A Quantified Epistemic Logic for Reasoning about Multi-Agent Systems*

F. Belardinelli
Department of Computing
Imperial College London
180 Queen's Gate
London SW7 2AZ
F.Belardinelli@imperial.ac.uk

A. Lomuscio
Department of Computing
Imperial College London
180 Queen's Gate
London SW7 2AZ
A.Lomuscio@imperial.ac.uk

Categories and Subject Descriptors

F.4.1 [**Theory of Computation**]: Mathematical Logic—*modal logic*; I.2.4 [**Artificial Intelligence**]: Knowledge Representation—*modal logic, predicate logic*

General Terms

Languages, Theory.

Keywords

Epistemic logic, First-order logic, Completeness.

ABSTRACT

We investigate quantified interpreted systems, a semantics for multi-agent systems in which agents can reason about individuals, their properties, and the relationships among them. We analyse a first-order epistemic language interpreted on this semantics and show soundness and completeness of $Q.S5_n$, an axiomatisation for these structures.

1. INTRODUCTION

Modal epistemic logic has been widely studied to reason about multi-agent systems (MAS), often in combination with temporal modalities. The typical language extends propositional logic by adding n modalities K_i representing the knowledge of agent i, as well as other modalities representing various mental states (explicit knowledge, beliefs, etc) and/or the flow of time. The use of modal propositional logic as a specification language requires little justification: it is a rather expressive language, well-understood from a theoretical point of view.

Still, it is hard to counterargue the remark, often raised from practitioners in Software Engineering, that quantification in specifications is so natural and convenient that it really should be brought explicitly into the language. Even when working with finite domains of individuals, without quantification one is forced to introduce ad-hoc propositions to emulate basic relations between individuals (as to express specifications like "the child of process p can send a message to all the processes that are allowed to invoke p"). In open MAS individuals may appear and disappear in the system at run-time, making the case for quantification even more compelling. Additionally, in a quantified modal language epistemic operators may be combined with quantifiers to express concepts such as *de re/de dicto* knowledge.

However, the use of first-order modal logic in MAS specifications is normally frown upon by theoreticians. Why should we use an undecidable language when a decidable one does the job already? Is the price that quantification brings in justified? While these objections are certainly sensible, we believe their strength has been increasingly weakened by recent progress in the area of MAS verification [8, 12, 13] by model checking. In the model checking approach [1] we do not check whether a formula representing a specification is satisfiable in some model based on the completeness class, but simply whether a formula is *true* on *the* model representing all possible evolutions of the system. While the former problem is undecidable for first-order modal logic (see [7]), the latter is decidable at least for some suitable fragments.

This paper takes inspiration from the considerations above and aims at making progress on the subject of first-order epistemic logic. The main contribution is the sound and complete axiomatisation of quantified interpreted systems (QIS) in Section 5. QIS are an extension to first order of Interpreted Systems semantics, the usual formalism for epistemic logic in MAS [4]. While we are not aware of completeness results of this nature on the subject, quantified modal logic (QML) has been discussed in MAS settings before. In [4] QML and its Kripke semantics are briefly discussed. In [10] the authors introduce a quantified logic of belief with doxastic modalities indexed to terms of a first-order language. In [14] a quantified temporal epistemic logic is discussed. On related subjects, in [2, 3] Cohen and Levesque develop a first-order logic of belief and action with quantification over agents, although the semantics is not given in terms of computationally grounded structures [15].

In all works above completeness is not tackled. This may be due to the technical difficulties associated with QML and the relatively poor status of the metatheoretical investigation in comparison with the propositional case. We hope the present contribution will be the first in a line of work in which a systematic analysis of these logics is provided.

*The present research was supported by the Royal Society.

AAMAS'07 May 14–18 2007, Honolulu, Hawai'i, USA.

2. SYSTEMS OF GLOBAL STATES

This paper is primarily concerned with the representation of knowledge in MAS, not their temporal evolution. Given this, we adopt the "static" perspective on the systems of global states [11], rather than their "dynamic" version [4]. More formally, consider a set L_i of local states $l_i, l'_i, \ldots$, for each agent $i \in A$, and a set L_e containing the states of the environment $l_e, l'_e, \ldots$; then define a system of global states as follows:

Definition 1. *A system of global states $\mathcal{S}$ - or SGS in short - is a pair $\langle S, D \rangle$ such that $S \subseteq L_e \times L_1 \times \ldots \times L_n$ and D is a non-empty domain of individuals.*

For $s = \langle l_e, l_1, \ldots, l_n \rangle \in S$, s_i is equal to l_i, for $i \in A$. We denote by $\mathcal{SGS}$ the class of the systems of global states.

Remarks: This definition of SGS is grounded on two assumptions. First, the domain D of individuals is the same for every agent i, so all the agents reason about the same objects. This choice is justified by the *external* account of knowledge adopted in the framework of interpreted systems. Second, the domain D is assumed to be the same for every global state, i.e. no individual appears nor disappears in moving from one state to another. Again, this is also consistent with the external account of knowledge. We discuss further options in Section 6. Finally, note that it can be the case that $A \subseteq D$. This means that the agents can reason about themselves, their properties and relationships.

3. SYNTAX AND SEMANTICS

First-order multi-modal formulas are defined on an alphabet containing the variables $x_1, x_2, \ldots$, the n-ary functors $f_1^n, f_2^n, \ldots$, and the n-ary predicative constants $P_1^n, P_2^n, \ldots$, for $n \in \mathbb{N}$, the identity $=$, the propositional connectives $\neg$ and $\rightarrow$, the universal quantifier $\forall$ and the epistemic operators K_i, for every $i \in A$. Terms and formulas in the language $\mathcal{L}_n$ are defined as follows:

$$t ::= x \mid f^k(t_1, \ldots, t_k)$$
$$\phi ::= P^k(t_1, \ldots, t_k) \mid t = t' \mid \neg\phi \mid \phi \rightarrow \psi \mid K_i\phi \mid \forall x\phi$$

Remarks: The symbols $\bot$, $\wedge$, $\vee$, $\leftrightarrow$ and $\exists$ are defined in the standard way. By $t[\vec{y}/\vec{t}]$ (resp. $\phi[\vec{y}/\vec{t}]$) we denote the term (resp. formula) obtained by simultaneously substituting some, possibly all, free occurrences of $y_1, \ldots, y_n$ in t (resp. ϕ) with $t_1, \ldots, t_n$, renaming bounded variables if necessary. By Var we denote the set of variables in $\mathcal{L}_n$.

We interpret the language $\mathcal{L}_n$ on a system of global states $\mathcal{S}$ by means of a function I mapping the syntactic features of $\mathcal{L}_n$ to the elements of $\mathcal{S}$.

Definition 2. *Given an SGS $\mathcal{S}$, a quantified interpreted system - or QIS in short - is a pair $\mathcal{P} = \langle \mathcal{S}, I \rangle$ such that $I(f^k)$ is a k-ary function from D^k to D; for every $s \in S$, $I(P^k, s)$ is a subset of D^k and $I(=, s)$ is the equality on D.*

Note that functors are interpreted rigidly. Let σ be an assignment, i.e. any function from Var to D, the valuation $I^\sigma(t)$ of a term t is inductively defined as follows:

$$I^\sigma(y) = \sigma(y)$$
$$I^\sigma(f^k(t_1, \ldots, t_k)) = I(f^k)(I^\sigma(t_1), \ldots, I^\sigma(t_k))$$

A variant $\sigma\binom{x}{a}$ of an assignment σ differs from σ at most on x and assigns element $a \in D$ to x.

Definition 3 (Satisfaction). *The satisfaction relation $\models$ for $\phi \in \mathcal{L}_n$, $s \in \mathcal{P}$ and assignment σ is defined as follows:*

$(\mathcal{P}^\sigma, s) \models P^k(t_1, \ldots, t_k)$ iff $\langle I^\sigma(t_1), \ldots, I^\sigma(t_k) \rangle \in I(P^k, s)$
$(\mathcal{P}^\sigma, s) \models t = t'$ iff $I^\sigma(t) = I^\sigma(t')$
$(\mathcal{P}^\sigma, s) \models \neg\psi$ iff $(\mathcal{P}^\sigma, s) \not\models \psi$
$(\mathcal{P}^\sigma, s) \models \phi \rightarrow \psi$ iff $(\mathcal{P}^\sigma, s) \not\models \phi$ or $(\mathcal{P}^\sigma, s) \models \psi$
$(\mathcal{P}^\sigma, s) \models K_i\psi$ iff $s_i = s'_i$ implies $(\mathcal{P}^\sigma, s') \models \psi$
$(\mathcal{P}^\sigma, s) \models \forall x\psi$ iff for every $a \in D, (\mathcal{P}^{\sigma\binom{x}{a}}, s) \models \psi$

The truth conditions for formulas containing $\bot$, $\wedge$, $\vee$, $\leftrightarrow$, $\exists$ are defined as usual. A formula ϕ in $\mathcal{L}_n$ is *true at a state s* iff it is satisfied at s by every assignment σ; ϕ is *valid on a QIS $\mathcal{P}$* iff it is true at every state in $\mathcal{P}$; ϕ is *valid on a SGS $\mathcal{S}$* iff it is valid on every QIS on $\mathcal{S}$; ϕ is *valid on a class $\mathcal{C}$ of SGS* iff it is valid on every SGS in $\mathcal{C}$.

4. EXPRESSIVENESS

Note that in the language $\mathcal{L}_n$ we can express an agent's knowledge of properties and relationships among individuals. Consider the following specifications:

1. agent a knows that for every process x, agent b knows that there exists a precondition y, which has to be fulfilled in order for x to start.

2. agent a knows that there exists an input x for which agent b does not know that every computation y on input x fails.

3. agent c knows that not every agent is identical to d; in particular, she knows that she is not identical to d.

These statements can be formalised as follows:

1. $K_a\forall x(Proc(x) \rightarrow K_b\exists y(Pre(y) \wedge (St(x) \rightarrow Fulfil(y))))$
2. $K_a\exists x(Input(x) \wedge \neg K_b\forall y(Comp(y) \rightarrow Fails(x, y)))$
3. $K_c\neg\forall x(Ag(x) \rightarrow x = agent\text{-}d) \wedge K_c(agent\text{-}c \neq agent\text{-}d)$

Clearly, in this framework one can model the knowledge agents have about themselves. In addition, we retain all the expressive power of propositional epistemic logic. Furthermore, we can now express the *de re/de dicto* distinction, that is, the difference between knowing something of someone and knowing that someone is something. For instance, when we use an informal specification to say that, as far as a security controller is concerned, every user is authorised to access the site, one could interpret this as (hence implement it!) either *de dicto*, i.e. descriptively:

a) the security controller knows that every user is authorised to access the site,

or *de re*, i.e. prescriptively:

b) for every user, the security controller knows that he is authorised to access the site.

These two readings express different concepts. While these cannot be easily separated by means of a propositional language, in $\mathcal{L}_n$ this is succintly done as follows:

a) $K_{SecCont}\forall x(Auth\text{-}user(x) \rightarrow Access(x))$

b) $\forall x(Auth\text{-}user(x) \rightarrow K_{SecCont}Access(x))$

The difference in meaning between the two specifications is clear. For instance, the security controller not granting access to an authorised user u is a violation of (b), but not of (a), if he does not regard u to be an authorised user.

5. AXIOMATISATION

The system $Q.S5_n$ on the language $\mathcal{L}_n$ is a first-order multi-modal version of the normal propositional system $S5$. Hereafter we list its axioms; note that $\Rightarrow$ is the inference relation between formulas.

DEFINITION 4. *The system $Q.S5_n$ on $\mathcal{L}_n$ contains the following schemes of axioms and inference rules:*

Taut	*every classic propositional tautology*
Dist	$K_i(\phi \to \psi) \to (K_i\phi \to K_i\psi)$
T	$K_i\phi \to \phi$
4	$K_i\phi \to K_iK_i\phi$
5	$\neg K_i\phi \to K_i\neg K_i\phi$
MP	$\phi \to \psi, \phi \Rightarrow \psi$
Nec	$\phi \Rightarrow K_i\phi$
Ex	$\forall x\phi \to \phi[x/t]$
Gen	$\phi \to \psi[x/t] \Rightarrow \phi \to \forall x\psi$*, x not free in* ϕ
Id	$t = t$
Func	$t = t' \to (t''[x/t] = t''[x/t'])$
Subst	$t = t' \to (\phi[x/t] \to \phi[x/t'])$

It is easy to check that every axiom in $Q.S5_n$ is valid on any system of global states and its rules preserve validity. We can also show that the axioms and inference rules in $Q.S5_n$ are sufficient to prove all the validities on $\mathcal{SGS}$. This result is obtained by extending the completeness proofs for first-order modal logic in [6, 9]. By combining together soundness and completeness we obtain the main result of this paper.

THEOREM 1. *A formula ϕ is valid on the class $\mathcal{SGS}$ of systems of global states iff ϕ is provable in $Q.S5_n$.*

6. CONCLUSIONS

It is clear that first-order modal formalisms offer expressivity advantages over propositional modal ones. However, the specialised literature has so far fallen short of a deep and systematic analysis of the machinery, even in the case of static epistemic logic.

In this paper we believe we have made a first attempt in this direction: the axiomatisation presented shows that the popular system $S5_n$ extends naturally to the first-order case. In carrying out this exercise we tried to remain as close as possible to the original epistemic logic's semantics of interpreted systems, so that fine grained specifications of MAS can be expressed as recent work on model checking interpreted systems demonstrates [5, 13].

Different extensions of the present framework seem pursuing. First of all, it seems interesting to relax the assumption on the domain of quantification and admit a different domain $D(s)$, for every state s. Further, we could assume a different domain of quantification $D_i(s)$ for every agent i in a state s. It would also be of interest to explore the completeness issues resulting from term-indexing of the epistemic operators [10]. In an orthogonal dimension to the above another significant extension would be to add temporal operators to the formalism. This would pave the way for an exploration of axiomatisations for temporal/epistemic logic for MAS.

7. REFERENCES

[1] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, Cambridge, Massachusetts, 1999.

[2] P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[3] P. R. Cohen and H. J. Levesque. Rational interaction as the basis for communication. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, Cambridge, MA, 1990.

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

[5] P. Gammie and R. van der Meyden. Mck: Model checking the logic of knowledge. In R. Alur and D. Peled, editors, *CAV*, volume 3114 of *Lecture Notes in Computer Science*, pages 479–483. Springer, 2004.

[6] J. Garson. Quantification in modal logic. In *Handbook of Philosophical Logic*, volume 2, pages 249–307. D. Gabbay and F. Guenthner Eds., 1984.

[7] I. M. Hodkinson, F. Wolter, and M. Zakharyaschev. Decidable fragment of first-order temporal logics. *Ann. Pure Appl. Logic*, 106(1-3):85–134, 2000.

[8] W. Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In M. Gini, T. Ishida, C. Castelfranchi, and W. L. Johnson, editors, *Proceedings of AAMAS'02*, pages 1167–1174. ACM Press, 2002.

[9] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, New York, 1996.

[10] A. Lomuscio and M. Colombetti. QLB: a quantified logic for belief. In J. Müller, M. Wooldridge, and N. Jennings, editors, *Proceedings of ATAL-96*, volume 1193 of *Lecture Notes in AI*. Springer-Verlag, Heidelberg, 1996.

[11] A. Lomuscio and M. Ryan. On the relation between interpreted systems and kripke models. In Springer lecture notes in AI, volume 1441, 1997.

[12] J.-J. C. Meyer and W. Hoek. *Epistemic Logic for AI and Computer Science*, volume 41 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 1995.

[13] F. Raimondi and A. Lomuscio. Automatic verification of multi-agent systems by model checking via OBDDs. *Journal of Applied Logic*, 2005. To appear in Special issue on Logic-based agent verification.

[14] M. Wooldridge. *The logical modelling of computational multi-agent systems*. PhD thesis, University of Manchester, Faculy of Technology, 1992.

[15] M. Wooldridge. Computationally grounded theories of agency. In E. Durfee, editor, *Proceedings of ICMAS, International Conference of Multi-Agent Systems*, pages 13–22. IEEE Press, 2000.