View-based explicit knowledge

R. Ramanujam

The Institute of Mathematical Sciences, C.I.T. campus, Chennai 600 113, India.
jam@imsc.ernet.in

In the information based definition of knowledge, an agent is said to know α at a state s if α is true in all states that look the same as s to that agent. However, in systems where an agent's view of the system is partial, even if a state s' may be logically indistinguishable from a state s, s' may not be visible from s. For instance, in a distributed system, all global states in which the agent's local state does not change look the same to that agent, but this set of global states may not be accessible because the agent may not even be aware of the existence of many agents in the network. We propose a logic of explicit knowledge built on agents' views and show it to be decidable.

1 A distinction

1.1 Basic logic of knowledge

During the last decade, there has been extensive research into the notion of ascribing knowledge to reasoning agents, particularly in the contexts of distributed systems, artificial intelligence and mathematical economics. The book [FHMV95] provides an excellent summary of dominant themes in this line of research.

Following Hintikka's logical treatment ([Hi62]), knowledge is studied as a modality in (typically) a propositional modal logic, with the semantics of formulas given in Kripke frames where the accessibility relation is an equivalence. Formally, we have a countable set of propositions $P = \{p_0, p_1, \ldots\}$ and the syntax of formulas is given by:

$$\Phi \ ::= \ p \in P \mid \neg \alpha \mid \alpha \vee \beta \mid \mathbf{K}\, \alpha$$

[1] This paper is a revised and expanded version of [R96b], prepared in honour of Rohit Parikh, on his $60th$ birthday.

Preprint submitted to Elsevier Preprint *26 February 1998*

The frames are Kripke structures $F = (W, \sim)$, where W is a set of possible worlds, and $\sim$ is an equivalence relation called the indistinguishability relation for the reasoning agent. Models are frames with valuations: $M = (F, V)$, where $F = (W, \sim)$ and $V : W \to 2^P$. The semantics of formulas is then given by:

- $M, w \models p$ iff $p \in V(w)$.
- $M, w \models \neg\alpha$ iff $M, w \not\models \alpha$.
- $M, w \models \alpha \vee \beta$ iff ($M, w \models \alpha$ or $M, w \models \beta$).
- $M, w \models \mathbf{K}\alpha$ iff $\forall w' \sim w,\ M, w' \models \alpha$.

The satisfiability and model checking problems for the logic are decidable, the set of valid formulas can be axiomatized ([FHMV95]).

1.2 *Ascribed knowledge*

An assertion $\mathbf{K}\alpha$ in this logic can be read as: 'the reasoning agent *knows* that the formula α holds'. What kind of knowledge is referred to here ? There are (at least) two ways of interpreting this statement (see [HM92] for a related discussion):

- The agent `possesses` this knowledge in some actual sense; for instance, the agent can provide evidence in support of the assertion α, or can answer some queries 'about' α.
- This is knowledge `ascribed` to the agent by an observer outside the system. The observer considers all the information available to the agent at the world state w and concludes that α is logically implied by such information, and therefore, can use the system as if the agent knew α.

Clearly, the indistinguishability relation in the frames captures the notion of information in the latter sense (by describing the agent's information partitions of the set of states), and thus the latter interpretation is more immediate. Limited as it may seem, this reading is still immensely useful, as we can meaningfully ascribe knowledge to thermostats, computer programs and other such entities in this manner ("the thermostat knows that the room is too warm, so it switches on the air conditioner", "the program knows that the shared resource which is in use will be released within t units, so it decides to wait until then", and so on). Such knowledge assertions can be useful in system verification.

1.3 *Limited visibility*

This reading is not useful in situations where an agent in a system needs to *compute* its knowledge on the basis of evidence available to it. There is a way in which explicit knowledge differs from implicit knowledge, one which is especially important in the context of distributed computing. A reasoner operating in an environment has only a limited view of the world, and her explicit knowledge is determined by the **visibility** of world states to her, whereas ascribed implicit knowledge of the agent depends on her behaviour in all possible worlds. In terms of the indistinguishability relation $\sim$ above, computing $\sim$ in itself may necessitate some effort. We can conceive of a situation where $s_0 \sim s_1 \sim s_2$, but at state s_0 only s_0 and s_1 are visible to a computing agent and not s_2. That is, an observer who has access to complete information about all the world states may declare that the agent would behave in the same way in all these three states, but the agent itself may require computation to realise this. If the agent makes the effort and computes the world state s_2, it may also realise that s_2 is in the same equivalence class as s_1.

Visibility can also seen as a manifestation of *resource boundedness*: a resource limited agent might need to reuse its resources, and thus computation rendering a state visible may entail losing information about other visible states making them no longer visible. Thus in the example above, as the agent computes the state s_2, it might 'forget' the information about s_0.

Thus, even among states within the same externally ascribed indistinguishability class, some may be 'farther' and some 'nearer' in terms of visibility to an agent. In the context of distributed systems, this happens routinely. A component of a system behaves in the same manner in all *global* system states in which its *local* state is the same. Hence all these states would be ascribed in the same equivalence class for that agent. However, this component would typically communicate only with immediate neighbours in the network, and may not even be aware of the existence of many other 'distant' components in the network. In this case, for the agent to compute its knowledge, the first task of computing these indistinguishable system states is well beyond its capabilities.

Moreover, at different system states, the agent may have different views, and this is also dependent on the computational resources available to the agent. For instance, a dynamic computational agent with bounded memory may forget events in the distant past. Even in static situations, computation may result in learning thereby enlarging visibility. This aspect of the distinction between ascribed and computed knowledge of view-limited reasoners seems to have been relatively less studied in the literature.

1.4 Algorithmic knowledge

In this context, we should mention one way in which implicit and explicit knowledge have been distinguished extensively in the literature, where again view limitations seem relevant. In the logic defined above, a reasoner knows every logical consequence of any fact she knows; further she knows all the valid formulas. Human reasoners simply do not exhibit such *logical omniscience*, and hence the logic above does not reflect reasoning about human knowledge. On the other hand, if we consider the reasoner above to be a computing agent, we come up with the problem of *resource boundedness*. The agent has only limited computational resources at its disposal, and exploring all the logical consequences of an assumption is expensive business. In particular, the validity problem for the propositional calculus is already co-NP-complete, and for the reasoner to know all tautologies requires, at the least, a co-NP machine. See [FHMV 95] for a detailed treatment of this issue. Clearly explicit knowledge should be resource bounded.

The papers of Rohit Parikh ([Pa 87a], [Pa 87b], [Pa 90] and [Pa 94]) provide all manners of criticism which include the two observations above and much more. For instance, he discusses the distinction between knowledge and information; the distinction between knowledge of a proposition and knowledge of the *sentence* denoting the proposition, and so on. In some sense, all these can be seen as further refinements of the distinction between knowledge explicitly available to an agent and implicit knowledge ascribed to the agent.

An **algorithm-based** notion of knowledge has been proposed in [FHMV 95] (Chapter 10) and [Pa 87a] to capture explicit knowledge. This works as follows: at any state s, when asked whether a formula ϕ holds, the reasoner evokes an algorithm available at s. If the algorithm returns 'Yes', we say that the reasoner explicitly (or algorithmically) knows ϕ. Depending on whether the algorithm gives only 'Yes/No' outputs or whether it also has the possibility of a '?' output meaning 'I cannot find out within the resources available to me', we get the notions studied in [FHMV 95] or [Pa 87a], respectively. [FHMV 95] restricts the algorithm to be local, in the sense that when $s \sim s'$, the algorithm invoked by the reasoner at s is the same as at s'.

Unfortunately, the framework is also very general and it is hard to study properties of this notion of knowledge. How do we relate different algorithms used by different reasoners ? When we iterate modalities to say that i knows that j knows that α holds, should i have access to j's algorithm as well ? If the system moves from a state s to state s', how should the algorithms available at the two states be related ? Further, the invocation of an algorithm is an extra-logical notion and as it is, we have no way of studying algorithmic knowledge *logically*.

On the other hand, the algorithm-based formulation of explicit knowledge is interesting particularly because it avoids many of the philosophical pitfalls which cause criticism. One way of studying algorithmic knowledge logically is to decide that all reasoners uniformly use the same algorithm, namely the *model checking algorithm* for the logic. This means that when α is valid, the reasoner knows α to be true but does not know that it is valid. This is because α, being a valid formula, happens also to be true in all the states being checked by the algorithm, and thus is explicitly known, but then so are many other non-valid formulas which simply happen to be true in the checked model. (See [HV 91] for more general arguments advocating model checking rather than theorem proving in the context of reasoning about systems.) The framework facilitates the study of knowledge based on probabilistic algorithms ([KNP 90]) and action based on knowledge ([HF 89], [R96a]). This circumvents many of the difficulties mentioned above, but a requirement that all reasoners use the same knowledge algorithm seems too strong. It is here that a view-based notion of knowledge helps, whereby we can have a uniform logical notion of all reasoners invoking the same model checking algorithm, but they run the algorithm on different sub-models.

1.5 This paper

In this paper, we approach the problem by asking: if knowledge amounts to model checking, *which* model is to be checked ? We hypothesize that due to limited visibility, the agent can see only a substructure of the given model, and hence checks for truth in that substructure. Thus an agent claims that ϕ holds at a state if ϕ holds in the substructure visible to the agent at that state. We can then say that the agent claims to know ϕ if ϕ holds in the substructure at all the visible indistinguishable states. This gives a framework in which a form of algorithmic knowledge is studied logically, and which takes into account the limited visibility of agents. Thus we have three basic notions: truth of a proposition, ascribed knowledge of propositions (determined by indistinguishability relations) and claimed knowledge (determined by views). The third notion is dependent on the agent's view, and hence when the view changes, knowledge claimed by the agent changes as well.

2 The logic

We now present a logic of view-based explicit knowledge. This is simply the usual logic of knowledge augmented with a new modality: $\mathbf{X}\alpha$ will denote that the formula α holds explicitly (or visibly) for the reasoner. For explicit knowledge, we will use formulas of the form $\mathbf{XK}\alpha$. (As [FHMV 95] observe in Chapter 10, this should be seen as a reasoner claiming to know α rather than knowing α.)

2.1 Syntax

Formally, we have a countable set of propositions $P = \{p_0, p_1, \ldots\}$ as before and the syntax of formulas is given by:

$$\Psi ::= p \in P \mid \neg \alpha \mid \alpha \vee \beta \mid \mathbf{K}\, \alpha \mid \mathbf{X}\alpha$$

The fixed formula $p_0 \vee \neg p_0$ is called $True$. Φ, the set of formulas introduced earlier, is a subset of Ψ and formulas in Φ are referred to as $\mathbf{X}$-free formulas. The other logical connectives are defined as usual. The dual of $\mathbf{K}\alpha$, namely $\neg\mathbf{K}\neg\alpha$, is denoted $\mathbf{L}\alpha$, and the dual of $\mathbf{X}\alpha$, namely $\neg\mathbf{X}\neg\alpha$, is denoted $\mathbf{Y}\alpha$.

The frames are Kripke equivalence structures augmented with view functions.

Definition 1 *A* **view frame** *is a tuple $F = (W, \sim, \nu)$, where W is a set of worlds, $\sim \subseteq W \times W$ is an equivalence on W, and $\nu : W \to 2^W$ is the* view function.

A model is a tuple $M = (F, V)$, where $F = (W, \sim, \nu)$ is a frame, and $V : W \to 2^P$ is the valuation. We denote the equivalence class of w under $\sim$ by $[w]$.

When $w' \in \nu(w)$, we say that w' is *visible* from w. Note that there is no assumption that every world state is visible to the reasoner at that state. (Indeed, we might well have $\nu(w) = \emptyset$ for some w). A world w is said to be *recognizable* if $w \in \nu(w)$. A frame F is said to be recognizable if every world in F is. While recognizability may be too strong an assumption to make about every world, it may seem unintuitive to have worlds w, w' where $w \in \nu(w')$ but $w \notin \nu(w)$. Thus, even if every world is not recognizable, we might at least want every visible world to be recognizable: call a frame F *weakly recognizable* iff for every w, if there exists w' such that $w \in \nu(w')$ then $w \in \nu(w)$. On the other hand, we will call a frame F *almost recognizable* iff for every w, if $\nu(w) \neq \emptyset$, then $w \in \nu(w)$. We will show that our logic cannot distinguish between general view models and those based on almost recognizable frames.

The standard Kripke frames for knowledge can now be seen as view frames with constant universal view function: $\forall w, \nu(w) = W$.

When we have $w \sim w'$ and $w' \notin \nu(w)$, we may wish to interpret this as w and w' being logically indistinguishable, but computationally distinguishable. (This happens in situations where discovering that w and w' constitute the same evidence requires computation).

Definition 2 *Given a view frame $F = (W, \sim, \nu)$ and $W' \subseteq W$, the* restriction *of F to W' is given by: $F\lceil W' \stackrel{\text{def}}{=} (W', \sim', \nu')$, where $\sim' \stackrel{\text{def}}{=} \sim \cap (W' \times W')$ and $\forall w \in W'$, $\nu'(w) \stackrel{\text{def}}{=} \nu(w) \cap W'$. For a model $M = (F, V)$, we define $M\lceil W' \stackrel{\text{def}}{=} (F\lceil W', V')$, where for $w \in W'$, $V'(w) \stackrel{\text{def}}{=} V(w)$.*

The intended meaning of $\mathbf{X}\alpha$ at w should be clear now: it simply asserts that α holds in the substructure induced by $\nu(w)$ at w. We assume that all other connectives have the same meanings as before.

$$M, w \models \mathbf{X}\alpha \ \text{ iff } \ w \in \nu(w) \text{ and } M\lceil \nu(w), w \models \alpha$$

As usual, we say α is satisfiable if there exists a view model $M = ((W, \sim, \nu), V)$ such that for some $w \in W$, we have $M, w \models \alpha$. In addition, we say α is *explicitly satisfiable* if there exists a model $M = ((W, \sim, \nu), V)$ such that for some $w \in W$, we have $W = \nu(w)$ and $M, w \models \alpha$. A formula is (explicitly) valid if and only if its negation is not (explicitly) satisfiable. Note that α is explicitly satisfiable iff $\mathbf{X}\alpha$ is satisfiable, and α is explicitly valid iff $\mathbf{Y}\alpha$ is valid.

The first thing to note is that the modality does not collapse — in the sense that there are formulas α such that $\alpha \equiv \mathbf{X}\alpha$ is not valid. To see this, consider the view frame given by: $W = \{w_1, w_2\}$, $\sim = W \times W$ and $\nu(w_j) = \{w_j\}$ for $j \in \{1, 2\}$. Let M be a model on this frame with the valuation: $V(w_1) = \emptyset$ and $V(w_2) = \{p\}$. Clearly, we see that $M, w_1 \models (\mathbf{L}p \wedge \neg\mathbf{XL}p)$, and that $M, w_2 \models (\neg\mathbf{K}p \wedge \mathbf{XK}p)$.

Secondly, consider the dual modality: $\mathbf{Y}\alpha$ is a weaker assertion – at w it only means that if $w \in \nu(w)$ then $M\lceil\nu(w), w \models \alpha$, and hence can be true vacuously. However, we can check that the implication $\mathbf{X}\alpha \supset \mathbf{Y}\alpha$ is valid and that the implication becomes an equivalence over models based on recognizable frames.

Thirdly, though $\mathbf{X}\alpha \supset \alpha$ is not valid, $\alpha \equiv \mathbf{X}\alpha$ is explicitly valid, as we would expect. Further, the equivalences $\mathbf{XX}\alpha \equiv \mathbf{X}\alpha \equiv \mathbf{XY}\alpha$ and $\mathbf{YX}\alpha \equiv \mathbf{Y}\alpha \equiv \mathbf{YY}\alpha$ are valid. Thus we have a modality which is not like a traditional S5 modality and somewhat different from a usual belief modality. To see this, note that $\mathbf{K}\alpha \supset \alpha$ and $\neg\mathbf{K}\alpha \supset \mathbf{K}\neg\mathbf{K}\alpha$ are valid formulas, but $\mathbf{X}\alpha \supset \alpha$ and $\neg\mathbf{X}\alpha \supset \mathbf{X}\neg\mathbf{X}\alpha$ are not. But $\mathbf{X}\neg\mathbf{X}\alpha \supset \mathbf{X}\neg\alpha$ is valid, but such an assertion is not defensible if we interpret $\mathbf{X}$ as a belief modality, since an agent may not believe that α holds and be aware of such a belief, but may yet not believe that the denial of α holds either.

Like the $\mathbf{K}$ and $\mathbf{L}$ modalities, the $\mathbf{X}$ and $\mathbf{Y}$ modalities also collapse at the second level. We get nontrivial second level modalities only by looking at *alternations* of the knowledge and explicitness modalities like $\mathbf{XK}$, $\mathbf{XL}$ etc. Indeed, a formula like $\mathbf{XLXLXK}p$ may not in general be collapsible to any smaller formula. Therefore, we will be interested in the `modal alternation depth` of a formula, which will be formally defined later.

2.3 Substructures

How do we determine the *size* of a finite view frame ? This is important, for questions like *model checking*, where we are given a view model M, an element w in M and a formula α, and asked whether $M, w \models \alpha$. When this is a decidable problem, we are interested in the complexity of the decision procedure in terms of the input size, namely the size of M and the length of α. For a finite frame $F = (W, \sim, \nu)$, we might think of $\mid W \mid$ as the size of F, but this can really be misleading.

Consider a structure $W = \{w_0, w_1, \ldots\}$, where for all j, $\nu(w_j) = \{w_j, w_{j+1}\}$ and some model M based on it. Then for any modal formula α, the truth of α at w_j $(j > 0)$ in the structure $M\lceil\nu(w_j)$ is in general independent of the truth of the same formula at the same state in the structure $M\lceil\nu(w_{j-1})$.

Thus, for any formula α and worlds w_1, w_2 such that $w_2 \in \nu(w_1)$, we need to see whether $M\lceil\nu(w_1), w_2 \models \alpha$ holds. But there is still worse to come: when this formula

α is of the form $\mathbf{X}\beta$, this requires us to look at the substructure obtained by restricting to $(\nu(w_1) \cap \nu(w_2))$ which need not be of the form $\nu(w)$ for any $w \in W$. Therefore, in the worst case, we should be prepared to look at $M\lceil S, w \models \alpha$ for arbitrary nonempty subsets S of W (such that $w \in S$). This makes the logic hard to decide.

Can we find formulas that force us to look at substructures in the manner described above ? For instance, consider the frame $F = (W, \sim, \nu)$, where $W = \{w_1, w_2, w_3, w_4\}$, $\sim = W \times W$, and ν is given by: $\nu(w_1) = \{w_1, w_2, w_3\}$, $\nu(w_2) = \{w_2, w_3, w_4\}$, $\nu(w_3) = \{w_1, w_3, w_4\}$ and $\nu(w_4) = \{w_4\}$. Consider the valuation $V(w_3) = \{p\}$, and for all $w \neq w_3$, $V(w) = \emptyset$. Let $\delta = \mathbf{XLXLXK}p$. It can be checked that $M, w_1 \models \delta$, and for any proper modal suffix δ' of δ, for all $w \in W$, $M, w \not\models \delta'$. Evaluating this formula at w_1 requires evaluation of suffixes $\mathbf{LXLXK}p$ at $M_1 = M\lceil\nu(w_1)$, $\mathbf{LXK}p$ at $M_2^j = M_1\lceil\nu(w_j), j \in \{2, 3\}$ and $\mathbf{XK}p$ at $M_3 = M_2^2\lceil\nu(w_3)$, respectively.

In general, given a frame $F = (W, \sim, \nu)$, where $|W| = n$, there may be as many as $O(n^n)$ substructures of interest for checking truth of a formula. This crucially depends on how nasty the view function ν of the frame happens to be, as well as the modal alternation depth of the formula to be checked. For formulas, we have already pointed out that modal alternation depth is a parameter of interest. Similarly, the depth to which we need to look at substructures of substructures matters as well. Formally, we have the following:

Definition 3 *Let $F = (W, \sim, \nu)$ be a frame. A* **view tree** *of depth d on F is a labelled tree $\tau_F = (\Theta, \rightarrow, \lambda)$ of depth d, where Θ is the set of tree nodes with the distinguished element θ_0 as the root, $\rightarrow \subseteq (\Theta \times W \times \Theta)$ is the set of edges labelled by elements of W, and $\lambda : \Theta \rightarrow 2^W$ labels each node by a subset of W, satisfying the following conditions:*

(i) $\lambda(\theta_0) = W$.
(ii) If there are tree edges $\theta \xrightarrow{w} \theta'$ and $\theta \xrightarrow{w} \theta''$, then $\theta' = \theta''$.
(iii) If $\theta \xrightarrow{w} \theta'$ then $w \in \lambda(\theta') = (\lambda(\theta) \cap \nu(w)) \neq \lambda(\theta)$.
(iv) If θ is a node of depth $k < d$ and $w \in (\lambda(\theta) \cap \nu(w)) \neq \lambda(\theta)$, then there exists a tree edge $\theta \xrightarrow{w} \theta'$.

Condition (iii) is referred to as a *coherence condition*; we will see a need to weaken this condition in the next section.

If $F = (W, \sim, \nu)$ is a frame, it induces a view tree on F: the root is labelled W, and whenever there is a node labelled S and $w \in (S \cap \nu(w)) \neq S$, there exists an edge to a node labelled $S \cap \nu(w)$. All view trees upto a specific depth on the same frame are isomorphic, so we can speak of *the* d depth view tree associated with a frame.

3 Decidability

For convenience, we assume from now on that the constant $True \in \Psi$. Before we proceed to look for decision procedures for satisfiability for this logic, some observations simplifying the problem can be made. In any structure $F = (W, \sim, \nu)$, when we are interested in satisfying a given formula at a world $w_0 \in W$, we can simply restrict ourselves to the substructure induced by the equivalence class of w_0 under $\sim$. For logics of knowledge where there is only one knower, the accessibility relation may as well be universal. In fact, this observation is used to show that satisfiability is in NP. It is no surprise that such universal models suffice for this logic as well.

Definition 4 *A view frame $F = (W, \sim, \nu)$ is called a* **universal** *view frame iff $\sim = W \times W$.*

Proposition 5 *If α is satisfiable, then it is satisfiable in a universal model.*

PROOF. Given a formula α_0 satisfied at w_0 in model $M = ((W, \sim, \nu), V)$, define the new model $M' = M\lceil[w_0]$. Note that M' has the structure $M' = ((W', \sim', \nu'), V')$, where $W' = [w_0]_\sim$. Thus for any $w \sim w_0$, $w \in W'$ and $w \sim' w_0$, and M' is a model based on a universal frame.

We now show by induction (on the structure of α) that, for all α, $\forall w \in [w_0]$, the following statement holds:

– $\forall S \subseteq W$ such that $w \in S$, $M\lceil S, w \models \alpha$ if and only if $M'\lceil(S \cap [w_0]), w \models \alpha$.

The cases when α is a boolean formula are routine. Let $S \subseteq W$ and $w \in (S \cap [w_0])$. Suppose $M\lceil S, w \models \mathbf{K}\beta$, then $(\forall w' \in S$, if $w \sim w'$ then $M\lceil S, w' \models \beta)$. Now $w \in [w_0]$, hence $w' \in [w_0]$, and by induction hypothesis, $M'\lceil(S \cap [w_0]), w' \models \beta$. Since this holds for all $w' \sim' w$, we have $M'\lceil(S \cap [w_0]), w \models \mathbf{K}\beta$, as required. The other direction is similar.

Now suppose α is of the form $\mathbf{X}\beta$ and that $M\lceil S, w \models \mathbf{X}\beta$. Hence $w \in \nu(w)$ and $M\lceil S', w \models \beta$, where $S' = S \cap \nu(w)$. We then have $w \in S'$ as well, and by induction hypothesis, $M\lceil(S' \cap [w_0]), w \models \beta$. Now, $S' \cap [w_0] = (S \cap \nu(w)) \cap [w_0] = (S \cap [w_0]) \cap \nu(w)$, and $w \in \nu(w)$. Therefore, $M'\lceil(S \cap [w_0]), w \models \mathbf{X}\beta$, as required. The other direction is again similar.

This completes the induction. Now, putting $S = W$, since $M, w_0 \models \alpha_0$, we get $M'\lceil[w_0], w_0 \models \alpha_0$. But $M'\lceil[w_0] = M'$. Thus α_0 is satisfiable in the universal model M'.

Proposition 6 *If α is satisfiable, then it is satisfiable in an almost recognizable model.*

PROOF. Suppose that $M = ((W, \sim, \nu), V)$ is a model and $w_0 \in W$ such that $M, w_0 \models \alpha_0$. Define $M' = ((W, \sim, \nu'), V)$, where ν' is defined as follows: if $w \in \nu(w)$, then $\nu'(w) = \nu(w)$, and otherwise, $\nu'(w) = \emptyset$. Note that the view tree τ_F of any depth d is absolutely identical to the view tree $\tau_{F'}$ of the same depth. We can then easily show that for any node θ in this (same) tree and $w \in S = \lambda(\theta)$, $M\lceil S, w \models \alpha$ iff $M'\lceil S, w \models \alpha$, for any formula α. In particular this holds for $w_0 \in W = \lambda(\theta_0)$ and α_0. Thus α_0 is satisfied in the almost recognizable model M'.

In universal frames, we can simply drop the indistinguishability relation $\sim$ and work with only the view function ν. Unfortunately the logic is still hard to analyse. In the examples considered earlier, the frames were all universal, and yet we needed to look at substructures. However, we will show that the logic enjoys a *bounded pseudo-model property*, so that we can limit the size of W when we search for models for any given formula. This bound will depend on the modal alternation depth of a formula, defined as follows:

Definition 7 *The* **modal alternation depth** *of a formula α, denoted $\mu(\alpha)$ is defined as follows:*

- $\mu(p) \stackrel{\text{def}}{=} 0$ *for* $p \in P$.
- $\mu(\neg\alpha) \stackrel{\text{def}}{=} \mu(\alpha)$.
- $\mu(\alpha \vee \beta) \stackrel{\text{def}}{=} max\{\mu(\alpha), \mu(\beta)\}$.
- $\mu(\mathbf{K}\alpha) \stackrel{\text{def}}{=} \mu(\alpha)$.
- $\mu(\mathbf{X}\alpha)$ *is defined inductively:*
  - $\mu(\mathbf{X}p) = 0$ *for* $p \in P$.
  - $\mu(\mathbf{X}\neg\beta) \stackrel{\text{def}}{=} \mu(\mathbf{X}\beta)$.
  - $\mu(\mathbf{X}(\beta_1 \vee \beta_2)) \stackrel{\text{def}}{=} max\{\mu(\mathbf{X}\beta_1), \mu(\mathbf{X}\beta_2)\}$.
  - $\mu(\mathbf{XK}\beta) \stackrel{\text{def}}{=} \mu(\beta) + 1$.
  - $\mu(\mathbf{XX}\beta) \stackrel{\text{def}}{=} \mu(\mathbf{X}\beta)$.

Note that for any α, $\mu(\alpha) \leq \mu(\mathbf{X}\alpha) \leq \mu(\alpha) + 1$. The following proposition suggests how μ proves useful.

Proposition 8 *Let $F = (W, \sim, \nu)$ be a frame and $\tau_F = (\Theta, \rightarrow, \lambda)$ be a view tree on F. Let $M = (F, V)$, $\theta \in \Theta$, $S = \lambda(\theta)$, $w \in (S \cap \nu(w))$ and $\mathbf{X}\alpha$ a formula such that $\mu(\mathbf{X}\alpha) = \mu(\alpha)$. Then $M\lceil S, w \models \mathbf{X}\alpha$ iff $M\lceil S, w \models \alpha$.*

PROOF. The proof is by induction on α. When $\alpha = p \in P$, the assertion follows from the fact that $V\lceil S = V\lceil(S \cap \nu(w))$. The boolean cases follow from the validity of identities $\mathbf{X}\neg\beta \equiv (\mathbf{X}True \wedge \neg\mathbf{X}\beta)$ and $\mathbf{X}(\beta_1 \vee \beta_2) \equiv (\mathbf{X}\beta_1 \vee \mathbf{X}\beta_2)$. α cannot be of the form $\mathbf{K}\beta$ since $\mu(\mathbf{XK}\beta) \neq \mu(\beta)$. If α is of the form $\mathbf{X}\beta$, the assertion follows from the validity of the identity $\mathbf{XX}\beta = \mathbf{X}\beta$.

Definition 9 *The* subformula closure *of a formula α, denoted $SF(\alpha)$ is defined in the usual manner: $SF'(\alpha)$ is the least set containing α and satisfying the following conditions — if $Op\ \beta \in SF'(\alpha)$ then $\beta \in SF'(\alpha)$, where $Op \in \{\neg, \mathbf{K}, \mathbf{X}\}$; if $\beta_1 \vee \beta_2 \in SF'(\alpha)$ then $\{\beta_1, \beta_2\} \subseteq SF'(\alpha)$; $\mathbf{X}True \in SF'(\alpha)$. We then define $SF(\alpha) \stackrel{\text{def}}{=} SF'(\alpha) \cup \{\neg\beta \mid \beta \in SF'(\alpha)\}$.*

Note that the size of $SF(\alpha)$ is linear in the length of α, when we treat $\neg\neg\beta$ to be the same as β. Note that for every $\beta \in SF(\alpha_0)$, $\mu(\beta) \leq \mu(\alpha_0)$.

3.1 Checking truth

We first consider the following problem: given a finite view model M, a world w in it, and a formula α, is it the case that $M, w \models \alpha$?

Theorem 10 *Given a finite model $M = (F, V)$ where $F = (W, \sim, \nu)$, $\mid W \mid = n$, $w_0 \in W$, and a formula α_0 of length m, checking whether $M, w_0 \models \alpha_0$ holds can be done in time $O(n^{d+1}.m)$, where $d = \mu(\alpha_0)$.*

PROOF. Assume M, w_0 and α_0 as in the statement of the theorem. Assume a presentation of F such that for any w_1, w_2, checking $w_1 \in \nu(w_2)$? or for a proposition p, checking $p \in V(w_1)$? is efficient. Let $\tau_F = (\Theta, \rightarrow, \lambda)$ be the associated view tree of depth d; note that $\mid \Theta \mid = N = O(n^d)$. Fix an enumeration of Θ, say $\theta_1, \theta_2, \ldots, \theta_N$ such that depth k nodes precede depth $k-1$ nodes in the enumeration. Fix an enumeration of $SF(\alpha_0)$ (hereafter referred to as SF for the rest of this proof) to be $\gamma_1, \gamma_2, \ldots, \gamma_m$ such that if γ_i precedes γ_j in the enumeration, then either $(\mu(\gamma_i) < \mu(\gamma_j))$, or $(\mu(\gamma_i) = \mu(\gamma_j)$ and $\mid \gamma_i \mid \leq \mid \gamma_j \mid)$. Consider the lexicographic ordering on pairs (θ, γ) defined by the enumeration of nodes in τ_F and formulas in SF.

We describe a labeling algorithm, that labels every pair (θ, w), where $\theta \in \Theta$ and $w \in \lambda(\theta)$ with $L(\theta, w) \subseteq SF$. Initially all the $L(\theta, w)$ are empty and we consider a formula for adding as follows: α_1 is considered for $L(\theta_1, w_1)$ earlier than α_2 for $L(\theta_2, w_2)$ iff (θ_1, α_1) precedes (θ_2, α_2) in the lexicographic ordering.

- Add $p \in P$ to $L(\theta, w)$ iff $p \in V(w)$ independent of S.
- Add $\neg\alpha$ to $L(\theta, w)$ iff $\alpha \notin L(\theta, w)$.
- Add $\alpha \vee \beta$ to $L(\theta, w)$ iff $\alpha \in L(\theta, w)$ or $\beta \in L(\theta, w)$.
- Add $\mathbf{K}\alpha$ to $L(\theta, w)$ iff for every $w' \in \lambda(\theta)$ such that $w \sim w'$, $\alpha \in L(\theta, w')$.
- Add $\mathbf{X}\alpha$ to $L(\theta, w)$ iff $w \in \nu(w)$ and
  - θ is a leaf node of τ_F and $\alpha \in L(\theta, w)$ or,
  - $\alpha \in L(\theta', w)$, where $\theta \xrightarrow{w} \theta'$ is an edge in τ_F.

Finally the algorithm terminates saying 'Yes' if $\alpha_0 \in L(W, w_0)$ and 'No' otherwise. A straightforward inductive argument establishes that for all $\theta \in \Theta$ and $w \in \lambda(\theta) = S$, and for every formula $\gamma \in SF$ such that $\mu(\gamma) \leq \mu(\alpha_0) - depth(\theta)$,

$$\gamma \in L(\theta, w) \text{ iff } M\lceil S, w \models \gamma.$$

In the only interesting case of the induction step when γ is of the form $\mathbf{X}\alpha$, note that θ is a leaf node of τ_F iff depth of θ is d or for every $w \in S = \lambda(\theta)$, if $w \in \nu(w)$ then $S = \nu(w)$. In the former case, $\mu(\gamma) = 0$, then $\mu(\mathbf{X}\alpha) = \mu(\alpha)$, hence by Proposition 8, $M\lceil S, w \models \mathbf{X}\alpha$ iff $M\lceil S, w \models \alpha$, and the induction hypothesis gives the required assertion. In the latter case, $S \cap \nu(w) = S$, hence $M\lceil S, w \models \mathbf{X}\alpha$ iff $M\lceil S, w \models \alpha$ iff (by induction hypothesis) $\alpha \in L(\theta, w)$ iff (by construction) $\mathbf{X}\alpha \in L(\theta, w)$. The case when S is not a leaf node also follows easily from the induction hypothesis and the construction.

Thus $M, w_0 \models \alpha_0$ iff $\alpha_0 \in L(\theta_0, w_0)$, since $M\lceil W = M$, thus establishing the correctness of the algorithm. It can be easily seen that the running time is $O(n^d.n.m)$, since the required checking at each node takes $O(n.m)$ time.

3.2 Satisfiability

Fix a formula α_0 and let SF denote $SF(\alpha_0)$. This is the formula whose satisfiability we will be concerned with:

Definition 11 *We say $A \subseteq SF$ is* an atom *iff it satisfies the following conditions:*

- $\forall \beta \in SF$, $\neg\beta \in A$ *iff* $\beta \notin A$.
- $\forall \beta_1 \vee \beta_2 \in SF$, $\beta_1 \vee \beta_2 \in A$ *iff* $\beta_1 \in A$ *or* $\beta_2 \in A$.
- $\forall \mathbf{K}\beta \in SF$, *if* $\mathbf{K}\beta \in A$, *then* $\beta \in A$.
- *If there is a formula* $\mathbf{X}\beta \in A$ *then* $\mathbf{X}True \in A$.
- *For every* $\mathbf{X}\beta \in SF$ *such that* $\mu(\mathbf{X}\beta) = \mu(\beta)$, $\mathbf{X}\beta \in A$ *iff* $\beta \in A$.

Let AT denote the set of all atoms contained in SF, and use $A, B, \ldots$ to denote atoms. We say $A \sim B$ iff for every formula of the form $\mathbf{K}\alpha$, $\mathbf{K}\alpha \in A$ iff $\mathbf{K}\alpha \in B$. (Though we use the same symbol $\sim$ for the equivalence relation on frames as well as for this relation, the usage will be clear from the context.) We say $A \preceq B$ iff for every $\mathbf{X}\alpha \in SF$, $\alpha \in A$ iff $\mathbf{X}\alpha \in B$.

Let $Q \subseteq 2^{AT}$. We call Q a *good* set iff for every $A, B \in Q$, $A \sim B$, and if $\mathbf{L}\beta \in A$ then for some $C \in Q$, $\beta \in C$. We say that Q is *good for α* if Q is good and there exists $A \in Q$ such that $\alpha \in A$.

Proposition 12 *An* $\mathbf{X}$*-free formula* $\alpha \in SF$ *is satisfiable iff* AT *contains a subset which is good for* α.

PROOF. For any formula α, if it is satisfiable, let $M = ((W, \sim, \nu), V)$ and $w \in W$ such that $M, w \models \alpha$. Consider the set $\{\mu(w') \mid w' \sim w\}$, where $\mu(w') = \{\beta \in SF \mid$

$M, w' \models \beta\}$. It is easy to see that each element of the set is an atom and that the set is good.

On the other hand, let $\alpha \in SF$ be $\mathbf{X}$-free and let $Q \subseteq AT$ be good with $A \in Q$, α in A. Define the frame $F = (Q, \sim, \nu)$, where for all $B \in Q$, $\nu(B) = Q$. Define $V(B) = B \cap P$. This gives a model M and by an easy induction on the structure of γ, we can show that for any $B \in Q$, $M, B \models \gamma$ iff $\gamma \in B$. We thus get $M, A \models \alpha$, hence α is satisfiable.

We now come to the main construction for checking satisfiability. The idea is (as usual) to build a finite model for any satisfiable formula α, in such a way that the size of the model is bounded by a function of $\mid \alpha \mid$. If this finite model is constructed in some canonical fashion, then checking for satisfiability amounts to an attempt at building this model; the attempt succeeds iff the formula is satisfiable. As it turns out, we will work with structures (called witness structures below) that do not quite define models, but act as pseudo-models from which we can pull out models.

Given α_0, we try to construct a model $((W, \sim, \nu), V)$. If α_0 was $\mathbf{X}$-free, we only need a subset Q of AT which is good for α_0; then we can set $\sim = (Q \times Q)$ and $\nu(w) = Q$ for all $w \in Q$. The idea is that $\alpha \in w$ represents the fact that α holds at w in the constructed model. When we have $\mathbf{X}$-formulas in SF, a formula like $\mathbf{L}p \wedge \mathbf{XK}\neg p$ is satisfiable, and it becomes necessary to have $\nu(w) \neq Q$. This creates a new problem: suppose $w' \in \nu(w) \subseteq Q$: different formulas are satisfied at w' in the model M being constructed and in $M\lceil\nu(w)$. Hence we cannot simply use sets of formulas as elements of Q. We need to build a view tree based on Q and for each node θ in the tree, we need a good subset such that $\alpha \in w' \in S = \lambda(\theta)$ represents the truth of α at w' in $M\lceil S$.

A **partial view tree** (of depth d) on $F = (W, \sim, \nu)$ is a labelled tree τ_F similar to a view tree, except that instead of the coherence condition, it satisfies a weaker condition:

- If $\theta \xrightarrow{w} \theta'$ then either ($\theta = \theta_0$ and $\lambda(\theta') = \nu(w) \neq W$) or ($w \in \lambda(\theta') \subseteq (\lambda(\theta) \cap \nu(w)) \neq \lambda(\theta)$).

Note that every view tree on a frame is a partial view tree on that frame as well. Also, nodes at depth 1 in partial view trees exactly specify ν, as in view trees. Thus for $w \in W$ such that $w \in \nu(w)$, either $\nu(w) = W = \lambda(\theta_0)$, or there is a depth 1 node θ_w in τ with $\lambda(\theta_w) = \nu(w)$. We use the notation θ_0 for the root, θ_w for this node (θ_0 or a unique depth 1 node), $\tau\lceil w$ to denote the subtree of τ rooted at θ_w, and $\Theta\lceil w$ for the set of nodes in this subtree.

When τ_F is a partial view tree on F of depth d, we say that τ_F has a (θ, w') `requirement at` w iff the node θ is in $\tau_F\lceil w$, depth of θ is $k < d$ and there is an edge $\theta \xrightarrow{w'} \theta'$ such that $(\lambda(\theta) \cap \nu(w')) \neq \lambda(\theta')$. Clearly, a partial view tree of depth d on a frame F is a view tree on F iff it has no requirements at any w in it.

Definition 13 *Let $F = (W, \sim, \nu)$ be a frame, and let $\tau_F = (\Theta, \rightarrow, \lambda)$ be a partial view tree on F of depth $d = \mu(\alpha_0)$. A* **witness structure** *on F for α_0 based on τ_F is a map $T : (\Theta \times W) \rightarrow AT$, defined for all pairs (θ, w) such that $w \in \lambda(\theta)$, satisfying the following conditions:*

(i) For every $\theta \in \Theta$ and $w \in \lambda(\theta)$, the following conditions hold:
  (a) $T(\theta, w) \cap P = T(\theta_0, w) \cap P$.
  (b) $\mathbf{X}True \in T(\theta, w)$ iff $w \in \nu(w)$.
  (c) $\{T(\theta, w') \mid w' \in \lambda(\theta)\}$ is a good set.
  (d) If $w \in \nu(w)$ then either $(T(\theta, w) \preceq T(\theta, w))$ or there is an edge $\theta \xrightarrow{w} \theta'$ such that $T(\theta', w) \preceq T(\theta, w)$.

(ii) There exists $w_0 \in W$ such that $\alpha_0 \in T(\theta_0, w_0)$.

When T is a witness structure on a frame F based on τ_F and τ_F has a requirement, then we usually also say that T has that requirement in F, and call T a **partial witness structure** in this case. Clearly, a witness structure T based on a view tree has no requirements, and in that case, we call T a **complete witness structure**.

Let T be a complete witness structure for α_0 on F. Define the model $M_T = (F, V_T)$, where $V_T(w) = T(\theta_0, w) \cap P$. The following lemma shows that $M_T, w_0 \models \alpha_0$, since $M_T \lceil (\lambda(\theta_0)) = M_T$.

Lemma 14 *Let $F = (W, \sim, \nu)$ be a frame, let $\tau_F = (\Theta, \rightarrow, \lambda)$ be a view tree on F of depth $d = \mu(\alpha_0)$ and let $T : (\Theta \times W) \rightarrow AT$ be a complete witness structure on F based on τ_F. Then for every $\theta \in \Theta$ and $w \in S = \lambda(\theta)$, for every formula $\alpha \in SF$ such that $\mu(\alpha) \leq (d - depth(\theta))$, $\alpha \in T(\theta, w)$ iff $M_T \lceil S, w \models \alpha$.*

PROOF. The proof is by induction on the structure of α. When $\alpha = p \in P$, $p \in T(\theta, w)$ iff $p \in V(w)$ iff $M, w \models p$ iff $M \lceil S, w \models p$ since $V(w) = (V \lceil S)(w)$, for $w \in S$. The cases when α is boolean are routine.

Now suppose that $\mathbf{K}\alpha \in T(\theta, w)$. To show that $M \lceil S, w \models \mathbf{K}\alpha$, consider $w' \in (S \cap [w])$; it suffices to show that $M \lceil S, w' \models \alpha$, which would follow by the induction hypothesis if $\alpha \in T(\theta, w')$. But $T(\theta, w) \sim T(\theta, w')$; hence $\mathbf{K}\alpha \in T(\theta, w')$ and since $T(\theta, w')$ is an atom, $\alpha \in T(\theta, w')$, as required.

Conversely, if $M \lceil S, w \models \mathbf{K}\alpha$ and $\mathbf{K}\alpha \notin T(\theta, w)$, we have $\mathbf{L}\neg\alpha \in T(\theta, w) \cap SF$. By condition 1(c) on witness structures and definition of a good subset of AT, there exists $w' \in (\lambda(\theta) \cap [w])$ such that $\neg\alpha \in T(\theta, w')$. Hence $\alpha \notin T(\theta, w')$. By induction hypothesis, $M \lceil S, w' \not\models \alpha$, contradicting the fact that $M \lceil S, w \models \mathbf{K}\alpha$, and we are done.

Now let $\mathbf{X}\alpha \in T(\theta, w)$. Then $\mathbf{X}True \in T(\theta, w)$ and by condition 1(b) on witness structures, $w \in \nu(w)$. Let $S' = (S \cap \nu(w))$. If $S' = S$, then there exists no edge labelled w out of θ, and by condition 1(d) on witness structures, $T(\theta, w) \preceq T(\theta, w)$, hence $\alpha \in T(\theta, w)$. By induction hypothesis, $M\lceil S, w \models \alpha$. But $S = (S \cap \nu(w))$, hence $M\lceil (S \cap \nu(w)), w \models \alpha$, so $M\lceil S, w \models \mathbf{X}\alpha$, as required. On the other hand, if $S' \neq S$, we have two cases: either $depth(\theta) = d$, or there exists an edge $\theta \xrightarrow{w} \theta'$. In the latter case, by condition 1(d), $T(\theta, w') \preceq T(\theta, w)$, and hence $\alpha \in T(\theta', w)$. Since $\lambda(\theta') = (S \cap \nu(w))$, by induction hypothesis we get $M\lceil (S \cap \nu(w)), w \models \alpha$, and again, $M\lceil S, w \models \mathbf{X}\alpha$, as required. When $depth(\theta) = d$, $\mu(\mathbf{X}\alpha) = 0 = \mu(\alpha)$ and since $T(\theta, w)$ is an atom, $\alpha \in T(\theta, w)$. By induction hypothesis, $M\lceil S, w \models \alpha$. But then, by Proposition 8, $M\lceil S, w \models \mathbf{X}\alpha$, as required.

On the other hand, when $M\lceil S, w \models \mathbf{X}\alpha$, then we can proceed in a similar manner to show that $\mathbf{X}\alpha \in T(\theta, w)$.

Thus, if we can build a view tree and a (complete) witness structure for a formula on it, then that formula is satisfiable. On the other hand, whenever we are given a model for a formula, we can build a partial view tree on a frame of bounded size and a (partial) witness structure for that formula on it.

Lemma 15 *If α_0 is satisfiable, then there exists a finite frame $F = (W, \sim, \nu)$ and a partial witness structure for α_0 on F such that $\mid W \mid \leq 2^{O(md)}$, where $m = \mid \alpha_0 \mid$ and $d = \mu(\alpha_0)$.*

PROOF. Suppose that α_0 is satisfiable. By Proposition 5, we may assume that it is satisfied in a universal model. Let $M = (F, V)$ be the given model where $F = (W, \nu)$ is a universal frame, and for some $w_0 \in W$, $M, w_0 \models \alpha_0$. Let $d = \mu(\alpha_0)$. By Proposition 6, we can also assume that for all $w \in W$, if $\nu(w) \neq \emptyset$, then $w \in \nu(w)$.

Let $\tau_F = (\Theta, \rightarrow, \lambda)$ denote the view tree associated with F of depth d. Consider the map $T : (\Theta \times W) \rightarrow 2^{AT}$ defined by: $T(\theta, w) \stackrel{\text{def}}{=} \{\gamma \in SF \mid M\lceil \lambda(\theta), w \models \gamma\}$. It is easy to check that for all $\theta \in \Theta$ and $w \in \lambda(\theta)$, $T(\theta, w)$ is an atom, and that T is a witness structure on τ_F. However, T is not the witness structure required to prove the lemma, as W may be infinite.

Consider the following equivalence relation on W: $w \equiv w'$ iff for all $\theta \in \Theta$ such that $depth(\theta) \leq d$ and $\{w, w'\} \subseteq \lambda(\theta)$, $T(\theta, w) = T(\theta, w')$. In particular, when $w_1 \equiv w_2$, $T(\theta_0, w_1) = T(\theta_0, w_2)$, and hence $V(w_1) \cap SF = V(w_2) \cap SF$. $\equiv$ gives us a quotiented structure. (Below, we use the notation $[w]$ to denote equivalence classes under $\equiv$; since we have dispensed with the use of $\sim$, this should cause no confusion.) Formally, $\overline{M} \stackrel{\text{def}}{=} (\overline{F}, \overline{V})$, where $\overline{F} = (\overline{W}, \overline{\nu})$ is defined as follows: $\overline{W} \stackrel{\text{def}}{=} \{[w] \mid w \in W\}$ and $\overline{\nu}([w]) \stackrel{\text{def}}{=} \{[w'] \mid w' \in \nu(w)\}$. $\overline{V}([w]) \stackrel{\text{def}}{=} V(w) \cap SF$. Note that $[w] \in \overline{\nu}([w])$ iff $w \in \nu(w)$, thanks to the assumption on the given model, that for any w', if $w' \notin \nu(w')$ then $\nu(w') = \emptyset$. Since the size of AT is $2^{O(m)}$, where $m = \mid \alpha_0 \mid$, it is easy to check that $\mid \overline{W} \mid \leq 2^{O(md)}$.

For $W' \subseteq W$, let $\overline{W'}$ denote the set $\{[w'] \mid w' \in W'\}$. Whenever $W_1 \subseteq W_2 \subseteq W$, we have $\overline{W_1} \subseteq \overline{W_2} \subseteq \overline{W}$ as well.

Let θ be a node in τ_F and let $S \subseteq \lambda(\theta)$. If $w \in (S \cap \nu(w))$, it is easy to see that $\overline{S} \cap \overline{\nu}([w]) \subseteq \overline{S \cap \nu(w)}$, though not necessarily the other way around. This is because, we can have $w_1 \in (S - \nu(w))$ and $w_2 \in (\nu(w) - S)$ such that $w_1 \equiv w_2$, and hence $[w_1] = [w_2] = x \in (\overline{S} \cap \overline{\nu}(w))$, but there need be no $w' \in (S \cap \nu(w))$ such that $[w'] = x$. Thus, in effect, when we 'quotient' τ_F, we get a partial view tree. We make this precise now.

Construct the tree $\overline{\tau} = (\Delta, \Rightarrow, \overline{\lambda})$ and a map $\phi : \Delta \to \Theta$ defined as follows: $\overline{\lambda}(\delta_0) = \overline{W}$ and $\phi(\lambda_0) = \theta_0$, where δ_0 is the root. Inductively assume that we have a node δ of depth $k < d$ such that $\phi(\delta) = \theta$, $X = \overline{\lambda}(\delta) = \overline{S}$, where $S = \lambda(\theta)$. Now, suppose there exists an edge $\theta \xrightarrow{w} \theta'$ in τ_F, but there is no edge labelled $[w]$ out of δ. We have $[w] \in (X \cap \overline{\nu}([w]))$. Add a node δ' in Δ, and an edge $\delta \stackrel{[w]}{\Rightarrow} \delta'$, with $\overline{\lambda}(\delta') = \overline{\lambda(\theta')}$. Set $\phi(\delta') = \theta'$.

We claim that $\overline{\tau}$ is a partial view tree on the frame $\overline{F}$. It is easy to see that the first three conditions defining partial view trees are satisfied. Now, suppose that we have a node δ of depth $k < d$ and $x \in (X \cap \overline{\nu}(x)) \neq X$, where $X = \overline{\lambda}(\delta)$. We need to show that there is an edge labelled x out of δ in $\overline{\tau}$. Let $\phi(\delta) = \theta$. It suffices to show that there exists an edge $\theta \xrightarrow{w} \theta'$ in τ_F, where $x = [w]$ and $w \in S$, by construction. Suppose there exists no such edge in τ_F. Note that $w \in \nu(w)$ since $x \in \overline{\nu}(x)$. We have $S \subseteq \nu(w)$, and hence $\overline{S} \subseteq \overline{\nu(w)} = \overline{\nu}([w])$. But $X = \overline{S}$, and we get $X \subseteq \overline{\nu}(x)$, contradicting our assumption that $(X \cap \overline{\nu}(x)) \neq X$. This establishes our claim.

Define $\overline{T} : (\Delta \times \overline{W}) \to 2^{AT}$ as follows: let $\delta \in \Delta$ and $x \in \overline{\lambda}(\delta)$. Let $\phi(\delta) = \theta$ and $S = \lambda(\theta)$. Then there exists $w \in S$ such that $x = [w]$. Define $\overline{T}(\delta, x) \stackrel{\text{def}}{=} T(\theta, w)$. This is well-defined, because for all $w' \in S$ such that $x = [w']$, $T(\theta, w) = T(\theta, w')$, by definition of $\equiv$.

We now claim that $\overline{T}$ is a witness structure on $\overline{F}$ for α_0 based on the partial view tree $\overline{\tau}$. This is proved easily: observe that when the set $\{T(\theta, w) \mid w \in \lambda(\theta)\}$ is good, then so is the set $\{\overline{T}(\delta, [w]) \mid [w] \in \overline{\lambda}(\delta)\}$ when $\phi(\delta) = \theta$. Further, when there is an edge $\theta \xrightarrow{w} \theta'$ in τ_F, we get an edge $\delta \xrightarrow{[w]} \delta'$ correspondingly in $\overline{\tau}$ and $\overline{T}(\delta', [w]) = T(\theta', w) \preceq T(\theta, w) = \overline{T}(\delta, [w])$.

We thus have a witness structure for α_0 of bounded size on a partial view tree, proving the lemma.

We next show that whenever we have a partial witness structure T for α_0 on a finite frame F, we can `extend` it to a (possibly infinite) frame F' with a complete witness structure T' on it. Once we show this, Lemmas 14 and 15 then imply that a formula is satisfiable iff there exists a partial witness structure for it on a bounded size frame. To prove this, we first need to make the notion of 'extending' frames precise.

Let T be a partial witness structure on a frame $F = (W, \nu)$ and $\tau_F = (\Theta, \rightarrow, \lambda)$, a partial view tree associated with it. Let T' be a partial witness structure on frame F' where $F' = (W', \nu')$ and $\tau_{F'} = (\Theta', \rightarrow', \lambda')$ is a partial view tree associated with F' such that $W \subseteq W'$. Let $w \in W$ such that $\nu(w) = \nu'(w)$. We say that T' **extends** T at w iff $T(\theta_0, w) = T'(\theta'_0, w)$ and if $w \in \nu(w)$, then $\tau_F \lceil w$ is isomorphic to $\tau_{F'} \lceil w$, given by a bijection $f : \Theta\lceil w \rightarrow \Theta'\lceil w$ such that for all $\theta \in \Theta\lceil w$, $\lambda(\theta) = \lambda'(f(\theta))$, and for all $w' \in \lambda(\theta)$, $T(\theta, w') = T'(f(\theta), w')$.

We now show how to extend partial view trees to view trees. The strategy is to pick one w, fulfill all requirements at w and then move to the next. A requirement in a partial view tree is typically an edge $\theta \xrightarrow{w'} \theta'$, in the subtree rooted at θ_w, where $S = (\lambda(\theta) \cap \nu(w')) - \lambda(\theta') \neq \emptyset$. Let $w_1 \in S$. We extend the frame to a larger one by adding a new element w_2 and replacing $\nu(w')$ by $(\nu(w') - \{w_1\}) \cup \{w_2\}$ in the new frame. If we do this for every element in S, in the resulting frame, this particular requirement does not exist. This construction is repeated at all levels in the subtree rooted at w, giving a new structure without any requirements at w.

Suppose that we have fulfilled all w_1-requirements. Then how can we be sure that while fulfilling some w_2-requirement, we will not generate new w_1-requirements ? Note that in the construction suggested above, while fulfilling requirements at w, a new requirement may be generated for w' only if $w' \in \nu(w)$. Therefore, once we fulfill all w_1-requirements, if we have that $w_1 \notin \nu(w_2)$ for any w_2 that may be considered later, we can ensure that no new w_1-requirements will arise.

The following proposition asserts that we can always "relabel" elements in a frame and retain a partial witness structure in such a way that a designated world w is not visible from any other world in the resulting frame, except possibly from itself.

Proposition 16 *Given a partial witness structure T on a finite frame $F = (W, \sim, \nu)$, and $w \in W$, there exists a partial witness structure T' on a frame $F' = (W', \sim', \nu')$ such that $W \subseteq W'$, $\nu(w) = \nu'(w)$, for all $w' \in (W' - \{w\})$, $w \notin \nu'(w')$ and T' extends T at w.*

PROOF. Fix w and let $w' \neq w$ such that $w \in \nu(w')$. Pick $x \notin W$ and set $W' = W \cup \{x\}$. Define $\nu'(x) = \emptyset$ if $\nu(w)$ is empty, and otherwise let $\nu'(x) = (\nu(w) - \{w\}) \cup \{x\}$. For $w'' \in W - \{w\}$, if $w \in \nu(w'')$ then define $\nu'(w'') = (\nu(w'') - \{w\}) \cup \{x\}$; otherwise set $\nu'(w'') = \nu(w'')$. Define $\sim' = \sim \cup \{(x, x)\} \cup \{(x, w''), (w'', x) \mid w \sim w''\}$.

We need to define T': construct $\tau_{F'} = (\Theta', \rightarrow', \lambda')$ as follows. The root θ'_0 is labelled W'. Define $T'(\theta'_0, w'') = T(\theta_0, w'')$ for $w'' \in W$ and $T'(\theta'_0, x) = T(\theta_0, w)$. Let $\theta'_1 \in \Theta'$. Then there exists $\theta_1 \in \Theta$ such that $\lambda(\theta_1) = \lambda'(\theta'_1)$ or $\lambda'(\theta'_1) = (\lambda(\theta_1) - \{w\}) \cup \{x\}$. Suppose there is an edge $\theta_1 \xrightarrow{w''} \theta_2$ in τ_F. Then add an edge $\theta'_1 \xrightarrow{y} \theta'_2$ where the label $y = w''$ if $w'' \neq w$, and the label is $y = x$, otherwise. Set $\lambda'(\theta'_2) = \lambda(\theta_2)$ or $\lambda'(\theta'_2) = (\lambda(\theta_2) - \{w\}) \cup \{x\}$, depending on the corresponding case with θ_1. Set $T'(\theta'_2, w') = T(\theta_2, w')$ if $w' \in W$

and $T'(\theta_2', x) = T(\theta_2, w)$, otherwise.

It can be checked that T' is a partial witness structure on F' as required.

Lemma 17 *Given a partial witness structure T on a finite universal frame $F = (W, \nu)$ and $w \in W$ such that T has some requirements at w, there exists a partial witness structure T' on a finite universal frame $F' = (W', \nu')$ such that $W \subseteq W'$, for all $w' \in (W - \nu(w)) \cup \{w\}$, $\nu'(w') = \nu(w')$, T' extends T at w', and T' has* no *requirements at w.*

PROOF. Suppose that T has some requirements at $\widehat{w}$. Fix an enumeration of nodes in $\tau_F \lceil \widehat{w}$ such that nodes at depth k precede nodes at depth $k+1$. Let (θ_1, w_1), ..., (θ_m, w_m) be an enumeration of requirements at $\widehat{w}$. We define a sequence of frames $F_0, F_1, \ldots, F_m$, where $F_j = (W_j, \nu_j)$ and maps $T_0, T_1, \ldots, T_j$ such that for $1 \leq j \leq m$, for every $w' \in (W - \nu(w)) \cup \{w\}$, $\nu_j(w') = \nu_{j-1}(w')$, T_j extends T_{j-1} at w' via the bijection f_j and there is no $(f_j(\theta_j), w_j)$ requirement at w in F_j.

Define F_0 by: $W_0 = W$, $\nu_0 = \nu$ and $T_0 = T$. Given F_{j-1} and T_{j-1}, $j > 0$, check if there is a (θ_j, w_j) requirement at $\widehat{w}$ in F_{j-1}. If there is no such requirement, set $F_j = F_{j-1}$ and $T_j = T_{j-1}$, with f_j being the identity map. Otherwise, consider the edge $\theta_j \xrightarrow{w_j} \theta'$ in F_{j-1} and the set $S = (\lambda_{j-1}(\theta_j) \cap \nu_{j-1}(w_j)) - \lambda_{j-1}(\theta')$. Since this is a requirement, $S \neq \emptyset$. Let $S = \{u_1, \ldots, u_k\}$. Pick a set $X = \{x_1, \ldots, x_k\}$ such that $X \cap W_{j-1} = \emptyset$, and define $g : S \to X$ by: $g(u_i) = x_i$, for $1 \leq i \leq k$. Define F_j by $W_j = W_{j-1} \cup X_j$, and ν_j is defined as follows: let $S' = \nu(w) - \{w\}$; for $w' \in (W - S')$, define $\nu_j(w') = \nu_{j-1}(w')$; for $w' \in S'$, define $\nu_j(w') = (\nu_{j-1}(w') - S) \cup \{g(u_l) \mid u_l \in (S \cap \nu_{j-1}(w'))\}$. For $w' \in X$, define $\nu_j(x_i) = (\nu_{j-1}(u_i) - S) \cup \{g(u_l) \mid u_l \in (S \cap \nu_{j-1}(u_i))\}$.

Now consider τ_j, and a map $\phi : \Theta_j \to \Theta_{j-1}$ defined in the obvious manner: the root θ_0^j is labelled W_j and $\phi(\theta_0^j) = \theta_0^{j-1}$. For all $w' \in W_{j-1}$, if $w' \in \nu(w')$, τ_j contains a subtree isomorphic to $\tau_{j-1} \lceil w'$ with labels changed as follows: given a node θ^j in $\tau_j \lceil w'$ whose image is θ^{j-1} in τ_{j-1}, if $\lambda_{j-1}(\theta^{j-1}) = R$, and $R \cap S' = \emptyset$, then $\lambda_j(\theta^j) = R$; in this case, set $T_j(\theta^j, w') = T_{j-1}(\theta^{j-1}, w')$, for $w' \in R$. Otherwise, $\lambda_j(\theta^j) = (R - S) \cup \{g(u_l) \mid u_l \in \lambda_{j-1}(\theta^{j-1})\}$. Then set $T_j(\theta^j, w') = T_{j-1}(\theta^{j-1}, w')$, for $w' \in (R - S)$, and set $T_j(\theta^j, x_l) = T_{j-1}(\theta^{j-1}, u_l)$. For x_i such that $x_i \in \nu_j(x_i)$, τ_j contains a subtree isomorphic to $\tau_{j-1} \lceil u_i$, with labels changed and T_j extended again exactly as above. It is easy to see that T_j is a partial witness structure on F_j based on τ_j, satisfying the inductive conditions, as required.

The construction gives the required partial witness structure T_m on F_m with no requirements at w.

Lemma 18 *If α_0 has a partial witness structure T on a finite universal frame $F = (W, \nu)$ such that $\alpha_0 \in T(\theta_0, w_0)$ for some $w_0 \in W$, then there exists a complete witness structure T' for α_0 on a universal frame F' such that T' extends T at w_0.*

PROOF. Fix T, $F = (W, \nu)$, $w_0 \in W$ and τ_F as above. Fix a countable set $U = W \cup \{u_0, u_1, \ldots\}$ such that for all j, $u_j \notin W$, and an enumeration of U such that elements of W precede elements in $U - W$ in the enumeration. We now define a sequence of finite universal frames $F_0, F_1, \ldots$, with partial witness structures T_i on F_i, $i \geq 0$, with the following properties: let $F_i = (W_i, \nu_i)$ and $S_i \subseteq W_i$ such that T_i has no requirement at any $w \in S_i$, $i \geq 0$.

- $W_i \subseteq W_{i+1}$ and $S_i \subseteq S_{i+1}$.
- For all $w \in S_i$, $\nu_{i+1}(w) = \nu_i(w)$ and T_{i+1} extends T_i at w.

The construction proceeds inductively as follows: define $W_0 = W$, $\nu_0 = \nu$ and $T_0 = T$.

Suppose that we are given T_i on F_i satisfying the inductive assumptions. If $W_i = S_i$, then set $F_{i+1} = F_i$ and $T_{i+1} = T_i$; otherwise, among elements in $W_i - S_i$, pick the least one in the enumeration, say w. By repeated applications of Proposition 16, we get T'_i on the finite universal frame $F'_i = (W'_i, \nu'_i)$ where $W_i \subseteq W'_i$, and for all $w' \in S_i$, we have $\nu'_i(w') = \nu_i(w')$, for all $w'' \in (W'_i - \{w'\})$, $w' \notin \nu'_i(w'')$ and T'_i extends T_i at w'. In particular, for all $w' \in S_i$, $w' \notin \nu'_i(w)$.

Note that T'_i has a requirement at w in F'_i as well. By Lemma 17, there exists a partial report T' on a finite universal frame F' which has no requirements at w, and in addition, for all $w' \in S_i$, $\nu'(w') = \nu'_i(w')$ and T' extends T'_i at w', as required. Set $F_{i+1} = F'$ and $T_{i+1} = T'$.

Such a construction ensures the following property: for all $w \in U$, there exists n_w such that for all $k \geq n_w$, $w \in W_k$, $\nu_k(w) = \nu_{n_w}(w)$ and T_k extends T_{n_w} at w.

Now define the universal frame $F_{lim} \stackrel{\text{def}}{=} (W_{lim}, \nu_{lim})$ as follows: $W_{lim} = \bigcup_i W_i$, and $\nu_{lim}(w) = \nu_{n_w}(w)$ for $w \in W_{lim}$. The tree $\tau_{lim} = (\Theta_{lim}, \rightarrow_{lim}, \lambda_{lim})$ is defined in the obvious manner: the root is labelled W_{lim} and for $w \in W_{lim}$, it has a subtree isomorphic to τ_{n_w}. It can be checked that τ_{lim} is a view tree on F_{lim}.

For any node $\theta \in \Theta_{lim}$, and $u \in \lambda_{lim}(\theta)$, if θ is the root, then $T_{lim}(\theta, u) = T_i(\theta_{0,i}, u)$, where i is the least index such that $u \in W_i$ and $\theta_{0,i}$ is the root in τ_i. Otherwise, $\theta \in \Theta_{lim}\lceil w$ for some $w \in W_{lim}$. Define $T_{lim}(\theta, u) = T_{n_w}(\theta, u)$.

It can be easily checked that T_{lim} is a complete witness structure on the frame F_{lim}, as required.

Lemmas 14, 15 and 18 togther show that a formula is satisfiable iff it has a partial witness structure.

Theorem 19 *The satisfiability of a formula α can be decided in nondeterministic time $2^{O(m.d^2)}$, where m is the length of α and $d = \mu(\alpha)$.*

PROOF. Fix the formula α_0 of length m whose satisfiability is to be decided. There are at most $2^{O(m)}$ atoms. Given a k-element set S, constructing a tree labelled by decreasing subsets of S, assigning an atom to each element in each node, checking that the structure is a view tree and that it gives a partial witness structure can be done by a nondeterministic Turing machine in time k^{d+1}. By Lemmas 14, 15 and 18, we know that α_0 is satisfiable iff there is a partial witness structure for it over a set whose size is at most $2^{O(m.d)}$. Thus, a nondeterministic machine can decide satisfiability of α_0 in time $2^{O(m)} + (2^{O(m.d)})^{d+1}$, which is time $2^{O(m.d^2)}$.

We have thus shown that the satisfiability and model checking problems for the logic are elementarily decidable. Moreover, note that the complexity of model checking is linear in the size of the formula, and if we confine our attention to formulas of bounded alternation depth, then it is also polynomial in the size of the model.

4 Monotone frames

In an example discussed earlier, we considered a frame where $w_0 \sim w_1 \sim w_2$ where $\nu(w_0) = \{w_0, w_1\}$ and $\nu(w_1) = \{w_1, w_2\}$. We might see this as the agent 'forgetting' the information about w_0 at w_1; for instance, this may be a case of reusing memory to compute the state w_2. We now consider the subclass of frames where this cannot happen, and in which visibility preserves logical indistinguishability.

Definition 20 *A view frame $F = (W, \sim, \nu)$ is said to be* **monotone** *if and only if it is weakly recognizable, and $\forall w, w_1, w_2 \in W$, if $\{w_1, w_2\} \subseteq \nu(w)$ and $w_1 \sim w_2$ then $\nu(w) \cap \nu(w_1) = \nu(w) \cap \nu(w_2)$.*

In monotone frames, once two logically indistinguishable worlds become visible from another, they both contain the same information, in the sense that at w, the agent considers the same world to be visible at w_1 as at w_2. The following properties of view functions in monotone frames make this clearer.

Proposition 21 *Let $F = (W, \sim, \nu)$ be a monotone frame.*

(i) $\forall w_1, w_2 \in W$, if $w_2 \in \nu(w_1)$ and $w_1 \sim w_2$ then $\nu(w_1) = \nu(w_2)$.
(ii) $\forall w, w_1, w_2 \in W$, if $\{w_1, w_2\} \subseteq \nu(w)$ and $w_1 \sim w_2$ then $\nu(w_1) = \nu(w_2)$.

PROOF. Assume that $w_2 \in \nu(w_1)$ and $w_1 \sim w_2$. By weak recognizability, $w_1 \in \nu(w_1)$. By definition of monotonicity, since $\{w_1, w_2\} \subseteq \nu(w_1)$ and $w_1 \sim w_2$, we get $\nu(w_1) = \nu(w_1) \cap \nu(w_2)$, that is, $\nu(w_1) \subseteq \nu(w_2)$. But then $w_1 \in \nu(w_2)$ as well, and again by weak recognizability, $w_2 \in \nu(w_2)$. Now we have $\{w_1, w_2\} \subseteq \nu(w_2)$ and as before we find that $\nu(w_2) \subseteq \nu(w_1)$. Thus, $\nu(w_1) = \nu(w_2)$, establishing the first statement. The definition of monotonicity and the hypotheses of (ii) together imply that $w_1 \in \nu(w_2)$ and $w_2 \in \nu(w_1)$, and the claim follows from (i).

Proposition 22 *If $F = (W, \sim, \nu)$ is a monotone frame and $W' \subseteq W$, then so also is $F\lceil W'$.*

We say α is *monotone* satisfiable if there exists a model $M = (F, V)$, where $F = (W, \sim, \nu)$ is a monotone frame and for some $w \in W$, we have $M, w \models \alpha$. A formula is monotone valid if and only if its negation is not monotone satisfiable.

Note that the implication $\mathbf{XL}\alpha \supset \mathbf{LX}\alpha$ is monotone valid, though not valid. To see this, suppose $M, w \models \mathbf{XL}\alpha$ in a monotone model M. Then, $w \in \nu(w)$ and $M\lceil\nu(w), w \models \mathbf{L}\alpha$. Let $w' \in \nu(w)$ and $w \sim w'$ such that $M\lceil\nu(w), w' \models \alpha$. By proposition above, $\nu(w) = \nu(w')$ and hence we have $M\lceil\nu(w'), w' \models \alpha$. That is, $M, w' \models \mathbf{X}\alpha$. Since $w \sim w'$, we get $M, w \models \mathbf{LX}\alpha$, as required.

4.1 Checking truth

The advantage of monotone frames is clearly seen by the implications of Proposition 21: consider a universal monotone frame $F = (W, \nu)$, and note that ν partitions W into sets $W_1, W_2, \ldots$ etc such that a world in W_i is visible from a world in W_j iff $i = j$. Thus the frame is split into a set of universal subframes. Further, observe that for every formula α and $w, w' \in W$ such that $w' \in \nu(w)$, since $\nu(w) = \nu(w')$,

$$M\lceil\nu(w), w' \models \alpha \text{ iff } M\lceil\nu(w'), w' \models \alpha.$$

Thus, given a frame with k worlds in it, there are at most $k + 1$ substructures which we need to work with when we wish to check truth. The benefits bestowed by these observations are immediate:

Theorem 23 *The monotone truth checking problem is solvable in time $O(n.k)$, where k is the size of the given structure and n is the length of the checked formula.*

PROOF. Suppose that we are given a monotone model $M = ((W, \sim, \nu), V)$, an element $w_0 \in W$ and a formula α_0. We confine our attention to the universal submodel defined by $[w_0] = \{w_0, w_1, \ldots, w_{k-1}\}$, say. Since M is monotone, ν partitions $[w_0]$. Assume that there is an enumeration of W listing the partitions contiguously, so that for $0 \leq j \leq k$, there exist j_1, j_2 such that $j_1 \leq j \leq j_2$ and $\{w_{j_1}, \ldots, w_{j_2}\}$ is one of the partitions induced by ν. Fix an enumeration of SF in order of non-decreasing length.

We label each world $w \in W$ by an ordered pair (L_1, L_2), both subsets of SF. Initially, for all w, set $w.L_1 = w.L_2 = V(w) \cap SF$. We now describe the labelling procedure for all $w_j.L_2$, $0 \leq j \leq k$, where we consider pairs of the form (γ, w_j) in the lexicographic order defined by the enumeration of SF and W fixed above:

- Add $\neg\alpha$ to $w_j.L_2$ only if $\alpha \notin w_j.L_2$.
- Add $\alpha \vee \beta$ to $w_j.L_2$ only if $\alpha \in w_j.L_2$ or $\beta \in w_j.L_2$.

– Add $\mathbf{X}\alpha$ to $w_j.L_2$ only if $\alpha \in w_j.L_2$.
– Add $\mathbf{K}\alpha$ to $w_j.L_2$ only if $\alpha \in w_i.L_2$ for all $i : j_1 \leq i \leq j_2$.

The labelling procedure for $w_j.L_1$, $0 \leq j \leq k$, is similar: the only changes are:

– Add $\mathbf{X}\alpha$ to $w_j.L_1$ only if $\alpha \in w_j.L_2$.
– Add $\mathbf{K}\alpha$ to $w_j.L_1$ only if $\alpha \in w_i.L_1$ for all $i : 0 \leq i \leq k$.

Finally, terminate saying 'Yes' if $\alpha_0 \in w_0.L_1$ and 'No', otherwise.

It is easy to check that for all $\alpha \in SF$ and $w \in W$, $M, w \models \alpha$ iff $\alpha \in w.L_1$, and for $w \in \nu(w)$, $M\lceil\nu(w), w \models \alpha$ iff $\alpha \in w.L_2$. Hence the algorithm is correct and runs in time $O(n.k)$ where k is the size of the given structure and n is the length of the checked formula.

4.2 Satisfiability

For proving decidability of the monotone satisfiability problem, we can again define witness structures and show that a formula is monotone satisfiable iff it has a witness structure. But the structures required are very simple now.

Fix a formula α_0, and let AT be the set of subsets of $SF(\alpha_0)$ which are atoms. A *monotine witness structure* for α_0 is a set of pairs $\Gamma = (A, B) \in (AT \times AT)$ satisfying the following conditions:

(i) For some (A_0, B_0) in Γ, $\alpha_0 \in A_0$.
(ii) If $\mathbf{L}\alpha \in A$, there exists $(A', B') \in \Gamma$ such that $A \sim A'$ and $\alpha \in A'$.
(iii) If $\mathbf{X}True \in A$, then
  (a) $A \cap P = B \cap P$.
  (b) $B \preceq B \preceq A$.
  (c) If $\mathbf{L}\alpha \in B$, there exists $(A', B') \in \Gamma$ such that $B \sim B'$ and $\alpha \in B'$.

Given such a witness structure for α_0, it is easy to construct a model M for α_0 as follows: $M = ((W, \sim, \nu), V)$, where $W = \Gamma$; $(A, B) \sim (A', B')$ iff $A \sim A'$; if $\mathbf{X}True \in A$, then $\nu((A, B)) \stackrel{\text{def}}{=} \{(A', B') \mid B \sim B'\}$, and $\nu((A, B)) \stackrel{\text{def}}{=} \emptyset$, otherwise. $V((A, B)) \stackrel{\text{def}}{=} A \cap P$. We can then check that for all $\alpha \in SF$, $M, (A, B) \models \alpha$ iff $\alpha \in A$, and when $\mathbf{X}True \in A$, $M\lceil\nu((A, B)), (A, B) \models \alpha$ iff $\alpha \in B$. This shows that $M, (A_0, B_0) \models \alpha_0$ and hence α_0 is satisfiable.

On the other hand, given a model $M = ((W, \sim, \nu), V)$ such that $M, w_0 \models \alpha_0$, constructing a witness structure for α_0 is also straightforward: for each world w form a pair (A_w, B_w), where $A_w = \{\gamma \in SF \mid M, w \models \gamma\}$; if $w \in \nu(w)$, then $B_w = \{\gamma \in SF \mid M\lceil\nu(w), w \models \gamma\}$ and $B_w = \emptyset$ otherwise.

Let $m =\mid \alpha_0 \mid$. There exists $W' \subseteq [w_0]$ such that $\mid W' \mid \leq m$, $w_0 \in W'$, and for every $w \in W'$ and $\mathbf{L}\beta \in SF$, if $M, w \models \mathbf{L}\beta$, then there exists $w' \in W'$ such that $M, w' \models \beta$.

Call this a good subset of $[w_0]$. For each $w \in (W' \cap \nu(w))$, let $g(w)$ denote a good subset of $\nu(w)$ containing w similarly, and of size at most m. Let $W'' = W' \cup \bigcup_{w \in W'} g(w)$. We can then show that $\Gamma = \{(A_w, B_w) \mid w \in W''\}$ is a monotone witness structure for α_0.

Thus, α is monotone satisfiable iff there is a monotone witness structure for α of size $O(\mid \alpha \mid^2)$. On the other hand a monotone witness structure of size k for α gives a model of size k for α as well. Therefore, we have:

Lemma 24 *α is monotone satisfiable iff it is satisfiable in a monotone model of size $O(\mid \alpha \mid^2)$.*

This lemma, alongwith the model checking procedure above, at once gives an NP decision procedure for monotone satisfiability.

Theorem 25 *The monotone satisfiability problem is NP-complete.*

4.3 Axiom system

We now present an axiomatization of monotone valid formulas. We present it in two layers in a mutually recursive manner, one layer for explicitly valid formulas, and another for monotone valid formulas. We use the notation $\vdash_X$ to denote formulas proved as theorems in the system for explicit validity and $\vdash$ for the theorems of the system for monotone validity. However, we will have rules for 'transferring' theorems from one system to the other, so we are always, in effect, speaking of the combined system.

Ax_M, The axiom system for monotone valid formulas

($A0$) All the substitutional instances of the tautologies of PC

($A1$) $\mathbf{K}(\alpha \supset \beta) \supset (\mathbf{K}\alpha \supset \mathbf{K}\beta)$

($A2$) $\mathbf{Y}(\alpha \supset \beta) \supset (\mathbf{Y}\alpha \supset \mathbf{Y}\beta)$

($A3$) $\mathbf{K}\alpha \supset \alpha$

($A4$) $\mathbf{K}\alpha \supset \mathbf{KK}\alpha$

($A5$) $\neg\mathbf{K}\alpha \supset \mathbf{K}\neg\mathbf{K}\alpha$

($A6$) $\mathbf{X}True \supset (\alpha \equiv \mathbf{X}\alpha)$ when $\mu(\mathbf{X}\alpha) = \mu(\alpha)$

($A7$) $\mathbf{X}\alpha \supset \mathbf{Y}\alpha$

($A8$) $\mathbf{XL}\alpha \supset \mathbf{LX}\alpha$

Inference rules

$$(MP_i)\ \frac{\alpha, \quad \alpha \supset \beta}{\beta} \qquad (KG)\ \frac{\alpha}{\mathbf{K}\alpha} \qquad (XG)\ \frac{\vdash_X \alpha}{\mathbf{Y}\alpha}$$

Ax_X, The axiom system for explicitly valid formulas

$(B0)\ \alpha \equiv \mathbf{X}\alpha$

Inference rules

$$(MP_X)\ \frac{\alpha, \quad \alpha \supset \beta}{\beta} \qquad (IE)\ \frac{\vdash \alpha}{\alpha}$$

A *derivation* of the combined system is a finite sequence of formulas where each line is of the form:

- $\vdash \alpha$ where α is an axiom of Ax_M or follows from formulas earlier in the sequence by an application of a rule of Ax_M, or
- $\vdash_X \alpha$ where α is an axiom of Ax_X or follows from formulas earlier in the sequence by an application of a rule of Ax_X.

A *thesis* or theorem of the combined system is a formula that occurs as the last line of a derivation. When it is of the form $\vdash \alpha$ ($\vdash_X \alpha$), we refer to it as a theorem of Ax_M (Ax_X).

The axiom schemes carry no surprises: schemes (A1) and (A2) are deductive closure axioms for the two modalities. Schemes (A3) through (A5) are the usual properties of the $\mathbf{K}$ operator. (A6) asserts Proposition 8, whereas (A7) asserts that the visible substructure is unique. (A8) is the characteristic axiom for monotonicity, specifying that ν partitions every $\sim$-equivalence class further. (B0) is the only crucial difference between the two systems, as it is not valid, but explicitly valid.

Note that every theorem of Ax_M is also a theorem of Ax_X. This allows us, for instance, to derive tautologies as theorems of Ax_X due to rule (IE), and '$\mathbf{X}$-versions' of tautologies thanks to rule (XG).

The following are some sample theses and derived rules of the combined system. The derivations are easy and hence omitted.

$(T1) \vdash (\mathbf{Y}\alpha \wedge \mathbf{X}\beta) \supset \mathbf{X}(\alpha \wedge \beta)$

$(T2) \vdash \mathbf{X}\alpha \equiv \mathbf{X}\mathbf{X}\alpha$

$(T3) \vdash \mathbf{X}\alpha \equiv \mathbf{X}\mathbf{Y}\alpha$

$(T4) \vdash \mathbf{Y}\alpha \equiv \mathbf{Y}\mathbf{X}\alpha$

$(T5) \vdash \mathbf{Y}\alpha \equiv \mathbf{Y}\mathbf{Y}\alpha$

$(T6) \vdash_X \mathbf{K}\alpha \supset \alpha$

$(T7) \vdash (\mathbf{X}\alpha \wedge \mathbf{X}\beta) \supset \mathbf{X}(\alpha \wedge \beta)$

$$(KG_X)\ \frac{\vdash_X \alpha}{\vdash_X \mathbf{K}\alpha}$$

As usual, we say a formula α is **consistent** if its negation is not a thesis of the combined system. Specifically, if $\not\vdash \neg\alpha$, we call α Ax_M-consistent, and if $\not\vdash_X \neg\alpha$, we call α Ax_X-consistent. Clearly, every Ax_X-consistent formula is Ax_M-consistent due to rule (IE), and whenever $\mathbf{X}\alpha$ is Ax_M-consistent, α is Ax_X-consistent, thanks to rule (XG).

A finite set of formulas A is consistent iff the conjunction of all formulas in it (denoted $\widehat{A}$) is consistent. We assume that $\widehat{\emptyset} = False$ by convention. A set of formulas Γ is consistent iff every finite subset of Γ is consistent. We will be interested in sets of Ax_M-consistent formulas as well as sets of Ax_X-consistent formulas. Since every set of the latter type is also one of the former type, when we simply speak of consistent sets, we mean Ax_M-consistent sets.

A set A is said to be *Maximal Consistent*, if whenever there is a formula α such that $A \cup \{\alpha\}$ is consistent, then $\alpha \in A$.

Theorem 26 *The combined system provides a sound and complete axiomatization of monotone valid formulas.*

PROOF. Soundness is straightforward, and we have already shown the validity of (A8), the characteristic axiom for monotonicity. To prove completeness, we show that every Ax_M-consistent formula is monotone satisfiable. (Strictly, we should simultaneously show that every Ax_X-consistent formula is explicitly satisfiable as well; but such a proof is easy to carry out on the same lines, and writing it explicitly only makes the presentation tedious, so we omit it.)

Fix an Ax_M-consistent formula α_0 and let SF denote the set SF_{α_0}. Let MAT denote the set of all maximal Ax_M-consistent subsets of SF and let XAT denote the set of all maximal Ax_X-consistent subsets of SF. Clearly, there exists $A_0 \in MAT$ such that $\alpha_0 \in A_0$.

Define $M = ((W, \sim, \nu), V)$ as follows: $W = \{(A, B) \mid A \in MAT$ such that $(\mathbf{X}True \notin A$ and $B = \emptyset)$ or $(\mathbf{X}True \in A$ and $\widehat{A} \wedge \mathbf{X}\widehat{B}$ is Ax_M-consistent$)\}$. We use w, w' etc to range over W and refer to the first and second components respectively by $w.1$ and $w.2$. (Note that $w_0 = (A_0, B_0) \in W$, where $\alpha_0 \in A_0$.) $\sim$ is defined by: $w \sim w'$ iff $\mathbf{K}\alpha \in w.1$ iff $\mathbf{K}\alpha \in w'.1$. $\nu(w) \stackrel{\text{def}}{=} \{w' \mid w'.2 \neq \emptyset$, and $\mathbf{K}\alpha \in w.2$ iff $\mathbf{K}\alpha \in w'.2\}$. $V(w) \stackrel{\text{def}}{=} w.1 \cap P$. It is easy to check that M is a monotone model.

The following series of observations shows how the pairs of atoms in W are related, and how pairs relate to other pairs in W. We omit some details, but these are standard in the study of modal logics, and can be found in texts like [HC84].

Proposition 27 *(i) For any $A \in XAT$, $\mathbf{X}\alpha \in A$ iff $\alpha \in A$.*

(ii) If $A \in MAT$ such that $\mathbf{X}True \in A$, then there exists $B \in XAT$ such that $(A, B) \in W$.

PROOF. The first assertion follows from axiom (B0). To prove the second, assume that $\mathbf{X}True \in A$, and note that $\vdash_X True \equiv (\widehat{B_1} \vee \ldots \vee \widehat{B_k})$, where $XAT = \{B_1, \ldots, B_k\}$. Hence $\vdash_M \mathbf{X}True \equiv \mathbf{X}(\widehat{B_1} \vee \ldots \vee \widehat{B_k})$. Hence, $\widehat{A} \wedge \mathbf{X}(\widehat{B_1} \vee \ldots \vee \widehat{B_k})$ is Ax_M-consistent; therefore, for some j, $\widehat{A} \wedge \mathbf{X}\widehat{B_j}$ is Ax_M-consistent. By definition of W, $(A, B) \in W$, as required.

Proposition 28 *Let $w \in W$ be of the form (A, B).*

(i) If $\mathbf{X}\alpha \in A$ then $\alpha \in B$.

(ii) If $\mathbf{L}\alpha \in A$ then there exists $(A', B') \sim (A, B)$ such that $\alpha \in A'$.

PROOF. Consider the first assertion: since $\widehat{A} \wedge \mathbf{X}\widehat{B}$ is Ax_M-consistent, and $\mathbf{X}\alpha \in A$, we can use thesis (T7) to show that $\mathbf{X}(\widehat{B} \wedge \alpha)$ is Ax_M-consistent, and by rule (XG), we have that $\widehat{B} \wedge \alpha$ is Ax_X-consistent. But $B \in XAT$, hence $\alpha \in B$, as required.

To prove the second assertion, suppose that $\mathbf{L}\alpha \in A$. Define the set $\Gamma = \{\mathbf{K}\beta \mid \mathbf{K}\beta \in A\} \cup \{\mathbf{L}\gamma \mid \mathbf{L}\gamma \in A\} \cup \{\alpha\}$. We can use axioms (A1) and (A3) through (A5) to show that this set is Ax_M-consistent. Hence there exists $A' \in MAT$ such that $\Gamma \subseteq A'$. If $\mathbf{X}True \in A'$, by Proposition 27 above, there exists $B' \in XAT$ such that $(A', B') \in W$. Otherwise, set $B' = \emptyset$, and again $(A', B') \in W$. By definition of $\sim$, $(A, B) \sim (A', B')$, and by construction, $\alpha \in A'$.

Proposition 29 *Let $w \in W$ be of the form (A, B), and let $\mathbf{X}True \in A$.*

(i) $A \cap P = B \cap P$.

(ii) If $\mathbf{L}\alpha \in B$ then there exists $(A', B') \in \nu((A, B))$ such that $\alpha \in B'$.

PROOF. The first assertion follows immediately from axiom (A6). Now suppose $\mathbf{L}\alpha \in B$. Let $\Gamma \stackrel{\text{def}}{=} \{\mathbf{K}\beta \mid \mathbf{K}\beta \in B\} \cup \{\mathbf{L}\gamma \mid \mathbf{L}\gamma \in B\} \cup \{\alpha\}$. It is easy to show that

$\vdash_X \widehat{B} \supset \mathbf{L}\widehat{\Gamma}$, and from this, that $\vdash (\widehat{A} \wedge \mathbf{X}\widehat{B}) \supset \mathbf{XL}\widehat{\Gamma}$. This implies, by axiom $(A8)$, that $\widehat{A} \wedge \mathbf{LX}\widehat{\Gamma}$ is Ax_M-consistent. We can now show, along the lines of Proposition 28, that there exists $(A', B') \in W$ such that $(A, B) \sim (A', B')$, and $\widehat{A'} \wedge \mathbf{X}\widehat{B'} \wedge \mathbf{X}\widehat{\Gamma}$ is Ax_M-consistent. Using thesis (T7) and (XG), we get that $\widehat{B'} \wedge \widehat{\Gamma}$ is Ax_X-consistent. But $B' \in XAT$, hence $\Gamma \subseteq B$. Therefore $(A', B') \in \nu((A, B))$ and $\alpha \in B'$, as required.

We now have all the ingredients to establish the following claim inductively. For every $w = (A, B) \in W$, for all $\gamma \in SF$,

– $\gamma \in A$ iff $M, w \models \gamma$, and
– If $w \in \nu(w)$, $\gamma \in B$ iff $M\lceil\nu(w), w \models \gamma$.

Now, since $\alpha_0 \in A_0$, where $w_0 = (A_0, B_0)$, we have $M, w_0 \models \alpha_0$ and we have a monotone model for α_0.

5 Discussion

We have argued that the distinction between logical distinguishability and visibility is one worth making and attempted to demonstrate that a logic of knowledge bearing such features is technically interesting. While this distinction has the sort of philosophical motivation presented in Section 1, there is also a more pragamatic reason for considering it, and that has to do with the way we model systems. One important reason for studying knowledge in the context of computing agents is that the agents' *actions* at any instant of time are based on the knowledge at that time. This interplay of knowledge, time and action is a major theme in the literature on knowledge in distributed systems (see [HF 89] and [FHMV 95]). However, knowledge based action must involve computed knowledge and it is here that we feel that the distinction made above is crucial, as the agent being resource limited, may have only limited visibility.

To study this, we need to enrich the frames to include time and action. Further, we have only talked about one reasoner so far, and this needs to be generalized to systems of many reasoners. We believe that the class of Knowledge Transition Systems introduced in [KR 94], [R 96a], augmented with view functions provide a suitable framework for such a study. Correspondingly the logic would also need to be enriched with modalities talking about reachability. The advantage of such a set-up is that resource limitations can be modelled in a natural way: for instance, we can consider systems where views are uniformly bounded to have at most k states; or insist that only the adjacent states (in the transition system graph) be visible at any state. We can then study interactions between agents one of whom does a one-step-look-ahead whereas the other always looks ahead two steps, and so on. This suggests that a *dynamic* or *temporal* logic of view-based knowledge may be worth studying.

An important issue that we have not satisfactorily addressed here is the precise `epistemic` status of the $\mathbf{X}$ modality. We have already observed that it is by itself neither like a knowledge modality nor a belief modality, in the sense originally discussed by Hintikka ([Hi62]). Since our logic has both the standard knowledge modality and this one, it is similar to logics where agents reason about knowledge of beliefs and so on, like in [KL88] and [MS93]. However, much of the complexity of the logic in [KL88] arises from systems of many reasoners, which we have not studied here. [MS93] reduces belief to a form of knowledge, whereas in our framework, belief comes in as an extra-logical notion in the form of views. The closest relationship to theories of belief seems to be in the fact that beliefs are based on evidence, and views can perhaps be seen as providing such evidence.

We have studied frames where the visibility of states is dependent on the world state. At the other end, we can think of *static* frames, where there is a fixed subset $W' \subseteq W$ such that $\forall w$, $\nu(w) = W'$ is constant. But then we do not need the $\mathbf{X}$ modality; a special proposition τ which when asserted at w means that $w \in W'$, would do the job. Construction of a meaning-preserving map from Ψ to Φ_τ (when interpreted over static frames) is a straightforward exercise.

For the basic logic of view-based explicit knowledge, we have only shown decidability here, and a complete axiom system has been given only for the monotone subclass of frames. In the general case, the following axiom is sound:

(G) $\mathbf{XL}\alpha \supset \mathbf{L}\alpha$ $\quad$ (α boolean)

We believe that a model construction can be carried out for consistent formulas of this system (where (A8) is replaced by (G)), along the lines of the construction in the proof of Theorem 19. However, the details seem tricky, and this is left for future work.

References

[FHMV95] Fagin, R., Halpern, J., Moses, Y. and Vardi, M., *Reasoning about knowledge*, M.I.T. Press, 1995.

[HF89] Halpern, J., and Fagin, R., "Modelling knowledge and action in distributed systems", *Distributed Computing*, vol 3, #4, 1989, 159-177.

[HMV94] Halpern, J., Moses, Y., and Vardi, M., "Algorithmic knowledge", *TARK V*, Theoretical Aspects of Reasoning about Knowledge, 1994, 255-266.

[HV91] Halpern, J., and Vardi, M., "Model checking vs theorem proving: a manifesto", in *Artificial Intelligence and Mathematical Theory of Computation*, (ed) V. Lifschitz, Acad Press, 1991, 151-176.

[Hi62] Hintikka, J., *Knowledge and Belief*, Ithaca, N.Y., Cornell University Press, 1962.

[HM92] Hoek, W. van der and Meyer, J.-J.C., "Making some issues of implicit knowledge explicit", *Int'l Journal of Foundations of Computer Science*, vol 3, #2, 1992, 193-223.

[HC84] Hughes, G.E., and Cresswell, M.J., *A Companion to Modal Logic*, London, Methuen, 1984.

[KNP90] Krasucki, P., Ndjatou, G. and Parikh, R., "Probabilistic knowledge and probabilistic common knowledge", *ISMIS 90*, International Symp. on Methodology for Intelligent Systems, 1990, 1-8.

[KR94] Krasucki, P., and Ramanujam, R., "Knowledge and the ordering of events in distributed systems", *TARK V*, Theoretical Aspects of Reasoning about Knowledge, 1994, 267-283.

[KL88] Kraus, S. and Lehmann, D., "Knowledge, belief and time", *Theoretical Computer Science*, vol 58, 1988, 155-174.

[MS93] Moses, Y. and Shoham, Y., "Belief as defeasible knowledge", *Artificial Intelligence*, vol 64, #2, 1993, 299-322.

[Pa87a] Parikh, R., "Knowledge and the problem of logical omniscience", *ISMIS 87*, International Symp. on Methodology for Intelligent Systems, 1987, 432-439.

[Pa87b] Parikh, R., "Some recent applications of knowledge" *FST and TCS 7*, Foundations of Software Technology and Theoretical Computer Science, 1987, LNCS #287, 528-539.

[Pa90] Parikh, R., "Recent trends in reasoning about Knowledge", *TARK III*, Theoretical Aspects of Reasoning about Knowledge, 1990, 3-10.

[Pa94] Parikh, R., "Logical omniscience", in *Logic and Computational Complexity*, Springer Lecture Notes in Computer Science 960, 22-29.

[R96a] Ramanujam, R., "Local knowledge assertions in a changing world", *TARK VI*, Theoretical Aspects of Rationality and Knowledge, 1996, pp. 1-17.

[R96b] Ramanujam, R., "A discussion on explicit knowledge", in *The Parikh Project: Seven papers in honour of Rohit*, ed. K. Segerburg, Uppsala Prints and Reprints in Philosophy, No. 18, 1996.