

A Complete First-Order Logic of Knowledge and Time

Francesco Belardinelli

Scuola Normale Superiore, Pisa

F.Belardinelli@sns.it

Alessio Lomuscio

Department of Computing

Imperial College London

A.Lomuscio@imperial.ac.uk

Keywords: Epistemic Logics: first-order temporal epistemic logic, message passing systems, completeness.

Abstract

We introduce and investigate quantified interpreted systems, a semantics to reason about knowledge and time in a first-order setting. We provide an axiomatisation, which we show to be sound and complete. We utilise the formalism to study message passing systems (Lamport 1978; Fagin et al 1995) in a first-order setting, and compare the results obtained to those available for the propositional case.

Introduction

The area of modal logic (Blackburn, van Benthem, and Wolter 2007; Chagrov and Zakharyashev 1997) has received considerable attention in artificial intelligence over the years. Research has pursued both fundamental theoretical investigations (completeness, decidability, complexity, etc), as well as the use of modal formalisms in specification and automatic system verification, as in model checking (Clarke, Grumberg, and Peled 1999).

Among the most well-known formalisms are propositional modal logics for reasoning about knowledge, or propositional epistemic logics (Fagin et al 1995; Meyer and Hoek 1995). The typical epistemic language extends propositional logic by adding n modalities K_i representing the knowledge of agent i in a group $A = \{1, \dots, n\}$ of agents. For expressiveness purposes, epistemic logic has been extended in several ways. In one direction, further modalities have been added to the formalism (distributed knowledge, common knowledge, belief, etc.) for representing the knowledge shared in a group of agents. In another one, the epistemic language has been enriched with temporal operators under the assumption of a given model of time (e.g., linear or branching, discrete or continuous, etc.). In all these lines of work there is a tension between extending the expressiveness of the language reflecting the system to be modeled and retaining some useful theoretical properties of the formalism, such as decidability.

This tension is still present in the exercise conducted here, where we aim at extending a combination of epistemic and

temporal logic to predicate level. We apply this result in the modeling of a class of computational structures normally referred to as message passing systems (Lamport 1978). We also show that known metatheoretical properties of message passing systems (Fagin et al 1995) become validities in the predicate logic here considered.

Our starting point is a number of results by Halpern, van der Meyden, and others regarding the combination of time and knowledge at propositional level (Fagin, Halpern, and Vardi 1992; Meyden 1994) together with studies by, among others, Hodkinson, Reynolds, Wolter, Zakharyashev for first-order temporal logic including both positive (Hodkinson, Wolter, and Zakharyashev 2000; Reynolds 1996; Wolter and Zakharyashev 2002) and negative results (Wolter 2000). In this note we also make use of our initial work in this direction (Belardinelli and Lomuscio 2007a; 2007b), where static (i.e., non-temporal) quantified epistemic logics were axiomatised.

Our motivation for the above comes from an interest in reasoning about reactive, autonomous distributed systems, or multi-agent systems (MAS), whose high-level properties may usefully be modeled by epistemic formalisms suitably extended to incorporate temporal logic. While temporal epistemic logics are well understood at propositional level (Fagin et al 1995; Meyer and Hoek 1995), their usefulness has been demonstrated in a number of applications (security and communication protocols, robotics), and model checking tools have been developed for them (Gammie and van der Meyden 2004; Raimondi and Lomuscio 2007; Dembiński et al 2003), still there is a growing need in web-services, security, as well as other areas, to extend these languages to first-order (see (Cohen and Dams 2007; Solanki, Cau, and Zedan 2006; Viganò 2007)). Moreover, a number of formalisms, including *BDI* logics (Rao and Georgeff 1991), the *KQML* framework (Cohen, and Levesque 1995), and *LORA* (Wooldridge 2000), have put forward agent theories that include the power of first-order quantification. However, most of these contributions do not address the issue of completeness, a core concern here.

In MAS applications the power of first-order logic is welcome every time agents' knowledge is concerned with:

- Relational statement, as in *agent i knows that message μ was sent by a to b* , or formally

$$K_i \langle P \rangle \text{Send}(a, b, \mu);$$

(where $\langle P \rangle$ is the diamond for past time);

- Functional dependency and identity: *at some future point agent i will know that message μ is the encryption of message μ' with key k* , formally

$$\langle F \rangle K_i(\mu = \text{enc}(k, \mu'));$$

- An infinite domain of individuals, or a finite domain whose cardinality cannot be bounded in advance: *agent i has to read an e-mail before deleting it*,

$$\forall \mu (\text{Delete}(i, \mu) \rightarrow \langle P \rangle \text{Read}(i, \mu));$$

- Quantification on agents (Lomuscio and Colombetti 1996): *the child of any process knows which process launched it*

$$\forall i K_{\text{child}(i)} \langle P \rangle \text{Launch}(i, \text{child}(i))$$

Furthermore, in the context of logics for knowledge it is known that epistemic modalities can be combined with quantifiers to express concepts such as knowledge *de re* and *de dicto* (Fitting, and Mendelsohn 1999; Hughes and Cresswell 1996). For instance, an agent i might know that every computation will eventually produce an output, thus having the *de dicto* knowledge expressed by the following specification:

$$\forall \text{comp } K_i \langle F \rangle \exists y \text{Output}(\text{comp}, y)$$

but she might not know the actual output of every computation. Therefore, the following *de re* specification:

$$\forall \text{comp } \exists y K_i \langle F \rangle \text{Output}(\text{comp}, y)$$

would not be satisfied. From the examples above we conclude that quantification can significantly extend the expressiveness of epistemic languages.

While the specifications above call for a first-order language, we need to consider why one should use an undecidable language when a decidable one (propositional temporal epistemic logic in our case) does a reasonable job already. Although this is a sensible objection, we should stress that in many practical applications, such as in model checking, we are typically not so much concerned with the validity problem but with satisfaction in a given model, which is often an easier problem, particularly for some classes of formulas. Additionally, recent research, including among others (Hodkinson, Wolter, and Zakharyashev 2000; Sturm, Wolter, and Zakharyashev 2000; 2002; Wolter and Zakharyashev 2001), has put forward useful decidable fragments of first-order modal logic, thereby opening the way for further extensions.

We approach the problem by introducing quantified interpreted systems, an extension to first-order of “standard” interpreted systems (Halpern, and Fagin 1989; Parikh and Ramanujam 1985), which are used to interpret a language for temporal epistemic logic including distributed knowledge. First, a sound and complete axiomatisation is presented. Second, message passing systems, a basic framework for reasoning about asynchronous systems (Lamport 1978) are analysed in the light of the novel formalism, and the results compared to the treatment in propositional logic.

A Quantified Temporal Epistemic Logic

In this section we extend to first-order the formalism of interpreted systems, a class of structures introduced to model the behaviour of multi-agent systems (Fagin et al 1995; Meyer and Hoek 1995). In what follows we assume a finite set $A = \{i_1, \dots, i_n\}$ of agents.

Syntax

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

Definition 1 *Terms and formulas in the language $\mathcal{L}_n$ are defined in the Backus-Naur form as follows:*

$$t ::= x \mid f^k(\vec{t})$$

$$\phi ::= P^k(\vec{t}) \mid t = t' \mid \neg \phi \mid \phi \rightarrow \psi \mid K_i \phi \mid D_G \phi \mid [F] \phi \mid [P] \phi \mid \forall x \phi$$

The formula $K_i \phi$ means “agent i knows ϕ ”, while $D_G \phi$ represents “ ϕ is distributed knowledge among the agents in G ”, and $[F] \phi$ (respectively $[P] \phi$) stands for “ ϕ will always be true” (respectively “ ϕ has always been true”). The symbols $\perp, \wedge, \vee, \leftrightarrow, \exists, \langle F \rangle$ (sometime in the future), $\langle P \rangle$ (sometime in the past) are defined as standard. The temporal operators $[F]^+$ (every future time including the present) and $[P]^+$ (every past time including the present) can be defined as $\phi \wedge [F] \phi$ and $\phi \wedge [P] \phi$ respectively.

We refer to 0-ary functors as *individual constants* $c_1, c_2, \dots$. A closed term v is a term where no variable appears; closed terms are either constants or terms obtained by applying functors to closed terms.

By $t[\vec{y}]$ (resp. $\phi[\vec{y}]$) we mean that $\vec{y} = y_1, \dots, y_n$ 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 substituting simultaneously some, possibly all, free occurrences of $\vec{y}$ in t (resp. ϕ) with $\vec{t} = t_1, \dots, t_n$, renaming bounded variables if necessary.

Quantified Interpreted Systems

Interpreted systems are widely used to model the behaviour of MAS, in this subsection we extend these structures to first-order. This extension can be performed in several ways, all leading to different results. For instance, we could introduce a domain of quantification for each agent and/or for each computational state (see (Belardinelli and Lomuscio 2007a; 2007b) for a discussion of the static case). In this paper we consider the simplest extension, obtained by adding a single quantification domain D common to all agents and states. We present further options in the conclusions.

More formally, for each agent $i \in A$ in a multi-agent system we introduce a set L_i of local states $l_i, l'_i, \dots$, and a set Act_i of actions $\alpha_i, \alpha'_i, \dots$. We consider local states and actions for the environment e as well. The set $\mathcal{S} \subseteq L_e \times L_1 \times \dots \times L_n$ contains all possible global states of the MAS, while $Act \subseteq Act_e \times Act_1 \times \dots \times Act_n$ is the set of

all possible joint actions. Note that some states may never be reached and some joint actions may never be performed. We also introduce a transition function $\tau : Act \rightarrow (\mathcal{S} \rightarrow \mathcal{S})$. Intuitively, $\tau(\alpha)(s) = s'$ encodes that the agents can access the global state s' from s by performing the joint action $\alpha \in Act$. The transition function τ defines the admissible evolutions of the MAS. We say that the global state s' is *reachable in one step* from s , or $s \prec s'$, iff there is $\alpha \in Act$ such that $\tau(\alpha)(s) = s'$; while s' is *reachable* from s iff $s \prec^+ s'$, where $\prec^+$ is the transitive closure of relation $\prec$.

To represent the temporal evolution of the MAS we consider the flow of time $T = \langle T, < \rangle$ defined as a weakly connected, strict partial order, i.e., T is a non-empty set and the relation $<$ on T is irreflexive, transitive and weakly connected: for n, n', n'' in T ,

- $n \not\prec n$
- $(n < n' \wedge n' < n'') \rightarrow (n < n'')$
- $(n < n' \wedge n < n'') \rightarrow (n' < n'' \vee n'' < n' \vee n' = n'')$
- $(n' < n \wedge n'' < n) \rightarrow (n' < n'' \vee n'' < n' \vee n' = n'')$

The relation $<$ can be thought of as the precedence relation on the set T of moments in time. A run r over $\langle \mathcal{S}, Act, \tau, T \rangle$, where $\mathcal{S}$, Act , τ , and T are defined as above, is a function from T to $\mathcal{S}$ such that $n < n'$ implies $r(n) \prec^+ r(n')$. Intuitively, a run represents a possible evolution of the MAS on the flow of time T .

We now define the quantified interpreted systems for the language $\mathcal{L}_n$ as follows:

Definition 2 A quantified interpreted system, or QIS, over $\langle \mathcal{S}, Act, \tau, T \rangle$ is a triple $\mathcal{P} = \langle R, D, I \rangle$ such that R is a non-empty set of runs over $\langle \mathcal{S}, Act, \tau, T \rangle$; D is a non-empty set of individuals; $I(f^k)$ is a k -ary function from D^k to D ; for $r \in R$, $n \in T$, $I(P^k, r, n)$ is a k -ary relation on D and $I(=, r, n)$ is the equality on D . We denote by QIS the class of all quantified interpreted systems.

Note that individual constants as well as functors in $\mathcal{L}_n$ are interpreted rigidly, that is, their interpretation is the same in every global state. Further, the present definition of quantified interpreted systems covers the most intuitive formalisations of time, as it includes $\mathbb{N}$, $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ with a notion of precedence among instants. Therefore, QIS are general enough to cover a wide range of cases, while still being interesting for applications.

Now we assign a meaning to the formulas of $\mathcal{L}_n$ in quantified interpreted systems. Following standard notation (Fagin et al 1995) a pair (r, m) is a *point* in $\mathcal{P}$. If $r(m) = \langle l_e, l_1, \dots, l_n \rangle$ is the global state at (r, m) , then $r_e(m) = l_e$ and $r_i(m) = l_i$ are the environment's and agent i 's local state at (r, m) respectively. We consider also the converse relation $>$ defined as $n > m$ iff $m < n$, and the partial order $\leq$ such that $n \leq m$ iff $n < m$ or $n = m$.

Let σ be an assignment from the variables in $\mathcal{L}_n$ to the individuals in D , the valuation $I^\sigma(t)$ of a term t is defined as $\sigma(y)$ for $t = y$, and $I^\sigma(t) = I(f^k)(I^\sigma(t_1), \dots, I^\sigma(t_k))$, for $t = f(\vec{t})$. A variant σ_a^x of an assignment σ assigns $a \in D$ to x and coincides with σ on all the other variables.

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

$$\begin{aligned} (\mathcal{P}^\sigma, r, m) \models P^k(\vec{t}) & \text{ iff } \langle I^\sigma(t_1), \dots, I^\sigma(t_k) \rangle \in I(P^k, r, m) \\ (\mathcal{P}^\sigma, r, m) \models t = t' & \text{ iff } I^\sigma(t) = I^\sigma(t') \\ (\mathcal{P}^\sigma, r, m) \models \neg\psi & \text{ iff } (\mathcal{P}^\sigma, r, m) \not\models \psi \\ (\mathcal{P}^\sigma, r, m) \models \psi \rightarrow \psi' & \text{ iff } (\mathcal{P}^\sigma, r, m) \not\models \psi \text{ or } (\mathcal{P}^\sigma, r, m) \models \psi' \\ (\mathcal{P}^\sigma, r, m) \models K_i\psi & \text{ iff } r_i(m) = r'_i(m') \text{ implies } (\mathcal{P}^\sigma, r', m') \models \psi \\ (\mathcal{P}^\sigma, r, m) \models D_G\psi & \text{ iff } r_i(m) = r'_i(m') \text{ for all } i \in G, \\ & \text{ implies } (\mathcal{P}^\sigma, r', m') \models \psi \\ (\mathcal{P}^\sigma, r, m) \models [F]\psi & \text{ iff } m < m' \text{ implies } (\mathcal{P}^\sigma, r, m') \models \psi \\ (\mathcal{P}^\sigma, r, m) \models [P]\psi & \text{ iff } m > m' \text{ implies } (\mathcal{P}^\sigma, r, m') \models \psi \\ (\mathcal{P}^\sigma, r, m) \models \forall x\psi & \text{ iff for all } a \in D, (\mathcal{P}^{\sigma_a^x}, r, m) \models \psi \end{aligned}$$

The truth conditions for $\perp$, $\wedge$, $\vee$, $\leftrightarrow$, $\exists$, $\langle F \rangle$, and $\langle P \rangle$ are defined from those above. In particular, the temporal operators $[F]^+$ and $[P]^+$ respect the intended semantics:

$$\begin{aligned} (\mathcal{P}^\sigma, r, m) \models [F]^+\psi & \text{ iff } m \leq m' \text{ implies } (\mathcal{P}^\sigma, r, m') \models \psi \\ (\mathcal{P}^\sigma, r, m) \models [P]^+\psi & \text{ iff } m \geq m' \text{ implies } (\mathcal{P}^\sigma, r, m') \models \psi \end{aligned}$$

A formula $\phi \in \mathcal{L}_n$ is said to be *true at a point* (r, m) iff it is satisfied at (r, m) by every σ ; ϕ is *valid on a QIS* $\mathcal{P}$ iff it is true at every point in $\mathcal{P}$; ϕ is *valid on a class* $\mathcal{C}$ of QIS iff it is valid on every QIS in $\mathcal{C}$.

The present definition of QIS is based on two assumptions. Firstly, the domain D of individuals is the same for every agent i , so all agents reason about the same objects. This choice is consistent with the *external account of knowledge* usually adopted in the framework of interpreted systems: if knowledge is ascribed to agents by an external observer, i.e., the specifier of the system, it seems natural to focus on the set of individuals assumed to exist by the observer. Secondly, 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 can be justified by the external account of knowledge: all individuals are supposed to be existing from the observer's viewpoint. However, either assumption can be relaxed to accommodate agent-indexed domains as well as individuals appearing and disappearing in the flow of time. We discuss further options in the conclusions. Finally, it can be the case that $A \subseteq D$: this means that the agents can reason about themselves, their properties, and relationships.

Expressiveness

Clearly, the language $\mathcal{L}_n$ is extremely expressive. We can use it to specify the temporal evolution of agents' knowledge, as well as the knowledge agents have of temporal facts about individuals. Both features are exemplified in the following specification: *agent i will know that someone sent him a message when he receives it*,

$$\forall j, \mu [F] (Rec(i, j, \mu) \rightarrow K_i \langle P \rangle Send(j, i, \mu)) \quad (1)$$

In $\mathcal{L}_n$ we can also express that *if agent i receives a message, then he will know that someone sent it to him*:

$$\forall \mu [F](\exists j Rec(i, j, \mu) \rightarrow K_i \exists j' \langle P \rangle Send(j', i, \mu)) \quad (2)$$

The latter specification is weaker than the former: (2) says nothing about the identity of the sender, while (1) requires that *the receiver knows the identity of the sender*. Further, we can express the fact that the existence of a sender is assumed only at the time the message is sent:

$$\forall \mu [F](\exists j Rec(i, j, \mu) \rightarrow K_i \langle P \rangle \exists j' Send(j', i, \mu))$$

In the section on message passing systems we provide further examples of the expressiveness of $\mathcal{L}_n$. Most importantly, we will show that this expressiveness is attained while retaining completeness.

We conclude this paragraph by considering some relevant validities on the class of QIS. Given that the domain of quantification is the same in every global state, both the Barcan formula and its converse are valid on the class of all QIS for all primitive modalities:

$$\begin{aligned} QIS &\models \forall x K_i \phi \leftrightarrow K_i \forall x \phi \\ QIS &\models \forall x D_G \phi \leftrightarrow D_G \forall x \phi \\ QIS &\models \forall x [F] \phi \leftrightarrow [F] \forall x \phi \\ QIS &\models \forall x [P] \phi \leftrightarrow [P] \forall x \phi \end{aligned}$$

Also, these validities are in line with the bird's eye approach usually adopted in epistemic logic. However, should we wish to do so, we can drop them by introducing quantified interpreted systems with varying domains.

For what concerns identity, the following principles hold:

$$\begin{aligned} QIS &\models t = t' \rightarrow K_i(t = t') & QIS &\models t \neq t' \rightarrow K_i(t \neq t') \\ QIS &\models t = t' \rightarrow D_G(t = t') & QIS &\models t \neq t' \rightarrow D_G(t \neq t') \\ QIS &\models t = t' \rightarrow [F](t = t') & QIS &\models t \neq t' \rightarrow [F](t \neq t') \\ QIS &\models t = t' \rightarrow [P](t = t') & QIS &\models t \neq t' \rightarrow [P](t \neq t') \end{aligned}$$

These validities, which hold because of rigid designation, are consistent with the external account of knowledge. However, should we require terms whose denotations depends on the epistemic states of agents, or change accordingly to the evolution of the MAS, we can consider introducing *flexible* terms in the language (Belardinelli and Lomuscio 2007b). In such an extended formalism none of the validities above holds whenever t and t' are flexible terms.

The System QKT.S5_n

In this section we provide a sound and complete axiomatisation of quantified interpreted systems. This result shows that, even though language $\mathcal{L}_n$ is highly expressive, QIS provide a perfectly adequate semantics for it. This also opens the possibility of developing automated verification methods for the formalism. We first prove the completeness of the first-order multi-modal system QKT.S5_n with respect to Kripke models. The proof presented here is an extension of (Gabbay, Hodkinson, and Reynolds 1993), where completeness of a first-order temporal language on weakly-connected partial orders was presented. Then, by means of a map from Kripke models to QIS, the completeness of QKT.S5_n with respect to QIS follows.

The system QKT.S5_n is a first-order multi-modal version of the propositional system S5 combined with a linear temporal logic. 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 QKT.S5_n. Note that $\Rightarrow$ is the inference relation between formulas, while $\Box$ is a placeholder for any primitive modality in $\mathcal{L}_n$ (both temporal and epistemic).

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

<i>Taut</i>	every instance of classic propositional tautologies
<i>MP</i>	$\phi \rightarrow \psi, \phi \Rightarrow \psi$
<i>Dist</i>	$\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$
<i>4</i>	$\Box\phi \rightarrow \Box\Box\phi$
<i>Nec</i>	$\phi \Rightarrow \Box\phi$
<i>T</i>	$K_i\phi \rightarrow \phi$
<i>5</i>	$\neg K_i\phi \rightarrow K_i\neg K_i\phi$
<i>D1</i>	$D_{\{i\}}\phi \leftrightarrow K_i\phi$
<i>D2</i>	$D_G\phi \rightarrow D_{G'}, \text{ for } G \subseteq G'$
<i>FP</i>	$\phi \rightarrow [F]\langle P \rangle\phi$
<i>PF</i>	$\phi \rightarrow [P]\langle F \rangle\phi$
<i>WConF</i>	$\langle P \rangle\langle F \rangle\phi \rightarrow (\langle P \rangle\phi \vee \phi \vee \langle F \rangle\phi)$
<i>WConP</i>	$\langle F \rangle\langle P \rangle\phi \rightarrow (\langle P \rangle\phi \vee \phi \vee \langle F \rangle\phi)$
<i>Ex</i>	$\forall x\phi \rightarrow \phi[x/t]$
<i>Gen</i>	$\phi \rightarrow \psi[x/t] \Rightarrow \phi \rightarrow \forall x\psi, \text{ where } x \text{ is not free in } \phi$
<i>Id</i>	$t = t$
<i>Func</i>	$t = t' \rightarrow (t''[x/t] = t''[x/t'])$
<i>Subst</i>	$t = t' \rightarrow (\phi[x/t] \rightarrow \phi[x/t'])$

By the definition above the operators K_i and D_G are S5 type modalities, while the future $[F]$ and past $[P]$ operators are axiomatised as linear-time modalities. To this we add the classic theory of quantification, consisting of postulates *Ex* and *Gen*, which are both sound in our interpretation as we are considering a unique domain of individuals. Finally, we have the axioms for identity.

We consider the standard definitions of *proof* and *theorem*: $\vdash \phi$ means that $\phi \in \mathcal{L}_n$ is a theorem in QKT.S5_n. A formula $\phi \in \mathcal{L}_n$ is *derivable* in QKT.S5_n from a set Δ of formulas, or $\Delta \vdash \phi$, iff $\vdash \phi_1 \wedge \dots \wedge \phi_n \rightarrow \phi$ for some $\phi_1, \dots, \phi_n \in \Delta$.

It can be easily checked that the axioms of QKT.S5_n are valid on every QIS and the inference rules preserve validity. As a consequence, we have the following soundness result:

Theorem 5 (Soundness) *The system QKT.S5_n is sound for the class QIS of quantified interpreted systems.*

Now we show that the axioms in QKT.S5_n are not only necessary, but also sufficient to prove all validities on QIS.

Kripke Models

Although quantified interpreted systems are useful for modeling MAS, for showing that QKT.S5_n is complete with respect to QIS we introduce an appropriate class of Kripke models (Blackburn, van Benthem, and Wolter 2007; Chagrov and Zakharyashev 1997), which are more suitable for theoretical investigations, namely, the completeness proof.

Definition 6 *A Kripke model, or K-model, for the language $\mathcal{L}_n$ is a tuple $\mathcal{M} = \langle W, \{\sim_i\}_{i \in A}, <, D, I \rangle$ such that W is a non-empty set; for $i \in A$, $\sim_i$ is an equivalence relation on W ; $<$ is a weakly connected, strict partial order on W ; D is a non-empty set of individuals; $I(f^k)$ is a k -ary function from D^k to D ; for $w \in W$, $I(P^k, w)$ is a k -ary relation on D , and $I(=, w)$ is the equality on D . The class of all Kripke models is denoted by $\mathcal{K}$.*

Further, the satisfaction relation $\models$ for an assignment σ is inductively defined as follows:

$$\begin{aligned} (\mathcal{M}^\sigma, w) &\models P^k(\vec{t}) && \text{iff } \langle I^\sigma(t_1), \dots, I^\sigma(t_k) \rangle \in I(P^k, w) \\ (\mathcal{M}^\sigma, w) &\models t = t' && \text{iff } I^\sigma(t) = I^\sigma(t') \\ (\mathcal{M}^\sigma, w) &\models \neg\psi && \text{iff } (\mathcal{M}^\sigma, w) \not\models \psi \end{aligned}$$

$$\begin{aligned}
(\mathcal{M}^\sigma, w) \models \psi \rightarrow \psi' & \text{ iff } (\mathcal{M}^\sigma, w) \not\models \psi \text{ or } (\mathcal{M}^\sigma, w) \models \psi' \\
(\mathcal{M}^\sigma, w) \models [F]\psi & \text{ iff } w < w' \text{ implies } (\mathcal{M}^\sigma, w') \models \psi \\
(\mathcal{M}^\sigma, w) \models [P]\psi & \text{ iff } w > w' \text{ implies } (\mathcal{M}^\sigma, w') \models \psi \\
(\mathcal{M}^\sigma, w) \models K_i\psi & \text{ iff } w \sim_i w' \text{ implies } (\mathcal{M}^\sigma, w') \models \psi \\
(\mathcal{M}^\sigma, w) \models D_G\psi & \text{ iff } (w, w') \in \bigcap_{i \in G} \sim_i \text{ implies } (\mathcal{M}^\sigma, w') \models \psi \\
(\mathcal{M}^\sigma, w) \models \forall x\psi & \text{ iff for all } a \in D, (\mathcal{M}^{\sigma(a)}, w) \models \psi
\end{aligned}$$

We formally compare Kripke models to quantified interpreted systems by means of a map $g : \mathcal{K} \rightarrow QIS$. Let $\mathcal{M} = \langle W, \{\sim_i\}_{i \in A}, <, D, I \rangle$ be a Kripke model. For every equivalence relation $\sim_i$, for $w \in W$, let the equivalence class $[w]_{\sim_i} = \{w' \mid w \sim_i w'\}$ be a local state for agent i ; while W is the set of local states for the environment. Let $\langle W, < \rangle$ be the irreflexive, transitive and weakly connected flow of time. Then define $g(\mathcal{M})$ as the triple $\langle R, D, I' \rangle$, where R contains the run r such that $r(w) = \langle w, [w]_{\sim_1}, \dots, [w]_{\sim_n} \rangle$ for $w \in W$, D is the same as in $\mathcal{M}$, and $I'(P^k, r, w) = I(P^k, w)$. The structure $g(\mathcal{M})$ is a QIS that satisfies the following result:

Lemma 7 For every $\phi \in \mathcal{L}_n$, $w \in W$,

$$(\mathcal{M}^\sigma, w) \models \phi \quad \text{iff} \quad (g(\mathcal{M})^\sigma, r, w) \models \phi$$

where r is the only run in $g(\mathcal{M})$. We refer to the appendix for a proof of this lemma.

Completeness

We show that the system QKT.S5_n is complete by extending to first-order the proof for the propositional system S5_n^D in (Fagin, Halpern, and Vardi 1992), together with the completeness proof for the first-order temporal logic discussed in (Gabbay, Hodkinson, and Reynolds 1993). The relevance of our result consists in showing that these two methods can be combined together to prove an original completeness result, as long as there is no interaction between epistemic and temporal modalities. Note that an independent completeness proof for S5_n^D appeared in (Meyer and Hoek 1992).

More formally, we show that if QKT.S5_n does not prove a formula $\phi \in \mathcal{L}_n$, then the canonical model $\mathcal{M}^{QKT.S5_n}$ for QKT.S5_n does not pseudo-validate ϕ . It is not guaranteed that pseudo-validity (as defined below) coincides with plain validity, but by results in (Fagin, Halpern, and Vardi 1992; Gabbay, Hodkinson, and Reynolds 1993) from $\mathcal{M}^{QKT.S5_n}$ we can obtain a K -model $\mathcal{M}^+$ such that $\mathcal{M}^{QKT.S5_n}$ pseudo-validates ϕ iff $\mathcal{M}^+ \models \phi$, and completeness follows.

In order to prove 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 $\mathcal{L}_n$,

- Λ is *consistent* iff $\Lambda \not\vdash \perp$;
- Λ is *maximal* iff for every $\phi \in \mathcal{L}_n$, $\phi \in \Lambda$ or $\neg\phi \in \Lambda$;
- Λ is *max-cons* iff Λ is consistent and maximal;
- Λ is *rich* iff $\exists x\phi \in \Lambda \Rightarrow \phi[x/c] \in \Lambda$, for some $c \in \mathcal{L}_n$;
- Λ is *saturated* iff Λ is max-cons and rich.

Assume that QKT.S5_n 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 8 (Saturation (Hughes and Cresswell 1996))

If Δ is a consistent set of formulas in $\mathcal{L}_n$, then it can

be extended to a saturated set Π of formulas on some expansion $\mathcal{L}_n^+$ obtained by adding an infinite enumerable set of new individual constants to $\mathcal{L}_n$.

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

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

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

If $QKT.S5_n \not\vdash \phi$, then by the saturation lemma there is a saturated set $w \supseteq \{\neg\phi\}$, so the set W of possible worlds is non-empty. Since T , 4 and 5 are axioms of QKT.S5_n, 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 $R_G \neq \bigcap_{i \in G} R_i$ (Fagin, Halpern, and Vardi 1992). On the other hand, the relation $<$ is transitive and weakly connected by axioms 4, $WConF$, $WConP$. By FP , PF the relation $w > w'$ defined as $\{\phi \mid [P]\phi \in w\} \subseteq w'$ is the converse of $<$. However, $<$ might not be irreflexive (Gabbay, Hodkinson, and Reynolds 1993).

These remarks give 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 $\mathcal{M}^{QKT.S5_n}$):

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

We state the *truth lemma* for the pseudo-satisfaction relation $\models^p$ and refer to (Fagin, Halpern, and Vardi 1992) for a proof.

Lemma 10 (Truth lemma) Let $w \in \mathcal{M}$, $\psi \in \mathcal{L}_n^+$, $\sigma(y_i) = [v_i]$,

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

We remarked that the canonical model $\mathcal{M}$ might not satisfy $\bigcap_{i \in G} R_i = R_G$. However, by applying the techniques in (Fagin, Halpern, and Vardi 1992) $\mathcal{M}$ can be unwound to get a K -model $\mathcal{M}'$ in such a way that $R_G = \bigcap_{i \in G} R_i$ and the same formulas hold. We refer to the appendix for a proof of the following lemma.

Lemma 11 For every $\psi \in \mathcal{L}_n^+$,

$$\mathcal{M}' \models \psi \quad \text{iff} \quad \mathcal{M} \models^p \psi$$

In conclusion, if QKT.S5_n $\not\vdash \phi$, then the canonical model $\mathcal{M}$ pseudo-satisfies $\neg\phi$ by lemma 10. By lemma 11 we obtain that the K -model $\mathcal{M}'$ does not validate ϕ .

Note that the relation $<'$ on W' might not be irreflexive, as $<$ on W is not such. However, we can apply the techniques in (Gabbay, Hodkinson, and Reynolds 1993) to construct an irreflexive K -model $\mathcal{M}^+$ from $\mathcal{M}'$ such that:

Lemma 12 For every $\psi \in \mathcal{L}_n^+$,

$$\mathcal{M}^+ \models \psi \quad \text{iff} \quad \mathcal{M}' \models \psi$$

Also in this case we refer to the appendix for a proof.

By lemma 12 we conclude that the K -model $\mathcal{M}^+$ falsifies the unprovable formula ϕ . Therefore, the following completeness result holds:

Theorem 13 (Completeness) The system $QKT.S5_n$ is complete for the class $\mathcal{K}$ of Kripke models.

In order to prove completeness for the class QIS consider the quantified interpreted system $g(\mathcal{M}^+)$. In lemma 7 we showed that $\mathcal{M}^+ \models \phi$ iff $g(\mathcal{M}^+) \models \phi$, hence $g(\mathcal{M}^+)$ satisfies $\neg\phi$. As a result, we have the following implications and a further completeness result:

$$QIS \models \phi \quad \Rightarrow \quad \mathcal{K} \models \phi \quad \Rightarrow \quad QKT.S5_n \vdash \phi$$

Theorem 14 (Completeness) The system $QKT.S5_n$ is complete for the class QIS of quantified interpreted systems.

By combining together the soundness and completeness theorems we can compare directly the axiomatisation $QKT.S5_n$ and QIS , so we state our main result:

Corollary 15 (Soundness and Completeness) A formula $\phi \in \mathcal{L}_n$ is valid on the class QIS of quantified interpreted systems iff ϕ is provable in $QKT.S5_n$.

Message Passing Systems as QIS

In this section we model message passing systems (Fagin et al 1995; Lamport 1978) in the framework of QIS. A message passing system (MPS) is a MAS in which the only external actions for the agents are message exchanges, specifically sending and receiving messages. This setting is common in the study of a variety of distributed systems, well beyond the realms of MAS and AI. Indeed, any synchronous or asynchronous networked system can be seen as an MPS.

The notion of time is crucial for the analysis of the ordering of events in MPS. As remarked in (Lamport 1978), a message μ can be said to have been sent (received) before message μ' if μ was sent (respectively received) at an earlier time than μ' . We can of course specify this condition in terms of an external global clock. However, maintaining synchronicity in a distributed system is known to be costly. An alternative is to study asynchronous MPS (or AMPS), where only internal clocks exist and agents can work at arbitrary rates relative to each other.

In what follows we show how both (synchronous) MPS and AMPS can be thought of as particular classes of QIS satisfying a finite number of specifications expressed in the first-order modal language $\mathcal{L}_n$. Further, we analyse in detail the agents' knowledge about the ordering of events in AMPS. Our main result consists in showing that the characterisation of AMPS at propositional level given as a metatheorem (specifically, in (Fagin et al 1995), Proposition 4.4.3) can naturally be cast as a formula in $\mathcal{L}_n$, which turns out to be a validity on the class of QIS we introduce. While the basic details are given below, we refer to (Fagin et al 1995), sections 4.4.5-6, for more details on MPS.

We introduce a set Act of actions $\alpha_1, \alpha_2, \dots$, and a set Msg of messages $\mu_1, \mu_2, \dots$. 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 $j \in A$, $\mu \in Msg$. 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 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 $s \in \mathcal{S}$ is a tuple $\langle l_e, l_1, \dots, l_n \rangle$, where $l_1, \dots, l_n$ are local states as above, and l_e contains all the events in $l_1, \dots, l_n$. In what follows we assume that the natural numbers $\mathbb{N}$ as the flow of time. This choice implies that we cannot provide a complete characterisation of MPS in this formalism, as first-order temporal logic on $\mathbb{N}$ is unaxiomatisable (Gabbay, Hodkinson, and Reynolds 1993). Still, we can express a number of interesting properties of MPS in the language $\mathcal{L}_n$.

A run r over $\langle \mathcal{S}, \mathbb{N} \rangle$ is a function from the natural numbers $\mathbb{N}$ to $\mathcal{S}$ such that:

- MP1 $r_i(m)$ is a history over Σ_i , Int_i and Msg ;
- MP2 for every event $rec(i, j, \mu)$ in $r_i(m)$ there exists a corresponding event $send(j, i, \mu)$ in $r_j(m)$.
- MP3 $r_i(0)$ is a sequence of length one (the initial state $init(i, \alpha)$), and $r_i(m+1)$ is either identical to $r_i(m)$ or results from appending an event to $r_i(m)$.

The last specification MP4 has only a simplifying purpose and does not restrict our analysis:

- MP4 All events in a given agent's history are distinct. An agent can never perform the same action twice in a given run.

By MP1 the local states of each agent records her initial state, the messages she has sent or received, as well as the internal actions she has taken. MP2 guarantees that any received message was actually sent, while MP3 specifies that at each step at most a single event occurs to any agent. Finally, MP4 is not essential, but it simplifies proofs as we do not have to distinguish different occurrences of the same action by, for example, time-stamping actions. We will use this constraint throughout the present section without explicitly mentioning it.

We now define message passing QIS (MPQIS) as a particular class of quantified interpreted systems $\mathcal{P} = \langle R, D, I \rangle$, where R is a non-empty set of runs satisfying the constraints MP1-4 above, D contains the agents in A , the actions in Act , the messages in Msg , and the events $e_1, e_2, \dots$, and I is an interpretation for $\mathcal{L}_n$. 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, $e_1, e_2, \dots$ are metaterms ranging over events; for instance, $\forall e \phi[e]$ is a shorthand for

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

where $\phi[t]$ means that the term t occurs in the formula ϕ .

We use the same notation for the objects in the model and the syntactic elements, the distinction will be clear by the context.

For the specification of MPS it is useful to introduce a predicative constant H for *happens* such that $(\mathcal{P}^\sigma, r, m) \models H(e, i)$ iff the event e occurs to agent i at time m in run r , i.e., $r_i(m)$ is the result of appending e to $r_i(m-1)$. We write $H(e)$ as a shorthand for $\exists i H(e, i)$. By definition of the environment's local state, $(\mathcal{P}^\sigma, r, m) \models H(e)$ iff e occurs at time m in run r . Also, we introduce the predicate $H'ed(e, i)$ for *happened* as $\langle P \rangle^+ H(e, i)$, and $H'ed(e) := \exists i H'ed(e, i)$. Finally, $Sent(i, j, \mu)$, $Recd(i, j, \mu)$, $Init(i, \alpha)$, and $Int(i, \alpha)$ are shorthands for $H'ed(send(i, j, \mu))$, $H'ed(rec(i, j, \mu))$, $H'ed(init(i, \alpha))$, and $H'ed(int(i, \alpha))$ respectively.

Let us now explore the range of specifications that can be expressed in the formalism. A property often required is *channel reliability*. We express this by stating that every sent message is eventually received. According to the definition of message passing QIS, it is possible that a message is lost during a run of the system. We can force channel reliability by requiring the following specification on MPQIS:

$$\forall i, j, \mu (Sent(i, j, \mu) \rightarrow \langle F \rangle^+ Recd(j, i, \mu))$$

Another relevant property of MPQIS concerns *authentication*: if agent i has received a message μ from agent j , then i knows that μ had actually been sent by j . This specification can be expressed as:

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

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

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

It is easy to show that by definition MPQIS satisfy authentication and perfect recall but not channel reliability.

We anticipated that the formalism of QIS is powerful enough for expressing the specifications MP1-4 in $\mathcal{L}_n$. Moreover, we can reason about the knowledge agents have of the ordering of events in asynchronous MPS. To show this, we define $Prec(e, e', i)$ as a shorthand for:

$$H'ed(e', i) \wedge H'ed(e, i) \wedge [P]^+ (H'ed(e', i) \rightarrow H'ed(e, i))$$

It follows that $(\mathcal{P}^\sigma, r, m) \models Prec(e, e', i)$ iff events e and e' both occur to agent i by round m of run r , and e occurs no later than e' in r . Also, the ordering $Prec(e, e')$ is defined as:

$$H'ed(e') \wedge H'ed(e) \wedge [P]^+ (H'ed(e') \rightarrow H'ed(e))$$

Note that in the propositional language of (Fagin et al 1995) $Prec(e, e')$ is assumed as a primitive proposition.

We can express that the events in a state $r(m)$ are partially ordered by specifying that $Prec(e, e')$ is a reflexive and transitive relation on the set of past events:

$$\forall e (H'ed(e) \rightarrow Prec(e, e)) \quad (3)$$

$$\forall e, e', e'' (Prec(e, e') \wedge Prec(e', e'') \rightarrow Prec(e, e'')) \quad (4)$$

Moreover, $Prec(e, e', i)$ can be defined as an anti-symmetric, linear, discrete order on the events in $r_i(m)$, where with each non-final point is associated an immediate successor, that is, it is also anti-symmetric and total:

$$\forall e, e' (Prec(e, e', i) \wedge Prec(e', e, i) \rightarrow (e = e')) \quad (5)$$

$$\forall e, e' (H'ed(e, i) \wedge H'ed(e', i) \rightarrow Prec(e, e', i) \vee Prec(e', e, i)) \quad (6)$$

and each non-final point has an immediate successor:

$$\forall e, e' (Prec(e, e', i) \rightarrow \exists e'' (Prec(e, e'', i) \wedge \neg \exists e''' (Prec(e, e''', i) \wedge Prec(e''', e'', i)))) \quad (7)$$

We define $LinDisc(Prec(e, e', i))$ as the conjunction of (3)-(7) above, expressing that the relation $Prec(e, e', i)$ is a linear, discrete order where every non terminal event has a successor. Also, we define the first event as the minimal one with respect to $Prec(e, e', i)$, that is,

$$Fst(e, i) ::= \forall e' (H'ed(e', i) \rightarrow Prec(e, e', i))$$

the first event is provably unique as the order on histories is total. We formally define the specifications MP1-4 as follows:

- MP1' $LinDisc(Prec(e, e', i)) \wedge \wedge \exists e (Fst(e, i) \wedge \exists \alpha (e = init(i, \alpha))) \wedge \wedge \forall e (H'ed(e, i) \wedge \neg Fst(e, i) \rightarrow \exists j, \alpha, \mu (e = int(i, \alpha) \vee \vee e = send(i, j, \mu) \vee e = rec(i, j, \mu)))$
- MP2' $\forall i, j, \mu (Recd(i, j, \mu) \rightarrow Sent(j, i, \mu))$
- MP3' $\langle P \rangle^+ ([P] \perp \wedge \exists e (H'ed(e, i) \wedge \exists \alpha (e = init(\alpha, i))) \wedge \wedge \forall e' (H'ed(e', i) \rightarrow e' = e)) \wedge \wedge \forall e (H'ed(e, i) \rightarrow ((\langle P \rangle H'ed(e, i) \vee \vee (H(e, i) \wedge \forall e' (H(e', i) \rightarrow e' = e))))$
- MP4' $H(e, i) \rightarrow ([P] \neg H(e, i) \wedge [F] \neg H(e, i))$

By MP1' the events in the local of agent i are a linear, discrete order, whose first element is an initial event, and whose following events are either send or receive events or internal events. According to MP2' each local state trivially satisfies MP2. By MP3' there is a moment (the starting point) when the only event in an agent's local state is the initial event, and for every event already happened, either it happened at some point strictly in the past, or it is the single event which happened in the last round. Finally, by MP4' each event happens only once in a given run, thus satisfying MP4. MP1'-4' are the basic specifications for MPQIS. We underline that these specifications are defined by means of only the predicative constant H .

As we pointed out above, synchronicity is a costly assumption in terms of computational resources in MPS. This remark prompts us to consider asynchronous MPS, where agents have no common clock. To make this informal definition precise, we follow once more (Fagin et al 1995). First, we say that a set V of histories is *prefix closed* if whenever $h \in V$, every non-empty prefix of h is in V as well. Then, we consider the following constraint for AMPQIS:

- MP5 The set R of runs in an AMPQIS includes *all* runs satisfying MP1-4 such that the local states of agent i belong to V_i , for some prefix closed set V_i of histories.

This constraint implies that at round m of a run r , each agent i considers possible that any other agent j has performed only a proper subset $r'_j(m)$ of the actions listed in $r_j(m)$.

We can now prove the main result of this section: Proposition 4.4.3 in (Fagin et al 1995) can be restated as a validity on the class of AMPQIS. We do not provide the full statement here, but we note that this metatheoretical result can be restated as a formula in the first-order modal language $\mathcal{L}_n$. We introduce a relation of *potential causality* between events, as first discussed in (Lamport 1978). 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 point (r, m) iff both e and e' occur by round m in the run r , and

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

Note that $\mapsto_G$ is a partial order on events, it is also anti-symmetric by MP4. We can say that two events e, e' are *concurrent* iff $e \not\mapsto_G e'$ and $e' \not\mapsto_G e$. Intuitively, the relation $\mapsto_G$ holds between events e and e' iff it is possible for event e to causally affect event e' . Two events are concurrent if neither can affect the other. We say that $(\mathcal{P}^\sigma, r, m) \models e \mapsto_G e'$ if $e \mapsto_G e'$ holds at (r, m) .

Now we prove that the potential causality relation $\mapsto_G$ is the closest we can come in AMPS to an ordering of events, that is, even if the agents in G could combine all their knowledge of the order $Prec(e, e')$ on events, they could not deduce any more about this ordering than is implied by the relation $\mapsto_G$. This is due to the fact that the delivery of messages can be arbitrarily delayed in AMPS, and the agents might be unaware of this because of asynchronicity. We refer to the appendix for a detailed proof.

Lemma 16 *The following validity holds in the class of AMPQIS satisfying the specifications MP1-5 above:*

$$AMPQIS \models \forall e, e' ((e \mapsto_G e') \leftrightarrow D_G Prec(e, e'))$$

By virtue of the analysis above we remark that the quantified language we have introduced has the power to express complex specifications, which identify metaproperties about the semantical class under discussion. In particular, by using language $\mathcal{L}_n$ we are able to formalise various constraints on MPS such as reliability, authentication and perfect recall. The traditional propositional specifications MP1-4 for MPS can be given formal counterparts MP1'-4' in $\mathcal{L}_n$, which can be shown valid on the corresponding semantical classes thereby signaling the general correctness of the approach.

Conclusions and Future Work

In this paper we analysed a quantified variant of interpreted systems and showed completeness for the axiomatisation QKT.S5_n involving temporal and epistemic modalities on

the first-order language $\mathcal{L}_n$. Retaining completeness seems noteworthy given the known difficulties of these formalisms.

Further, we used this formalism to reason about message passing systems, a mainstream framework to reason about asynchronous systems. In particular, we compared the results obtained at first-order with what was already known at propositional level, and observed that some properties in the latter setting become formal validities in the former.

Still, further work seems to be needed in this line of research. First, it seems interesting to relax the assumption on the domain of quantification, and admit a different domain $D_i(s)$ for each agent i and for each global state s . In such a framework we should check how to modify the completeness proof for QKT.S5_n to accommodate varying domains.

Moreover, we aim at extending the temporal fragment of our language with the *next* $\bigcirc$ and *until* U operators. Completeness results are available for various *monodic* fragments of such a language (Wolter and Zakharyashev 2002), and for the fragment with $\bigcirc$ over the rational numbers (Meyden 1994). It is yet to be checked whether these results extend to first-order languages with epistemic operators as well. Also, we would like to analyse relevant classes of QIS, such as *synchronous* QIS and QIS with *perfect recall*. We have sound and complete axiomatisations for these structures at propositional level (Fagin et al 1995), but it is not clear whether these results extend to first-order.

Acknowledgements

The research described in this paper was partly supported by the EC Framework 6 funded project CONTRACT (IST Project Number 034418), and by the Scuola Normale Superiore, Pisa. The authors would like to thank the reviewers for useful comments on an earlier version of this paper.

References

- Belardinelli, F., and Lomuscio, A. 2008. A Complete Quantified Epistemic Logic for Reasoning about Message Passing Systems. in *Proceedings of the 8th International Workshop on Computational Logic in Multi-Agent Systems (CLIMA VIII)*, 258-273.
- Belardinelli, F., Lomuscio, A. 2007. Quantified Epistemic Logics with Flexible Terms. *LORI workshop on Logic, Rationality and Interaction*, Beijing, 5-9 August, 2007.
- Blackburn, P., van Benthem, J., and Wolter, F. 2007. *Handbook of Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Elsevier.
- Chagrov, A., and Zakharyashev, M. 1997. *Modal Logic*, volume 35 of *Oxford Logic Guides*. Oxford: Clarendon Press.
- Clarke, E. M., Grumberg, O., and Peled, D. A. 1999. *Model Checking*. Cambridge, Massachusetts, The MIT Press.
- Cohen, P., Levesque H. 1995. Communicative Actions for Artificial Agents. *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*.
- Dembinski, P. et al. 2003. Verics: A Tool for Verifying Timed Automata and Estelle Specifications. In *TACAS*, 278-283, Springer.
- Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge*. Cambridge: MIT Press.
- Fagin, R., Halpern, J. Y., and Vardi, M. Y. 1992. What can machines know? On the properties of knowledge in distributed systems. *Journal of the ACM* 39(2):328-376.

Fitting, M., Mendelsohn, R. 1999. *First-order Modal Logic*. Kluwer, Dordrecht.

Gabbay, D. M., Hodkinson, I. M., and Reynolds, M. A. 1993. *Temporal Logic, Mathematical Foundations and Computational Aspects, Volume 1*. Oxford University Press.

Gammie, P., and van der Meyden, R. 2004. MCK: Model checking the logic of knowledge. In *Proceedings of CAV'04*, 479–483. Springer.

Halpern, J., Fagin, R. 1989. Modelling knowledge and action in distributed systems. *Distributed Computing*, 3(4):159–179.

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

Hughes, G. E., and Cresswell, M. J. 1966. *A New Introduction to Modal Logic*. New York: Routledge.

Lamport, L. 1978. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565.

Cohen, M., and Dams, M. 2007. A complete axiomatisation of knowledge and cryptography. In *LICS*.

Meyden, R., and Wong, K. 2003. Complete axiomatizations for reasoning about knowledge and branching time. *Studia Logica*, 75(1):93–123.

Meyden, R. 1994. Axioms for knowledge and time in distributed systems with perfect recall. In *LICS*, 448–457.

Meyer, J.-J. C., and Hoek, W. 1992. Making some issues of implicit knowledge explicit. in *Int. J. of Foundations of Computer Science*, 3 (2), pp. 193–223.

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

Lomuscio, A., and Colombetti, M. 1996. QLB: a quantified logic for belief. In *Proceedings of ATAL-96*, Springer-Verlag.

Parikh, R., and Ramanujam, R. 1985. Distributed processes and the logic of knowledge. In *Logic of Programs*, 256–268. Springer.

Raimondi, F., and Lomuscio, A. 2007. Automatic verification of multi-agent systems by model checking via OBDDs. *Journal of Applied Logic*.

Rao, A., and Georgeff, M. 1991. Modeling Rational Agents within a BDI-Architectures. *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*.

Reynolds, M. 1996. Axiomatizing first-order temporal logic: until and since over linear time. *Studia Logica*, 57(2/3):279–302.

Solanki, M., Cau, A., and Zedan, H. 2006. ASDL: a wide spectrum language for designing web services. In *WWW*, 687–696. ACM.

Sturm, H., Wolter, F., and Zakharyashev, M. 2000. Monodic epistemic predicate logic. In *JELIA*, 329–344. Springer.

Sturm, H., Wolter, F., and Zakharyashev, M. 2002. Common knowledge and quantification. *Economic Theory*, 19:157–186.

Viganó, F. 2007. A framework for model checking institutions. In *MoChart IV*, volume 4428 of *LNCS*. Springer.

Wolter, F. 2000. First Order Common Knowledge Logics. *Studia Logica*, 65, 249–271.

Wolter, F., and Zakharyashev, M. 2001. Decidable fragments of first-order modal logics. *J. Symb. Log.* 66(3):1415–1438.

Wolter, F., and Zakharyashev, M. 2002. Axiomatizing the monodic fragment of first-order temporal logic. *Ann. Pure Appl. Logic* 118(1-2):133–145.

Wooldridge, M. 2000. *Reasoning about Rational Agents*. MIT Press.

Appendix

Lemma 7 For every $\phi \in \mathcal{L}_n$, $w \in W$,

$$(\mathcal{M}^\sigma, w) \models \phi \quad \text{iff} \quad (g(\mathcal{M})^\sigma, r, w) \models \phi$$

Proof. The proof of this lemma is by induction on the length of the formula ϕ . The base of induction for $\phi = P^k(\vec{t})$ or $\phi = (t = t')$ follows by definition of the interpretation I' in $g(\mathcal{M})$. The inductive cases for the propositional connectives are straightforward.

For $\phi = K_i\psi$, $(\mathcal{M}^\sigma, w) \models \phi$ iff for all $w' \sim_i w$, $(\mathcal{M}^\sigma, w') \models \psi$, iff for $r_i(w') = r_i(w)$, $(g(\mathcal{M})^\sigma, r, w) \models \phi$, by definition of r and induction hypothesis, iff $(g(\mathcal{M})^\sigma, r, w) \models \phi$.

The inductive cases for the other modal operators can be shown similarly.

Lemma 11 For every $\psi \in \mathcal{L}_n^+$,

$$\mathcal{M}' \models \psi \quad \text{iff} \quad \mathcal{M} \models^p \psi$$

Proof. We first show that if the canonical model $\mathcal{M}$ pseudo-validates $\psi \in \mathcal{L}_n$, then there is a tree-like structure $\mathcal{M}^*$ which pseudo-validates ϕ as well. Then, from $\mathcal{M}^*$ we can obtain a K -model $\mathcal{M}'$ satisfying lemma 11.

In order to define $\mathcal{M}^*$ we need few more definitions. Let w, w' be worlds in W , a *path from w to w'* is a sequence $\langle w_1, l_1, w_2, l_2, \dots, l_{k-1}, w_k \rangle$ such that (1) $w = w_1$ and $w' = w_k$; (2) $w_1, \dots, w_k \in W$; (3) each l_j is either an agent or a set of agents; (4) $\langle w_j, w_{j+1} \rangle \in R_{l_j}$.

The *reduction* of a path $\langle w_1, l_1, w_2, l_2, \dots, l_{k-1}, w_k \rangle$ is obtained by replacing each maximal consecutive subsequence $\langle w_q, i_q, w_{q+1}, i_{q+1}, \dots, i_{r-1}, w_r \rangle$ where $i_q = i_{q+1} = \dots = i_{r-1}$ by $\langle w_q, i_q, w_r \rangle$. A path is said to be *reduced* if it is equal to its reduction.

Given the canonical model $\mathcal{M} = \langle W, R, <, D, I \rangle$, we define a structure $\mathcal{M}^* = \langle W^*, R^*, <^*, D, I^* \rangle$ and a surjective function $h : W^* \rightarrow W$ such that (i) $\mathcal{M}^*$ is a tree, that is, for $w, w' \in W^*$ there is at most one reduced path from w to w' ; (ii) $w R_i^* w'$ implies $h(w) R_i h(w')$; (iii) $w R_G^* w'$ implies $h(w) R_G h(w')$; (iv) $w <^* w'$ implies $h(w) < h(w')$; (v) $\langle a_1, \dots, a_k \rangle \in I^*(P^k, w)$ iff $\langle a_1, \dots, a_k \rangle \in I(P^k, h(w))$.

We define W^* by induction. Let W_1^* be W , and define W_{k+1}^* as the set of worlds $v_{w,l,w'}$ such that $w \in W_k^*$, $w' \in W$ and l is an agent or group of agents. Let $W^* = \bigcup_{k \in \mathbb{N}} W_k^*$, then define $h : W^* \rightarrow W$ by letting $h(w) = w$, for $w \in W_1^*$ and $h(v_{w,l,w'}) = w'$, for $w \in W_k^*$. Further, R_l^* is the reflexive, symmetric and transitive closure of the relation defined for $w, w' \in W^*$ if $w' = v_{w,l,w''}$, for some $w'' \in W$, and $h(w) R_l h(w')$; while $<^*$ is the relation defined for $w, w' \in W^*$ if $h(w) < h(w')$. Finally, $I^*(P^k, w) = I(P^k, h(w))$. By results in (Fagin, Halpern, and Vardi 1992) $\mathcal{M}^*$ and h satisfy (i)-(v) above. In particular, we can show the following:

Proposition 17 For $w \in W^*$, $\psi \in \mathcal{L}_n^+$,

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

Finally, we make use of the structure $\mathcal{M}^*$ to define a K -model $\mathcal{M}'$ such that lemma 11 holds. 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, transitive and symmetric, R'_i is an equivalence relation. We state the following result about $\mathcal{M}'$ and refer to (Fagin, Halpern, and Vardi 1992) for further details.

Proposition 18 For $w \in W'$, $\psi \in \mathcal{L}_n^+$,

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

In conclusion, The canonical model $\mathcal{M}$ pseudo-validates $\psi \in \mathcal{L}_n$ if and only if $\mathcal{M}^*$ pseudo-validates ψ by proposition 17, iff by proposition 18 the K -model $\mathcal{M}'$ validates ψ .

Lemma 12 For every $\psi \in \mathcal{L}_n^+$,

$$\mathcal{M}^+ \models \psi \quad \text{iff} \quad \mathcal{M}' \models \psi$$

Proof. Let $W^{ir} = \{w \in W' \mid w \not<' w\}$ be the set of irreflexive worlds in $\mathcal{M}'$ and define the equivalence relation $\approx$ on $W^r = \{w \in W' \mid w <' w\}$ as $w_1 \approx w_2$ iff $w_1 <' w_2$ and $w_2 <' w_1$. For every $\approx$ -equivalence class a , define a map $a(\cdot)$ from the reals $\mathbb{R}$ onto a such that for every $w \in a$, $p \in \mathbb{R}$ there are $s, t \in \mathbb{R}$ and

- $s < p < t$;
- $a(s) = w = a(t)$.

This can be done as every $\approx$ -equivalence class contains at most 2^{N_0} saturated sets of formulas.

Further, for $w \in W^{ir}$ we set $\{w\}(0) = w$. Now we define the K -model $\mathcal{M}^+$, where $W^+ = \{(\{w\}, 0) \mid w \in W^{ir}\} \cup \{(a, p) \mid a \text{ is a } \approx\text{-equivalence class, } p \in \mathbb{R}\}$ is the set of possible worlds. The order $<^+$ on W^+ is such that $(a, p) <^+ (b, s)$ iff

- $a \neq b$ and there are $w_a \in a, w_b \in b$ and $w_a <' w_b$; or
- $a = b$ and $p < s$.

The relation $<^+$ is a weakly connected, strict partial order on W^+ , in particular $<^+$ is irreflexive. Also, the relation R_i^+ on W^+ such that $(a, p) R_i^+ (b, s)$ iff $a(p) R'_i b(s)$ is an equivalence relation as R'_i is such. Finally, the domain D^+ is equal to D' , and I^+ is such that $\langle u_1, \dots, u_k \rangle \in I^+(P^k, (a, p))$ iff $\langle u_1, \dots, u_k \rangle \in I'(P^k, a(p))$.

It is straightforward to check that $(\mathcal{M}^{+\sigma}, (a, p)) \models \psi$ iff $(\mathcal{M}'^\sigma, a(p)) \models \psi$, so the lemma follows.

Lemma 16 The following validity holds in the class of $AMPQIS$ satisfying the specifications $MP1$ -5 above:

$$AMPQIS \models \forall e, e' ((e \mapsto_G e') \leftrightarrow D_G \text{Prec}(e, e'))$$

Proof. $\Rightarrow$ Assume $(\mathcal{P}^\sigma, r, m) \models e \mapsto_G e'$. If e' is a receive event and e is the corresponding send event, then $r_i(m) = r'_i(m')$ for all $i \in G$ implies $(\mathcal{P}^\sigma, r', m') \models H'ed(e) \wedge H'ed(e') \wedge [P]^+(H'ed(e') \rightarrow H'ed(e))$. In fact, for all $m'' \leq m'$, $(\mathcal{P}^\sigma, r', m'') \models \text{Recd}(i, j, \mu) \rightarrow \text{Sent}(j, i, \mu)$ by $MP2'$. Thus, $(\mathcal{P}^\sigma, r, m) \models D_G \text{Prec}(e, e')$.

If e, e' are both in $r_i(m)$ and either $e = e'$ or e comes earlier than e' in $r_i(m)$, then $r'_i(m') = r_i(m)$ implies $(\mathcal{P}^\sigma, r', m') \models H'ed(e) \wedge H'ed(e') \wedge [P]^+(H'ed(e') \rightarrow H'ed(e))$, then $(\mathcal{P}^\sigma, r, m) \models K_i \text{Prec}(e, e')$. By $D1$ and $D2$, $(\mathcal{P}^\sigma, r, m) \models D_G \text{Prec}(e, e')$.

If there exists some e'' such that $e \mapsto_G e''$ and $e'' \mapsto_G e'$, then without loss of generality we assume that $e \mapsto_G e''$ and $e'' \mapsto_G e'$ for either case 1 or 2 above, in both cases $(\mathcal{P}^\sigma, r, m) \models D_G \text{Prec}(e, e'') \wedge D_G \text{Prec}(e'', e')$. This means that $r_i(m) = r'_i(m')$ for all $i \in G$ implies $(\mathcal{P}^\sigma, r', m') \models [P]^+(H'ed(e'') \rightarrow H'ed(e)) \wedge [P]^+(H'ed(e') \rightarrow H'ed(e''))$. By distributivity and transitivity, $(\mathcal{P}^\sigma, r', m') \models [P]^+(H'ed(e') \rightarrow H'ed(e))$. Thus, $(\mathcal{P}^\sigma, r, m) \models D_G \text{Prec}(e, e')$.

$\Leftarrow$ Assume that $(\mathcal{P}^\sigma, r, m) \models H(e) \wedge H(e')$ but $(\mathcal{P}^\sigma, r, m) \not\models e \mapsto_G e'$. The events e, e' must be distinct. Moreover, if they both appear in $r_i(m)$, for some i , by hypothesis there must be some $m' < m$ such that $(\mathcal{P}^\sigma, r, m') \models H(e') \wedge \neg H(e)$. Thus, $(\mathcal{P}^\sigma, r, m) \not\models D_G \text{Prec}(e, e')$.

If e and e' appear in the local states of distinct agents i, j , then consider the minimal $m_{e'}$ such that $e' \in r_i(m_{e'})$. If $e \notin r_j(m_{e'})$ we are done. Otherwise, consider the minimal m_e such that $e \in r_j(m_e)$. We define a run r' such that $e \mapsto_G e''$ implies that e'' occurs $(m_{e'} - m_e) + 1$ round later in r' than in r . Specifically, for each agent k , if there is no e'' such that $e \mapsto_G e''$ then $r'_k(m) = r_k(m)$ for every m . Otherwise, let $m_{e''}$ be the minimal round such that $e'' \in r_k(m_{e''})$, then define r' as follows:

$$r'_k(m) = \begin{cases} r_k(m) & \text{for } m < m_{e''} \\ r_k(m_{e''} - 1) & \text{for } m_{e''} \leq m \leq m_{e''} + (m_{e'} - m_e) \\ r_k(m - (m_{e'} - m_e)) & \text{for } m_{e''} + (m_{e'} - m_e) < m \end{cases}$$

We can show that r' is well defined, and $r' \in \mathcal{P}$ by $MP5$. Finally, for all $i \in G$, $r_i(m) = r'_i(m + (m_{e'} - m_e) + 1)$ and $(\mathcal{P}^\sigma, r', m_{e'}) \models H(e') \wedge \neg H(e)$. Thus, $(\mathcal{P}^\sigma, r, m) \not\models D_G \text{Prec}(e, e')$.



2008

This paper was published in the *Proceedings, Eleventh International Conference on Principles of Knowledge Representation and Reasoning*, edited by Gerhard Brewka and Jérôme Lang (Menlo Park, California: AAAI Press).

The 2008 KR conference was held in Sydney, Australia, 16–19 September, 2008.