A Complete Quantified Epistemic Logic for Reasoning about Message Passing Systems*

F. Belardinelli and A. Lomuscio

Department of Computing
Imperial College London, UK
{F.Belardinelli,A.Lomuscio}@imperial.ac.uk

Abstract. We investigate quantified interpreted systems, a semantics to model multi-agent systems in which 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 analyse a first-order epistemic language interpreted on this semantics and show 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 modelling message passing systems, a typical 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 [6, 17]. The use of modal propositional logic as a specification language is so routine to require 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 often forced to introduce ad-hoc propositions to emulate basic relations among 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"). Not always quantification is simply syntactic sugar: certain expressivity needs do require infinite domains (e.g., see section 4 below). Further, epistemic modalities can be combined with quantifiers to express concepts such as knowledge *de re*/*de dicto* [8].

* 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.

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? 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 [9, 18]. In the model checking approach [5] the verification problem is tackled not by checking theoremhood but simply model satisfaction. In other words, we do not check whether a formula representing a specification is a theorem in some logic (or, given completeness, whether a formula is satisfiable), 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, the latter is decidable at least in some suitable fragment.

This paper takes inspiration from the considerations above and aims to make progress on the subject of first-order epistemic logic. The main contribution of the paper is the axiomatisation presented 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 [6].

While completeness results for quantified modal logic are customarily proved with respect to Kripke semantics [8, 13], we should state clearly that QML has been discussed in a MAS setting before. In [6] quantified epistemic logic is briefly discussed, along with its Kripke semantics and some significant validities; in [15] the authors introduce a quantified logic of belief, in which the doxastic modalities are indexed to terms of a first-order language; in [1] 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 $\mathcal{L}_n^D$ and interpret it on quantified interpreted systems, a valued version of systems of global states. In section 4 we exemplify syntax and semantics by describing a formal model for message-passing systems and discuss some specification patterns in $\mathcal{L}_n^D$. 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 [6, 10], Kripke frames are best employed to get a deeper understanding of the formal properties of these systems [3, 4]. Technically, we extend the corresponding propositional structures to the first-order. This extension is entirely not trivial, as there are many ways of performing it: for instance, we have to choose between a single or several quantifying domains for each agent and/or for each computational state, not to mention domains of intensional objects [2]. 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 $A = \{1, \dots, n\}$.

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 [16], rather than the "dynamic" version [6]. 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 $l_i, l'_i, \dots$, for each agent $i \in A$, and a set L_e containing the local states of the environment $l_e, l'_e, \dots$. 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 $S \subseteq L_e \times L_1 \times \dots \times L_n$ 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 the 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 the 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 the 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, notably completeness investigations. We work with frames with equivalence relations, so we take the following definition:

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

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

2.3 Maps between $\mathcal{SGS}$ and $\mathcal{F}_E$

We explore the relationship between these structures by means of two maps f and g from $\mathcal{SGS}$ to $\mathcal{F}_E$ 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 equivalence 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 consequence, every sound and complete axiomatisation of the equivalence frames is also an axiomatisation of the systems of global states.

We start with the map $f : \mathcal{SGS} \to \mathcal{F}_E$. 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 an equivalence frame.

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

Proof. If $\mathcal{F} = \langle W, \sim_1, \ldots, \sim_n, D\rangle$ is an equivalence 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. Moreover, $w \sim_i w'$ 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. $\square$

By Lemma 1 we will show in section 5 that a sound and complete axiomatisation of equivalence frames is adequate also with respect to SGSs.

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 the identity. The language $\mathcal{L}_n^D$ is interpreted on models based on equivalence frames. Finally, we present the quantified interpreted systems, a valued version of the systems of global states.

3.1 Syntax

Our first-order multi-modal formulas are defined on an alphabet containing 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 $=$, 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 $G \subseteq A$. Terms and formulas in the language $\mathcal{L}_n^D$ are defined as follows:

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

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. 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 define validity on quantified interpreted systems in terms of validity on Kripke models.

Definition 3 (model). *A Kripke model $\mathcal{M}$ - or simply a model - based on an equivalence 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.

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 follows:

$$\begin{aligned} I^\sigma(y) &= \sigma(y) \\ I^\sigma(f^k(t_1, \ldots, t_k)) &= I(f^k)(I^\sigma(t_1), \ldots, I^\sigma(t_k)) \end{aligned}$$

In particular, the valuation $I^\sigma(d)$ of constant d is $I(d) \in D$. 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 define the truth conditions for the formulas in $\mathcal{L}_n^D$.

Definition 4 (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:*

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

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

Let Δ be a set of formulas in $\mathcal{L}_n^D$, we say that $\mathcal{M}$ is a model *for* Δ if every formula in Δ is valid on $\mathcal{M}$. Moreover, $\mathcal{F}$ is a frame *for* Δ if every model based on $\mathcal{F}$ is a model for Δ.

Now we have all the preliminary definitions to introduce the quantified interpreted systems (QIS).

Definition 5 (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 $f(\mathcal{S})$.*

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 $(\mathcal{P}^\sigma, s) \models \phi$ if $(\mathcal{P}_f^\sigma, s) \models \phi$. In particular,

$$\begin{array}{l}
(\mathcal{P}^\sigma, \langle l_e, l_1, \ldots, l_n\rangle) \models P^k(\vec{t}) \text{ if } \langle I^\sigma(t_1), \ldots, I^\sigma(t_k)\rangle \in I(P^k, \langle l_e, l_1, \ldots, l_n\rangle) \\
(\mathcal{P}^\sigma, \langle l_e, l_1, \ldots, l_n\rangle) \models K_i\psi \text{ if } l_i = l'_i \text{ implies } (\mathcal{P}^\sigma, \langle l'_e, l'_1, \ldots, l'_n\rangle) \models \psi \\
(\mathcal{P}^\sigma, \langle l_e, l_1, \ldots, l_n\rangle) \models D_G\psi \text{ if } l_i = l'_i \text{ for all } i \in G \text{ implies } (\mathcal{P}^\sigma, \langle l'_e, l'_1, \ldots, l'_n\rangle) \models \psi
\end{array}$$

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

Definition 6 (Validity on QIS). *If ϕ is a formula in $\mathcal{L}_n^D$ and $\mathcal{P}$ is a quantified interpreted systems, then $\mathcal{P} \models \phi$* if *$\mathcal{P}_f \models \phi$.*

Thus, we can reason about a multi-agent system by using the expressiveness of QISs, but rely on Kripke models to prove properties of the system. By the definition of validity on QISs, if $\phi \in \mathcal{L}_n^D$ is a validity on the class $\mathcal{F}_E$ of equivalence frames, then ϕ holds on the class of SGSs.

3.3 Some Validities

We briefly explore the semantics of QISs by considering the traditional Barcan formulas [8]. 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 QISs, i.e., they hold in every quantified interpreted system (for a proof see the reference above):

$$\mathcal{QIS} \models \forall x K_i \phi \to K_i \forall x \phi \qquad BF_i$$
$$\mathcal{QIS} \models K_i \forall x \phi \to \forall x K_i \phi \qquad CBF_i$$
$$\mathcal{QIS} \models \forall x D_G \phi \to D_G \forall x \phi \qquad BF_G$$
$$\mathcal{QIS} \models D_G \forall x \phi \to \forall x D_G \phi \qquad CBF_G$$

These validities are in line with the bird's eye approach usually adopted in epistemic logic. By BF_i if agent i knows that a is ϕ for each individual a, then she knows that all the individuals are ϕ, even if she has not to be aware of this fact. 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_i if agent i knows that all the individuals are ϕ, then she knows that a is ϕ, for each individual $a \in D$. Similar considerations apply to BF_G and CBF_G. Once again we underline that the external account of knowledge applies both to *de re* and *de dicto* modalities, i.e., whether the quantifier is outside or inside the scope of a knowledge operator. We have also generalised versions of the Barcan formula and its converse, for arbitrary strings of epistemic operators:

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

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 either. They say that agents can generalise and particularise not only their direct knowledge, but also the knowledge they have of other agents' knowledge, when this is considered from the viewpoint of an external observer.

4 Message-Passing QIS

In this section we show how to model a message-passing system (MPS) in the framework of QISs. An MPS is a multi-agent system where the most relevant actions are sending and receiving messages. In an MPS the local state of an agent i contains information about its initial state, the messages it has sent and received, and the internal actions it has taken. For the formal presentation of message-passing systems we refer to [6], par. 4.4.5–6, although here we do not explicitly consider the temporal evolution of MPSs. The main result of this section consists in showing that Proposition 4.4.3 in [6] can be reformulated as a validity on the class of QISs modelling MPSs.

More formally, for every agent $i \in A$ we introduce an initial event $init(i)$, a set INT_i of internal actions $\alpha_1, \alpha_2, \dots$, and a set MSG of messages $\mu_1, \mu_2, \dots$ A local state l_i for agent i is a sequence of events whose first element is $init(i)$ and whose following elements are events of the form $send(i, j, \mu)$, $rec(i, j, \mu)$ or $int(i, \alpha)$, for $j \in A$, $\mu \in MSG$ and $\alpha \in INT_i$, describing the actions performed by i. Intuitively, $send(i, j, \mu)$ represents the event *agent i sends message μ to j*, while $rec(i, j, \mu)$ represents the event *agent i receives message μ from j*, and $int(i, \alpha)$ represents the event *agent i performs internal action α*.

A global state s is an n-tuple $\langle l_e, l_1, \dots, l_n \rangle$, where l_e contains all the events in $l_1, \dots, l_n$. We now 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 $i \in A$. The message-passing QISs (MPQISs) we consider are linearly ordered, that is, for every $s, s' \in \mathcal{P}$, either $s \leq s'$ or $s' \leq s$, and contain the initial global state $\langle init(e), init(1), \ldots, init(n)\rangle$. These constraints correspond to the assumption that every MPQIS models the evolution of a single MPS: starting from the initial state, the MPQIS contains all the states reachable during the execution of the MPS. The temporal evolution of an MPS can be represented as a sequence $s_0, s_1, \ldots$ of global states such that $s_0 = \langle init(e), init(1), \ldots, init(n)\rangle$, 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 MPQIS represents various temporal evolutions differing on the number of idle steps. Finally, in each MPQIS the domain D of individuals comprises all agents in A, the messages in MSG, the actions in the various INT_i, and the events $e_1, e_2, \ldots$

We assume that the language $\mathcal{L}_n^D$ 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 $\forall i, j, \mu, \alpha(\phi[send(i, j, \mu)] \wedge \phi[rec(i, j, \mu)] \wedge \phi[init(i)] \wedge \phi[int(i, \alpha)])$. 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 language. In $\mathcal{L}_n^D$ we can define events by formulas which are provably valid in every MPQIS (the existence of a unique individual $\exists!$ can be defined by means of $=$):

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

The first formula expresses the fact that every event is either a *send* or *receive* event, where the sender is different from the receiver, or an initial event, or an internal action. Thus, it cannot be the case that $e = send(i, j, \mu) = send(i', j', \mu')$, for distinct agents and messages. The second formula says that every *send* or *receive* event, initial event, 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 MPQISs validates these specifications.

It is more interesting to consider specifications involving epistemic operators. In [6], p. 132, the authors list three constraints on MPSs, the third one involves runs in an SGS. Nonetheless, we can reformulate the first two without introducing runs:

MP1 every $l_i(s)$ contains only events over $init(i)$, INT_i and MSG;
MP2 for every event $rec(i, j, \mu)$ in $l_i(s)$ there exists a corresponding event $send(j, i, \mu)$ in $l_j(s)$.

We formalise these specifications in the language of MPQISs. First, we introduce a predicative constant H for 'happened' such that $(\mathcal{P}^\sigma, s) \models H(e)$ iff e is an event in s. The formulas $Send(i, j, \mu)$, $Rec(i, j, \mu)$, $Init(i)$, and $Int(i, \alpha)$ are shorthands for $H(send(i, j, \mu))$, $H(rec(i, j, \mu))$, $H(init(i))$, and $H(int(i, \alpha))$ respectively. Now we formalise our specification as follows:

MP1' $\forall e(K_i H(e) \to \exists \alpha, j, \mu(e = init(i) \vee e = int(i, \alpha) \vee e = send(i, j, \mu) \vee e = rec(i, j, \mu)))$

MP2' $\forall j, \mu(Rec(i, j, \mu) \to K_i Send(j, i, \mu))$

If $e \in l_i(s)$ then $(\mathcal{P}^\sigma, s) \models K_i H(e)$. By MP1' $(\mathcal{P}^\sigma, s) \models e = init(i) \vee \exists \alpha, j, \mu(e = int(i, \alpha) \vee e = send(i, j, \mu) \vee e = rec(i, j, \mu))$, which means that if event e belongs to the local state of agent i, then it is either i's initial event, or it is an internal action of i, or it is a *send* or *receive* event of i, that is, MP1 holds. Further, if event $rec(i, j, \mu)$ appears in $l_i(s)$, then by MP2' $(\mathcal{P}^\sigma, s) \models K_i Send(j, i, \mu)$, which means that in particular $send(j, i, \mu) \in l_j(s)$, that is, MP2 holds.

It is easy to check that MP1' holds in the class of all MPQISs by the way they are defined, while MP2' in general can fail. Moreover, in [6] the authors single out the *reliable* MPSs, where every message sent is eventually received. Modified from [6], an MPS is reliable iff it satisfies the specification below:

MP4 for all agents i, j and all states s, if $send(i, j, \mu)$ is in $l_i(s)$, then there exists a s' such that $rec(j, i, \mu)$ is in $l_j(s')$.

We formalise this specification as follows:

MP4' $\forall j, \mu(Send(i, j, \mu) \to \neg K_i \neg Rec(j, i, \mu))$

In fact, if $send(i, j, \mu)$ is in $l_i(s)$, by MP4' $(\mathcal{P}^\sigma, s) \models \neg K_i \neg Rec(j, i, \mu)$, this means that there exists a global state s' such that $(\mathcal{P}^\sigma, s') \models Rec(j, i, \mu)$, that is, $rec(j, i, \mu) \in l_j(s')$. Thus, MP4 holds. Note that MP4' is far stronger than MP4 as the former requires that the local states of agent i in s and s' are identical.

We now prove the main result of this section, that is, Proposition 4.4.3 in [6] can be restated as a validity on the class of MPQISs satisfying MP1, MP2 and two simplifying assumptions. As to the former, in presenting the MPSs in [6], the authors model the local state of an agent as a sequence of sets of events. Then they introduce the semplifying assumption MP5, according to which the sets of events are actually singletons. Since we defined the local states as tuples of events, MP5 is already satisfied in the present framework. As regards the latter assumption:

MP6 All the events in a given agent's local state are distinct.

Also for MP6 we can find a formula in $\mathcal{L}_n^D$ whose validity guarantees that this specification holds. We say that $(\mathcal{P}^\sigma, s) \models Prec(e, e')$ iff $(\mathcal{P}^\sigma, s) \models H(e) \wedge H(e')$ and for every $s' \leq s$, $(\mathcal{P}^\sigma, s') \models H(e') \to H(e)$. Intuitively, $Prec(e, e')$ means that event e appears no later than e'; while $Succ(e, e')$ stands for that event e' happens immediately after event e.

MP6' $\forall e, e'(e \neq e' \wedge K_i H(e) \wedge K_i H(e') \to K_i(\neg Succ(e, e) \wedge \neg(Prec(e, e') \wedge Prec(e', e)))$

If it is not the case that MP6, then we have two occurrences of event e in $l_i(s)$. If these are not separated by any event e', then $Succ(e, e)$ holds and we have a contradiction. If they are separated by some other event e', then $(\mathcal{P}^\sigma, s) \models Prec(e, e') \wedge Prec(e', e)$, that is, MP6' fails.

Henceforth we consider only MPQISs satisfying the specifications above, except MP4. We define a notion of *potential causality* between events, which is intended to capture the intuition that event e might have caused event e'. Fix a subset G of A, the relation $\mapsto$ holds between events e, e' iff:

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

Note that $\mapsto$ is anti-symmetric because of MP6. We say that $(\mathcal{P}^\sigma, s) \models e \mapsto e'$ if $e \mapsto e'$ (we use the same notation for semantic and syntactic elements).

Now we prove that the potential causality relation $\mapsto$ respects the order *Prec* of events by showing that the following validity holds in the class of MPQISs satisfying the specifications above. Note that this is the right to left implication of Proposition 4.4.3 in [6]:

$$MPQIS \models \forall e, e'(H(e) \land H(e') \to ((e \mapsto e') \to D_G Prec(e, e')))$$

Proof. Assume that $(\mathcal{P}^\sigma, s) \models H(e) \land H(e') \land e \mapsto e'$. If e' is a *receive* event $rec(i, j, \mu)$ and e is the corresponding *send* event $send(j, i, \mu)$, then $(s, s') \in \bigcap_{i \in G} \sim_i$ implies $(\mathcal{P}^\sigma, s') \models H(e) \land H(e')$ and for $s'' \leq s'$, $(\mathcal{P}^\sigma, s'') \models H(e') \to H(e)$ 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 $(\mathcal{P}^\sigma, s) \models K_i(H(e) \land H(e'))$ and for $s' \leq s$, $(\mathcal{P}^\sigma, s') \models K_i(H(e') \to H(e))$. Also in this case $(\mathcal{P}^\sigma, s) \models D_G Prec(e, e')$.

Finally, if there exists some event e'' such that $e \mapsto e''$ and $e'' \mapsto e'$, then $(\mathcal{P}^\sigma, s) \models H(e'') \to D_G Prec(e, e''), H(e'') \to D_G Prec(e'', e')$. Without loss of generality we can assume that $e'' \mapsto e'$ for either case 1 or 2 above, in both cases $(\mathcal{P}^\sigma, s) \models H(e'')$. Therefore, for every s', $(s, s') \in \bigcap_{i \in G} \sim_i$ implies $(\mathcal{P}^\sigma, s') \models H(e) \land H(e')$ and for $s'' \leq s'$, $(\mathcal{P}^\sigma, s'') \models H(e'') \to H(e) \land H(e') \to H(e'')$. By transitivity $(\mathcal{P}^\sigma, s'') \models H(e') \to H(e)$. Thus, $(\mathcal{P}^\sigma, s) \models D_G Prec(e, e')$. □

The example of the message-passing systems analysed in this section clearly shows the advantages of first-order modal languages in comparison with propositional ones. We were able to formalize in $\mathcal{L}_n^D$ various constraints on MPSs. Most important, Proposition 4.4.3 in [6] turned out out be a validity on the class of QISs modelling MPSs.

5 Axiomatisation

In this section we provide a sound and complete axiomatisation of the validities on the 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 $\mathcal{QIS}$. Technically, we first prove the completeness of the first-order multi-modal system $Q.S5_n^D$ with respect to equivalence frames. Then, by Lemma 1 the completeness of $Q.S5_n^D$ with respect to $\mathcal{SGS}$ follows.

In [14] Kripke proved the completeness of monomodal $Q.S5$ without distributed knowledge (see also [8, 13]). The novelty of this section consists in showing that the techniques in [7] for propositional $S5^D_n$ can be rather straightforwardly extended to the first-order for proving the completeness of $Q.S5^D_n$.

5.1 System $Q.S5^D_n$

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

Definition 7. *The system $Q.S5^D_n$ on $\mathcal{L}^D_n$ contains the following schemes of axioms and inference rules:*

Taut	*every classic propositional tautology*
MP	$\phi \to \psi, \phi \Rightarrow \psi$
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	$\neg K_i\phi \to K_i\neg K_i\phi$
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	$D_G\phi \to D_GD_G\phi$
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	$\phi \to \psi[x/t] \Rightarrow \phi \to \forall x\psi$, *$x$ not free in ϕ*
Id	$t = t$
Func	$t = t' \to (t''[x/t] = t''[x/t'])$
Subst	$t = t' \to (\phi[x/t] \to \phi[x/t'])$, *for atomic ϕ*
K_iId	$t = t' \to K_i(t = t')$
K_iDif	$t \neq t' \to K_i(t \neq t')$

We take the standard definitions of *proof* and *theorem*: $\vdash \phi$ stands for formula $\phi \in \mathcal{L}^D_n$ is a theorem in $Q.S5^D_n$. Moreover, $\phi \in \mathcal{L}^D_n$ is *derivable* in $Q.S5^D_n$ from a set Δ of formulas in $\mathcal{L}^D_n$ - $\Delta \vdash \phi$ 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 every equivalence frame $\mathcal{F}$ is a frame for $Q.S5^D_n$. As a consequence, we have the following soundness result.

Lemma 2 (Soundness). *The system $Q.S5^D_n$ is sound with respect to the class $\mathcal{F}_E$ of equivalence frames.*

By this lemma and the definition of validity on SGSs, the following implications hold:

$$Q.S5^D_n \vdash \phi, \textit{ then } \mathcal{F}_E \models \phi, \textit{ then } \mathcal{SGS} \models \phi$$

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 [8] 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 the first-order the proof for the propositional system $S5_n^D$ in [7]. 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 the results in [7] we can obtain from $\mathcal{M}^{Q.S5_n^D}$ an equivalence model $\mathcal{M}'$ such that $\mathcal{M}^{Q.S5_n^D}$ pseudo-validates ϕ iff $\mathcal{M}' \models \phi$. Thus completeness follows.

To show the first part of the result we rely on two lemmas: the *saturation lemma* and the *truth lemma*. In order to state these partial results we need the following definitions: let Λ be a set of formulas in $\mathcal{L}_n^D$,

Λ is *consistent*	iff $\Lambda \nvdash \bot$;
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.

The worlds $w, w', \dots$ in the canonical model are saturated sets of formulas on the language $\mathcal{L}_n^{D+}$, obtained by expanding $\mathcal{L}_n^D$ with an infinite denumerable set of new constants. Moreover, for closed terms $v, v' \in \mathcal{L}_n^+$, we define $v \sim_w v'$ iff $v = v' \in w$. This defines an equivalence relation; we write $[v]_w = \{v' | v = v' \in w\}$ for the equivalence class of v in w. Since the accessibility relation in $\mathcal{M}^{Q.S5_n^D}$ will be defined by wR_iw' iff $\{\phi | \Box\phi \in w\} \subseteq w'$, by axioms K_iId, K_iDif and $Func$ we can show that the definition of $[v]_w$ is independent from v and from w - i.e. wR_iw' implies $[v]_w = [v]_{w'}$ - so we simply write $[v]$.

The following result follows from Henkin's and Lindenbaum's lemmas, a proof can be found in [13].

Lemma 3. *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 $\mathcal{L}_n^{D+}$ of $\mathcal{L}_n^D$.*

If we assume that $Q.S5_n^D \not\vdash \phi$, the lemma above guarantees that the set W of possible worlds in the canonical model, defined as follows, is non-empty.

Definition 8 (Canonical model). *The canonical model $\mathcal{M}^{Q.S5_n^D}$ for $Q.S5_n^D$ on the language $\mathcal{L}_n^D$, with an expansion $\mathcal{L}_n^{D+}$, is a 4-tuple $\langle W, R, D, I\rangle$ such that*

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

Since T, 4 and 5 are all axioms of $Q.S5_n^D$, it is easy to show that the various R_i and R_G are equivalence relations. Moreover, from $D1$ 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 the introduction of the pseudo-satisfaction relation $\models^p$, defined as $\models$ but for the distributed knowledge operator D (in what follows we simply write $\mathcal{M}$ for $\mathcal{M}^{Q.S5_n^D}$):

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

Now we can prove the *truth lemma* for the pseudo-satisfaction relation $\models^p$. To obtain such a result we first 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 $w \in \mathcal{M}$, $\phi \in \mathcal{L}_n^{D+}$, for $\sigma(y_i) = [v_i]$,*

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

Proof. The proof is by induction on the structure of $\phi \in \mathcal{L}_n^{D+}$.
$\phi = P^k(t_1, \dots, t_k)$. By the definitions of pseudo-satisfaction and canonical interpretation $(\mathcal{M}^\sigma, w) \models^p P^k(t_1[\vec{y}], \dots, t_k[\vec{y}])$ iff $\langle I^\sigma(t_1[\vec{y}]), \dots, I^\sigma(t_k[\vec{y}])\rangle \in I(P^k, w)$ iff $\langle [t_1[\vec{y}/\vec{v}]], \dots, [t_k[\vec{y}/\vec{v}]]\rangle \in I(P^k, w)$ iff $P^k(t_1[\vec{y}/\vec{v}], \dots, t_k[\vec{y}/\vec{v}]) \in w$.
$\phi = \neg\psi, \psi \to \psi', \forall x\psi$. 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 inductive 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 [8, 13] 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 may 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 [7]. 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 path from w to w', and no path from w back to itself, (ii) wR_i^*w' implies $h(w)R_ih(w')$ and wR_G^*w' 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. $\langle w_j, w_{j+1}\rangle \in R^*_{i_j}$.

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 $w, w' \in W^*$ if $w' = v_{w,i,w''}$ for some $w'' \in W$, and $h(w)R_ih(w')$. Finally, define $I^*(P^k, w) = I(P^k, h(w))$. It can be checked that $\mathcal{M}^*$ and h satisfy (i)-(iii) above, we omit the proof for reasons of space and refer to [7]. In particular, we can show what follows:

Lemma 5. *For every $w \in W^*$, $\phi \in \mathcal{L}_n^D$, $(\mathcal{M}^{*\ \sigma}, w) \models^p \phi$ 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 wR_i^*w' and $(\mathcal{M}^{*\ \sigma}, w') \not\models^p \psi$. This means that $h(w)R_ih(w')$ and $(\mathcal{M}^\sigma, h(w')) \not\models^p \psi$ by inductive 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 $h(w)R_iw'$ and $(\mathcal{M}^\sigma, w') \not\models^p \psi$. By construction $v_{w,i,w'} \in W^*$, $h(v_{w,i,w'}) = w'$ and $wR_i^*v_{w,i,w'}$. By the inductive 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 model $\mathcal{M}'$ that does not validate any 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 an equivalence frame. Further, we can prove the following result:

Lemma 6. *For every w, $\phi \in \mathcal{L}_n^D$, $(\mathcal{M}'^\sigma, w) \models \phi$ 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 [7] 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 inductive 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 $\mathcal{F}_E$ of equivalence 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 model $\mathcal{M} = \langle \mathcal{F}, I\rangle$ based on an equivalence 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 QISs, $\mathcal{P} \models \phi$ 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:

$$\mathcal{SGS} \models \phi, \textit{ then } \mathcal{F}_E \models \phi, \textit{ then } Q.S5_n^D \vdash \phi$$

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 ϕ 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 modal 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 [9, 18].

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 MPQISs 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 [15].

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 [11].

References

1. T. Ågotnes, W. van der Hoek, M. Wooldridge. Quantified Coalition Logic. *Proceedings of IJCAI*, 1181–1186, Hyderabad, India, January 6-12, 2007.
2. F. Belardinelli, A. Lomuscio. Quantified Epistemic Logics with Flexible Terms. *LORI workshop on Logic, Rationality and Interaction*, Beijing, 5-9 August, 2007.
3. P. Blackburn, M. de Rijke, Y. Venema. *Modal Logic.* Cambridge University Press, Cambridge, 2001.
4. A. Chagrov, M. Zakharyaschev. *Modal Logic.* Oxford University Press, 1997.
5. E. Clarke, O. Grumberg, D. Peled. *Model Checking.* MIT Press, Cambridge, 1999.
6. R. Fagin, J. Halpern, Y. Moses, M. Vardi. *Reasoning about Knowledge.* MIT Press, Cambridge, 1995.
7. R. Fagin, J. Halpern, M. Vardi. What can machines know? on the properties of knowledge in distributed systems. *J. ACM*, 39(2):328–376, 1992.
8. M. Fitting, R. Mendelsohn. *First-order Modal Logic.* Kluwer, Dordrecht, 1999.
9. P. Gammie, R. van der Meyden. MCK: Model checking the logic of knowledge. In *CAV*, volume 3114 of *LNCS*, pages 479–483. Springer, 2004.
10. J. Halpern, R. Fagin. Modelling knowledge and action in distributed systems. *Distributed Computing*, 3(4):159–179, 1989.
11. I. Hodkinson, F. Wolter, M. Zakharyaschev. Decidable fragment of first-order temporal logics. *Ann. Pure Appl. Logic*, 106(1-3):85–134, 2000.
12. W. van der Hoek, M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of AAMAS'02*, pages 1167–1174. ACM Press, 2002.
13. G. Hughes, M. Cresswell. *A New Introduction to Modal Logic.* Routledge, 1996.
14. S. Kripke. A Completeness Theorem in Modal Logic. *J. Sym. Log.*, 24, 1-14, 1959.
15. A. Lomuscio, M. Colombetti. QLB: a quantified logic for belief. In *Proceedings of ATAL-96*, volume 1193 of *Lecture Notes in AI.* Springer-Verlag, Heidelberg, 1996.
16. A. Lomuscio, M. Ryan. On the relation between interpreted systems and kripke models. In *Proceedings of AI97*, volume 1441. Springer lecture notes in AI, 1997.
17. J.-J. Ch. Meyer, W. van der Hoek. *Epistemic Logic for AI and Computer Science.* Cambridge University Press, 1995.
18. F. Raimondi, A. Lomuscio. Automatic verification of multi-agent systems by model checking via OBDDs. *Journal of Applied Logic*, 5(2), 235–251, 2007.