

Decidability of Propositionally Quantified Logics of Knowledge

Tim French

Murdoch University and the University of Western Australia
Perth, W.A.,
Australia.
tim@csse.uwa.edu.au

Abstract. Logics of knowledge have important applications for reasoning about security protocols and multi-agent systems. We extend the semantics for the logic of necessity with local propositional quantification $\mathcal{L}_{(\Box, \exists, \exists_1, \dots, \exists_k)}$ introduced in [4] to allow reasoning about knowledge in more general (non-hierarchical) systems. We show that these new semantics preserve the properties of knowledge in a multi-agent system, give a significant and useful increase in expressivity and most importantly, have a decidable satisfiability problem. The new semantics interpret propositional (local and non-local) quantification with respect to bisimulations, and the satisfiability problem is shown to be solvable via an embedding into the temporal logic, QCTL.

1 Introduction

Logics of knowledge [5] have important applications for reasoning about security protocols and multi-agent systems. Such a logic allows you to formalize what facts (represented by propositions) are known by which agents or whether one agent knows if another agent knows (or considers possible) some fact. Formalizations of knowledge also allow us to represent such notions as common knowledge (every agent knows, and every agent knows every agent knows, and so on), and distributed knowledge (what a group of agents could infer if they shared their knowledge). Propositional quantification in modal logics [6],[9] has often been considered as a way of increasing expressivity.

Recently several extensions have been investigated which allow quantification of propositions. This allows us to reason about an agent's knowledge independent of individual propositions. Being able to quantify over propositions (so that they represent arbitrary facts) allows us to examine and verify protocols independent of the context. However the resulting increase in expressive power extends much further and, in the most general case, the language has been shown to be highly undecidable. This problem has been addressed in several ways, including by weakening the properties of knowledge and by enforcing various semantic structures on the knowledge of agents.

In this paper we present a third approach. Rather than restricting any semantic properties of the system, we generalize the notion of propositional quantification to apply to bisimulations of the model. We show that interpretations of knowledge, common knowledge and distributed knowledge are unaffected by this change. Furthermore the resulting language will be shown to be decidable. This is significant as it allows us to express a wide range of second order properties which are of practical value (for example “agent 1 knows strictly more than agent 2 and 3 combined”), whilst avoiding any undecidable second-order properties. To the author’s knowledge, this combination of expressivity and decidability is not present in any previous languages.

This paper will be organized as follows. We will first introduce the basic logic of knowledge, its syntax and semantics. We will then generalize this to a logic of local propositions [3]. The new notion of quantification will be introduced and we will show that interpretations of knowledge are invariant under this change. We will discuss several formalizations of these semantics and show they are equivalent. We will also explore the expressive power of the language. The next part of the paper will present a proof of decidability for the language. This proof will require us to examine the temporal logic QCTL (since the satisfiability problem is reduced to the satisfiability problem for QCTL). This language is shown to be decidable in [8]. We will prove that this reduction is both sound and correct.

2 The Language

The language L_k^{CD} , for some natural number k , can express properties of knowledge for k agents and is built from the following abstract syntax:

$$\alpha := x \in \mathcal{V} \mid \alpha_1 \wedge \alpha_2 \mid \neg\alpha \mid K_i\alpha \mid C_G\alpha \mid D_G\alpha$$

where $i \in I = \{1, \dots, K\}$ is the label of an agent, and $G \subseteq I$ represents a set of agents. The interpretations of these formulas are as follows:

- $x \in \mathcal{V}$ are the propositional variables. These are used to represent properties of the system.
- \wedge and \neg are the standard boolean operators.
- $K_i\alpha$ represents the statement, “agent i knows α is true”.
- $C_G\alpha$ represents the statement, “it is common knowledge to all agents in G that α is true”
- $D_G\alpha$ represents the statement, “if all the agents in G were to combine their knowledge, they would be able to deduce α ”.

For details on common semantics for logics of knowledge see [5] We assume that each agent has a set of possible states, and each of these states correspond to a partition of the set of possible worlds. An agent, i , cannot distinguish two

worlds if the agent's local states in those two worlds are the same (the worlds are i -local).

We suppose that each agent, i , has a set of local states \mathcal{L}_i . Rather than using a distinct label for each world, we let each world be a tuple of agent's local states, along with the set of propositions that are true at that world. We say $w = (w_0, w_1, \dots, w_k)$ is a world where $w_0 \subseteq \mathcal{V}$ and for each i , $w_i \in \mathcal{L}_i$. The model M, w is simply a set of worlds, with one world specified as "this world", (i.e. $M \subseteq \wp(\mathcal{V}) \times \mathcal{L}_1 \times \dots \times \mathcal{L}_k$) and we define the satisfiability of a formula at a given world, w , as follows:

$$M, w \models_K x \iff x \in w_0 \quad (1)$$

$$M, w \models_K \neg\alpha \iff M, w \not\models_K \alpha \quad (2)$$

$$M, w \models_K \alpha \wedge \beta \iff M, w \models_K \alpha \text{ and } M, w \models_K \beta \quad (3)$$

$$M, w \models_K K_i\alpha \iff M, w' \models_K \alpha \text{ for all } w' \text{ where } w_i = w'_i \quad (4)$$

$$M, w \models_K C_G\alpha \iff M, w' \models_K \alpha \text{ for all } w' \text{ where } w \cong_G w' \quad (5)$$

$$M, w \models_K D_G\alpha \iff M, w' \models_K \alpha \text{ for all } w' \text{ where } \forall i \in G, w_i = w'_i, \quad (6)$$

where $w \cong_G w'$ is the smallest relation recursively defined by

$$w \cong_G w' \iff w = w' \text{ or } \exists u \in M, \exists i \in G \text{ such that } w_i = u_i \text{ and } u \cong_G w'.$$

We say α is a validity of L_k^{CD} if $M, w \models \alpha$ for all models M, w , and α is satisfiable if $\neg\alpha$ is not a validity. We will examine the semantic interpretation for these operators in more detail in the following section.

3 Semantics for Propositional Quantification

We are now ready to introduce propositional quantification into the language of local propositions. We could simply add propositional quantification to L_k^{CD} , but this would unnecessarily complicate the language. We use the notion of local quantifiers to express the properties of knowledge. These were introduced in [3].

The logic of quantified local propositions, (QL_k^1) is as follows:

$$\alpha := x \in \mathcal{V} \mid \alpha_1 \wedge \alpha_2 \mid \neg\alpha \mid \exists x\alpha \mid \exists_i x\alpha \mid \Box\alpha.$$

The new operators are existential propositional quantification (\exists), existential i -local propositional quantification (\exists_i) and necessity (\Box). The formula $\exists x\alpha$ states "there is an interpretation of the variable x that makes α true"; the formula $\exists_i x\alpha$ states "there is some interpretation of the variable x that is consistent with agent i 's local state, and which makes α true"; and the formula $\Box\alpha$ states "in all worlds α is true". We use $\exists\{x_1, \dots, x_n\}\alpha$ as an abbreviation for $\exists x_1 \dots \exists x_n \alpha$

¹ While the language is identical to $\mathcal{L}_{(\Box, \exists, \exists_1, \dots, \exists_k)}$ of [3], we will use QL_k so as to distinguish the semantics.

This language appears to be very different from the language given above, however we will show that every formula of L_k^{CD} is expressible in QL_k . To give the formal semantic interpretations of the operators \Box , \exists and \exists_i we will require the following definition. *Note the definition has been changed from the original paper, where it was incorrectly reported*

Definition 1. *Given some model M, w we say a model M', w' is an X -variant of M if there is some relation $B \subseteq M \times M'$ such that*

1. $(w, w') \in B$, and
2. for all $(u, u') \in B$
 - (a) $u_0 \setminus X = u'_0 \setminus X$,
 - (b) $\forall v \in M, \exists(v, v') \in B$ such that $\forall i \in I(u_i = v_i \iff u'_i = v'_i)$,
 - (c) $\forall v' \in M', \exists(v, v') \in B$ such that $\forall i \in I(u_i = v_i \iff u'_i = v'_i)$

Given $i \in I$ we say M', w' is an i -local X -variant of M, w if M', w' is an X -variant of M, w and for all $x \in X$, for all $u' \in M'$, $x \in u'_0$ if and only if for all $v' \in M'$ with $v'_i = u'_i$ we have $x \in v'_0$. If M', w' is a \emptyset -variant of M, w we say that M', w' is bisimilar to M, w .

Essentially, we allow quantification to not only apply to the given model, but also to models that are bisimilar to the given model. Since the semantic description of states forces all possibility relations in a model to be equivalence relations, we can assume the basic axioms of knowledge are preserved. However we should note that we are using a restricted set of bisimulations, since every model is bisimilar to a tree. We are now able to give the semantic interpretations for the new operators.

$$M, w \models_L \Box \alpha \iff M, w' \models_L \alpha \text{ for all } w' \in M$$

$$M, w \models_L \exists_i x \alpha \iff M', w' \models_L \alpha \text{ where } M \text{ is an } i\text{-local } \{x\}\text{-variant of } M, w$$

$$M, w \models_L \exists x \alpha \iff M', w' \models_L \alpha \text{ where } M \text{ is an } \{x\}\text{-variant of } M, w$$

Alternative semantic interpretations have been offered for logics of knowledge with propositional quantification. In [3], two alternatives based on local propositions were considered. The first, referred to as the *strong* semantics considered propositional quantification over a fixed set of worlds. That is

$$M, w \models_S \exists x \alpha \iff M', w' \models_S \alpha \tag{7}$$

where for all $w \in M$ there is some $w' \in M'$ such that $w_0 \setminus \{x\} = w'_0 \setminus \{x\}$ and $w_i = w'_i$ for $i > 0$. These semantics were shown to be too strong, being expressively equivalent to full second order logic.

In the same paper, a set of *weak* semantics were also considered. We will not go into the formal definitions here, but the basic idea was to restrict the

interpretation of atoms over sets of worlds. This interpretation was shown to be unable to fully express the knowledge operators required by logics of knowledge.

Finally in [4], a restricted version of the strong semantics were considered where the knowledge relations formed linear hierarchy. That is, for any pair of agents, one knew strictly more than the other. This language was shown to be decidable, and in the same paper a complete axiomatization was given.

The main advantage of the semantics presented here is that we are now able to reason about non-linear hierarchies of knowledge, without losing decidability or any of the necessary expressiveness.

4 Definability

The logic QL_k combines the power of logics of knowledge with propositional quantification. We will first show that our intuitions of propositional quantification are preserved. For example, if α is a validity, $\forall x\alpha$ should also be a validity, (this follows trivially from the definition), and if the atom, x , does not appear in α then $\alpha \rightarrow \forall x\alpha$ should be a validity. This follows from the following lemma.

Lemma 1. *Suppose that (M, w) is a model, and (N, v) is a model bisimilar to (M, w) . Then for all formulas α ,*

$$M, w \models_L \alpha \iff N, v \models_L \alpha. \quad (8)$$

Proof. This can be shown by induction over the complexity of formulas. By definition 1, we know $w_0 = v_0$, so the propositional case is trivial. Likewise the inductive steps for \neg and \wedge are trivial. The inductive steps for the quantifiers \exists and \exists_i follow directly from the fact that bisimilarity is an equivalence relation. This leaves the \Box operator.

By the induction hypothesis we suppose for all bisimilar models $M, w \models_L \alpha$ if and only if $N, t \models_L \alpha$. The result follows from the conditions of definition 1 which requires the relation to be defined for every element of M and every element of N . If there were some $u \in M$ such that $M, u \not\models_L \alpha$, then by the induction hypothesis there must be some $v \in N$ where $N, v \not\models_L \alpha$, and vice-versa. Therefore $M, w \models_L \Box\alpha$ if and only if $N, t \models_L \Box\alpha$ and the proof is complete.

It follows that the quintessential properties of propositional quantification are retained in QL_k . We will now show that the standard knowledge operators can be expressed in QL_k , and retain their semantic interpretation. As was shown in [3], the formulas $K_i(\alpha)$, $C_G(\alpha)$, and $D_G(\alpha)$ can be expressed in terms of these local quantifier and the \Box operator. Let x, x_1, \dots, x_k be variables that do not

appear in α .

$$K_i(\alpha) = \exists_i x(x \wedge \Box(x \rightarrow \alpha)) \quad (9)$$

$$C_G(\alpha) = \exists x(x \wedge \bigwedge_{i \in G} \exists_i y \Box(y \leftrightarrow x) \wedge \Box(x \rightarrow \alpha)) \quad (10)$$

$$D_G(\alpha) = [\exists_i x_i]_{i \in G} \left[\left(\bigwedge_{i \in G} x_i \right) \wedge \Box \left(\left(\bigwedge_{i \in G} x_i \right) \rightarrow \alpha \right) \right] \quad (11)$$

We will now show that as with the semantics discussed in [3], the interpretation of knowledge is retained.

Lemma 2. *For every model M, w ,*

$$M, w \models_K K_i \alpha \iff M, w \models_L K_i \alpha \quad (12)$$

$$M, w \models_K C_G \alpha \iff M, w \models_L C_G \alpha \quad (13)$$

$$M, w \models_K D_G \alpha \iff M, w \models_L D_G \alpha \quad (14)$$

Proof. Suppose that $M, w \models_K K_i \alpha$. Then for all worlds $v \in M$, where v is i -local to w , $M, v \models_K \alpha$. Let x be true at exactly these worlds, v . Then x is i -local so $M, w \models \exists_i x(x \wedge \Box(x \rightarrow \alpha))$.

Now suppose that $M, w \models_L K_i(\alpha)$. Then there is some i -local x -variant, (N, t) , of (M, s) such that $N, t \models_L x \wedge \Box(x \rightarrow \alpha)$. Therefore for every world $u \in N$ that is i -local to t , we must have $x \in u_0$ and hence, $N, u \models \alpha$. From Lemma 1, it follows that for all $v \in M$, i -local to w , we must have $M, v \models \alpha$ (since we assume that x is not a variable of α). Therefore $M, w \models K_i \alpha$, and we have shown the equivalence (12) holds.

We will omit the proofs of the following remaining two equivalences, as they are similar.

We will now briefly consider the expressive power of our logic. From the above proof, anything that is expressible in L_k^{CD} is expressible in QL_k . For examples of the expressive power of L_k^{CD} , see [5]. One of the significant advantages of QL_k is the power to reason about arbitrary hierarchies of knowledge. In the standard logics of knowledge we can express concepts, such as “If agent i knows p is true, then agent j knows p is true”, (i.e. $K_i p \rightarrow K_j p$). In QL_k we can express such properties independent of the proposition p , so the hierarchy applies to all formulas expressible in the language, (i.e. $\forall x(K_i x \rightarrow K_j x)$). Furthermore, we can reason about non-linear hierarchies of knowledge. The formula

$$\exists x(K_1 x \wedge \neg D_{2,3} x) \wedge \exists x(K_2 x \wedge \neg D_{1,3} x) \wedge \exists x(K_3 x \wedge \neg D_{1,2} x) \quad (15)$$

expresses the property that there are three agents, and if any two agents were to combine their knowledge, they would still not know strictly more than the third agent. Such a property could be important for ensuring the security of shared network resources, or the fairness of three player games.

5 Decidability

We can now present the proof of decidability for the language. We will do this via a translation into the temporal logic QCTL [1],[2]. This is because the decidability of QCTL is a very complicated result [8], [7], and we would like to avoid repeating it. The proof of decidability for QCTL involves a new kind of tree automata (amorphous Street tree automata) which act on ω -trees. (In fact, they act on the set of bisimulations of an ω -tree). There is a deterministic translation from a formula α , of QCTL to an amorphous automata which accepts exactly the models of α . Therefore the emptiness of the automata is equivalent to the unsatisfiability of α . Since the emptiness of an amorphous automata can be determined, this solves the satisfiability for QCTL.

The decidability process described above is particularly appropriate for QCTL since trees are a natural model for branching temporal logics. However the models of QL_k involve multiple equivalence relations interacting in an arbitrary manner, so it is not immediately apparent how we could define such a translation. The answer comes from the expressive power of QCTL. Every model of QL_k can be “untangled” along the relations to give an ω -tree. We will find that QCTL has the expressive power to interpret formulas on this tree with respect to the original structure. That is every relation can be syntactically “entangled” back into an equivalence relation. This way we simplify the semantic specification of a formula, but increase the complexity of the syntax to preserve validities in the language.

We will now briefly describe the logic QCTL. The syntax is given as

$$\alpha ::= x \mid \neg\alpha \mid \alpha_1 \vee \alpha_2 \mid AX\alpha \mid AG\alpha \mid \exists x\alpha \quad (16)$$

The formula $AX\alpha$, states that alpha is true at the next moment of time for all possible futures, $AG\alpha$ states that α is true for all moments of time from here on, regardless of future. The abbreviations $\wedge, \rightarrow, \leftrightarrow$ are defined as usual, and we define $EX\alpha$ to be $\neg AX\neg\alpha$, $EG\alpha$ to be $\neg AF\neg\alpha$ and $EF\alpha$ to be $\neg AG\neg\alpha$. To give the semantics for QCTL we define \mathcal{V} -labeled Kripke frames:

Definition 2. A Kripke frame is a tuple (S, R) where

1. S is a nonempty set of states, or moments.
2. $R \subseteq S^2$ is a total binary relation.

A \mathcal{V} -labeled Kripke frame is a Kripke frame with a valuation $\pi : S \rightarrow \wp(\mathcal{V})$.

Let $T = (S, R, \pi)$ be a \mathcal{V} -labeled Kripke frame, and let R^* be the reflexive, transitive closure of R . We interpret a formula α of QCTL with respect to a \mathcal{V} -labeled Kripke frame T and a state $s \in S$. We write $T, s \models_C \alpha$ where:

$$T, s \models_C x \iff x \in \pi(s)$$

$$\begin{aligned}
T, s \models_C \neg\alpha &\iff T, s \not\models_C \alpha \\
T, s \models_C \alpha \vee \beta &\iff T, s \models_C \alpha \text{ or } T, s \models_C \beta \\
T, s \models_C AX\alpha &\iff \text{for all } t \text{ where } (s, t) \in R, T, t \models_C \alpha \\
T, s \models_C AG\alpha &\iff \text{for all } t \text{ where } (s, t) \in R^*, T, t \models_C \alpha \\
T, s \models_C \exists x\alpha &\iff \text{there is some } x\text{-variant } T', s' \text{ of } T, s \text{ where } T', s' \models_C \alpha.
\end{aligned}$$

For the semantic interpretation of $\exists x\alpha$, we require the following definition:

Definition 3. *Given $X \subseteq \mathcal{V}$, $(T, s_0) = (S, R, \pi, s_0)$ is an X -variant of $(T', s') = (S', R', \pi', s'_0)$ if there exists some relation $B \subseteq S \times S'$ with $(s_0, s'_0) \in B$ and for all $(s, s') \in B$:*

1. $\pi(s) \setminus X = \pi'(s') \setminus X$.
2. For all $t \in S$ such that $(s, t) \in R$, there exists $t' \in S'$ with $(s', t') \in R'$ such that $(t, t') \in B$.
3. For all $t' \in S'$ such that $(s', t') \in R'$, there exists $t \in S$ with $(s, t) \in R$ such that $(t, t') \in B$.

We say α is a validity if for all models, (T, s) , we have $T, s \models_C \alpha$, and α is satisfiable if $\neg\alpha$ is not a validity. It is important to note that the truth of formulas of QCTL with respect to a given model is invariant under bisimulation. As all QCTL models are bisimilar to tree, from now on we will consider all models of QCTL to be trees where the relation R acts as the parent-child relation.

The definitions of x -variant are quite similar for QCTL and QL_k . They both use the basic idea of a bisimulations [10] [11] to access the power of propositional quantification whilst abstracting out the non-modal properties of the model. For more details on QCTL see [7].

Our approach will involve translating the satisfiability problem for a formula of QL_k into the satisfiability problem for a formula of QCTL. We must first describe the translation, and secondly show that it preserves satisfiability and unsatisfiability. Let \mathcal{M}_L be the set of models for QL_k and let \mathcal{M}_T be the set of models for QCTL.

Definition 4. *Given some formula α , we define translation $\phi : \mathcal{M}_L \longrightarrow \mathcal{M}_T$ as follows: Let $I_\alpha = \{x_1, \dots, x_k\}$ be a set of propositions not appearing in α . Given $(M, s) \in \mathcal{M}_L$ Let $(M, s)^\phi$ be the model (S, R, π, s) where*

- $S = \{sw \mid w \in M^*\} \subseteq M^*$.
- $\pi(s) = s_0$
- For all $tu \in S$ where $t \in S$ and $u \in M$, for all $v \in M$, $\pi(tuv) = (v_0 \setminus I_\alpha) \cup \{x_i \mid u_i = v_i\}$.
- $R = \{(t, tu) \mid t \in S, u \in M\}$.

The QCTL model $(M, s)^\phi$ is similar to the model (M, s) , except where (M, s) has k different relations, $(M, s)^\phi$ only has the one. To compensate for this each node is labeled with the propositions I_α , which indicate which local states a node shares with its parent.

We will now define a syntactic translation $*$ on QL_k , such that α is satisfiable for QL_k if and only if α^* is satisfiable for QCTL. Essentially what we are doing is untangling the semantics of QL_k to be represented by a tree, and entangling the syntax of QL_k to retain the necessary properties of knowledge.

Definition 5. We define the translation $*$: $QL_k \rightarrow QCTL$ in several stages. Let α be some formula of QL_k and we define $\mathcal{V}_\alpha = \{y_\beta | \beta \subseteq \alpha\}$ where for each sub-formula $\beta \subseteq \alpha$, y_β is some variable not appearing in either α or I_α . For each sub-formula β of α we define β' recursively as follows:

$$\begin{aligned}
x' &= AG(y_x \leftrightarrow (x)) \\
(\neg\gamma)' &= AG(y_\beta \leftrightarrow \neg y_\gamma) \wedge \gamma' \\
(\gamma_1 \wedge \gamma_2)' &= AG(y_\beta \leftrightarrow (y_{\gamma_1} \wedge y_{\gamma_2})) \wedge \gamma'_1 \wedge \gamma'_2 \\
(\Box\gamma)' &= AG(y_\gamma) \leftrightarrow AG(y_\beta) \wedge EF(y_\beta) \rightarrow AG(y_\gamma) \wedge \gamma' \\
(\exists x\gamma)' &= \exists x(AG(y_\beta \leftrightarrow y_\gamma) \wedge \gamma') \\
(\exists_i x\gamma)' &= \exists x(\gamma' \wedge AG((x \rightarrow AX(x_i \rightarrow x)) \wedge (EX(x_i \wedge x) \rightarrow x) \wedge (y_\beta \rightarrow y_\gamma)))
\end{aligned}$$

Finally let $\alpha^* = \exists \mathcal{V}_\alpha(\alpha' \wedge y_\alpha)$.

This definition implements the semantic interpretations for each operator. The states at which any sub-formula, β , is true is marked with the proposition y_β and any sub-formula referring to β is then interpreted with respect to that proposition. This allows us to implement the semantic interpretation for QL_k . Particularly, the operators of QCTL are naturally anti-symmetric, irreflexive, and transitive, so the concept of an equivalence relation is enforced in the translation of $\exists_i x\beta$, which considers all states in the model which can be reached by passing only through states with x_i in their label.

Lemma 3. $M, s \models_L \alpha \implies (M, s)^\phi \models_C \alpha^*$.

Proof. We first must show that for all $\beta \subseteq \alpha$, the formula β' is satisfiable for QCTL. This can be shown by induction. This is trivial in the case that $\beta \in \mathcal{V}$, so suppose $\beta = \mathcal{O}\gamma$, where \mathcal{O} is some operator, and γ' is satisfiable. Then β' simply dictates that the interpretation of y_β is defined with respect to the semantic interpretation of \mathcal{O} and y_γ . Since propositional quantification is interpreted with respect to bisimulations for QCTL, it follows that α' is not only satisfiable, but also $\exists \mathcal{V}_\alpha \alpha'$ is a validity.

Let $(M, s)^\phi = (S, R, \pi, s)$. We must show that if $(M, s)^\phi \models_L \alpha'$, then for all $\beta \subseteq \alpha$, $M, u \models_L \beta$ if and only if $y_\beta \in \pi(wu)$ (where $w \in M^*$, $u \in M$). Again,

this can be shown by a simple induction over the complexity of α where the base case ($\beta \in \mathcal{V}$) is trivial, as are the cases for \neg , \wedge and $\exists x$.

The leaves the operators \Box and \exists_i , so suppose that for all models, M, s , for all $u \in M$, $M, t \models_L \beta$ then $y_\beta \in \pi(wu)$. In the case of $\Box\beta$, $M, s \models_L \Box\beta$ if and only if β is true at every world in M . By the induction hypothesis it follows that $y_\beta \in \pi(u)$ for all $u \in M$. As the definition of $(\Box\beta)'$ defines the interpretation of $y_{\Box\beta}$ with respect to the entire tree it follows that if $M, u \models_L \Box\beta$ then $y_{\Box\beta} \in \pi(wu)$.

In the case of $\exists_i x\beta$, $M, s \models_L \exists_i x\beta$ if and only if there is i -local x -variant, (N, t) , of (M, s) such that $N, t \models_L \beta$. We state without proof (though it is easy to show) that if (N, t) is an i -local x -variant of (M, s) , then $(N, t)^\phi$ is an x -variant of $(M, s)^\phi$. From the induction hypothesis $(N, t)^\phi \models_C y_\beta$. The definition of ϕ will ensure that all worlds which share an i -local state will be mapped to a subtree, upon which x_i is always true. The definition of $(\exists_i x\beta)'$ will ensure that the interpretation of x will not vary on this subtree. Therefore $(N, t)^\phi$ is an x -variant of $(M, s)^\phi$ which satisfies

$$(\gamma' \wedge \text{AG}((x \rightarrow \text{AX}(x_i \rightarrow x)) \wedge (\text{EX}(x_i \wedge x) \rightarrow x) \wedge (y_\beta \rightarrow y_\gamma))), \quad (17)$$

and it follows from the definition of an x -variant that for all $u \in M$, if $M, u \models_L \exists_i x\beta$, then $y_{\exists_i x\beta} \in \pi(wu)$.

This completes the induction, so it follows that if $M, s \models_L \alpha$, then $(M, s)^\phi \models_C \alpha'$ implies that $y_\alpha \in \pi(s)$ completing the proof.

We now know that if α is satisfiable, then α^* is also satisfiable. To prove that the satisfiability problem for QL_k is decidable, we must also show that if α^* are satisfiable in QCTL, then α is satisfiable in QL_k . To do this we will show that from any QCTL model that satisfies α^* we can generate a QL_k model that satisfies α . We first define a map from \mathcal{M}_T to \mathcal{M}_L .

Definition 6. We define the translation $\kappa^\alpha : \mathcal{M}_T \rightarrow \mathcal{M}_L$ as follows. We suppose that there are a set of atoms I_α not appearing in α , that agree with the set of atoms used in the translation $*$, and $T, s = (S, R, \pi, s)$ is a model of QCTL.

We define a function $\kappa : S \rightarrow \wp(\mathcal{V}) \times S^k$ recursively by

- $\kappa(s) = (\pi(s), s, \dots, s)$
- If $\kappa(t) = (a, \tau_1, \dots, \tau_k)$, and $(t, u) \in R$, then $\kappa(u) = (\pi(u), \sigma_1, \dots, \sigma_k)$ where for all $i \in I$,
 - if $x_i \in \pi(u)$ then $\sigma_i = \tau_i$
 - if $x_i \notin \pi(u)$ then $\sigma_i = u$.

We let $\kappa(S)$ be the range of κ , and define $\kappa^\alpha(T, s) = (\kappa(S), \kappa(s))$.

This translation essentially builds up a QL_k model by taking the reflexive, symmetric, transitive closure of the relations represented by the atoms in I_α .

Lemma 4. $T, s \models_C \alpha^* \implies \kappa^\alpha(T, s) \models_L \alpha$.

Proof. Suppose that $(T, s) \models_C \alpha^*$ and $\mathcal{V}_\alpha = \{y_\beta | \beta \subseteq \alpha\}$. Then there is a \mathcal{V}_α -variant, T', s' of T, s such that $T', s' \models_C \alpha' \wedge y_\alpha$. We state the following fact without proof:

If (M, s) and (N, t) are bisimilar models of QCTL, then $\kappa(M, s)$ is bisimilar to $\kappa(N, t)$ with respect to QL_k .

This is easy to see from the construction of κ . Since the atoms of \mathcal{V}_α do not appear in α , it is enough to show that

$$T, s \models_C \alpha' \wedge y_\alpha \implies \kappa(T, s) \models_L \alpha. \quad (18)$$

This is shown by induction over the complexity of α . Specifically, we will show for all models (T, s) , for all $t \in S$, if $T, s \models_C \alpha'$ and $y_\beta \in \pi(t)$, then $\kappa(S, t) \models_L \beta$. This is trivial for $\beta \in \mathcal{V}$ (the base case), and the inductive steps for \neg , \wedge and $\exists x$ are easy to show (noting that $\alpha' \rightarrow \beta'$ is a validity for all $\beta \subseteq \alpha$).

So suppose that for all models, (T, s) , for all $t \in S$, if $T, s \models_C \alpha'$ and $y_\beta \in \pi(t)$, then $\kappa(S), \kappa(t) \models_L \beta$. For the necessity operator, suppose that $y \square_\beta \in \pi(t)$ for some t . By the definition of β' , we see that $y \square_\beta \in \pi(t)$ for all $t \in S$. By the induction hypothesis we must have $\kappa(S), \kappa(t) \models_L \beta$.

For the local quantifier, \exists_i , suppose that $y_{\exists_i x \beta} \in \pi(t)$ for some t . Therefore there is some x -variant, (T', s') , of (T, s) such that $y_\beta \in \pi(t')$, where $(t, t') \in B$, for some bisimulation, B , defined by the x -variant. By the definition of $(\exists_i x \beta)'$, we see that the interpretation of x cannot change from one node to its successor if x_i is true at it's successor. From the definition of κ , we see that x will be i -local in $\kappa(M', s')$. Applying the induction hypothesis we have $\kappa(S'), \kappa(t') \models_L \beta$, and by applying the fact stated above, $\kappa(S), \kappa(t) \models_L \exists_i x \beta$.

This completes the induction and (18) follows.

Theorem 1. *The satisfiability problem for QL_k is decidable.*

Proof. This follows trivially from the above lemmas. Given any formula α , satisfiable in QL_k , we can compute the formula α^* , and demonstrate the satisfiability in QCTL, via Lemma 3 and the decision procedure for QCTL. If α is not satisfiable in QL_k by Lemma 4 it is enough to show that α^* is not satisfiable in QCTL.

6 Conclusion

In this paper we have presented an effective interpretation for logics of knowledge with propositional quantification. We have briefly looked at the expressive power

of such a logic, and shown how we can express a wide range of useful second-order properties. Most importantly, we have achieved this gain in expressivity without sacrificing the decidability of the logic. The decidability of QL_k was shown via a translation to QCTL. From this translation we can see that the decision process presented is infeasible for automated reasoning. We also note that the language QCTL has a polynomial translation into the language QL_k , and the satisfiability problem for QCTL is known to be non-elementary [8], i.e. the decision procedure is optimal. This does not mean that the logic is of no use. Its expressive power, and the fact that it is decidable, suggest that it is valuable for formalizing the specification of systems. Automating reasoning could be effectively applied to sub-languages, and axiomatizations could be found to assist reasoning about systems and verification of protocols.

One of the main avenues for future work is to find an axiomatization of QL_k working from the foundation laid in [4]. The other area for future work is applying the above decision procedure to a more general class of logics. The process of semantic untangling and syntactic entangling appears to have potential for generating a large class of propositionally quantified multi-modal logics. Furthermore, interpreting propositional quantification with respect to bisimulations will maintain the natural abstractions of the modal logics.

References

1. E. Clarke and E. Emerson. Synthesis of synchronization skeletons for branching time temporal logic. In *Proc. IBM Workshop on Logic of Programs, Yorktown Heights, NY*, pages 52–71. Springer, Berlin, 1981.
2. E. Emerson and A. Sistla. Deciding full branching time logic. *Information and Control*, 61:175 – 201, 1984.
3. K. Englehardt, R. van der Meyden, and Y. Moses. Knowledge and the logic of local propositions. In *Theoretical Aspects of Rationality and Knowledge, Proceedings of the Seventh Conference*, pages 29–41, 1998.
4. K. Englehardt, R. van der Meyden, and K. Su. Modal logics with a linear hierarchy of local propositional quantifiers. In *Proceedings of AiML 2002*, to appear.
5. R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about knowledge*. MIT Press, 1995.
6. K. Fine. Propositional quantifiers in modal logic. *Theoria*, 36:336–346, 1970.
7. T. French. Decidability of quantified propositional branching time logics. In *Proceedings of the 14th Australian Joint Conference on Artificial Intelligence*, pages 165–176, 2001.
8. T. French. *The Decidability of Modal Logics with Bisimulation Quantifiers*. PhD thesis, Murdoch University, in preparation.
9. D. Kaplan. S5 with quantifiable propositional variables. *Journal of Symbolic Logic*, 35:355, 1970.
10. R. Milner. A calculus of communicating systems. *Lecture Notes in Computer Science*, 92, 1980.
11. David Park. Concurrency and automata on infinite sequences. *Lecture Notes in Computer Science*, 104:167–183, 1981.