed in the context of PDL.

For an arbitrary $\phi \in \mathcal{L}$, let $\sim \phi = \psi$ whenever $\phi = \neg \psi$ and $\sim \phi = \neg \phi$ otherwise.

For an arbitrary $\phi \in \mathcal{L}$, let $\bar{k}_i \phi = \phi$ whenever $\phi = K_i \psi$ or $\phi = \neg K_i \psi$ and $\bar{k}_i \phi = K_i \phi$ otherwise.

Definition 5.4. The (*Fischer-Ladner*) *closure* of ϕ , written $\mathcal{FL}(\phi)$, is defined as a set of formulas such that:

- $\phi, \langle - \rangle_a p, \langle - \rangle_a \top \in \mathcal{FL}(\phi)$,
- if $\psi \in \mathcal{FL}(\phi)$, then $\sim \psi \in \mathcal{FL}(\phi)$, $\bar{k}_i \psi$ and any subformula of ψ is in $\mathcal{FL}(\phi)$,
- if $\langle - \rangle_a \psi \in \mathcal{FL}(\phi)$ or $\langle + \rangle_a \psi \in \mathcal{FL}(\phi)$, then $\langle - \rangle \psi, \langle + \rangle \psi \in \mathcal{FL}(\phi)$,
- if $\diamond^* \psi \in \mathcal{FL}(\phi)$, then $\diamond \diamond^* \psi \in \mathcal{FL}(\phi)$.

The following lemma is immediate but important to state because we have to ensure that we always have finite sets of formulas when we construct models out of sets of formulas.

Lemma 5.5. For any $\phi \in \mathcal{L}$, $\mathcal{FL}(\phi)$ is finite.

In what follows we fix a consistent formula $\theta \in \mathcal{L}$ and we construct a finite model for θ . This means that we construct an epistemic frame $\mathcal{E}_\theta = (\mathcal{F}_\theta, (\approx_i)_{i \in \mathcal{I}})$, a valuation $\rho : \text{supp}(\mathcal{F}_\theta) \Rightarrow 2^{\mathcal{P}}$ and we will identify a state $s \in \text{supp}(\mathcal{F}_\theta)$ such that $s \models \theta$.

Let Ω_θ be the set of $\mathcal{FL}(\theta)$ -maximally consistent sets. Because $\mathcal{FL}(\theta)$ is finite, Ω_θ and any $\Gamma \in \Omega_\theta$ are finite sets. In the construction of the model we will use Ω_θ as the support set for \mathcal{F}_θ . The transitions on Ω_θ are defined as follows. For each $a \in \mathcal{A}$, let $\xrightarrow{a} \subseteq \Omega_\theta \times \Omega_\theta$ be defined by

$$\Gamma \xrightarrow{a} \Gamma' \text{ iff for any } \psi \in \mathcal{L}, [+]_a \psi \in [\Gamma] \text{ implies } \psi \in [\Gamma'].$$

Now we prove a few properties of these transitions that will be important for the rest of the proof.

Lemma 5.6. For arbitrary $\Gamma, \Gamma' \in \Omega_\theta$ the following are equivalent

1. for any $\phi \in \mathcal{L}$, $[+]_a \phi \in [\Gamma]$ implies $\phi \in [\Gamma']$,
2. for any $\phi \in \mathcal{L}$, $[-]_a \phi \in [\Gamma']$ implies $\phi \in [\Gamma]$.

Proof. (1) implies (2): Suppose that $[-]_a \phi \in [\Gamma']$. Then, $\vdash \Gamma' \Rightarrow [-]_a \phi$ and using axiom (AB1), $\vdash \langle + \rangle_a \Gamma' \Rightarrow \phi$. If we prove that $\langle + \rangle_a \Gamma' \in [\Gamma]$, then $\phi \in [\Gamma]$ and the proof is done. Observe that $\langle + \rangle_a \top \in [\Gamma]$ because otherwise $\neg \langle + \rangle_a \top \in [\Gamma]$ implying $[+]_a \perp \in [\Gamma]$ and from the hypothesis we obtain $\perp \in [\Gamma']$ - impossible. Hence, $\langle + \rangle_a \top \in [\Gamma]$ and if $\langle + \rangle_a \Gamma' \notin [\Gamma]$, from Lemma 5.2 instantiated to $\Box^x = [+]_a$, we obtain that $\{\psi \mid [+]_a \psi \in [\Gamma]\} \cup \{\neg \Gamma'\}$ is consistent. But this is impossible because, from the hypothesis, $\{\psi \mid [+]_a \psi \in [\Gamma]\} \subseteq [\Gamma']$.

(2) implies (1) Suppose that $[+]_a \phi \in [\Gamma]$. Then, $\vdash \Gamma \Rightarrow [+]_a \phi$ implying $\vdash \langle - \rangle_a \Gamma \Rightarrow \langle - \rangle_a [+]_a \phi$. Now (AB2) guarantees that $\vdash \langle - \rangle_a \Gamma \Rightarrow \phi$. In any normal modal logic we have that $\vdash (\Box \psi \wedge \Diamond \top) \Rightarrow \Diamond \psi$. We use this with the previous formula and we obtain $\vdash ([-]_a \Gamma \wedge \langle - \rangle_a \top) \Rightarrow \phi$.

Note that $\langle - \rangle_a \top \in \Gamma'$ because otherwise $[-]_a \perp \in \Gamma'$ and, from the hypothesis we obtain that $\perp \in [\Gamma]$ - impossible. Now, if we prove that $[-]_a \Gamma \in [\Gamma']$, then $\phi \in [\Gamma']$ and the proof is done. Now note that $[-]_a \Gamma \notin [\Gamma']$ implies, using Lemma 5.2 instantiated with $\Box^x = [-]_a$, that $\{\psi \mid [-]_a \psi \in [\Gamma']\} \cup \{\neg \Gamma\}$ is consistent. But this is impossible because, from the hypothesis, $\{\psi \mid [-]_a \psi \in [\Gamma']\} \subseteq [\Gamma]$. ■

This lemma tells us that we can define the transitions either using $[+]$ or $[-]$.

Lemma 5.7. For arbitrary $\Gamma \in \Omega_\theta$ and $[+]_a \phi \in \mathcal{FL}(\theta)$,

1. $[+]_a \phi \in \Gamma$ iff for any $\Gamma' \in \Omega_\theta$, $\Gamma \xrightarrow{a} \Gamma' \Rightarrow \phi \in \Gamma'$;

2. $\langle + \rangle_a \phi \in \Gamma$ iff there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma \xrightarrow{a} \Gamma', \phi \in \Gamma'$;
3. $\langle - \rangle_a \phi \in \Gamma$ iff for any $\Gamma' \in \Omega_\theta$ such that $\Gamma' \xrightarrow{a} \Gamma, \phi \in \Gamma'$;
4. $\langle - \rangle_a \phi \in \Gamma$ iff there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \xrightarrow{a} \Gamma, \phi \in \Gamma'$.

Proof. 1. (\Rightarrow): From the definition of \xrightarrow{a} .

(\Leftarrow): Let ϕ be such that $\phi \in [\Gamma']$ for each $\Gamma' \in \Omega_\theta$ with $\Gamma \xrightarrow{a} \Gamma'$. We need to prove that $[+]_a \phi \in [\Gamma]$. Note that a formula that is in $[\Gamma]$ and also in $\mathcal{FL}(\theta)$ is automatically in Γ .

Let $\Delta = \{\Gamma' \in \Omega_\theta \mid \Gamma \xrightarrow{a} \Gamma'\}$ and let $\delta = \bigvee_{\Gamma' \in \Delta} \Gamma'$. Obviously, $\vdash \delta \Rightarrow \phi$ implying $\vdash [+]_a \delta \Rightarrow [+]_a \phi$. Now, if we prove that $[+]_a \delta \in [\Gamma]$, the proof is done.

Suppose that $[+]_a \delta \notin [\Gamma]$. Lemma 5.2 implies that $\Lambda \cup \{-\delta\}$ is consistent, where $\Lambda = \{\psi \mid [+]_a \psi \in [\Gamma]\}$. But $[+]_a \psi \in [\Gamma]$ implies $\psi \in \Gamma'$ for each $\Gamma' \in \Delta$ and this proves that $\Lambda \cup \{-\delta\}$ cannot be consistent.

(2) is the De Morgan dual of (1).

(3) and (4) are proved in the same way as (1) and (2). ■

We draw the reader's attention to a minor subtlety in the proof because it recurs in several later proofs. We showed that a formula in $\mathcal{FL}(\theta)$, say ϕ , is in the *deductive closure* of a maximally consistent subset, say Γ , of $\mathcal{FL}(\theta)$, in other words we showed that $\phi \in [\Gamma]$. From the fact that ϕ is itself in $\mathcal{FL}(\theta)$ we were able to deduce that ϕ is in Γ itself precisely because Γ is maximal consistent as a subset of $\mathcal{FL}(\theta)$.

We now need to establish the analogous results for the starred modalities. In what follows, let $\rightarrow = \bigcup_{a \in \mathcal{A}} \xrightarrow{a}$ and \rightarrow^* be its reflexive-transitive closure. This means that $\Gamma \rightarrow^* \Gamma'$ if there exists a sequence $\Gamma_1, \dots, \Gamma_k \in \Omega_\theta$ such that

$$\Gamma = \Gamma_1 \rightarrow \Gamma_2 \rightarrow \dots \rightarrow \Gamma_{k-1} \rightarrow \Gamma_k = \Gamma';$$

Because \rightarrow^* is reflexive, k can be 1.

Lemma 5.8. For arbitrary $\Gamma, \Gamma' \in \Omega_\theta$ the following are equivalent

1. for any $\phi \in \mathcal{L}$, $[+]^* \phi \in [\Gamma]$ implies $\phi \in [\Gamma']$,
2. for any $\phi \in \mathcal{L}$, $[-]^* \phi \in [\Gamma']$ implies $\phi \in [\Gamma]$,
3. $\Gamma \rightarrow^* \Gamma'$.

Proof. (1) \implies (3): Let $\Delta = \{\Lambda \in \Omega_\theta \mid \Gamma \rightarrow^* \Lambda\}$ and $\delta = \bigvee_{\Lambda \in \Delta} \Lambda$.

By construction, if $[+] \phi \in [\Lambda]$ for some $\Lambda \in \Delta$, there exists $\Lambda' \in \Delta$ such that $\phi \in [\Lambda']$. This entails $\vdash \delta \Rightarrow [+] \delta$ which guarantees that $\vdash [+]^*(\delta \Rightarrow [+] \delta)$. Using axiom (C4), we obtain $\vdash \delta \Rightarrow [+]^* \delta$. But $\Gamma \in \delta$ (because \rightarrow^* is reflexive),

consequently $\vdash \Gamma \Rightarrow \delta$. From here and the previous we derive $\vdash \Gamma \Rightarrow [+]*\delta$ implying $[+]*\delta \in [\Gamma]$. Now using 1., $\delta \in [\Gamma']$ implying $\Gamma' \in \Delta$.

(3) \Rightarrow (1): Suppose that $\Gamma = \Gamma_1 \rightarrow \dots \rightarrow \Gamma_k = \Gamma'$ and $[+]*\phi \in [\Gamma]$. Axiom (C3) guarantees that $\phi \in [\Gamma_1]$ and $[+][+]*\phi \in [\Gamma_1]$. Hence $[+]*\phi \in [\Gamma_2]$ from the definition of \rightarrow . The same argument can be repeated for the k cases eventually giving $[+]*\phi \in [\Gamma_k] = [\Gamma']$ which implies, using axiom (C3), $\phi \in [\Gamma']$.

(2) \Leftrightarrow (3): It is proved in the same way using the axioms (D1) and (D2) in instances of Lemma 5.1 and (D3), (D4) respectively. \blacksquare

Lemma 5.9. For arbitrary $\Gamma \in \Omega_\theta$ and $[+]*\phi \in \mathcal{FL}(\theta)$,

1. $[+]*\phi \in \Gamma$ iff for any $\Gamma' \in \Omega_\theta$ such that $\Gamma \rightarrow^* \Gamma', \phi \in \Gamma'$;
2. $\langle + \rangle^*\phi \in \Gamma$ iff there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma \rightarrow^* \Gamma', \phi \in \Gamma'$;
3. $[-]*\phi \in \Gamma$ iff for any $\Gamma' \in \Omega_\theta$ such that $\Gamma' \rightarrow^* \Gamma, \phi \in \Gamma'$;
4. $\langle - \rangle^*\phi \in \Gamma$ iff there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \rightarrow^* \Gamma, \phi \in \Gamma'$.

Proof. (1) \Rightarrow : From Lemma 5.8.

(\Leftarrow): Let ϕ be such that $\phi \in [\Gamma']$ for each $\Gamma' \in \Omega_\theta$ with $\Gamma \rightarrow^* \Gamma'$. We need to prove that $[+]*\phi \in [\Gamma]$.

Let $\Delta = \{\Gamma' \in \Omega_\theta \mid \Gamma \rightarrow^* \Gamma'\}$ and let $\delta = \bigvee_{\Gamma' \in \Delta} \Gamma'$. Obviously, $\vdash \delta \Rightarrow \phi$ implying $\vdash [+]*\delta \Rightarrow [+]*\phi$. Now, if we prove that $[+]*\delta \in [\Gamma]$, the proof is done.

Suppose that $[+]*\delta \notin [\Gamma]$. Lemma 5.2 implies that $\Lambda \cup \{\neg\delta\}$ is consistent, where $\Lambda = \{\psi \mid [+]*\psi \in [\Gamma]\}$. But $[+]*\psi \in [\Gamma]$ implies $\psi \in \Gamma'$ for each $\Gamma' \in \Delta$ and this proves that $\Lambda \cup \{\neg\delta\}$ cannot be consistent.

(2) is equivalent to (1).

(3) and (4) are proved in the same way. \blacksquare

Now we can proceed with our construction of the model for θ . We start by showing that $(\Omega_\theta, \xrightarrow{a})_{a \in \mathcal{A}}$ is a forest. For this we need to verify that the past is unique and that the graphs have no loops. The precise statement is given in the following theorem.

Theorem 5.10. If $f \in \mathcal{L}$ is consistent, then $\mathcal{F}_\theta = (\Omega_\theta, \xrightarrow{a})_{a \in \mathcal{A}}$ is a forest over \mathcal{A} .

The proof of this theorem is broken down into two lemmas.

Lemma 5.11. For arbitrary $\Gamma, \Gamma_1, \Gamma_2 \in \Omega_\theta$, if $\Gamma_1 \xrightarrow{a} \Gamma$ and $\Gamma_2 \xrightarrow{b} \Gamma$, then $a = b$ and $\Gamma_1 = \Gamma_2$.

Proof. To prove that $a = b$ it is sufficient to observe that $\langle - \rangle_a \top \wedge \langle - \rangle_b \top$ is inconsistent, result that is a direct consequence of axiom (B3).

Now, from $\Gamma_1 \xrightarrow{a} \Gamma$ and $\Gamma_2 \xrightarrow{a} \Gamma$ we prove that $\Gamma_1 = \Gamma_2$. Suppose that there exists $\phi \in \mathcal{FL}(\theta)$ s.t. $\phi \in \Gamma_1$ and $\neg\phi \in \Gamma_2$. Then, from axiom (AB1) we obtain that $[+]_a \langle - \rangle_a \phi \in [\Gamma_1]$ and $[+]_a \langle - \rangle_a \neg\phi \in [\Gamma_2]$. Now $\Gamma_1 \xrightarrow{a} \Gamma$ guarantees that $\langle - \rangle_a \phi \in [\Gamma]$ while $\Gamma_2 \xrightarrow{a} \Gamma$ guarantees that $\langle - \rangle_a \neg\phi \in [\Gamma]$. Further, using axiom (B4) we obtain that $[-]\phi, [-]\neg\phi \in [\Gamma]$ implying $[-]\perp \in [\Gamma]$. On the other hand, $\langle - \rangle_a \phi \in [\Gamma]$ implies $\langle - \rangle_a \top \in [\Gamma]$ which is equivalent to $\neg[-]\perp \in [\Gamma]$ - contradicts the consistency of $[\Gamma]$. ■

Now we prove that in the graph $(\Omega_\theta, \xrightarrow{a})_{a \in \mathcal{A}}$ there are no backwards infinite sequences; this will conclude the proof that $(\Omega_\theta, \xrightarrow{a})_{a \in \mathcal{A}}$ is a forest over \mathcal{A} .

Lemma 5.12. There exists no infinite sequence $\Gamma_1, \dots, \Gamma_k, \dots \in \Omega_\theta$ such that

$$\dots \Gamma_k \rightarrow \Gamma_{k-1} \rightarrow \dots \rightarrow \Gamma_1 = \Gamma.$$

Proof. Suppose that there exists such a sequence. Axiom (BD1) guarantees that $\langle - \rangle^* [-]\perp \in [\Gamma]$ and using Lemma 5.9 we obtain that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \rightarrow^* \Gamma$ and $[-]\perp \in \Gamma'$. Lemma 5.11 guarantees that Γ' is one of the elements of our sequence, hence $\neg \langle - \rangle \top \in \Gamma'$. But this implies that there exists no $\Gamma'' \in \Omega_\theta$ such that $\Gamma'' \rightarrow^* \Gamma'$, this contradiction establishes the result. ■

To complete the construction of the model for θ we need to define the indistinguishability relations on Ω_θ that will eventually organize our forest as an epistemic frame.

For each $i \in \mathcal{I}$, let $\approx_i \subseteq \Omega_\theta \times \Omega_\theta$ be defined as follows:

$$\Gamma \approx_i \Gamma' \text{ iff for any } \phi \in \mathcal{L}, K_i \phi \in [\Gamma] \text{ iff } K_i \phi \in [\Gamma'].$$

By construction, \approx_i is an equivalence relation. Now, to finalize our construction, we must prove that for each $i \in \mathcal{I}$, \approx_i preserves the branching structure of \mathcal{F}_θ and finally that we have an epistemic frame.

Theorem 5.13. $\mathcal{E}_\theta = (\mathcal{F}_\theta, (\approx_i)_{i \in \mathcal{I}})$, where $\mathcal{F}_\theta = (\Omega_\theta, \xrightarrow{a})_{a \in \mathcal{A}}$ and \approx_i are defined as before, is an epistemic frame.

The proof is broken into a number of lemmas. The first lemma that we need is the following.

Lemma 5.14. For arbitrary $\Gamma, \Gamma' \in \Omega_\theta$, if for any ϕ , $K_i \phi \in [\Gamma]$ implies $\phi \in [\Gamma']$, then for any ϕ , $K_i \phi \in [\Gamma]$ implies $K_i \phi \in [\Gamma']$.

Proof. Suppose that for any ϕ , $K_i \phi \in [\Gamma]$ implies $\phi \in [\Gamma']$ and let $K_i \psi \in [\Gamma]$. From our hypothesis we obtain that if $K_i \psi \notin [\Gamma']$, then $K_i K_i \psi \notin [\Gamma]$. From the axioms (E3) and (E4), $\vdash K_i \psi \leftrightarrow K_i K_i \psi$. Hence, $K_i \psi \notin [\Gamma]$, this contradiction completes the proof. ■

Now we can prove that for each $i \in \mathcal{I}$, \approx_i preserves the backwards branching structure of \mathcal{F}_θ .

Theorem 5.15. For arbitrary $\Gamma, \Gamma' \in \Omega_\theta$, if $\Gamma \approx_i \Gamma'$ and there exists $\Gamma_0 \in \Omega_\theta$ such that $\Gamma_0 \xrightarrow{a} \Gamma$, then there exists $\Gamma'_0 \in \Omega_\theta$ such that $\Gamma'_0 \xrightarrow{a} \Gamma'$ and $\Gamma'_0 \approx_i \Gamma_0$.

Proof. Because $\vdash \top$, using (E2) we obtain $\vdash K_i \top$. Because $K_i \top \in [\Gamma_0]$, we obtain that $\langle - \rangle_a K_i \top \in [\Gamma]$ and axiom (BE1) implies $K_i \langle - \rangle_a \top \in [\Gamma]$. Now, from $\Gamma \approx_i \Gamma'$, $\langle - \rangle_a \top \in [\Gamma']$. From Lemma 5.7 we obtain that there exists $\Gamma'_0 \in \Omega_\theta$ such that $\Gamma'_0 \rightarrow \Gamma'$.

We prove now that $\Gamma'_0 \approx \Gamma_0$. Suppose that $K_i \phi \in [\Gamma_0]$. Then, $\langle - \rangle_a K_i \phi \in [\Gamma]$ and axiom (BE1) implies $K_i \langle - \rangle_a \phi \in [\Gamma]$. Now from $\Gamma \approx_i \Gamma'$, $\langle - \rangle_a \phi \in [\Gamma']$. Now axiom (B4) implies $[-] \phi \in [\Gamma']$ and because $\Gamma'_0 \rightarrow \Gamma'$, Lemma 5.7 implies $\phi \in [\Gamma'_0]$.

Hence, $K_i \phi \in [\Gamma_0]$ implies $\phi \in [\Gamma'_0]$ and Lemma 5.14 concludes that $K_i \phi \in [\Gamma_0]$ implies $K_i \phi \in [\Gamma'_0]$. Similarly can be proved that $K_i \phi \in [\Gamma'_0]$ implies $K_i \phi \in [\Gamma_0]$. ■

Lemma 5.14 also establishes the next result that is needed for the proof of the theorem.

Lemma 5.16. For arbitrary $\Gamma \in \Omega_\theta$ and $K_i \phi \in \mathcal{FL}(\theta)$,

$$K_i \phi \in \Gamma \text{ iff for any } \Gamma' \in \Omega_\theta \text{ such that } \Gamma \approx_i \Gamma', \phi \in \Gamma'$$

Proof. (\Rightarrow) This follows directly from Lemma 5.14.

(\Leftarrow) Let ϕ be such that $K_i \phi \in \mathcal{FL}(\theta)$ and $\phi \in \Gamma'$ for each $\Gamma' \in \Omega_\theta$ with $\Gamma \approx_i \Gamma'$. We need to prove that $K_i \phi \in \Gamma$.

Let $\Delta = \{\Gamma' \in \Omega_\theta \mid \Gamma \approx_i \Gamma'\}$, let $\Lambda = \{f_1, \dots, f_n\} = \bigcap_{\Gamma' \in \Delta} \Gamma'$ and let $F = f_1 \wedge \dots \wedge f_n$. Then $\vdash F \Rightarrow \phi$ implying $\vdash K_i F \Rightarrow K_i \phi$. Consequently, if we prove that $K_i F \in [\Gamma]$, the proof is done.

Suppose that $K_i F \notin [\Gamma]$. Then, there exists $f_t \in \Lambda$ such that $K_i f_t \notin \Gamma$. Then, $\neg K_i f_t \in \Gamma$ and axiom (E5) implies $K_i \neg K_i f_t \in [\Gamma]$. The definition of \approx_i guarantees that for any $\Gamma' \in \Delta$, $K_i \neg K_i f_t \in [\Gamma']$ and axiom (E3) entails that for any $\Gamma' \in \Delta$, $\neg K_i f_t \in \Gamma'$. Hence, $\vdash F \Rightarrow \neg K_i f_t$ which is equivalent to $\vdash K_i f_t \Rightarrow \neg F$. But $\vdash F \Rightarrow f_t$ implying $\vdash K_i F \Rightarrow K_i f_t$. Consequently, $\vdash K_i F \Rightarrow \neg F$. But from axiom (E3), $\vdash K_i F \Rightarrow F$, implying $\vdash \neg K_i F$. But Λ is consistent and $K_i F \notin [\Lambda]$, then a similar argument with the one used in Lemma 5.2 (notice that K_i is a normal modal operator of type \Box) shows that $\Lambda \cup \{\neg F\}$ is consistent, which is impossible. ■

This completes the proof of the theorem.

We are now ready to complete the construction of the model of θ . \mathcal{E}_θ is the epistemic frame of the model and we define a valuation $\rho_\theta : \Omega_\theta \rightarrow 2^{\mathcal{P}}$ by $\rho_\theta(\Gamma) = \{p \in \mathcal{P} \mid p \in \Gamma\}$. With this definition we prove the Truth Lemma.

Lemma 5.17 (Truth Lemma). If $\theta \in \mathcal{L}$ is consistent, \mathcal{E}_θ and ρ_θ are defined as before, then for any $\phi \in \mathcal{FL}(\theta)$ and $\Gamma \in \Omega_\theta$,

$$\phi \in \Gamma \text{ iff } \Gamma \models \phi.$$

Proof. Induction on ϕ .

[**The case** $\phi = p \in \mathcal{P}$:] from definition of $Prop_\theta$.

[**The case** $\phi = \neg\psi$:] (\implies) Suppose that $\Gamma \not\models \neg\psi$. Then $\Gamma \models \psi$ and from the inductive hypothesis, $\psi \in \Gamma$, hence $\phi \notin \Gamma$.

(\impliedby) Suppose that $\Gamma \models \neg\psi$ and $\neg\psi \notin \Gamma$. Then, $\psi \in \Gamma$ and the inductive hypothesis guarantees that $\Gamma \models \psi$ - contradiction.

[**The case** $\phi = \phi_1 \wedge \phi_2$:] $\phi_1 \wedge \phi_2 \in \Gamma$ iff $\phi_1, \phi_2 \in \Gamma$ which is equivalent, using the inductive hypothesis, to $[\Gamma \models \phi_1 \text{ and } \Gamma \models \phi_2]$, equivalent to $\Gamma \models \phi_1 \wedge \phi_2$.

[**The case** $\phi = \langle + \rangle_a \psi$:] (\implies) If $\langle + \rangle_a \psi \in \Gamma$, Lemma 5.7 implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma \xrightarrow{a} \Gamma'$ and $\psi \in \Gamma'$. From the inductive hypothesis, $\Gamma' \models \psi$, implying $\Gamma \models \phi$.

(\impliedby) $\Gamma \models \langle + \rangle_a \psi$ implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma \xrightarrow{a} \Gamma'$ and $\Gamma' \models \psi$. From the inductive hypothesis, $\psi \in \Gamma'$ and Lemma 5.7 implies $\langle + \rangle_a \psi \in \Gamma$.

[**The case** $\phi = \langle - \rangle_a \psi$:] (\implies) If $\langle - \rangle_a \psi \in \Gamma$, Lemma 5.7 implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \xrightarrow{a} \Gamma$ and $\psi \in \Gamma'$. From the inductive hypothesis, $\Gamma' \models \psi$, implying $\Gamma \models \phi$.

(\impliedby) $\Gamma \models \langle - \rangle_a \psi$ implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \xrightarrow{a} \Gamma$ and $\Gamma' \models \psi$. From the inductive hypothesis, $\psi \in \Gamma'$ and Lemma 5.7 implies $\langle - \rangle_a \psi \in \Gamma$.

[**The case** $\phi = \langle + \rangle^* \psi$:] (\implies) If $\langle + \rangle^* \psi \in \Gamma$, Lemma 5.9 implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma \rightarrow^* \Gamma'$ and $\psi \in \Gamma'$. From the inductive hypothesis, $\Gamma' \models \psi$, implying $\Gamma \models \phi$.

(\impliedby) $\Gamma \models \langle + \rangle^* \psi$ implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma \rightarrow^* \Gamma'$ and $\Gamma' \models \psi$. From the inductive hypothesis, $\psi \in \Gamma'$ and Lemma 5.9 implies $\langle + \rangle^* \psi \in \Gamma$.

[**The case** $\phi = \langle - \rangle^* \psi$:] (\implies) If $\langle - \rangle^* \psi \in \Gamma$, Lemma 5.9 implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \rightarrow^* \Gamma$ and $\psi \in \Gamma'$. From the inductive hypothesis, $\Gamma' \models \psi$, implying $\Gamma \models \phi$.

(\impliedby) $\Gamma \models \langle - \rangle^* \psi$ implies that there exists $\Gamma' \in \Omega_\theta$ such that $\Gamma' \rightarrow^* \Gamma$ and $\Gamma' \models \psi$. From the inductive hypothesis, $\psi \in \Gamma'$ and Lemma 5.9 implies $\langle - \rangle^* \psi \in \Gamma$.

[**The case** $\phi = K_i \psi$:] (\implies) If $K_i \psi \in \Gamma$, Lemma 5.16 implies that for any $\Gamma' \in \Omega_\theta$ such that $\Gamma \approx_i \Gamma'$, $\psi \in \Gamma'$. From the inductive hypothesis, $\Gamma' \models \psi$, implying $\Gamma \models \phi$.

(\impliedby) $\Gamma \models K_i \psi$ implies that for any $\Gamma' \in \Omega_\theta$ such that $\Gamma \approx_i \Gamma'$, $\Gamma' \models \psi$. From the inductive hypothesis, $\psi \in \Gamma'$ and Lemma 5.16 implies $K_i \psi \in \Gamma$. \blacksquare

A direct consequence of Truth Lemma is the finite model property.

Theorem 5.18 (Finite model property). For any consistent formula $\phi \in \mathcal{L}$ there exists a finite model. Moreover, the size of the model is bound by the structure of ϕ .

The finite model property in this context has two important consequences: the weak completeness of the axiomatic system and the decidability of the satisfiability problem.

Theorem 5.19 (Weak completeness). The axiomatic system of \mathcal{L} is complete, i.e., for any $\phi \in \mathcal{L}$,

$$\models \phi \text{ implies } \vdash \phi.$$

Proof. The proof is based on the fact that any consistent formula has a model. We wish to show that $\models \phi$ implies $\vdash \phi$. Now we have shown that if ϕ is consistent it has a model. Clearly then, if $\neg\phi$ is consistent there is a model of $\neg\phi$. The last statement is equivalent to saying that if $\not\vdash \phi$ then $\neg\phi$ is satisfiable. If $\neg\phi$ is satisfiable it follows that not every model models ϕ , i.e. $\not\models \phi$. Thus we have $\not\vdash \phi$ implies $\not\models \phi$, or taking the contrapositive, $\models \phi$ implies $\vdash \phi$. ■

Observe that in the previous construction, the size of Ω_θ depends on the number and type of operators that θ contains. In what follows we refer to the cardinality $|\Omega_\theta|$ of Ω_θ as *the size of θ* .

The satisfiability problem is the problem of deciding, given an arbitrary formula $\phi \in \mathcal{L}$, if ϕ has at least one model. The finite model property entails that the satisfiability problem for our logic is decidable.

Theorem 5.20 (Decidability). The satisfiability problem for \mathcal{L} is decidable.

Proof. We have proved that θ has at least one model iff it is consistent. And if θ is consistent we have proved that it has a model of size $|\Omega_\theta| \in \mathbb{N}$. But the class of models of size $k \in \mathbb{N}$ is finite. Consequently, we can decide in a finite number of steps if θ does or does not have a model by checking all the models of the appropriate sizes. ■

6 Conclusions and Related Work

There seems to be a mysterious divide between concurrency theory, which is primarily a European enterprise, and distributed systems theory which is intensively explored in the United States, Israel and a few other places. This is unfortunate because the two have much to learn from each other. Concurrency theorists can learn sophisticated new tools like algebraic topology and deeper problems whereas the distributed systems community could learn about, for example, compositional reasoning. Epistemic logic is one of the areas where the distributed systems community got an early start [FHMV95] in the mid 1980s whereas the concurrency theory community is only just starting to use these ideas. This schizophrenia is manifested even in the work of individuals! For example, the third author of the present paper worked on common knowledge in asynchronous distributed systems in the late 1980s [PT88, PT92] and later on concurrency theory [SRP91] without making the connection. The present work is intended to make epistemic logic more readily accessible to the concurrency theory community by providing a combination of epistemic logic with

the Hennessy-Milner logic that the concurrency community is accustomed to using.

The ground breaking paper of Halpern and Moses [HM84, HM90] showed the importance of common knowledge as a way of formalizing agreement protocols in distributed systems. Very quickly variants of common knowledge were developed [NT87, PT88] and many new applications were explored [NT90]. Extensions to probability [HT89] and zero-knowledge protocols [HMT88] quickly followed. The textbook of Fagin et al. [FHMV95] made these ideas widely accessible and stimulated even more interest and activity. There are numerous recent papers by Halpern and his collaborators, Parikh and his collaborators and students, van Benthem and the Amsterdam school and by several other authors as well. Applications of epistemic concepts range across game theory, economics, spatial reasoning and even social systems.

In the concurrency theory community there is very little work on this topic. Two striking examples are a recent paper by Chadha, Delaune and Kremer [CDK09] and one by Dechesne, Mousavi and Orzan [DMO07]. The former paper defines an epistemic logic for a particular process calculus, a variant of the π -calculus and uses it to reason about epistemic situations. The latter paper explores the connection between operational semantics and epistemic logic and is closer in spirit to our work which is couched in terms of labelled transition systems. Neither of these paper really integrate Hennessy-Milner logic and epistemic logic. A recent paper by Knight et al. [CKPP12] uses a rudimentary epistemic logic to capture epistemic strategies for games on concurrent processes. A recent paper by Pacuit and Simon [PS11] develops a PDL-style logic for reasoning about protocols. They also prove a completeness theorem for their logic; it is perhaps the closest in spirit to our work.

Acknowledgments

We have benefitted from helpful discussions on epistemic logic with Samson Abramsky, Alexandru Baltag, Adam Brandenburger, Kostas Chatzikokolakis, Valentin Goranko, Joe Halpern, Eric Pacuit, Rohit Parikh, Caitlin Phillips, Doina Precup, R. Ramanujam and Mehrnoosh Sadrzadeh. This research was supported by a grant from NSERC and from an EPSRC grant that allowed Prakash Panangaden to spend a sabbatical year at Oxford. We are very grateful to Alexandra Silva and Robert Constable for the invitation to submit this article to Dexter Kozen's Festschrift.

References

- [BdRV01] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.

- [CDK09] Rohit Chadha, Stéphanie Delaune, and Steve Kremer. Epistemic logic for the applied pi calculus. In David Lee, Antónia Lopes, and Arnd Poetzsch-Heffter, editors, *Proceedings of IFIP International Conference on Formal Techniques for Distributed Systems (FMOODS/FORTE'09)*, Lecture Notes in Computer Science, pages 182–197. Springer, June 2009.
- [CKPP12] Konstantinos Chatzikokolakis, Sophia Knight, Prakash Panangaden, and Catuscia Palamidessi. Epistemic strategies and games on concurrent processes. *ACM Transactions on Computational Logic*, 2012. in press.
- [DMO07] Francien Dechesne, MohammadReza Mousavi, and Simona Orzan. Operational and epistemic approaches to protocol analysis: Bridging the gap. In *Proceedings of the 14th International Conference on Logic for Programming Artificial Intelligence and Reasoning (LPAR'07)*, volume 4790 of *Lecture Notes in Artificial Intelligence*, pages 226–241. Springer-Verlag, 2007.
- [FHMV95] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [FL79] Michael Fischer and Richard Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 1979.
- [Hin62] J. Hintikka. *Knowledge and Belief*. Cornell University Press, 1962.
- [HM84] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. In *Proceedings of the Third A.C.M. Symposium on Principles of Distributed Computing*, pages 50–61, 1984. A revised version appears as IBM Research Report RJ 4421, Aug., 1987.
- [HM85] M. Hennessy and R. Milner. Algebraic laws for nondeterminism and concurrency. *Journal of the ACM*, 32(1):137–162, 1985.
- [HM90] J. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *JACM*, 37(3):549–587, 1990.
- [HMT88] Joseph Y. Halpern, Yoram Moses, and Mark R. Tuttle. A knowledge-based analysis of zero knowledge. In *Proceedings of the 20th ACM Symposium on Theory of Computing*, pages 132–147, 1988.
- [HT89] Joseph Y. Halpern and Mark R. Tuttle. Knowledge, probability and adversaries. In *Proceedings Of The Eighth Annual ACM Symposium On Principles of Distributed Computing*, pages 103–118. ACM, 1989.

- [KP81] Dexter Kozen and Rohit Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.
- [Kri63] S. Kripke. Semantical analysis of modal logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.
- [Mil89] R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [Mil90] R. Milner. *Handbook of Theoretical Computer Science: Volume B*, chapter Operational and Algebraic Semantics of Concurrent Processes, pages 1201–1242. MIT Press, 1990.
- [NT87] G. Neiger and S. Toueg. Substituting for real time and common knowledge in asynchronous distributed systems. In *Proceedings of the Sixth A.C.M. Symposium on Principles of Distributed Computing*, pages 281–293, 1987.
- [NT90] Gil Neiger and Mark R. Tuttle. Common knowledge and consistent simultaneous coordination. In J. van Leeuwen and N. Santoro, editors, *Proceedings of the Fourth International Workshop on Distributed Algorithms*, volume 486 of *Lecture Notes In Computer Science*, pages 334–352. Springer-Verlag, 1990.
- [Pop94] Sally Popkorn. *First Steps in Modal Logic*. Cambridge University Press, 1994.
- [PS11] Eric Pacuit and Sunil Simon. Reasoning with protocols under imperfect information. *The Review of Symbolic Logic*, 4(3):412–444, September 2011.
- [PT88] P. Panangaden and K. E. Taylor. Concurrent common knowledge. In *Proceedings of the Seventh Annual ACM Symposium on Principles of Distributed Computing*, pages 197–209, 1988.
- [PT92] P. Panangaden and K. E. Taylor. Concurrent common knowledge. *Distributed Computing*, 6:73–93, 1992.
- [SRP91] V. A. Saraswat, M. Rinard, and P. Panangaden. Semantic foundations of concurrent constraint programming. In *Proceedings of the Eighteenth Annual ACM Symposium on Principles of Programming Languages*, 1991.
- [vDvdHK08] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi. *Dynamic Epistemic Logic*. Number 337 in Synthese Library. Springer-Verlag, 2008.