A Complete Quantified Epistemic Logic for Reasoning about Message Passing Systems

Francesco Belardinelli[1] and Alessio Lomuscio[2]

[1] Scuola Normale Superiore, Pisa
belardinelli@sns.it
[2] Department of Computing
Imperial College London, UK
A.Lomuscio@imperial.ac.uk

Abstract. We investigate quantified interpreted systems, a semantics to model multi-agent systems in which the agents can reason about individuals, their properties, and relationships among them. The semantics naturally extends interpreted systems to first-order by introducing a domain of individuals. We present a first-order epistemic language interpreted on this semantics and prove soundness and completeness of the quantified modal system $QS5_n^D$, an axiomatisation for these structures. Finally, we exemplify the use of the logic by modeling message passing systems, a relevant class of interpreted systems analysed in epistemic logic.

1 Introduction

Modal epistemic logic has been widely studied in multi-agent systems (MAS) both on its own and in combination with other modalities, very often temporal ones. The typical language extends propositional logic by adding n modalities K_i representing the knowledge of agent i, as well as other modalities representing different mental states for the agents (distributed and common knowledge, beliefs, etc) and/or the temporal flow of time [10,25].

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 by 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 among individuals. Not always quantification is simply syntactic sugar: certain expressivity needs do require infinite domains (e.g., see section 4 below).

In multi-agent systems the power of first-order logic is required every time agents reason about:

– relational statement, as in *agent i knows that message μ has been sent by a to b*, or formally

$$K_i Sent(a, b, \mu)$$

F. Sadri and K. Satoh (Eds.): CLIMA VIII, LNAI 5056, pp. 248–267, 2008.

– functional dependency or identity: *agent i knows that message μ is the encryption of message μ' with key k*, formally

$$K_i(\mu = enc(k, \mu'))$$

– an infinite domain of individuals, or a finite domain whose cardinality cannot be bounded in advance: *agent i has to read all e-mails before deleting them*, or formally

$$\forall \mu(Delete(i, \mu) \rightarrow Read(i, \mu))$$

– quantification on agents' indexes [21]: *all agents knows..., at least one agent knows...*

$$\forall i K_i \ldots, \exists i K_i \ldots$$

Further, epistemic modalities can be combined with quantifiers to express concepts such as knowledge *de re*/*de dicto* [12].

Irrespective of the above, the use of first-order modal logic in MAS specifications is normally frowned upon by theoreticians. Why should we use an undecidable language when a decidable one does the job reasonably well already? Is the price that quantification brings in justified? While these objections are certainly sensible, we believe that their strength has been increasingly weakened by recent progress in the verification of MAS by model checking [13,28,26,23]. In the model checking approach [8] the decision problem is tackled not by checking validity but simply model satisfaction. In other words, we do not check whether a formula representing a specification is satisfiable, but simply whether it is true on the model representing all possible evolutions of the system. While the former problem is undecidable for first-order modal logic, the latter is decidable at least in some suitable fragments, such as the *monodic* fragments studied in [16,32,33,35]. Moreover, we have specification languages supporting first-order interval temporal logic [29,30]. Recently, first-order modal logic has been applied to the analysis of security protocols [1,5,9]. Finally, we have some preliminary works on first-order model checking [31,34].

This paper takes inspiration from the considerations above and aims at making progress on the subject of first-order epistemic logic. The main contribution of the paper is the axiomatisation in section 5, where a sound and complete system for quantified interpreted systems (QIS) is presented. We argue that QIS are the natural extension to first-order of Interpreted Systems semantics, the usual formalism for epistemic logic in MAS [10,25].

While completeness results for quantified modal logic are customarily proved with respect to Kripke semantics [12,18], we should state clearly that QML has been discussed in MAS settings before. In [10] quantified epistemic logic is briefly discussed, along with its Kripke semantics and some significant validities; in [21] the authors introduce a quantified logic of belief, in which the doxastic modalities are indexed to terms of a first-order language; in [2] a limited form of quantification is added to Coalition Logic. However, in most of the 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 this contribution will be the first in a line of work in which a systematic analysis of these logics is provided.

Scheme of the paper. In section 2 we present two classes of first-order structures: systems of global states and Kripke frames. In section 3 we introduce the first-order modal language $\mathcal{L}_n^D$ which is interpreted on quantified interpreted systems, a valued version of the systems of global states. In section 4 we exemplify syntax and semantics by describing three formal models for multi-agent systems and discuss some specification patterns. In section 5 we introduce the first-order modal system $Q.S5_n^D$, and prove the main result of this paper: $Q.S5_n^D$ is a sound and complete axiomatisation of the validities in the structures of global states. Finally, section 6 outlines some extensions of the present formalism.

2 Systems of Global States and Kripke Frames

In this section we introduce the systems of global states and Kripke frames in a first-order setting. While the first ones are used in computer science to model the behaviour of MAS [10,14,25], Kripke frames are best employed to get a deeper understanding of the formal properties of these systems [6,7]. Technically, we extend the corresponding propositional structures to first-order. This extension is not trivial, as there are many ways of performing it: for instance, we can choose a single domain of quantification or several domains for each agent and/or for each computational state, not to mention domains of intensional objects [4]. In this paper we consider the simplest construction, where we have just a single quantification domain D common to all the agents and states, which contains all possible objects. We leave other options for further work. In what follows we assume a set of agents $A = \{1, \ldots, n\}$.

2.1 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 [22], rather than the "dynamic" version [10]. So, while we assume that the states of the system result from the evolution given by protocols and transitions, for the time being we do not consider them explicitly. 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 local states of the environment $l_e, l_e', \ldots$. We define a system of global states as follows:

Definition 1 (SGS). *A system of global states $\mathcal{S}$ is a couple $\langle S, D\rangle$ such that $S \subseteq L_e \times L_1 \times \ldots \times L_n$ is a non-empty set of global states, and D is a non-empty domain of individuals. $\mathcal{SGS}$ is the class of the systems of global states.*

This definition of SGS is based on two assumptions. First, the domain D of individuals is the same for every agent i, so that all agents effectively reason about the same objects. This choice is justified by the *external account of knowledge*

usually adopted in the framework of interpreted systems. If knowledge is ascribed to agents by an external observer, it seems natural to focus on a unique set of individuals: the ones assumed to exist by the external observer. 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. This also is consistent with the external account of knowledge: all individuals are supposed to be existing from the observer's viewpoint. We discuss further options in section 6. Finally, it can be the case that $A \subseteq D$. This means that the agents can reason about themselves, their properties, and relationships.

2.2 Kripke Frames

While Kripke frames are less intuitive than interpreted systems to model MAS, they are more convenient for the purpose of formal analysis, namely completeness investigations. We work with frames with equivalence relations, so we take the following definition:

Definition 2. *A Kripke frame $\mathcal{F}$ is a $n+2$-tuple $\langle W, \sim_1, \ldots, \sim_n, D\rangle$ such that W is a non-empty set; for $i \in A$, $\sim_i$ is an equivalence relation on W; D is a non-empty set of individuals. $\mathcal{K}$ is the class of all Kripke frames.*

Now we have systems of global states modelling MAS and Kripke frames. In order to axiomatise $\mathcal{SGS}$, it is useful to map SGS into Kripke frames.

2.3 Maps between $\mathcal{SGS}$ and $\mathcal{K}$

We explore the relationship between these structures by means of two maps f and g from $\mathcal{SGS}$ to $\mathcal{K}$ and viceversa. We show that every SGS $\mathcal{S}$ is isomorphic to $g(f(\mathcal{S}))$, that is, there is a one-to-one correspondence onto the sets of global states and the domains of individuals. Further, we prove that every Kripke frame $\mathcal{F} = \langle W, \sim_1, \ldots, \sim_n, D\rangle$ is isomorphic to $f(g(\mathcal{F})) = \langle W', \sim'_1, \ldots, \sim'_n, D'\rangle$, that is, there are bijections between W and W' and between D and D'; in addition $w \sim_i w'$ iff $(f \circ g)(w) \sim'_i (f \circ g)(w')$. As a result, every sound and complete axiomatisation of Kripke frames is also an axiomatisation of SGS.

We start with the map $f : \mathcal{SGS} \to \mathcal{K}$. Let $\mathcal{S} = \langle S, D\rangle$ be an SGS, define $f(\mathcal{S})$ as the $n+2$-tuple $\langle S, \sim_1, \ldots, \sim_n, D\rangle$, where S is the set of possible states and D is the domain of individuals. Moreover, for each $i \in A$, the relation $\sim_i$ on S such that $\langle l_e, l_1, \ldots, l_n\rangle \sim_i \langle l'_e, l'_1, \ldots, l'_n\rangle$ iff $l_i = l'_i$ is an equivalence relation. So $f(\mathcal{S})$ is a Kripke frame.

For the converse map $g : \mathcal{K} \to \mathcal{SGS}$, let $\mathcal{F} = \langle W, \sim_1, \ldots, \sim_n, D\rangle$ be a Kripke frame. For every epistemic state $w \in W$, let the equivalence class $[w]_{\sim_i} = \{w' | w \sim_i w'\}$ be a local state for agent i, and W is the set of local states for the environment. Define $g(\mathcal{F}) = \langle S, D\rangle$, where S contains all the $n+1$-tuples $\langle w, [w]_{\sim_1}, \ldots, [w]_{\sim_n}\rangle$, for $w \in W$, while D is defined as above. The structure $g(\mathcal{F})$ is trivially an SGS.

We prove that the composition of the two maps gives isomorphic structures.

Lemma 1. *Every Kripke frame $\mathcal{F}$ is isomorphic to $f(g(\mathcal{F}))$.*

Proof. If $\mathcal{F} = \langle W, \sim_1, \ldots, \sim_n, D\rangle$ is a Kripke frame, then $f(g(\mathcal{F})) = \langle W', \sim'_1, \ldots, \sim'_n, D\rangle$ is such that W' is the set of n+1-tuples $\langle w, [w]_{\sim_1}, \ldots, [w]_{\sim_n}\rangle$, for $w \in W$. The composition $f \circ g$ is a bijection between W and W': it is one-to-one as if $w, w' \in W'$ and $w = w'$, then in particular the first components of w and w' are equal. It is onto as the first component w_1 of $w \in W'$ is such that $w_1 \in W$ and $f(g(w_1)) = w$. Also, the identity on D is a bijection. Finally, $w \sim_i w'$ iff $[w]_{\sim_i} = [w']_{\sim_i}$ iff $\langle w, [w]_{\sim_1}, \ldots, [w]_{\sim_n}\rangle \sim'_i \langle w', [w']_{\sim_1}, \ldots, [w']_{\sim_n}\rangle$. Thus, the two structures are isomorphic. □

By lemma 1 we will show in section 5 that a sound and complete axiomatisation of Kripke frames is adequate also with respect to SGS.

3 Syntax and Semantics

In this section we introduce the first-order multi-modal language $\mathcal{L}_n^D$ containing individual variables and constants, as well as quantifiers, n epistemic operators, the distributed knowledge operator, and identity. The language $\mathcal{L}_n^D$ is interpreted on models based on Kripke frames. Finally, we present the quantified interpreted systems, a valued version of the systems of global states.

3.1 Syntax

The first-order multi-modal language $\mathcal{L}_n^D$ contains individual variables $x_1, x_2, \ldots$, n-ary functors $f_1^n, f_2^n, \ldots$ and n-ary predicative letters $P_1^n, P_2^n, \ldots$, for $n \in \mathbb{N}$, the identity predicate $=$, the propositional connectives $\neg$ and $\rightarrow$, the universal quantifier $\forall$, the epistemic operators K_i, for $i \in A$, and the distributed knowledge operator D_G, for $G \subseteq A$.

Definition 3. *Terms and formulas in the language $\mathcal{L}_n^D$ are defined in the Backus-Naur form as follows:*

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

Intuitively, the formula $K_i\phi$ means that *agent i knows* ϕ, while $D_G\phi$ is read as *ϕ is distributed knowledge among the agents in G*. The symbols $\bot$, $\wedge$, $\vee$, $\leftrightarrow$ and $\exists$ are defined by means of the other logical constants; we refer to the 0-ary functors as individual constants $c_1, c_2, \ldots$ A closed term v is a term where no variable appears, the closed terms are only constants and terms obtained by applying functors to closed terms. Finally, by $t[\vec{y}]$ (resp. $\phi[\vec{y}]$) we mean that $\vec{y} = y_1, \ldots, y_n$ are all the free variables in t (resp. ϕ); while $t[\vec{y}/\vec{t}]$ (resp. $\phi[\vec{y}/\vec{t}]$) denotes the term (resp. formula) obtained by simultaneously substituting some, possibly all, free occurrences of $\vec{y}$ in t (resp. ϕ) with $\vec{t} = t_1, \ldots, t_n$, renaming bounded variables if necessary.

3.2 Semantics

In order to assign a meaning to the formulas in $\mathcal{L}_n^D$ we make use of Kripke models. We then define validity on quantified interpreted systems in terms of validity on Kripke models.

Definition 4 (model). *A Kripke model $\mathcal{M}$ - or simply K-model - based on a Kripke frame $\mathcal{F}$, is a couple $\langle \mathcal{F}, I \rangle$ where I is an interpretation such that:*

- *if f^k is a k-ary functor, then $I(f^k)$ is a function from D^k to D;*
- *if P^k is a k-ary predicative letter and $w \in W$, then $I(P^k, w)$ is a k-ary relation on D, i.e. $I(P^k, w) \subseteq D^k$;*
- *the interpretation $I(=, w)$ of the identity $=$ in w is the equality on D.*

Note that function symbols are interpreted rigidly, that is, for every $w, w' \in W$ the interpretation of a functor f^k in w is the same as the interpretation of f^k in w'. Given that our approach is the one of the external observer, rigid designators seem appropriate.

Now let σ be an assignment, i.e., any function from the set of variables in $\mathcal{L}_n^D$ to the domain D, the valuation $I^\sigma(t)$ of a term t is defined as $\sigma(y)$ for $t = y$, and $I^\sigma(t) = I(f^k)(I^\sigma(t_1), \ldots, I^\sigma(t_k))$, for $t = f^k(t_1, \ldots, t_k)$. In particular, the valuation $I^\sigma(d)$ of a constant d is an individual $I(d)$ in D. The variant $\sigma\binom{x}{a}$ of the assignment σ differs from σ at most on x and assigns element $a \in D$ to x. Now we are able to define the truth conditions for the formulas in $\mathcal{L}_n^D$.

Definition 5 (Satisfaction). *The satisfaction relation $\models$ for a formula $\phi \in \mathcal{L}_n^D$, a world $w \in \mathcal{M}$ and an assignment σ is inductively defined as follows:*

$$
\begin{array}{lll}
(\mathcal{M}^\sigma, w) \models P^k(\vec{t}) & \text{iff} & \langle I^\sigma(t_1), \ldots, I^\sigma(t_k) \rangle \in I(P^k, w) \\
(\mathcal{M}^\sigma, w) \models t = t' & \text{iff} & I^\sigma(t) = I^\sigma(t') \\
(\mathcal{M}^\sigma, w) \models \neg\psi & \text{iff} & (\mathcal{M}^\sigma, w) \not\models \psi \\
(\mathcal{M}^\sigma, w) \models \psi \rightarrow \psi' & \text{iff} & (\mathcal{M}^\sigma, w) \not\models \psi \text{ or } (\mathcal{M}^\sigma, w) \models \psi' \\
(\mathcal{M}^\sigma, w) \models K_i\psi & \text{iff} & \text{for } w' \in W,\ w \sim_i w' \text{ implies } (\mathcal{M}^\sigma, w') \models \psi \\
(\mathcal{M}^\sigma, w) \models D_G\psi & \text{iff} & \text{for } w' \in W,\ (w, w') \in \bigcap_{i \in G} \sim_i \text{ implies } (\mathcal{M}^\sigma, w') \models \psi \\
(\mathcal{M}^\sigma, w) \models \forall x\psi & \text{iff} & \text{for all } a \in D, (\mathcal{M}^{\sigma\binom{x}{a}}, w) \models \psi
\end{array}
$$

The truth conditions for the formulas containing the symbols $\bot$ $\wedge$, $\vee$, $\leftrightarrow$ and $\exists$ are standardly defined from those above. Further, a formula ϕ in $\mathcal{L}_n^D$ is said to be *true at a world w* iff it is satisfied at w by every assignment σ; ϕ is *valid on a model $\mathcal{M}$* iff it is true at every world in $\mathcal{M}$; ϕ is *valid on a frame $\mathcal{F}$* iff it is valid on every model on $\mathcal{F}$; ϕ is *valid on a class $\mathcal{C}$ of frames* iff it is valid on every frame in $\mathcal{C}$.

Now we have all preliminary definitions to introduce quantified interpreted systems.

Definition 6 (QIS). *A quantified interpreted systems $\mathcal{P}$ based on an SGS $\mathcal{S}$, is a couple $\langle \mathcal{S}, I \rangle$ such that I is an interpretation of $\mathcal{L}_n^D$ in the Kripke frame $f(\mathcal{S})$.*

The notions of satisfaction, truth and validity are defined as above, i.e., let $\mathcal{P}_f = \langle f(\mathcal{S}), I\rangle$ be the Kripke model for the quantified interpreted system $\mathcal{P} = \langle \mathcal{S}, I\rangle$, then $(\mathcal{P}^\sigma, s) \models \phi$ iff $(\mathcal{P}^\sigma_f, s) \models \phi$. In particular, the present definition of satisfaction agrees with the usual definition for interpreted systems:

$$
\begin{array}{lll}
(\mathcal{P}^\sigma, s) \models P^k(t_1, \ldots, t_k) & \text{iff} & \langle I^\sigma(t_1), \ldots, I^\sigma(t_k)\rangle \in I(P^k, s) \\
(\mathcal{P}^\sigma, s) \models K_i\psi & \text{iff} & l_i(s) = l_i(s') \text{ implies } (\mathcal{P}^\sigma, s') \models \psi \\
(\mathcal{P}^\sigma, s) \models D_G\psi & \text{iff} & l_i(s) = l_i(s') \text{ for all } i \in G \text{ implies } (\mathcal{P}^\sigma, s') \models \psi \\
(\mathcal{P}^\sigma, s) \models \forall x\psi & \text{iff} & \text{for all } a \in D,\ (\mathcal{P}^{\sigma\binom{a}{x}}, s') \models \psi
\end{array}
$$

Moreover, a formula $\phi \in \mathcal{L}_n^D$ is valid on a quantified interpreted systems $\mathcal{P}$ iff ϕ is valid on $\mathcal{P}_f$, or more formally:

Definition 7 (Validity on QIS). *If ϕ is a formula in $\mathcal{L}_n^D$ and $\mathcal{P}$ is a quantified interpreted systems, then $\mathcal{P} \models \phi$ iff $\mathcal{P}_f \models \phi$.*

Thus, we can reason about a multi-agent system by using the expressiveness of QIS, but rely on Kripke models to prove formal properties of the system.

3.3 Some Validities

Clearly, the language $\mathcal{L}_n^D$ is very expressive. We can specify the knowledge agents have of facts about individuals, as in the following specification: *agent i knows that someone sent him a message when he receives it*,

$$\forall j, \mu(Recd(i, j, \mu) \rightarrow K_i Sent(j, i, \mu)) \tag{1}$$

This specification can be expressed also in some propositional modal languages. However, in $\mathcal{L}_n^D$ we can make more subtle distinctions as in *if agent i receives a message, then he knows that someone sent it to him*:

$$\forall \mu(\exists j Recd(i, j, \mu) \rightarrow K_i\ \exists j' Sent(j', i, \mu)) \tag{2}$$

The latter specification is weaker than the former, as (2) says nothing about the identity of the sender, while (1) requires that the receiver knows the identity of the sender.

We briefly explore the semantics of QIS by considering the traditional Barcan formulas [12]. There has been much discussion on these principles and their soundness in epistemic contexts. Given that the domain of quantification is the same for every global state, both the Barcan formula and its converse are valid on the class $\mathcal{QIS}$ of all QIS, i.e., they hold in every quantified interpreted system:

$$
\begin{array}{lll}
\mathcal{QIS} \models \forall x K_i\phi \rightarrow K_i \forall x\phi & \mathcal{QIS} \models \forall x D_G\phi \rightarrow D_G \forall x\phi & BF \\
\mathcal{QIS} \models K_i \forall x\phi \rightarrow \forall x K_i\phi & \mathcal{QIS} \models D_G \forall x\phi \rightarrow \forall x D_G\phi & CBF
\end{array}
$$

We remarked that these formulas are direct consequences of having fixed domains, which were justified by the external account of knowledge usually adopted in epistemic logic: the domain of quantification consists of the individuals considered by the designer. Also, we deem these validities in line with the bird's

eye approach of epistemic logic. By BF if agent i knows that a is ϕ for each individual $a \in D$, then she knows that all the individuals are ϕ. In fact, in any epistemic alternative considered by agent i *at most* the individuals she currently considers (i.e., those in D) are present. In other words, agents are assumed to be able to generalise their knowledge, at least when this is considered from an external point of view. By CBF if agent i knows that all the individuals are ϕ, then she knows that a is ϕ, for each individual $a \in D$. This happens because in any epistemic alternative considered by agent i *at least* the individuals she currently considers (again those in D) are present. In other words, agents are assumed to be able to particularise their knowledge.

We stress the fact that in this interpretation the formula $K_i \forall x \phi$ does not mean that *agent i knows that all the individuals she considers are* ϕ, but rather *agent i knows that all the individuals (considered by the external designer) are* ϕ. Further, following the external account of knowledge typical in epistemic logic, the truth of $K_i \forall x \phi$ does not imply that agent i has to be aware of all the individuals considered by the designer. As it is the case in propositional epistemic logic, the formula $\forall x \phi$ expresses the knowledge attributed by the external observer to agent i, rather than the explicit knowledge possessed by i.

We have also generalised versions of the Barcan formula and its converse, for arbitrary strings of epistemic operators:

$$\mathcal{QIS} \models \forall x E_{j_1} \dots E_{j_m} \phi \rightarrow E_{j_1} \dots E_{j_m} \forall x \phi \qquad BF_{j_1,\dots,j_m}$$
$$\mathcal{QIS} \models E_{j_1} \dots E_{j_m} \forall x \phi \rightarrow \forall x E_{j_1} \dots E_{j_m} \phi \qquad CBF_{j_1,\dots,j_m}$$

where each E_{j_k} is either K_i or D_G. Even if these principles seem quite strong, by considering an external notion of knowledge they do not appear problematic.

For what concerns identity, the following principles hold:

$$\mathcal{QIS} \models (t = t') \rightarrow K_i(t = t') \qquad \mathcal{QIS} \models (t \neq t') \rightarrow K_i(t \neq t')$$
$$\mathcal{QIS} \models (t = t') \rightarrow D_G(t = t') \qquad \mathcal{QIS} \models (t \neq t') \rightarrow D_G(t \neq t')$$

These validities result from rigid designation and require further explanation. Suppose message μ' is the encryption of message μ with key k, i.e. $\mu' = enc(k, \mu)$, then by the principles above any agent i should know this identity, that is, for each i, we have $K_i(\mu' = enc(k, \mu))$. But this seems to imply that we cannot represent encryption in this formalism. However, if we assume the *de re* interpretation of modality, we can reconcile encryption and the validities above. In fact, if μ' and $enc(k, \mu)$ are one and the same message, then any agent i knows that this message is identical to itself, which is the *de re* interpretation of $K_i(\mu' = enc(k, \mu))$. Still, agent i may not have explicit, *de dicto* knowledge of the fact that message μ' is obtained by encrypting μ with key k.

4 Message-Passing QIS

In this section we show how to model message-passing systems [10] in the framework of QIS. A m.p. system is a multi-agent system where agents communicate by exchanging messages, so the most relevant events are sending and receiving

messages. The formalism of message passing systems is useful to model a wide range of MAS. For instance, a network of computers, such as the Internet, can be seen as a m.p. system. In general, any multi-agent system is a m.p. system if the message transmission delay is not negligible. In a m.p. system the local state of each agent contains information about its initial state, the messages it has sent or received, as well as the internal actions it has taken.

In what follows we show that m.p. systems can be defined as a particular class of SGS satisfying a finite number of specifications in the first-order modal language $\mathcal{L}_n^D$. Our main result consists in showing that Proposition 4.4.3 in [10], concerning the knowledge of the ordering of events in m.p. systems, can be restated as a validity on the class of QIS modeling m.p. systems. Thus, the formalism of QIS is powerful enough to deal with the theory of m.p. systems. Throughout the section we refer to [10], par. 4.4.5-6, for the details of m.p. systems.

More formally, we introduce a set Act of actions $\alpha_1, \alpha_2, \ldots$, and a set MSG of messages $\mu_1, \mu_2, \ldots$ For each agent $i \in A$, we consider a set Σ_i of initial events $init(i, \alpha)$, and a set INT_i of internal events $int(i, \alpha)$. We define the local state l_i for agent i as a *history* over Σ_i, INT_i and MSG, that is, a sequence of events whose first element is in Σ_i, and whose following elements either belong to INT_i or are events of the form $send(i, j, \mu)$, $rec(i, j, \mu)$ for $j \in A$, $\mu \in MSG$ and $\alpha \in Act$. Intuitively, $init(i, \alpha)$ represents the event where *agent i performs the initial action α*, $send(i, j, \mu)$ represents the event where *agent i sends message μ to j*, while the intended meaning of $rec(i, j, \mu)$ is that *agent i receives message μ from j*. Finally, $int(i, \alpha)$ means that *agent i performs the internal action α*.

A global state s is a tuple $\langle l_e, l_1, \ldots, l_n \rangle$, where $l_1, \ldots, l_n$ are local states as above and l_e contains all the events in $l_1, \ldots, l_n$. We define a reflexive, transitive and anti-symmetric relation $\leq$ on the local states of agent i such that $l_i \leq l_i'$ iff l_i is a prefix of l_i'. This order extends to global states, so that $s \leq s'$ iff $l_i \leq l_i'$ for every $i \in A$. We can define message passing systems as a special class of quantified interpreted systems by considering the class of QIS $\mathcal{P} = \langle S, D, I \rangle$ where S is a non-empty set of global states; the domain D of individuals includes all agents in A, the messages in MSG, the actions in Act, and the events $e_1, e_2, \ldots$; I is the interpretation for $\mathcal{L}_n^D$. Intuitively, each m.p. QIS models the evolution of a m.p. system: starting from an initial state, the m.p. QIS contains the states reachable during the execution of the m.p. system. The temporal evolution of a m.p. QIS can be represented as a sequence $s_0, s_1, \ldots$ of global states such that $s_0 = \langle init(e, \alpha_e), init(1, \alpha_1), \ldots, init(n, \alpha_n) \rangle$, and for every $n \in \mathbb{N}$, either s_{n+1} is identical to s_n or there is an i such that $l_i(s_n) \leq l_i(s_{n+1})$ but $l_i(s_n) \neq l_i(s_{n+1})$. Note that a single m.p. QIS can contain several temporal evolutions of the same m.p. system.

We assume that our language has terms and predicative letters for representing the objects in the domain D and the relations among them. In particular, $e_1, e_2, \ldots$ are metaterms ranging over events: we write $\forall e \phi[e]$ as a shorthand for

$$\forall i, j, \mu, \alpha \; \phi[send(i, j, \mu)] \wedge \phi[rec(i, j, \mu)] \wedge \phi[init(i, \alpha)] \wedge \phi[int(i, \alpha)]$$

In fact, any event is either a send or a receive event, or an initial action or an internal action. We use the same notation for the objects in the model and the syntactic elements, as the ones mirror the others; the distinction will be made clear by the context. We immediately give some examples of the expressiveness of our formalism. In $\mathcal{L}_n^D$ we can define events by formulas which are provably valid in every m.p. QIS (the existence of a unique individual $\exists!$ can be defined by means of $=$):

$$\forall e \exists! i, j, \mu, \alpha \; (i \neq j) \wedge \; (e = send(i, j, \mu) \vee e = rec(i, j, \mu) \vee e = init(i, \alpha) \vee e = int(i, \alpha))$$

$$\begin{aligned}\forall i, j, \mu, \alpha \exists! e_1, e_2, e_3, e_4 (send(i, j, \mu) = e_1 \wedge rec(i, j, \mu) = e_2 \wedge init(i, \alpha) = e_3 \wedge int(i, \alpha) = e_4 \wedge \\ \wedge e_1 \neq e_2 \wedge e_1 \neq e_3 \wedge e_1 \neq e_4 \wedge e_2 \neq e_3 \wedge e_2 \neq e_4 \wedge e_3 \neq e_4).\end{aligned}$$

The first specification expresses the fact that every event is either a *send* or a *receive* event, where the sender is different from the receiver, or an initial action, or an internal action. The second specification says that every *send* or *receive* event, initial action, and internal action are distinct events. Thus, we cannot have $send(i, j, \mu) = e = rec(i', j', \mu')$. It is easy to check that our definition of m.p. QIS validates these specifications.

In [10], p. 132 the authors list three constraints on m.p. systems, the third one involves the notion of *run* on an SGS. Nonetheless, we can restate the first two without introducing runs:

MP1 a local state l_i for agent i is a *history* over Σ_i, INT_i and MSG;
MP2 for every event $rec(i, j, \mu)$ in $l_i(s)$ there exists an event $send(j, i, \mu)$ in $l_j(s)$.

Further, the following simplifying assumption is considered.

* all events in a given agent's local state are distinct, an agent can never perform the same action twice.

Note that this does not restrict our analysis as identical actions can be timestamped. We show how to formalise these specifications in the language $\mathcal{L}_n^D$ of m.p. QIS. First, we introduce a predicate H for *happened* such that $(\mathcal{P}^\sigma, s) \models H(e, i)$ iff e is an event in $l_i(s)$. The formula $H(e)$ is a shorthand for $\exists i H(e, i)$. By the definition of m.p. system, we can show that $(\mathcal{P}^\sigma, s) \models H(e)$ iff e is an event in s. Further, we define an order $Prec$ on events as follows:

$$\begin{aligned}(\mathcal{P}^\sigma, s) \models Prec(e, e', i) \text{ iff } & (\mathcal{P}^\sigma, s) \models H(e, i) \wedge H(e', i) \text{ and} \\ & \text{for all } s' \leq s, \; (\mathcal{P}^\sigma, s') \models H(e', i) \rightarrow H(e, i).\end{aligned}$$

The definition of $Prec(e, e')$ is similar, with $H(e)$ instead of $H(e, i)$. We can force the events in a global state s to be partially ordered by specifying that $Prec(e, e')$ is a reflexive and transitive relation on the set of past events:

$$H(e) \rightarrow Prec(e, e) \tag{3}$$

$$Prec(e, e') \wedge Prec(e', e'') \rightarrow Prec(e, e'') \tag{4}$$

As an example, we show that (4) holds. Suppose that $(\mathcal{P}^\sigma, s) \models Prec(e, e') \wedge Prec(e', e'')$. This means that $(\mathcal{P}^\sigma, s) \models \exists i H(e, i) \wedge \exists j H(e'', j)$. Moreover, we

have that for all $s' \leq s$, $(\mathcal{P}^\sigma, s') \models \exists i H(e', i) \rightarrow \exists i H(e, i)$ and $(\mathcal{P}^\sigma, s') \models \exists j H(e'', j) \rightarrow \exists j H(e', j)$. By renaming bounded variables and the transitivity of implication, we obtain that for all $s' \leq s$, $(\mathcal{P}^\sigma, s') \models \exists i H(e'', i) \rightarrow \exists i H(e, i)$. As a result, $(\mathcal{P}^\sigma, s) \models Prec(e, e'')$. Further, $Prec(e, e', i)$ can be defined as a linear order on the events in l_i, i.e., it is also anti-symmetric and total:

$$Prec(e, e', i) \wedge Prec(e', e, i) \rightarrow (e = e') \tag{5}$$

$$H(e, i) \wedge H(e', i) \rightarrow Prec(e, e', i) \vee Prec(e', e, i) \tag{6}$$

We define $Linear(Prec(e, e', i))$ as the conjunction of (3)–(6) above, expressing the fact that the relation $Prec(e, e', i)$ is linear. Also, we define the first event as the minimal one with respect to $Prec(e, e', i)$, that is,

$$Fst(e, i) ::= H(e, i) \wedge \forall e'(e' \neq e \rightarrow (H(e', i) \rightarrow \neg Prec(e', e, i)))$$

Finally, the formulas $Sent(i, j, \mu)$, $Recd(i, j, \mu)$, $Init(i, \alpha)$, and $Int(i, \alpha)$ are shorthands for $H(send(i, j, \mu))$, $H(rec(i, j, \mu))$, $H(init(i, \alpha))$, and $H(int(i, \alpha))$ respectively. Now we can formalise the specifications MP1-2 and * as follows:

MP1* $Linear(Prec(e, e', i)) \wedge$
$\wedge \exists! e(Fst(e, i) \wedge \exists \alpha(e = init(i, \alpha))) \wedge \forall e(H(e, i) \wedge \neg Fst(e, i) \rightarrow$
$\rightarrow \exists j, \alpha, \mu(e = int(i, \alpha) \vee e = send(i, j, \mu) \vee e = rec(i, j, \mu)))$

MP2' $\forall i, j, \mu(Recd(i, j, \mu) \rightarrow Sent(j, i, \mu))$

MP1* forces the local state of any agent i to satisfy MP1 and *; while by MP2' specification MP2 is satisfied. MP1*-2 are the basic specifications for m.p. QIS, but we can consider further constraints on message passing system. A property often required in the framework of m.p. systems is *channel reliability*. Modified from [10], a m.p. system is reliable if every sent message is eventually received, or more formally:

MP4 if $send(i, j, \mu)$ is in $l_i(s)$, then there exists a global state s' such that $rec(j, i, \mu)$ is in $l_j(s')$.

In $\mathcal{L}_n^D$ we can formalise this specification as follows:

MP4' $\forall j, \mu(Sent(i, j, \mu) \rightarrow \neg K_i \neg Recd(j, i, \mu))$

In fact, if $send(i, j, \mu)$ is in $l_i(s)$, by MP4' $(\mathcal{P}^\sigma, s) \models \neg K_i \neg Recd(j, i, \mu)$, this means that there exists a global state s' such that $(\mathcal{P}^\sigma, s') \models Recd(j, i, \mu)$, that is, $rec(j, i, \mu) \in l_j(s')$. Thus, MP4 holds. Note that MP4' is stronger than MP4 as the former requires that the local states of agent i in s and s' are identical. This can be considered a limit of our epistemic language, due to the lack of temporal operators. Further, a relevant property of m.p. systems concerns *authentication*: if agent i receives a message μ from agent j, then i knows that μ was actually sent by j. This specification can be expressed as

$$\forall j, \mu(Recd(i, j, \mu) \rightarrow K_i Sent(j, i, \mu))$$

Finally, we may require that agents have *perfect recall*, that is, they know everything that has happened to them:

$$\forall e(H(e,i) \rightarrow K_i H(e,i))$$

It is easy to show that by the way they are defined, m.p. QIS satisfy authentication and perfect recall but not channel reliability. We remark that all the specifications introduced are defined by means of only the predicative constants $H(e,i)$ and $Prec(e,e',i)$.

We now prove the main result of this section: Proposition 4.4.3 in [10] can be restated as a validity on the class of m.p. QIS satisfying MP1, MP2, and *. We do not give the full statement, but we note that this metatheoretical result can be restated as a formula in the first-order modal language $\mathcal{L}_n^D$. First, we introduce a relation $\mapsto_G$ of *potential causality* between events, as discussed in [20]. This relation is intended to capture the intuition that event e might have caused event e'.

Fix a subset G of A, the relation $\mapsto_G$ holds between events e, e' at a state s iff both e and e' appears in s, and

1. for some $i, j \in G$, e' is a *receive* event and e is the corresponding *send* event;
2. for some $i \in G$, events e, e' are both in $l_i(s)$ and either $e = e'$ or e comes earlier than e' in $l_i(s)$;
3. for some event e'', we have that $e \mapsto_G e''$ and $e'' \mapsto_G e'$ hold at s.

Note that $\mapsto_G$ is a partial order on events. We say that $(\mathcal{P}^\sigma, s) \models e \mapsto_G e'$ if $e \mapsto_G e'$ hold at s (we use the same notation for semantic and syntactic elements).

Now we prove that the potential causality relation $\mapsto_G$ respects the order $Prec$ of events by showing that the following validity holds in the class of m.p. QIS. This means that if event e is the "cause" of event e', then it is distributed knowledge among the agents that e happened before e'. Note that this is the right to left implication of Proposition 4.4.3 in [10]:

$$m.p.\ QIS \models \forall e, e'((e \mapsto_G e') \rightarrow D_G Prec(e,e'))$$

Proof. Assume that $(\mathcal{P}^\sigma, s) \models e \mapsto_G e'$. If e' is a *receive* event and e is the corresponding *send* event, then $l_i(s) = l_i(s')$ for all $i \in G$ implies $(\mathcal{P}^\sigma, s') \models H(e) \wedge H(e')$, and for $s'' \leq s'$, $(\mathcal{P}^\sigma, s'') \models Recd(i,j,\mu) \rightarrow Sent(j,i,\mu)$ by MP2'. Thus, $(\mathcal{P}^\sigma, s) \models D_G Prec(e,e')$.

If e, e' are both in $l_i(s)$ and either $e = e'$ or e comes earlier than e' in $l_i(s)$, then $l_i(s) = l_i(s')$ implies that $(\mathcal{P}^\sigma, s') \models H(e) \wedge H(e')$, and for $s'' \leq s'$, $(\mathcal{P}^\sigma, s'') \models H(e') \rightarrow H(e)$. It follows that $(\mathcal{P}^\sigma, s) \models K_i Prec(e,e')$, and by $D1$, $D2$, $(\mathcal{P}^\sigma, s) \models D_G Prec(e,e')$.

Finally, if there exists some event e'' such that $e \mapsto_G e''$ and $e'' \mapsto_G e'$, we can assume without loss of generality that this happens because we are in either the first or second case above. In both cases we have that $(\mathcal{P}^\sigma, s) \models D_G Prec(e,e'') \wedge D_G Prec(e'',e')$. Therefore, for every s', $l_i(s) = l_i(s')$ for all $i \in G$

implies $(\mathcal{P}^\sigma, s') \models H(e) \wedge H(e')$, and for $s'' \leq s'$, $(\mathcal{P}^\sigma, s'') \models H(e'') \rightarrow H(e) \wedge H(e') \rightarrow H(e'')$. By transitivity, $(\mathcal{P}^\sigma, s'') \models H(e') \rightarrow H(e)$. Thus, $(\mathcal{P}^\sigma, s) \models D_G Prec(e, e')$. □

The analysis of message-passing systems carried out in this section clearly shows the advantages of first-order modal formalisms in comparison with propositional ones. By means of language $\mathcal{L}_n^D$ we are able to formalise various constraints on m.p. systems, thereby signaling the general correctness of the approach. Most importantly, the right to left implication of Proposition 4.4.3 in [10] turned out out be a validity on the class of QIS modelling m.p. systems.

In the second part of this paper we will show that this expressivity gain is obtained while still retaining completeness of the logical formalism.

5 Axiomatisation

In this section we provide a sound and complete axiomatisation of systems of global states. Note that while it is customary in modal logic to axiomatise unvalued structures (hence our choice of $\mathcal{SGS}$), the same result applies to $\mathcal{QIS}$ as well. Technically, we first prove the completeness of the first-order multi-modal system $Q.S5_n^D$ with respect to Kripke frames. Then, by lemma 1 the completeness of $Q.S5_n^D$ with respect to $\mathcal{SGS}$ follows.

In [19] Kripke proved the completeness of monomodal $Q.S5$ (see also [12,18]). The novelty of this section consists in showing that the techniques in [11] for the completeness of propositional $S5_n^D$ can be straightforwardly extended to first-order for proving the completeness of $Q.S5_n^D$. Also, note that an independent completeness proof for $S5_n^D$ appeared in [24].

5.1 The System $Q.S5_n^D$

The system $Q.S5_n^D$ on the language $\mathcal{L}_n^D$ is a first-order multi-modal version of the propositional system $S5$. Although tableaux proof systems and natural deduction calculi are more suitable for automated theorem proving, Hilbert-style systems are easier to handle for the completeness proof. Hereafter we list the postulates of $Q.S5_n^D$; note that $\Rightarrow$ is the inference relation between formulas.

Definition 8. *The system $Q.S5_n^D$ on $\mathcal{L}_n^D$ contains the following schemes of axioms and inference rules:*

Taut	*every instance of classic propositional tautologies*
MP	$\phi \rightarrow \psi, \phi \Rightarrow \psi$
Dist	$K_i(\phi \rightarrow \psi) \rightarrow (K_i\phi \rightarrow K_i\psi)$
T	$K_i\phi \rightarrow \phi$
4	$K_i\phi \rightarrow K_iK_i\phi$
5	$\neg K_i\phi \rightarrow K_i\neg K_i\phi$
Nec	$\phi \Rightarrow K_i\phi$

Dist	$D_G(\phi \to \psi) \to (D_G\phi \to D_G\psi)$
T	$D_G\phi \to \phi$
4	$D_G\phi \to D_G D_G\phi$
5	$\neg D_G\phi \to D_G \neg D_G\phi$
D1	$D_{\{i\}}\phi \leftrightarrow K_i\phi$
D2	$D_G\phi \to D_{G'}$, *for* $G \subseteq G'$
Nec	$\phi \Rightarrow D_G\phi$
Ex	$\forall x\phi \to \phi[x/t]$
Gen	$\phi \to \psi[x/t] \Rightarrow \phi \to \forall x\psi$, *where* x *is not free in* ϕ
Id	$t = t$
Func	$t = t' \to (t''[x/t] = t''[x/t'])$
Subst	$t = t' \to (\phi[x/t] \to \phi[x/t'])$

We consider the standard definitions of *proof* and *theorem*: $\vdash \phi$ means that $\phi \in \mathcal{L}_n^D$ is a theorem in $Q.S5_n^D$. Moreover, $\phi \in \mathcal{L}_n^D$ is *derivable* in $Q.S5_n^D$ from a set Δ of formulas in $\mathcal{L}_n^D$ - $\Delta \vdash \phi$ in short - iff there are $\phi_1, \ldots, \phi_n \in \Delta$ such that $\vdash \phi_1 \wedge \ldots \wedge \phi_n \to \phi$. It is easy to check that the axioms of $Q.S5_n^D$ are valid on every Kripke frame and the inference rules preserve validity. As a consequence, we have the following soundness result.

Lemma 2 (Soundness). *The system* $Q.S5_n^D$ *is sound with respect to the class* $\mathcal{K}$ *of Kripke frames.*

By lemma 2 and the definition of validity on SGS, these implications hold:

$$Q.S5_n^D \vdash \phi \quad \Rightarrow \quad \mathcal{K} \models \phi \quad \Rightarrow \quad \mathcal{SGS} \models \phi$$

Thus, we have soundness also for the systems of global states.

Corollary 1 (Soundness). *The system* $Q.S5_n^D$ *is sound with respect to the class* $\mathcal{SGS}$ *of systems of global states.*

In the next paragraph we show that the axioms in $Q.S5_n^D$ are not only necessary, but also sufficient to prove all the validities on $\mathcal{SGS}$. In conclusion we show that the converse of the Barcan formula is provable in $Q.S5_n^D$. For a proof of BF, we refer to [12], p.138.

1. $\forall x\phi \to \phi$	Ex
2. $K_i(\forall x\phi \to \phi)$	from 1 by Nec
3. $K_i(\forall x\phi \to \phi) \to (K_i\forall x\phi \to K_i\phi)$	$Dist$
4. $K_i\forall x\phi \to K_i\phi$	from 2, 3 by MP
5. $K_i\forall x\phi \to \forall x K_i\phi$	from 4 by Gen

5.2 Completeness

We prove the completeness of $Q.S5_n^D$ by extending to first-order the proof for the propositional system $S5_n^D$ in [11]. The novelty of our result consists in showing that this method can be straightforwardly applied to first-order Kripke frames.

Specifically, we show that if $Q.S5_n^D$ does not prove a formula $\phi \in \mathcal{L}_n^D$, then the canonical model $\mathcal{M}^{Q.S5_n^D}$ for $Q.S5_n^D$ does not pseudo-validate ϕ. It is not guaranteed that the notion of pseudo-validity (to be defined below) coincides with plain validity, but by results in [11] we can obtain from $\mathcal{M}^{Q.S5_n^D}$ a Kripke model $\mathcal{M}'$ such that $\mathcal{M}^{Q.S5_n^D}$ pseudo-validates ϕ iff $\mathcal{M}' \models \phi$. Thus completeness follows.

In order to show the first part of the completeness result we rely on two lemmas: the *saturation lemma* and the *truth lemma*, whose statements require the following definitions: let Λ be a set of formulas in $\mathcal{L}_n^D$,

Λ is *consistent*	iff	for every $\phi \in \Lambda$, $\nvdash \neg\phi$;
Λ is *maximal*	iff	for every $\phi \in \mathcal{L}_n^D$, $\phi \in \Lambda$ or $\neg\phi \in \Lambda$;
Λ is *max-cons*	iff	Λ is consistent and maximal;
Λ is *rich*	iff	$\exists x\phi \in \Lambda$ implies $\phi[x/d] \in \Lambda$, for some constant $d \in \mathcal{L}_n^D$;
Λ is *saturated*	iff	Λ is max-cons and rich.

Assume that $Q.S5_n^D$ does not prove ϕ, then the set $\{\neg\phi\}$ is consistent, and by the saturation lemma below $\{\neg\phi\}$ can be extended to a saturated set:

Lemma 3 (Saturation [18]). *If Δ is a consistent set of formulas in $\mathcal{L}_n^D$, then it can be extended to a saturated set Π of formulas on some expansion $\mathcal{L}_n^{D+}$ obtained by adding an infinite set of new individual constants to $\mathcal{L}_n^D$.*

Now we introduce the canonical model for $Q.S5_n^D$. Note that $\wp^+(A)$ is the set of non-empty sets of agents.

Definition 9 (Canonical model). *The canonical model for $Q.S5_n^D$ on the language $\mathcal{L}_n^D$, with an expansion $\mathcal{L}_n^{D+}$, is a tuple $\mathcal{M}^{Q.S5_n^D} = \langle W, \{R_j\}_{j \in A \cup \wp^+(A)}, D, I\rangle$ such that*

- *W is the set of saturated sets of formulas in $\mathcal{L}_n^{D+}$;*
- *for $i \in A$, $w, w' \in W$, wR_iw' iff $\{\phi | K_i\phi \in w\} \subseteq w'$;*
- *for $G \subseteq A$, $w, w' \in W$, wR_Gw' iff $\{\phi | D_G\phi \in w\} \subseteq w'$;*
- *D is the set of equivalence classes $[v] = \{v' | v = v' \in w\}$, for every closed term $v \in \mathcal{L}_n^{D+}$;*
- *$I(f^k)([v_1], \ldots, [v_k]) = [f^k(v_1, \ldots, v_k)]$;*
- *$\langle [v_1], \ldots, [v_k]\rangle \in I(P^k, w)$ iff $P^k(v_1, \ldots, v_k) \in w$.*

If we assume that $Q.S5_n^D \nvdash \phi$, by the saturation lemma there exists a saturated set $w \supseteq \{\neg\phi\}$, so the set W of possible worlds is non-empty. Further, by definition of R_i and R_G, and axioms *Func* and *Subst*, we can show that the definition of $[v]$ is independent from w, so D is well defined. Since T, 4 and 5 are all axioms of $Q.S5_n^D$, the various R_i and R_G are equivalence relations. Moreover, from $D1$ and $D2$ it follows that $R_{\{i\}}$ is equal to R_i, and $R_G \subseteq \bigcap_{i \in G} R_i$. However, in general it is not the case that $R_G = \bigcap_{i \in G} R_i$. This remark gives the *rationale*

for introducing the pseudo-satisfaction relation $\models^p$, defined as $\models$ but for the distributed knowledge operator D_G (in what follows we simply write $\mathcal{M}$ for $\mathcal{M}^{Q.S5^D_n}$):

$$(\mathcal{M}^\sigma, w) \models^p D_G\psi \text{ iff for every } w' \in W,\ wR_Gw' \text{ implies } (\mathcal{M}^\sigma, w') \models^p \psi$$

Now we can prove the *truth lemma* for the pseudo-satisfaction relation $\models^p$. In order to obtain this result we observe that for an assignment σ such that $\sigma(y_i) = [v_i]$, for $1 \leq i \leq n$, we have that $I^\sigma(t[\vec{y}]) = [t[\vec{y}/\vec{v}]]$.

Lemma 4 (Truth lemma). *For every $w \in \mathcal{M}$, $\phi \in \mathcal{L}^{D+}_n$, for $\sigma(y_i) = [v_i]$,*

$$(\mathcal{M}^\sigma, w) \models^p \phi[\vec{y}] \textit{ iff } \phi[\vec{y}/\vec{v}] \in w$$

Proof. The proof is by induction on the structure of $\phi \in \mathcal{L}^{D+}_n$.

$\phi = P^k(t_1, \ldots, t_k)$. By the definitions of pseudo-satisfaction and canonical interpretation $(\mathcal{M}^\sigma, w) \models^p P^k(t_1[\vec{y}], \ldots, t_k[\vec{y}])$ iff $\langle I^\sigma(t_1[\vec{y}]), \ldots, I^\sigma(t_k[\vec{y}])\rangle \in I(P^k, w)$ iff $\langle [t_1[\vec{y}/\vec{v}]], \ldots, [t_k[\vec{y}/\vec{v}]]\rangle \in I(P^k, w)$ iff $P^k(t_1[\vec{y}/\vec{v}], \ldots, t_k[\vec{y}/\vec{v}]) \in w$.

$\phi = \neg\psi, \psi \to \psi', \forall x\psi$. The cases for the propositional connectives follows by the maximality and consistency of the worlds in the canonical model; whereas for the universal quantifier, the induction step is proved by the richness of w.

$\phi = K_i\psi$. $\Leftarrow$ Assume that $K_i\psi[\vec{y}/\vec{v}] \in w$ and wR_iw'. By definition of R_i, $\psi[\vec{y}/\vec{v}] \in w'$ and by the induction hypothesis $(\mathcal{M}^\sigma, w') \models^p \psi[\vec{y}]$. Therefore $(\mathcal{M}^\sigma, w) \models^p K_i\psi[\vec{y}]$.

$\Rightarrow$ Assume that $K_i\psi[\vec{y}/\vec{v}] \notin w$. Note that the set $\{\phi | K_i\phi \in w\} \cup \{\neg\psi[\vec{y}/\vec{v}]\}$ is consistent. By standard techniques [12,18] we can extend it to a saturated set w' such that $\{\phi | K_i\phi \in w\} \cup \{\neg\psi[\vec{y}/\vec{v}]\} \subseteq w'$. This means that wR_iw' and $(\mathcal{M}^\sigma, w') \models^p \neg\psi[\vec{y}]$ by the induction hypothesis. Hence $(\mathcal{M}^\sigma, w) \not\models^p K_i\psi[\vec{y}]$.

$\phi = D_G\psi$. Similar to the previous case. □

We remarked that the canonical model might not satisfy $\bigcap_{i\in G} R_i = R_G$. However, it can be unwound to get a structure $\mathcal{M}'$ in such a way that the same formulas are valid [11]. More formally, given the canonical model $\mathcal{M} = \langle W, R, D, I\rangle$, there is another structure $\mathcal{M}^* = \langle W^*, R^*, D, I^*\rangle$ and a surjective function $h : W^* \to W$ such that (i) $\mathcal{M}^*$ is a tree, that is, for all $w, w' \in W^*$, there is at most one *reduced* path from w to w', (ii) wR^*_iw' implies $h(w)R_ih(w')$ and wR^*_Gw' implies $h(w)R_Gh(w')$, and (iii) $\langle a_1, \ldots, a_k\rangle \in I^*(P^k, w)$ iff $\langle a_1, \ldots, a_k\rangle \in I(P^k, h(w))$.

In order to define $\mathcal{M}^*$ and h we need more definitions. Let w, w' be worlds in W, a *path from w to w'* is a sequence $\langle w_1, i_1, w_2, i_2, \ldots, i_{k-1}, w_k\rangle$ such that:

1. $w = w_1$ and $w' = w_k$;
2. $w_1, \ldots, w_k \in W$;
3. each i_j is either an agent or a set of agents;
4. $\langle w_j, w_{j+1}\rangle \in R^*_{i_j}$.

The *reduction* of a path $\langle w_1, i_1, w_2, i_2, \ldots, i_{k-1}, w_k\rangle$ is obtained by replacing each maximal consecutive subsequence $\langle w_q, i_q, w_{q+1}, i_{q+1}, \ldots, i_{r-1}, w_r\rangle$ where

$i_q = i_{q+1} = \ldots = i_{r-1}$ by $\langle w_q, i_q, w_r \rangle$. A path is said to be *reduced* is it is equal to its reduction.

We define W^* by induction. Let W_1^* be W, and define W_{k+1}^* as the set of worlds $v_{w,i,w'}$ such that $w \in W_k^*$, $w' \in W$ and i is an agent or group of agents. Let $W^* = \bigcup_{k \in \mathbb{N}} W_k^*$, then define $h : W^* \to W$ by letting $h(w) = w$, for $w \in W_1^*$ and $h(v_{w,i,w'}) = w'$, for $w \in W_k^*$. Further, R_i^* is the reflexive, transitive and symmetric closure of the relation defined for $w, w' \in W^*$ iff $w' = v_{w,i,w''}$ for some $w'' \in W$, and $h(w) R_i h(w')$. Finally, $I^*(P^k, w) = I(P^k, h(w))$. It can be checked that $\mathcal{M}^*$ and h satisfy (i)-(iii) above, we omit the proof for reasons of space and refer to [11] for the details. In particular, we can show what follows:

Lemma 5. *For $w \in W^*$, $\phi \in \mathcal{L}_n^D$,*

$$(\mathcal{M}^{*\sigma}, w) \models^p \phi \text{ iff } (\mathcal{M}^{\sigma}, h(w)) \models^p \phi$$

Proof. The proof is by induction on the length of ϕ. If ϕ is an atomic formula, then the coimplication follows by the definition of I^*. The cases for the propositional connectives and the universal quantifier are straightforward.

$\phi = K_i\psi$. $\Leftarrow$ Suppose that $(\mathcal{M}^{*\sigma}, w) \not\models^p K_i\psi$, then there is a world $w' \in W^*$ such that $w R_i^* w'$ and $(\mathcal{M}^{*\sigma}, w') \not\models^p \psi$. This means that $h(w) R_i h(w')$ and $(\mathcal{M}^{\sigma}, h(w')) \not\models^p \psi$ by induction hypothesis. Thus $(\mathcal{M}^{\sigma}, h(w)) \not\models^p K_i\psi$.

$\Rightarrow$ If $(\mathcal{M}^{\sigma}, h(w)) \not\models^p K_i\psi$, then there is a world $w' \in W$ such that $h(w) R_i w'$ and $(\mathcal{M}^{\sigma}, w') \not\models^p \psi$. By construction $v_{w,i,w'} \in W^*$, $h(v_{w,i,w'}) = w'$ and $w R_i^* v_{w,i,w'}$. By induction hypothesis $(\mathcal{M}^{*\sigma}, v_{w,i,w'}) \not\models^p \psi$, hence $(\mathcal{M}^{*\sigma}, w) \not\models^p K_i\psi$.

$\phi = D_G\psi$. Similar to the previous case. □

Now we make use of the structure $\mathcal{M}^*$ to define a Kripke model $\mathcal{M}'$ that does not validate the unprovable formula $\phi \in \mathcal{L}_n^D$. Define $\mathcal{M}' = \langle W', R', D', I' \rangle$ as follows:

- $W' = W^*$, $D' = D$ and $I' = I^*$;
- R_i' is the transitive closure of $R_i^* \cup \bigcup_{i \in G} R_G^*$.

Since the various R_i^* and R_G^* are reflexive and symmetric, it follows that R_i' is an equivalence relation, and therefore $\mathcal{M}'$ is based on a Kripke frame. Further, we can prove the following result:

Lemma 6. *For $\phi \in \mathcal{L}_n^D$,*

$$(\mathcal{M}'^{\sigma}, w) \models \phi \text{ iff } (\mathcal{M}^{*\sigma}, w) \models^p \phi$$

Proof. Also this proof is by induction on the length of ϕ. If ϕ is an atomic formula, then the coimplication follows because $I' = I^*$. The cases for the propositional connectives are straightforward.

For $\phi = K_i\psi$ or $\phi = D_G\psi$, the inductive step goes as in the propositional case; we refer to [11] for a detailed proof.

$\phi = \forall x \psi$. If $(\mathcal{M}'^{\sigma}, w) \models \phi$, then for all $a \in D'$, $(\mathcal{M}'^{\sigma\binom{a}{x}}, w) \models \psi$. By induction hypothesis $(\mathcal{M}^{*\sigma\binom{a}{x}}, w) \models^p \psi$, and since $D' = D$, $(\mathcal{M}^{*\sigma}, w) \models^p \phi$. □

In conclusion, if $\phi \in \mathcal{L}_n^D$ is not provable in $Q.S5_n^D$, then the canonical model $\mathcal{M}$ pseudo-satisfies $\neg\phi$ by lemma 4. By lemma 5 also $\mathcal{M}^*$ pseudo-satisfies $\neg\phi$, and by the last result above $\mathcal{M}'$ does not validate ϕ. Thus, we state the following completeness result.

Theorem 1 (Completeness). *The system $Q.S5_n^D$ is complete with respect to the class $\mathcal{K}$ of Kripke frames.*

As a consequence, we have completeness also with respect to the systems of global states. In fact, if $\nvdash \phi$ then by Theorem 1 there exists a K-model $\mathcal{M} = \langle \mathcal{F}, I \rangle$, based on a Kripke frame $\mathcal{F}$, which falsifies ϕ. In order to prove that $\mathcal{SGS} \not\models \phi$ we have to find a quantified interpreted system $\mathcal{P}$ falsifying ϕ. Define $\mathcal{P}$ as $\langle g(\mathcal{F}), I \rangle$: by the definition of validity in QIS, $\mathcal{P} \models \phi$ iff $\mathcal{P}_f = \langle f(g(\mathcal{F})), I \rangle$ models ϕ, but by lemma 1 $f(g(\mathcal{F}))$ is isomorphic to $\mathcal{F}$. Hence $\mathcal{P} \not\models \phi$.

As a result, we have the following implications and a further completeness result:

$$\mathcal{SGS} \models \phi \quad \Rightarrow \quad \mathcal{K} \models \phi \quad \Rightarrow \quad Q.S5_n^D \vdash \phi$$

Corollary 2 (Completeness). *The system $Q.S5_n^D$ is complete with respect to the class $\mathcal{SGS}$ of systems of global states.*

By combining together the soundness and completeness theorems we compare directly the axiomatisation $Q.S5_n^D$ and the systems of global states, so we state our main result:

Corollary 3 (Soundness and Completeness). *A formula $\phi \in \mathcal{L}_n^D$ is valid on the class $\mathcal{SGS}$ of systems of global states iff ϕ is provable in $Q.S5_n^D$.*

6 Conclusions

As we argued in the Introduction, first-order modal formalisms offer expressivity advantages over propositional ones. But the cited explorations already carried out on this subject in MAS and, more in general, in knowledge representation and Artificial Intelligence, have 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, even if limited to the static case, shows that the popular system $S5_n^D$ extends naturally to first-order. In carrying out this exercise we tried to remain as close as possible to the original semantics of interpreted systems, so that fine grained specifications of MAS may be expressed, as recent work on model checking interpreted systems demonstrates [13,28].

Different extensions of the present framework seem worth pursuing. First of all, it seems interesting to relax the assumption on the domain of quantification and admit a different domain $d(w)$ for every state w. Further, we could assume a different domain of quantification $d_a(w)$ for each agent a in a state w. In this case quantification would be agent-indexed, i.e. we would be using a different quantifier $\forall_a$ for every agent $a \in A$. In such an extended framework we should

check whether the validities on m.p. QIS in section 4 still hold, and how to modify the completeness proof for $Q.S5_n^D$. Also, it would be of interest to explore the completeness issues resulting from term-indexing epistemic operators as in [21].

In an orthogonal dimension to the above, another significant extension would be to add temporal operators to the formalism. This would open the way for an exploration of axiomatisations for temporal/epistemic logic for MAS. While as reported in the Introduction we are not so concerned with the satisfiability problem, in doing so attention will have to be paid to the results in [16].

Acknowledgments

The research described in this paper was supported by the EC Framework 6 funded project CONTRACT (IST Project Number 034418). The authors would like to thank Wiebe van der Hoek for useful comments on an earlier version of this paper, as well as CLIMA's reviewers for their suggestions.

References

1. Armando, A., et al.: The Avispa tool for the automated validation of internet security protocols and applications. In: Etessami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 281–285. Springer, Heidelberg (2005)
2. Ågotnes, T., van der Hoek, W., Wooldridge, M.: Quantified Coalition Logic. In: Proceedings of IJCAI, Hyderabad, India, January 6-12, 2007, pp. 1181–1186 (2007)
3. Belardinelli, F., Lomuscio, A.: A Quantified Epistemic Logic to reason about Multi-Agent Systems. In: Proceedings of AAMAS 2007, Honolulu, Hawaii (2007)
4. Belardinelli, F., Lomuscio, A.: Quantified Epistemic Logics with Flexible Terms. In: LORI workshop on Logic, Rationality and Interaction, Beijing, August 5-9 (2007)
5. Bieber, P.: A logic of communication in hostile environments. In: CSFW, pp. 14–22 (1990)
6. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge, UP (2001)
7. Chagrov, A., Zakharyaschev, M.: Modal Logic. Oxford University Press, Oxford (1997)
8. Clarke, E., Grumberg, O., Peled, D.: Model Checking. MIT Press, Cambridge (1999)
9. Dams, M., Cohen, M.: A complete axiomatisation of knowledge and cryptography. In: LICS (2007)
10. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. MIT Press, Cambridge (1995)
11. Fagin, R., Halpern, J., Vardi, M.: What can machines know? on the properties of knowledge in distributed systems. J. ACM 39(2), 328–376 (1992)
12. Fitting, M., Mendelsohn, R.: First-order Modal Logic. Kluwer, Dordrecht (1999)
13. Gammie, P., van der Meyden, R.: MCK: Model checking the logic of knowledge. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 479–483. Springer, Heidelberg (2004)
14. Halpern, J., Fagin, R.: Modelling knowledge and action in distributed systems. Distributed Computing 3(4), 159–179 (1989)
15. Garson, J.: Quantification in modal logic. Handbook of Philosophical Logic, vol. 2, pp. 249–307 (1984)

16. Hodkinson, I., Wolter, F., Zakharyaschev, M.: Decidable fragment of first-order temporal logics. Ann. Pure Appl. Logic 106(1-3), 85–134 (2000)
17. van der Hoek, W., Wooldridge, M.: Tractable multiagent planning for epistemic goals. In: Proceedings of AAMAS 2002, pp. 1167–1174. ACM Press, New York (2002)
18. Hughes, G., Cresswell, M.: A New Introduction to Modal Logic. Routledge (1996)
19. Kripke, S.: A Completeness Theorem in Modal Logic. J. Sym. Log. 24, 1–14 (1959)
20. Lamport, L.: Time, Clocks, and the Ordering of Events in a Distributed System. Communication of the ACM 21(7), 558–565 (1978)
21. Lomuscio, A., Colombetti, M.: QLB: a quantified logic for belief. In: Jennings, N.R., Wooldridge, M.J., Müller, J.P. (eds.) ECAI-WS 1996 and ATAL 1996. LNCS, vol. 1193. Springer, Heidelberg (1996)
22. Lomuscio, A., Ryan, M.: On the relation between interpreted systems and kripke models. In: Wobcke, W., Pagnucco, M., Zhang, C. (eds.) Agents and Multi-Agent Systems Formalisms, Methodologies, and Applications. LNCS (LNAI), vol. 1441. Springer, Heidelberg (1997)
23. Lomuscio, A., Penczek, W., Wozna, B.: Bounded model checking knowledge and real time. Artificial Intelligence (to appear)
24. Meyer, J.-J., van der Hoek, W.: Making some issues of implicit knowledge explicit. Int. J. of Foundations of Computer Science 3(2), 193–223 (1992)
25. Meyer, J.-J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science. Cambridge University Press, Cambridge (1995)
26. Penczek, W., Lomuscio, A.: Verifying Epistemic Properties of multi-agent systems via bounded model checking. Fund. Inform. 55(2), 167–185 (2003)
27. Quine, W.v.O.: Quantifiers and Propositional Attitudes. Journal of Philosophy 53, 177–187 (1956)
28. Raimondi, F., Lomuscio, A.: Automatic verification of multi-agent systems by model checking via OBDDs. Journal of Applied Logic 5(2), 235–251 (2007)
29. Solanki, M.: A Compositional Framework for the Specification, Verification and Runtime Validation of Reactive Web Service. PhD thesis (2005)
30. Solanki, M., Cau, A., Zedan, H.: ASDL: a wide spectrum language for designing web services. In: WWW, pp. 687–696. ACM, New York (2006)
31. Spoletini, P.: Verification of Temporal Specification via Model Checking. PhD thesis, Politecnico di Milano, Dipartimento di Elettronica e Informatica (2006)
32. Sturm, H., Wolter, F., Zakharyaschev, M.: Monodic epistemic predicate logic. In: Brewka, G., Moniz Pereira, L., Ojeda-Aciego, M., de Guzmán, I.P. (eds.) JELIA 2000. LNCS (LNAI), vol. 1919, pp. 329–344. Springer, Heidelberg (2000)
33. Sturm, H., Wolter, F., Zakharyaschev, M.: Common knowledge and quantification. Economic Theory 19, 157–186 (2002)
34. Viganó, F.: A Framework for Model Checking Institutions. In: Edelkamp, S., Lomuscio, A. (eds.) MoChArt IV. LNCS (LNAI), vol. 4428, pp. 129–145. Springer, Heidelberg (2007)
35. Wolter, F., Zakharyaschev, P.: Decidable fragments of first-order modal logics. J. Symb. Log. 66(3), 1415–1438 (2001)