Interactions between Time and Knowledge in a First-order Logic for Multi-Agent Systems

F. Belardinelli and **A. Lomuscio**

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

Abstract

We investigate a class of first-order temporal epistemic logics for the specification of multi-agent systems. We consider well-known properties of multi-agent systems including perfect recall, synchronicity, no learning, unique initial state, and define natural correspondences between these and quantified interpreted systems. Our findings identify several monodic fragments of first-order temporal epistemic logic that we prove to be both sound and complete with respect to the corresponding classes of quantified interpreted systems. The results show that interaction axioms for propositional temporal epistemic logic can be lifted to the monodic fragment.

Introduction

First-order modal logics for reasoning about knowledge and time have attracted increasing interest from logicians and researchers in AI, both as regards their theoretical properties (completeness, decidability, complexity) (Gabbay et al. 2003; Hodkinson et al. 2003; Sturm, Wolter, and Zakharyaschev 2000), and their applications to multi-agent systems (Cohen and Levesque 1995; Rao and Georgeff 1991; Wooldridge 2000).

In this paper we introduce several classes of *quantified interpreted systems* (Belardinelli and Lomuscio 2008; 2009a; 2009b) that are suitable for modeling the interaction between temporal and epistemic modalities at the first order. Specifically, we analyse systems with perfect recall, no learning, synchronicity and a unique initial state (Fagin et al. 1995). For all these we present sound and complete axiomatisations of the set of *monodic* validities, where at most one free variable appears in the scope of any modal operator (Hodkinson, Wolter, and Zakharyaschev 2000).

Our starting point for this contribution consists of results on the axiomatisability (Sturm, Wolter, and Zakharyaschev 2000; Wolter and Zakharyaschev 2002), decidability (Hodkinson, Wolter, and Zakharyaschev 2000; Wolter and Zakharyaschev 2001), and complexity (Hodkinson 2006; Hodkinson et al. 2003) of first-order modal logics, together with completeness results for propositional temporal epistemic logics (Halpern, Meyden, and Vardi 2003; Halpern

and Moses 1992). Specifically, we prove the completeness of our first-order temporal epistemic logics via a *quasimodel* construction, which has previously been used in (Hodkinson, Wolter, and Zakharyaschev 2000; 2002) to prove decidability for *monodic* fragments of first-order temporal logic (FoTL) with respect to linear and branching flows of time (Hodkinson, Wolter, and Zakharyaschev 2000; 2002). Quasimodels have also been applied to first-order temporal and epistemic logic in (Sturm, Wolter, and Zakharyaschev 2000; Wolter and Zakharyaschev 2002). In (Wolter and Zakharyaschev 2002) the authors present a complete axiomatisation for the monodic fragment of FoTL on the naturals. In (Sturm, Wolter, and Zakharyaschev 2000) we have a similar result for a variety of first-order epistemic logics with common knowledge. However, the interaction between temporal and epistemic modalities at the first order has not yet been taken into account, nor has the interpreted systems semantics (Fagin et al. 1995; Parikh and Ramanujam 1985): both of these are relevant for applications to multi-agent systems (MAS).

In this paper we also make use of the results in (Halpern, Meyden, and Vardi 2003; Halpern and Moses 1992) on the completeness of propositional temporal epistemic logics. In particular, in (Halpern, Meyden, and Vardi 2003) the authors provide a framework for proving completeness on semantics similar to those here considered. We combine their approach with the quasimodel technique to prove completeness for monodic fragments of first-order temporal epistemic logic.

This contribution is motivated by an interest in first-order temporal epistemic formalisms to model high-level properties of multi-agent systems. Recent papers witness an increasing need in web-services, security, communication protocols, as well as other areas, to extend the expressive power of temporal epistemic languages to the first order (see for instance (Deutsch, Sui, and Vianu 2004; Hallé and Villemaire 2009)). As a preliminary contribution to this project in (Belardinelli and Lomuscio 2008) we introduced quantified interpreted systems (QIS) to model a first-order temporal epistemic formalism. These investigations were further pursued in (Belardinelli and Lomuscio 2009a), which explicitly assumes linear-time operators and the natural numbers as the flow of time. Neither contribution considers the interaction between time and knowledge, which is addressed here.

Quantified Interpreted Systems

We extend interpreted systems to the first order by endowing each structure with a domain of individuals. Preliminary investigations in "static" quantified interpreted systems, where no account of evolution for the system is given, have appeared in (Belardinelli and Lomuscio 2009b). Fully-fledged QIS on a language with temporal modalities have been introduced in (Belardinelli and Lomuscio 2008; 2009a). We follow the definition of QIS there provided, but differently from these contributions, we also consider the interaction between temporal and epistemic modalities.

Given a set $A = \{1, \ldots, m\}$ of agents, the first-order temporal epistemic language $\mathcal{L}_m$ contains individual variables $x_1, x_2, \ldots$, individual constants $c_1, c_2, \ldots$, n-ary predicative letters $P_1^n, P_2^n, \ldots$, for $n \in \mathbb{N}$, the connectives $\neg$ and $\rightarrow$, the quantifier $\forall$, the temporal operators $\bigcirc$ and $\mathcal{U}$, and the epistemic operator K_i for each agent $i \in A$. The only terms $t_1, t_2, \ldots$ in $\mathcal{L}_m$ are individual variables and constants.

Definition 1 *Formulas in $\mathcal{L}_m$ are defined as follows:*

$$\phi ::= P^k(t_1, \ldots, t_k) \mid \neg\psi \mid \psi \rightarrow \psi' \mid \forall x\psi \mid \bigcirc\psi \mid \psi\mathcal{U}\psi' \mid K_i\psi$$

The formulas $\bigcirc\phi$ and $\phi\mathcal{U}\phi'$ are read as "*at the next step ϕ*" and "*eventually ϕ' and until then ϕ*"; $K_i\phi$ represents "*agent i knows ϕ*". We define the symbols $\wedge$, $\vee$, $\leftrightarrow$, $\exists$, G ("always in the future"), F ("some time in the future") as standard; $\bar{K}_i\phi$ is short for $\neg K_i\neg\phi$. By $\phi[\vec{y}]$ we mean that $\vec{y} = y_1, \ldots, y_n$ are all the free variables in ϕ; while $\phi[\vec{y}/\vec{t}]$ is the formula obtained by substituting simultaneously some, possibly all, free occurrences of $\vec{y}$ in ϕ with $\vec{t} = t_1, \ldots, t_n$ and renaming bounded variables.

To introduce quantified interpreted systems we assume a set L_i of local states $l_i, l_i', \ldots$, a set Act_i of actions $a_i, a_i', \ldots$, and a protocol $P_i : L_i \rightarrow 2^{Act_i}$ from local states to non-empty sets of actions for each agent $i \in A$ in a multi-agent system. We consider local states, actions, and a protocol for the environment e as well. The set $S \subseteq L_e \times L_1 \times \ldots \times L_m$ contains the global states of the MAS, $Act \subseteq Act_e \times Act_1 \times \ldots \times Act_m$ is the set of joint actions, while $P = (P_e, P_1, \ldots, P_m)$ is the joint protocol. We also introduce a transition function $\tau : Act \rightarrow (S \rightarrow S)$ such that $\tau(a)(s) = s'$ only if $a \in P(s)$. Intuitively, $\tau(a)(s) = s'$ encodes that the system moves from state s to state s' if agents perform the joint action a. We say that the global state s' is *reachable in one step* from s, or $s \sqsubset s'$, if there is $a \in Act$ such that $\tau(a)(s) = s'$. To represent the temporal evolution of the MAS we consider the flow of time $\mathbb{N}$ of the naturals numbers. A *run* is any function $r : \mathbb{N} \rightarrow S$ such that $r(n) \sqsubset r(n+1)$. Intuitively, a run represents a possible evolution of the MAS according to the transition function τ. Finally, we define the quantified interpreted systems for the language $\mathcal{L}_m$ as follows:

Definition 2 (QIS) *A quantified interpreted system is a triple $\mathcal{P} = \langle \mathcal{R}, \mathcal{D}, I \rangle$ such that (i) $\mathcal{R}$ is a non-empty set of runs; (ii) $\mathcal{D}$ is a non-empty set of individuals; (iii) I is an interpretation of $\mathcal{L}_m$ such that $I(c) \in \mathcal{D}$, and for $r \in \mathcal{R}$, $n \in \mathbb{N}$, $I(P^k, r, n)$ is a k-ary relation on $\mathcal{D}$.*

Following standard notation (Fagin et al. 1995) a pair (r, n) is a *point* in $\mathcal{P}$. If $r(n) = \langle l_e, l_1, \ldots, l_m \rangle$ is the global state at point (r, n) then $r_e(n) = l_e$ and $r_i(n) = l_i$ are the environment's and agent i's local state at (r, n) respectively. Further, for $i \in A$ the equivalence relation $\sim_i$ is defined such that $(r, n) \sim_i (r', n')$ if $r_i(n) = r_i'(n')$.

In this paper we consider the following classes of QIS.

Definition 3 • *A QIS $\mathcal{P}$ is* synchronous *if for every agent $i \in A$, $(r, n) \sim_i (r', n')$ implies $n = n'$.*

• *A QIS satisfies* perfect recall *if for all points $(r, n) \sim_i (r', n')$, if $n > 0$ then either $(r, n-1) \sim_i (r', n')$ or there is $l < n'$ such that $(r, n-1) \sim_i (r', l)$ and for all $l < k \leq n'$ we have $(r, n) \sim_i (r', k)$.*

• *A QIS satisfies* no learning *if for all points $(r, n) \sim_i (r', n')$ either $(r, n+1) \sim_i (r', n')$ or there is $l > n'$ such that $(r, n+1) \sim_i (r', l)$ and for all $l > k \geq n'$ we have $(r, n) \sim_i (r', k)$.*

• *A QIS has a* unique initial state *if for all $r, r' \in \mathcal{R}$, $r(0) = r'(0)$.*

These conditions have extensively been discussed in the literature (Halpern, Meyden, and Vardi 2003) together with equivalent formulations. Intuitively, a QIS is synchronous if time is part of the local state of each agent. A QIS satisfies perfect recall if an agent's local state registers everything that has happened to her. No learning is dual to perfect recall. Finally, a QIS has a unique initial state if all runs start from the same global state.

By $\mathcal{QIS}_m$ we denote the class of QIS with m agents; the superscripts $sync, pr, nl, uis$ denote specific subclasses of $\mathcal{QIS}_m$ satisfying the respective constraints. For instance, $\mathcal{QIS}_m^{sync, uis}$ is the class of synchronous QIS with m agents and a unique initial state.

We now assign a meaning to the formulas of $\mathcal{L}_m$ in quantified interpreted systems. Let σ be an assignment from the variables to the individuals in $\mathcal{D}$, the valuation $I^\sigma(t)$ of a term t is defined as $\sigma(y)$ for $t = y$, and $I^\sigma(t) = I(c)$ for $t = c$. A variant $\sigma\binom{x}{a}$ of an assignment σ assigns $a \in \mathcal{D}$ to x and coincides with σ on all the other variables.

Definition 4 *The satisfaction relation $\models$ for $\phi \in \mathcal{L}_m$, $(r, n) \in \mathcal{P}$, and an assignment σ is defined as follows:*

$(\mathcal{P}^\sigma, r, n) \models P^k(\vec{t})$ if $\langle I^\sigma(t_1), \ldots, I^\sigma(t_k) \rangle \in I(P^k, r, n)$

$(\mathcal{P}^\sigma, r, n) \models \neg\psi$ if $(\mathcal{P}^\sigma, r, n) \not\models \psi$

$(\mathcal{P}^\sigma, r, n) \models \psi \rightarrow \psi'$ if $(\mathcal{P}^\sigma, r, n) \not\models \psi$ or $(\mathcal{P}^\sigma, r, n) \models \psi'$

$(\mathcal{P}^\sigma, r, n) \models \forall x\psi$ if for all $a \in \mathcal{D}$, $(\mathcal{P}^{\sigma\binom{x}{a}}, r, n) \models \psi$

$(\mathcal{P}^\sigma, r, n) \models \bigcirc\psi$ if $(\mathcal{P}^\sigma, r, n+1) \models \psi$

$(\mathcal{P}^\sigma, r, n) \models \psi\mathcal{U}\psi'$ if there is $n' \geq n$ such that $(\mathcal{P}^\sigma, r, n') \models \psi'$ and $n \leq n'' < n'$ implies $(\mathcal{P}^\sigma, r, n'') \models \psi$

$(\mathcal{P}^\sigma, r, n) \models K_i\psi$ if $(r, n) \sim_i (r', n')$ implies $(\mathcal{P}^\sigma, r', n') \models \psi$

The truth conditions for $\wedge$, $\vee$, $\leftrightarrow$, $\exists$, G and F are defined from those above. A formula $\phi \in \mathcal{L}_m$ is *true at a point* (r, n) if it is satisfied at (r, n) by every σ; ϕ is *valid on a QIS $\mathcal{P}$* if it is true at every point in $\mathcal{P}$; ϕ is *valid on a class $\mathcal{C}$ of QIS* if it is valid on every QIS in $\mathcal{C}$.

By considering all subsets of $\{sync, pr, nl, uis\}$ we obtain 16 subclasses of $\mathcal{QIS}_m$ for any $m \in \mathbb{N}$. Not all of them are independent nor axiomatisable. Some of these are not axiomatisable already at the propositional level (Halpern and Moses 1992; Halpern and Vardi 1989). In Table 1 we group together the classes of QIS that share the same set of

validities on $\mathcal{L}_m$ for $m > 1$. The proofs of these equivalences can be obtained similarly to the propositional case.

$$\boxed{\begin{array}{l} \mathcal{QIS}_m, \mathcal{QIS}_m^{sync}, \mathcal{QIS}_m^{uis}, \mathcal{QIS}_m^{sync,uis} \\ \mathcal{QIS}_m^{sync,pr}, \mathcal{QIS}_m^{sync,pr,uis} \\ \mathcal{QIS}_m^{pr}, \mathcal{QIS}_m^{pr,uis} \\ \mathcal{QIS}_m^{nl} \\ \mathcal{QIS}_m^{sync,nl} \\ \mathcal{QIS}_m^{nl,pr} \\ \mathcal{QIS}_m^{nl,pr,uis} \\ \mathcal{QIS}_m^{nl,uis} \\ \mathcal{QIS}_m^{sync,nl,pr} \\ \mathcal{QIS}_m^{sync,nl,uis}, \mathcal{QIS}_m^{sync,nl,pr,uis} \end{array}}$$

Table 1: Equivalences among classes of QIS.

By following (Halpern, Meyden, and Vardi 2003) we remark that the sets of propositional validities on all classes above are axiomatisable but $\mathcal{QIS}_m^{nl,pr,uis}$ and $\mathcal{QIS}_m^{nl,uis}$. Also, notice that $\mathcal{QIS}_1^{nl,uis}$ is equivalent to $\mathcal{QIS}_1^{nl}$ and $\mathcal{QIS}_1^{nl,pr,uis}$ is equivalent to $\mathcal{QIS}_1^{nl,pr}$. Thus, for $m = 1$ the sets of propositional validities on $\mathcal{QIS}_m^{nl,pr,uis}$ and $\mathcal{QIS}_m^{nl,uis}$ are nonetheless axiomatisable.

In the next section we show that the known axiomatisability result at the propositional level can be extended to the monodic fragment of the language $\mathcal{L}_m$ defined as follows:

Definition 5 *The monodic fragment $\mathcal{L}_m^1$ is the set of formulas $\phi \in \mathcal{L}_m$ such that any subformula of ϕ of the form $K_i\psi$, $\bigcirc\psi$ or $\psi_1\mathcal{U}\psi_2$ contains at most one free variable.*

In other words the monodic fragment of $\mathcal{L}_m$ contains formulas such as

$$\forall y \left(Resource(y) \rightarrow K_i \left(\forall z Available(y,z) \mathcal{U} \exists x Request(x,y) \right) \right)$$

The monodic fragments of a number of first-order modal logics have been thoroughly investigated (Hodkinson, Wolter, and Zakharyaschev 2000; Hodkinson et al. 2003; Wolter and Zakharyaschev 2001; 2002). In the case of $\mathcal{L}_m$ this fragment is quite expressive as it contains all *de dicto* formulas, i.e., formulas where no free variable appears in the scope of any modal operator.

Axiomatisations

In this section we present sound and complete axiomatisations of the sets of monodic validities for the classes of quantified interpreted systems defined in the previous section. We begin by introducing the basic system QKT_m that extends to the first order the epistemic logic S5 combined with the linear temporal logic LTL.

Definition 6 *The system QKT_m contains the following schemes of axioms and rules, where ϕ, ψ and χ are formulas in $\mathcal{L}_m^1$.*

Taut	*classic propositional tautologies*
MP	$\phi \rightarrow \psi, \phi \Rightarrow \psi$
K	$\bigcirc(\phi \rightarrow \psi) \rightarrow (\bigcirc\phi \rightarrow \bigcirc\psi)$
T1	$\bigcirc\neg\phi \leftrightarrow \neg\bigcirc\phi$
T2	$\phi\mathcal{U}\psi \leftrightarrow \psi \vee (\phi \wedge \bigcirc(\phi\mathcal{U}\psi))$
Nec	$\phi \Rightarrow \bigcirc\phi$
T3	$\chi \rightarrow \neg\psi \wedge \bigcirc\chi \Rightarrow \chi \rightarrow \neg(\phi\mathcal{U}\psi)$
K	$K_i(\phi \rightarrow \psi) \rightarrow (K_i\phi \rightarrow K_i\psi)$
T	$K_i\phi \rightarrow \phi$
4	$K_i\phi \rightarrow K_iK_i\phi$
5	$\neg K_i\phi \rightarrow K_i\neg K_i\phi$
Nec	$\phi \Rightarrow K_i\phi$
BF	$\bigcirc\forall x\phi \leftrightarrow \forall x\bigcirc\phi$
BF	$K_i\forall x\phi \leftrightarrow \forall x K_i\phi$
Ex	$\forall x\phi \rightarrow \phi[x/t]$
Gen	$\phi \rightarrow \psi[x/t] \Rightarrow \phi \rightarrow \forall x\psi$, *for x not free in ϕ*

The epistemic operator K_i is an $S5$ modality, while the next $\bigcirc$ and until $\mathcal{U}$ operators are axiomatised as linear-time modalities. To this we add the classic postulates Ex and Gen for quantification. Note that both are sound as we are considering a unique domain $\mathcal{D}$ of individuals in our structures. We consider the standard definitions of *proof* and *theorem*: $\vdash \phi$ means that $\phi \in \mathcal{L}_m^1$ is a theorem in QKT_m.

In this paper we focus on the schemes of axioms in Table 2 that specify the interaction between time and knowledge.

KT1	$K_i\bigcirc\phi \rightarrow \bigcirc K_i\phi$
KT2	$K_i\phi \wedge \bigcirc(K_i\psi \wedge \neg K_i\chi) \rightarrow \bar{K}_i((K_i\phi)\mathcal{U}((K_i\psi)\mathcal{U}\neg\chi))$
KT3	$(K_i\phi)\mathcal{U}K_i\psi \rightarrow K_i((K_i\phi)\mathcal{U}K_i\psi)$
KT4	$\bigcirc K_i\phi \rightarrow K_i\bigcirc\phi$
KT5	$K_i\phi \leftrightarrow K_j\phi$

Table 2: the axioms KT1-KT5.

We use $1, \ldots, 5$ as superscripts to denote the systems obtained by adding to QKT_m any combination of KT1-5. For instance, the system $\text{QKT}_m^{2,3}$ extends QKT_m with the axioms KT2 and KT3.

It is easy to check that the axioms of QKT_m are valid on every QIS and the inference rules preserve validity. On the other hand, the axioms KT1-5 are valid only on specific classes of QIS as stated in the following theorem.

Theorem 7 *The systems in the first column are sound for the corresponding classes of QIS in the second column.*

System	QIS
QKT_m	$\mathcal{QIS}_m, \mathcal{QIS}_m^{sync}, \mathcal{QIS}_m^{uis}, \mathcal{QIS}_m^{sync,uis}$
QKT_m^1	$\mathcal{QIS}_m^{sync,pr}, \mathcal{QIS}_m^{sync,pr,uis}$
QKT_m^2	$\mathcal{QIS}_m^{pr}, \mathcal{QIS}_m^{pr,uis}$
QKT_m^3	$\mathcal{QIS}_m^{nl}$
QKT_m^4	$\mathcal{QIS}_m^{sync,nl}$
$\text{QKT}_m^{2,3}$	$\mathcal{QIS}_m^{nl,pr}$
$\text{QKT}_1^{2,3}$	$\mathcal{QIS}_1^{nl,pr,uis}$
QKT_1^3	$\mathcal{QIS}_1^{nl,uis}$
$\text{QKT}_m^{1,4}$	$\mathcal{QIS}_m^{sync,nl,pr}$
$\text{QKT}_m^{1,4,5}$	$\mathcal{QIS}_m^{sync,nl,uis}, \mathcal{QIS}_m^{sync,nl,pr,uis}$

We now show that the systems in Theorem 7 are not only sound but also complete for the corresponding classes of QIS. For proving these results we need to introduce Kripke models as generalizations of quantified interpreted systems.

Kripke Models

To prove the completeness of the systems above we first introduce an appropriate class of Kripke models as a generalization of QIS, and prove completeness for these models. Then we apply a map between Kripke models and QIS to obtain the desired result.

Definition 8 *A Kripke model for $\mathcal{L}_m$ is a tuple $\mathcal{M} = \langle \mathcal{R}, \{\sim_i\}_{i \in A}, \mathcal{D}, I \rangle$ such that (i) $\mathcal{R}$ is a non-empty set of indexes $r, r', \ldots$; (ii) for $i \in A$, $\sim_i$ is an equivalence relation on the set of points (r, n) for $r \in \mathcal{R}$ and $n \in \mathbb{N}$; (iii) the elements $\mathcal{D}$ and I are defined as for QIS.*

Kripke models can be seen as abstractions of QIS where no details are given about the inner structure of points. The clauses for the satisfaction relation $(\mathcal{M}^\sigma, (r, n)) \models \phi$ are straightforwardly defined from those for QIS, as well as the notions of truth and validity. For instance, we have

$$(\mathcal{M}^\sigma, (r, n)) \models K_i \psi \text{ if } (r, n) \sim_i (r', n') \Rightarrow (\mathcal{M}^\sigma, (r', n')) \models \psi$$

We will consider Kripke models satisfying synchronicity, perfect recall, no learning, or with a unique initial state. The definition of these subclasses can be derived directly from Definition 3. For instance, a Kripke model satisfies *perfect recall* if for all points $(r, n) \sim_i (r', n')$, if $n > 0$ then either $(r, n-1) \sim_i (r', n')$ or there is $l < n'$ such that $(r, n-1) \sim_i (r', l)$ and for all $l < k \leq n'$ we have $(r, n) \sim_i (r', k)$.

Now let $\mathcal{K}_m$ be the class of Kripke models with m agents; in the following we adopt the same naming conventions as for QIS. For instance, $\mathcal{K}_m^{sync,uis}$ is the class of synchronous Kripke models with a unique initial state

We compare Kripke models and quantified interpreted systems by means of a map $g : \mathcal{K}_m \to \mathcal{QIS}_m$. Let $\mathcal{M} = \langle \mathcal{R}, \{\sim_i\}_{i \in A}, \mathcal{D}, I \rangle$ be a Kripke model. For every equivalence relation $\sim_i$, for $(r, n) \in \mathcal{M}$, let the equivalence class $[(r, n)]_{\sim_i} = \{(r', n') \mid (r, n) \sim_i (r', n')\}$ be a local state for agent i, while each (r, n) is a local state for the environment. Then define $g(\mathcal{M})$ as the tuple $\langle \mathcal{R}', \mathcal{D}, I' \rangle$ where $\mathcal{R}'$ contains the runs $\mathbf{r}_r$ for $r \in \mathcal{R}$ such that $\mathbf{r}_r(n) = \langle (r, n), [(r, n)]_{\sim_1}, \ldots, [(r, n)]_{\sim_m} \rangle$, $\mathcal{D}$ is the same as in $\mathcal{M}$, and $I'(P^k, \mathbf{r}_r, n) = I(P^k, r, n)$. The structure $g(\mathcal{M})$ is a QIS that satisfies the following result:

Lemma 9 *For every $\phi \in \mathcal{L}_m$ and $n \in \mathbb{N}$,*

$$(\mathcal{M}^\sigma, (r, n)) \models \phi \quad \textit{iff} \quad (g(\mathcal{M})^\sigma, \mathbf{r}_r, n) \models \phi$$

This lemma is proved by induction on the length of ϕ. Note that if $\mathcal{M}$ satisfies any of synchronicity, perfect recall, no learning, or unique initial state, then also $g(\mathcal{M})$ satisfies the corresponding condition. Thus, g defines a map from each of the 16 subclasses of $\mathcal{K}_m$ outlined above to the corresponding subclass of $\mathcal{QIS}_m$.

For reasoning about the monodic fragment of $\mathcal{L}_m$ when we have no learning and perfect recall we need to introduce the following class of "monodic friendly" Kripke models.

Definition 10 (mf-model) *A monodic friendly Kripke model for $\mathcal{L}_m$ is a tuple $\mathcal{M} = \langle \mathcal{R}, \{\sim_{i,a}\}_{i \in A, a \in \mathcal{D}}, \mathcal{D}, I \rangle$ such that (i) the elements $\mathcal{R}$, $\mathcal{D}$ and I are defined as for Krikpe models; (ii) for $i \in A$, $a \in \mathcal{D}$, $\sim_{i,a}$ is an equivalence relation on the set of points in $\mathcal{M}$.*

We can define synchronicity, perfect recall, no learning, and having a unique initial state also for mf-models by specifying Definition 3 for each relation $\sim_{i,a}$. For instance, a mf-model satisfies *perfect recall* if for all points $(r, n) \sim_{i,a} (r', n')$, if $n > 0$ then either $(r, n-1) \sim_{i,a} (r', n')$ or there is $l < n'$ such that $(r, n-1) \sim_{i,a} (r', l)$ and for all $l < k \leq n'$ we have $(r, n) \sim_{i,a} (r', k)$.

As regards the subclasses of the class $\mathcal{MF}_m$ of all mf-models with m agents, we adopt the same naming conventions as for QIS and Kripke models. Also notice that Kripke models are isomorphic to the mf-models such that for all $i \in A$, $a, b \in \mathcal{D}$, $\sim_{i,a}$ is equal to $\sim_{i,b}$.

Finally, the satisfaction relation $\models$ for $\phi \in \mathcal{L}_m^1$ in a mf-model $\mathcal{M}$ is defined as for Kripke models, but for the epistemic operator:

$$(\mathcal{M}^\sigma, (r, n)) \models K_i \psi[y] \text{ if}$$
$$(r, n) \sim_{i,\sigma(y)} (r', n') \Rightarrow (\mathcal{M}^\sigma, (r', n')) \models \psi$$

where at most y appears free in ψ.

We can now prove the following lemma, which will be useful in the completeness proof for systems encompassing either perfect recall or no learning.

Lemma 11 *For every $\phi \in \mathcal{L}_m^1$ and for every subset $\mathfrak{x}$ of $\{sync, pr, nl, uis\}$,*

$$\mathcal{K}_m^{\mathfrak{x}} \models \phi \quad \textit{iff} \quad \mathcal{MF}_m^{\mathfrak{x}} \models \phi$$

Proof sketch. The implication from right to left is immediate by the remark above. For the converse, assume that $\mathcal{M}$ is a mf-model such that $(\mathcal{M}^\sigma, (r, n)) \not\models \phi$ for some assignment σ, $r \in \mathcal{R}$ and $n \in \mathbb{N}$. We show how to build a Kripke model $\mathcal{M}' = \langle \mathcal{R}', \{\sim_i'\}_{i \in A}, \mathcal{D}', I' \rangle$ such that $(\mathcal{M}'^\sigma, (r', n')) \not\models \phi$ for some $r' \in \mathcal{R}'$ and $n' \in \mathbb{N}$. Let $\mathcal{R}' = \mathcal{R}$ and $\mathcal{D}' = \mathcal{D}$. In order to define each $\sim_i'$ for $i \in A$ we reason as follows. Suppose that $(\mathcal{M}^\sigma, (r, n)) \models K_i \psi[x]$ and $(r, n) \sim_{i,\sigma(x)} (r', n')$, then $(r, n) \sim_i' (r', n')$. However, if $(\mathcal{M}^\sigma, (r, n)) \models K_i \chi[y]$, $(r, n) \sim_{i,\sigma(y)} (r'', n'')$ and $(\mathcal{M}^\sigma, (r'', n'')) \not\models \psi[x]$, we have that $(r, n) \sim_i' (r'', n'')$ and $(\mathcal{M}^\sigma, (r'', n'')) \not\models \psi[x]$ against the fact that $(\mathcal{M}^\sigma, (r, n)) \models K_i \psi[x]$. So, we have to define the interpretation I' so that $(\mathcal{M}^\sigma, (r'', n'')) \models \psi[x]$. This is possible as we are considering only the monodic fragment $\mathcal{L}_m^1$ of $\mathcal{L}_m$. We repeat this process for all points reachable from (r, n) via any epistemic relation.

The Completeness Proof

We outline the main steps of the completeness proof, which is based on a quasimodel construction (Gabbay et al. 2003). Intuitively, a quasimodel for a monodic formula ϕ is a relational structure whose points are sets of sets of subformulas of ϕ. Each set of sets of subformulas describes a "possible state of affairs", and contains sets of subformulas defining the individuals in the point.

Given a formula $\phi \in \mathcal{L}_n^1$ we denote by $sub_{\bigcirc \neg} \phi$ the set $sub\phi \cup \{\neg \psi \mid \psi \in sub\phi\} \cup \{\bigcirc \psi \mid \psi \in sub\phi\} \cup \{\bigcirc \neg \psi \mid \psi \in sub\phi\}$ where $sub\phi$ is the set of subformulas of ϕ. Further, let $sub_n \phi$ be the subset of $sub_{\bigcirc \neg} \phi$ containing formulas with at most n free variables, and let x be a variable not occurring in ϕ, we define $sub_x \phi = \{\psi[y/x] \mid \psi[y] \in sub_1 \phi\}$.

Clearly, x is the only free variable in $sub_x\phi$. By $con\phi$ we denote the set of all constants occurring in ϕ.

Definition 12 *For $k \in \mathbb{N}$ we define the closures $cl_k\phi$ and $cl_{k,i}\phi$ by mutual recursion. Let $cl_0\phi = sub_x\phi$ and for $k \geq 1$, $cl_k\phi = \bigcup_{i \in A} cl_{k,i}\phi$. For $k \geq 0$, $i \in A$, $cl_{k,i}\phi = cl_k\phi \cup \{K_i(\psi_1 \vee \ldots \vee \psi_n), \neg K_i(\psi_1 \vee \ldots \vee \psi_n) \mid \psi_1, \ldots, \psi_n \in cl_k\phi\}$.*

Following (Halpern, Meyden, and Vardi 2003) we define $ad(\phi)$ as the greatest number of alternations of distinct K_i modalities along any branch in ϕ's parse tree. Further, an *index* is any finite sequence $\iota = i_1, \ldots, i_k$ of agents such that $i_n \neq i_{n+1}$; the length of ι is denoted by $|\iota|$. Also, $\iota \sharp i$ is the absorptive concatenation of index ι and i such that $\iota \sharp i = \iota$ if $i_k = i$. Finally, we write $K_\iota \psi$ for $K_{i_1} \ldots K_{i_k} \psi$.

Definition 13 *Let ι be an index such that $|\iota| \leq ad(\phi)$. If ι is the empty sequence ϵ then $cl_\iota\phi = cl_{ad(\phi)}\phi$. If $\iota = \iota' \sharp i$, then $cl_\iota\phi = cl_{k,i}\phi$ for $k = ad(\phi) - |\iota|$. A ι-type $\mathfrak{t}$ for ϕ is any maximal and consistent subset of $cl_\iota\phi$.*

Two ι-types $\mathfrak{t}$, $\mathfrak{t}'$ are said to *agree* if they contains the same closed formulas, i.e., formulas where no free variable appears. Given a ι-type $\mathfrak{t}$ for ϕ and a constant $c \in con\phi$, $\mathfrak{t}^c$ is an *indexed type* for ϕ.

Definition 14 *A ι-state candidate for ϕ is a pair $\mathfrak{C} = \langle T, T^{con} \rangle$ such that (i) T is a set of ι-types for ϕ that agree; and (ii) T^{con} is a set containing for each $c \in con\phi$ an indexed type $\mathfrak{t}^c$ such that $\mathfrak{t} \in T$. A ι-point for ϕ is a pair $\mathfrak{P} = \langle \mathfrak{C}, \mathfrak{t} \rangle$ such that (i) $\mathfrak{C}$ is a ι-state candidate for ϕ; and (ii) $\mathfrak{t} \in \mathfrak{C}$ is a ι-type.*

Note that by abuse of notation, we call *points* both the states (r, n) and the pairs $\mathfrak{P} = \langle \mathfrak{C}, \mathfrak{t} \rangle$. This is to be consistent with our references (Fagin et al. 1995; Halpern, Meyden, and Vardi 2003); the context will disambiguate.

Given a ι-state candidate $\mathfrak{C} = \langle T, T^{con} \rangle$ and a point $\mathfrak{P} = \langle \mathfrak{C}, \mathfrak{t} \rangle$ we define the formulas $\alpha_\mathfrak{C}$ and $\beta_\mathfrak{P}$ as follows:

$$\alpha_\mathfrak{C} := \bigwedge_{\mathfrak{t} \in T} \exists x \mathfrak{t}[x] \wedge \forall x \bigvee_{\mathfrak{t} \in T} \mathfrak{t}[x] \wedge \bigwedge_{\mathfrak{t}^c \in T^{con}} \mathfrak{t}[x/c]$$

$$\beta_\mathfrak{P} := \alpha_\mathfrak{C} \wedge \mathfrak{t}$$

A ι-state candidate $\mathfrak{C}$ is *consistent* if the formula $\alpha_\mathfrak{C}$ is consistent; similarly for points. Consistent state candidates represent the states of our quasimodels. We now define the relations of *suitability* that constitute the relational part of quasimodels.

Definition 15 • *A ι_1-type $\mathfrak{t}_1$ and a ι_2-type $\mathfrak{t}_2$ are $\bigcirc$-suitable, or $\mathfrak{t}_1 \Rightarrow \mathfrak{t}_2$, if $\iota_1 = \iota_2$ and $\mathfrak{t}_1 \wedge \bigcirc \mathfrak{t}_2$ is consistent. They are i-suitable, or $\mathfrak{t}_1 \approx_i \mathfrak{t}_2$, if $\iota_1 \sharp i = \iota_2 \sharp i$ and $\mathfrak{t}_1 \wedge \bar{K}_i \mathfrak{t}_2$ is consistent.*

• *A ι_1-state candidate $\mathfrak{C}_1$ and a ι_2-state candidate $\mathfrak{C}_2$ are $\bigcirc$-suitable, or $\mathfrak{C}_1 \Rightarrow \mathfrak{C}_2$, if $\iota_1 = \iota_2$ and $\alpha_{\mathfrak{C}_1} \wedge \bigcirc \alpha_{\mathfrak{C}_2}$ is consistent. They are i-suitable if $\iota_1 \sharp i = \iota_2 \sharp i$ and $\alpha_{\mathfrak{C}_1} \wedge \bar{K}_i \alpha_{\mathfrak{C}_2}$ is consistent.*

• *A ι_1-point $\mathfrak{P}_1$ and a ι_2-point $\mathfrak{P}_2$ are $\bigcirc$-suitable, or $\mathfrak{P}_1 \Rightarrow \mathfrak{P}_2$, if $\iota_1 = \iota_2$ and $\beta_{\mathfrak{P}_1} \wedge \bigcirc \beta_{\mathfrak{P}_2}$ is consistent. They are i-suitable if $\iota_1 \sharp i = \iota_2 \sharp i$ and $\beta_{\mathfrak{P}_1} \wedge \bar{K}_i \beta_{\mathfrak{P}_2}$ is consistent.*

Furthermore, for $c \in con\phi$, $\mathfrak{C}_1 \Rightarrow^c \mathfrak{C}_2$ if $\mathfrak{t}_1^c \in \mathfrak{C}_1$, $\mathfrak{t}_2^c \in \mathfrak{C}_2$ and $\langle \mathfrak{C}_1, \mathfrak{t}_1^c \rangle \Rightarrow \langle \mathfrak{C}_2, \mathfrak{t}_1^c \rangle$.

We now present the frame underlying the quasimodel for ϕ.

Definition 16 *A frame $\mathcal{F}$ is a tuple $\langle \mathcal{R}, \{\sim_{i,a}\}_{i \in A, a \in \mathcal{D}}, \mathcal{D}, \mathfrak{f} \rangle$ where (i) $\mathcal{R}$, $\{\sim_{i,a}\}_{i \in A, a \in \mathcal{D}}$ and $\mathcal{D}$ are defined as for mf-models; (ii) $\mathfrak{f}$ is a partial function associating to each point (r, n) a consistent state candidate $\mathfrak{f}(r, n) = \mathfrak{C}_{r,n}$ such that (a) the domain of $\mathfrak{f}$ is not empty; (b) if $\mathfrak{f}$ is defined on (r, n) then it is defined on $(r, n+1)$; (c) if $\mathfrak{f}$ is defined on (r, n) and $(r, n) \sim_{i,a} (r', n')$ then $\mathfrak{f}$ is defined on (r', n').*

Next, we provide the definition of *objects*, which correspond to the *runs* in (Gabbay et al. 2003). We choose this name to avoid confusion with the runs in QIS.

Definition 17 *For $a \in \mathcal{D}$, an object in $\mathcal{F}$ is a map ρ_a associating with every $(r, n) \in Dom(\mathfrak{f})$ a type $\rho_a(r, n) \in T_{r,n}$ such that:*

1. *$\rho_a(r, n) \Rightarrow \rho_a(r, n+1)$; and if $(r, n) \sim_{i,a} (r', n')$ then $\rho_a(r, n) \approx_i \rho_a(r', n')$*

2. *$\chi \mathcal{U} \psi \in \rho_a(r, n)$ iff there is $n' \geq n$ such that $\psi \in \rho_a(r, n')$ and $\chi \in \rho_a(r, n'')$ for all $n \leq n'' < n'$;*

3. *if $\rho_a(r, n) \approx_i \mathfrak{t}$ are ι-types then for some (r', n'), $(r, n) \sim_{i,a} (r', n')$ and $\rho_a(r', n') = \mathfrak{t}$.*

   An object$^+$ satisfies (1), (2) above and (3') instead of (3).

3' *if $\rho_a(r, n)$ is a ι-type, $\mathfrak{t}$ is a $\iota \sharp i$-type, and $\rho_a(r, n) \approx_i \mathfrak{t}$ then for some $(r', n') \sim_{i,a} (r, n)$, $\rho_a(r', n') = \mathfrak{t}$.*

Now we have all the elements to give the definition of quasimodel.

Definition 18 *A quasimodel for ϕ is a tuple $\mathfrak{Q} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a frame, and*

1. *$\phi \in \mathfrak{t}$ for some $\mathfrak{t} \in T_{r,n}$ and $T_{r,n} \in \mathfrak{C}_{r,n}$*

2. *$\mathfrak{C}_{r,n} \Rightarrow \mathfrak{C}_{r,n+1}$; and if $(r, n) \sim_{i,\rho} (r', n')$ then $\rho(r, n) \approx_i \rho(r', n')$*

3. *for every $\mathfrak{t} \in T_{r,n}$ there exists an object $\rho \in \mathcal{O}$ such that $\rho(r, n) = \mathfrak{t}$*

4. *for every $c \in con\phi$, the function ρ^c such that $\rho^c(r, n) = \mathfrak{t}^c \in T_{r,n}^{con}$ is an object in $\mathcal{O}$.*

A *quasimodel$^+$* is defined as a quasimodel in which clauses (3) and (4) refer to objects$^+$ rather than objects. We can define quasimodels satisfying perfect recall, no learning, synchronicity, or unique initial state by assuming the corresponding condition on the frame.

We now state the main result of this section, that is, satisfability in quasimodels implies satisfability in mf-models.

Theorem 19 *If there is a quasimodel (respectively quasimodel$^+$) $\mathfrak{Q}$ for a monodic formula $\phi \in \mathcal{L}_m^1$ then ϕ is satisfiable in a monodic friendly Kripke model.*

Proof sketch. The proof is inspired by those for Lemmas 11.72 and 12.9 in (Gabbay et al. 2003).

First, for every monodic formula $\psi \in \mathcal{L}_m^1$ of the form $K_i\chi$, $\bigcirc\chi$ or $\chi_1 \mathcal{U} \chi_2$ we consider a k-ary predicate P_ψ^k for k

equal to 0 or 1. The formula $P_\psi^k(x)$ is called the *surrogate* of ψ. Given a formula $\phi \in \mathcal{L}_m^1$ we denote by $\overline{\phi}$ the formula obtained from ϕ by substituting all its modal subformulas which are not within the scope of another modal operator by their surrogates.

Since every state candidate $\mathfrak{C}$ in the quasimodel $\mathfrak{Q}$ is consistent and the system QKT_m is based on first-order logic, the formula $\overline{\alpha}_\mathfrak{C}$ is consistent with first-order (non-modal) logic. By completeness of first-order logic, there is a first-order structure $\mathcal{I} = \langle I, \mathcal{D} \rangle$, where $\mathcal{D}$ is a non-empty set of individuals and I is an interpretation on $\mathcal{D}$, which satisfies $\overline{\alpha}_\mathfrak{C}$, that is, $I^\sigma \models \overline{\alpha}_\mathfrak{C}$ for some assignment σ to $\mathcal{D}$.

Now, we consider a cardinal number $\kappa \geq \aleph_0$ greater than the cardinality of the set $\mathcal{O}$ of all objects in $\mathfrak{Q}$, and define $\mathcal{D} = \{\langle \rho, \xi \rangle \mid \rho \in \mathcal{O}, \xi < \kappa\}$. By Löwenheim-Skolem theorem we can assume without loss of generality that $\mathcal{D}$ is the domain of the first-order structure $\mathcal{I}_{r,n} = \langle I_{r,n}, \mathcal{D} \rangle$ satisfying $\overline{\alpha}_{\mathfrak{C}_{r,n}}$, that is, all structures $\mathcal{I}_{r,n}$ share a common domain $\mathcal{D}$, and for every $\mathfrak{t} \in T_{r,n}$, $\langle \rho, \xi \rangle \in \mathcal{D}$, we have that $\rho(r,n) = \mathfrak{t}$ iff $I_{r,n}^\sigma \models \overline{\mathfrak{t}}[x]$ for $\sigma(x) = \langle \rho, \xi \rangle$. Moreover, $I_{r,n}(c) = \langle \rho, 0 \rangle$ for every $c \in con\phi$.

We define the mf-model $\mathcal{M}$ as the triple $\langle \mathcal{R}, \{\sim_{i,a}\}_{i \in A, a \in \mathcal{D}}, \mathcal{D}, I \rangle$ such that $\mathcal{R}$ is the set of runs in the quasimodel $\mathfrak{Q}$, for $a = \langle \rho, \xi \rangle \in \mathcal{D}$, $\sim_{i,a}$ is equal to $\sim_{i,\rho}$, $\mathcal{D}$ is defined as above, and the interpretation I is obtained by gluing together the various $I_{r,n}$. We can now prove the following result for $\mathcal{M}$.

Remark 20 *If $\mathcal{M}$ is obtained from a quasimodel $\mathfrak{Q}$ as described above, then for every $\psi \in sub_x\phi$*

$$I_{r,n}^\sigma \models \overline{\psi} \quad iff \quad (\mathcal{M}^\sigma, (r,n)) \models \psi$$

Moreover, if $\mathfrak{Q}$ is a quasimodel$^+$, $\mathfrak{f}(r,n)$ is a ι-state candidate, and $ad(K_\iota\psi) \leq d$ then

$$I_{r,n}^\sigma \models \overline{\psi} \quad iff \quad (\mathcal{M}^\sigma, (r,n)) \models \psi$$

Furthermore, if $\mathfrak{Q}$ satisfies any of perfect recall, no learning, synchronicity, or unique initial state, then the mf-model $\mathcal{M}$ obtained from $\mathfrak{Q}$ satisfies the corresponding constraints.

Dealing with each System

In this section we consider the completeness proof for each system in Theorem 7. In particular, we show that if $\phi \in \mathcal{L}_m^1$ is consistent with respect to a system S, then we can build a quasimodel (or a quasimodel$^+$ in some cases) for ϕ based on a frame for S. We only present the main steps of the construction. Notice that in the following sections the symbol $\vdash$ represents provability in the appropriate system S.

The Classes $\mathcal{QIS}_m$, $\mathcal{QIS}_m^{sync}$, $\mathcal{QIS}_m^{uis}$, $\mathcal{QIS}_m^{sync,uis}$

We start the completeness proof for the basic system QKT_m with the following definition.

Definition 21 *Let a $\Rightarrow$-sequence be a possibly infinite sequence $\mathfrak{C}_0 \Rightarrow \mathfrak{C}_1 \Rightarrow \dots$ of state candidates.*
A $\Rightarrow$-sequence is acceptable *if*

(i) for all $k \geq 0$ if $\psi\mathcal{U}\chi \in \mathfrak{t}_k \in \mathfrak{C}_k$ then there is $n \geq k$ such that $\chi \in \mathfrak{t}_n \in \mathfrak{C}_n$ and $\psi \in \mathfrak{t}_m \in \mathfrak{C}_m$ for all $n \leq m < k$;

(ii) for all $k \geq 0$ if $\psi\mathcal{U}\chi \in \mathfrak{t}_k^c \in \mathfrak{C}_k$ then there is $n \geq k$ such that $\chi \in \mathfrak{t}_n^c \in \mathfrak{C}_n$, $\psi \in \mathfrak{t}_m^c \in \mathfrak{C}_m$ for all $n \leq m < k$ and $\mathfrak{C}_k \Rightarrow^c \mathfrak{C}_{k+1} \Rightarrow^c \dots \Rightarrow^c \mathfrak{C}_m$.

The following lemmas entail the completeness result.

Lemma 22 *For any consistent $\phi \in \mathcal{L}_m^1$ there is a consistent ϵ-state candidate $\mathfrak{C} = \langle T, T^{con} \rangle$ for ϕ such that $\phi \in \mathfrak{t}$ for some $\mathfrak{t} \in T$.*

Lemma 23 *Every $\Rightarrow$-sequence of state candidates can be extended to an infinite acceptable $\Rightarrow$-sequence.*

These lemmas are proved similarly to Claims 11.75-76 in (Gabbay et al. 2003).

If $\phi \in \mathcal{L}_m^1$ is consistent then by Lemma 22 there is a consistent ϵ-state candidate $\mathfrak{C} = \langle T, T^{con} \rangle$ such that ϕ belongs to some type $\mathfrak{t} \in T$. So, by Lemma 23 the set of infinite acceptable $\Rightarrow$-sequences is non-empty. Now let $\mathfrak{x}$ be a new object. A sequence $\mathfrak{x}, \dots, \mathfrak{x}, \mathfrak{C}_n, \mathfrak{C}_{n+1}, \dots$ is *acceptable from n* if it starts with n copies of $\mathfrak{x}$ and $\mathfrak{C}_n, \mathfrak{C}_{n+1}, \dots$ is an acceptable $\Rightarrow$-sequence. Let $\mathcal{R}$ be the set of all such acceptable $\Rightarrow$-sequences, and for each (r,k) define the partial function $\mathfrak{f}$ as $\mathfrak{f}(r,k) = \mathfrak{C}_k$ if r is the $\Rightarrow$-sequence $\mathfrak{x}, \dots, \mathfrak{x}, \mathfrak{C}_n, \mathfrak{C}_{n+1}, \dots$ acceptable from n and $k \geq n$, undefined otherwise. Finally, let $\mathcal{O}$ be the set of all objects on the frame $\mathcal{F} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\mathcal{F}$ is synchronous. We can prove the following result.

Lemma 24 *The tuple $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a quasimodel for ϕ.*

The completeness of QKT_m with respect to the classes $\mathcal{QIS}$ and $\mathcal{QIS}^{sync}$ directly follows from Theorem 19. To prove completeness for $\mathcal{QIS}^{uis}$ and $\mathcal{QIS}^{sync,uis}$ we use the next result.

Remark 25 *Suppose $\mathfrak{y}$ is a subset of $\{pr, sync\}$. If $\phi \in \mathcal{L}_m$ is satisfiable in $\mathcal{QIS}^\mathfrak{y}$ then it is also satisfiable in $\mathcal{QIS}^{\mathfrak{y},uis}$.*

The Classes $\mathcal{QIS}_m^{pr}$, $\mathcal{QIS}_m^{pr,uis}$

The completeness proof for QKT_m^2 with respect to $\mathcal{QIS}^{pr}$ and $\mathcal{QIS}^{pr,uis}$ relies on the following lemma.

Lemma 26 *For ι-points $\mathfrak{P}_1 = \langle \mathfrak{C}_1, \mathfrak{t}_1 \rangle$, $\mathfrak{P}_2 = \langle \mathfrak{C}_2, \mathfrak{t}_2 \rangle$ and $\iota\sharp i$-type $\mathfrak{t}_2'$, if $\mathfrak{P}_1 \Rightarrow \mathfrak{P}_2$ and $\mathfrak{t}_2 \approx_i \mathfrak{t}_2'$ then there is a $\iota\sharp i$-type $\mathfrak{P}_2' = \langle \mathfrak{C}_2', \mathfrak{t}_2' \rangle$ such that either (a) $\mathfrak{t}_1 \approx_i \mathfrak{t}_2'$ or (b) there is a $\iota\sharp i$-type $\mathfrak{P}_1' = \langle \mathfrak{C}_1', \mathfrak{t}_1' \rangle$ such that $\mathfrak{t}_1 \approx_i \mathfrak{t}_1'$ and a $\Rightarrow$-sequence of $\iota\sharp i$-points $\mathfrak{S}_0 \Rightarrow \dots \Rightarrow \mathfrak{S}_n = \mathfrak{P}_2'$ such that $\mathfrak{S}_k = \langle \mathfrak{D}_k, \mathfrak{s}_k \rangle$ and $\mathfrak{s}_k \approx_i \mathfrak{t}_2$ for $k \leq n$, and $\mathfrak{P}_1' \Rightarrow \mathfrak{S}_0$.*

For any consistent $\phi \in \mathcal{L}_m^1$ we define a quasimodel$^+$ for ϕ to establish the completeness of QKT^2 with respect to $\mathcal{QIS}^{pr}$. Let $\mathcal{R}$ be the set of all acceptable $\Rightarrow$-sequences, and define $\mathfrak{f}$ such that $\mathfrak{f}(r,k) = \mathfrak{C}_k$ if r is the $\Rightarrow$-sequence $\mathfrak{C}_0, \mathfrak{C}_1, \dots$. Finally, let $\mathcal{O}$ be the set of all objects$^+$ on the frame $\mathcal{F} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\mathcal{F}$ satisfies perfect recall. We can now show the following lemma.

Lemma 27 *The tuple $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a quasimodel$^+$ for ϕ.*

However, we need to ensure that the set $\mathcal{O}$ of objects$^+$ is non-empty. In particular, we need the following lemma to show that clause (3') is satisfied.

Lemma 28 *if $\rho(r,n) \in \mathfrak{f}(r,n)$ is a ι-type, $\mathfrak{t}$ is a $\iota\sharp i$-type and $\rho(r,n) \approx_i \mathfrak{t}$ then for some $(r',n') \sim_{i,\rho} (r,n)$, $\mathfrak{t} \in \mathfrak{f}(r',n')$.*

Proof sketch. This proof is similar to the one for Lemma 5.6 in (Halpern, Meyden, and Vardi 2003); it proceeds by induction on n. For $n = 0$ we define a consistent $\iota\sharp i$-state candidate $\mathfrak{D} = \{\mathfrak{s} \mid \mathfrak{s} \in \mathfrak{f}(r,0), \mathfrak{s} \neq \rho(r,0)\} \cup \{\mathfrak{t}\}$. By Lemma 23 $\mathfrak{D}$ can be extended to a $\Rightarrow$-acceptable sequence r' such that $\rho(r',0) = \mathfrak{t}$. Finally, $(r',n') \sim_{i,\rho} (r,n)$ and $\mathfrak{t} \in \mathfrak{f}(r',n')$.

For the inductive step assume that $\mathfrak{f}(r,n-1) \Rightarrow \mathfrak{f}(r,n)$ and $\rho(r,n) \approx_i \mathfrak{t}$. By Lemma 26 either (a) $\rho(r,n-1) \approx_i \mathfrak{t}$ or (b) there is a $\iota\sharp i$-type $\mathfrak{P}' = \langle \mathfrak{C}', \mathfrak{t}' \rangle$ such that $\rho(r,n-1) \approx_i \mathfrak{t}'$ and a $\Rightarrow$-sequence of $\iota\sharp i$-points $\mathfrak{S}_0 \Rightarrow \ldots \Rightarrow \mathfrak{S}_l = \langle \mathfrak{D}, \mathfrak{t} \rangle$ such that $\mathfrak{S}_k = \langle \mathfrak{D}_k, \mathfrak{s}_k \rangle$ and $\mathfrak{s}_k \approx_i \rho(r,n)$ for $k \leq l$ and $\mathfrak{P}' \Rightarrow \mathfrak{S}_0$. If we apply the induction hypothesis in case (a) then we obtain that for some $(r',n') \sim_{i,\rho} (r,n-1)$, $\mathfrak{t} \in \mathfrak{f}(r',n')$ and $\rho(r,n-1) \approx_i \rho(r,n)$. Thus, also $(r',n') \sim_{i,\rho} (r,n)$. In case (b) by induction hypothesis we have that for some $(r',n') \sim_{i,\rho} (r,n-1)$, $\mathfrak{t}' \in \mathfrak{f}(r',n')$. Now assume that run r' is derived from the $\Rightarrow$-acceptable sequence $\mathfrak{C}_0, \mathfrak{C}_1, \ldots$, and let r'' be the run derived from the sequence with initial segment $\mathfrak{C}_0, \ldots, \mathfrak{C}_{n'}, \mathfrak{D}_0, \ldots, \mathfrak{D}_l$ by Lemma 23. By construction $\mathfrak{f}(r'', n'+l+1) = \mathfrak{D}$ and $\rho(r,n) \approx_i \rho(r'', n'+l+1) = \mathfrak{t}$. Hence, $(r,n) \sim_{i,\rho} (r'', n'+l+1)$.

This completes the proof for $\mathcal{QIS}_m^{pr}$. The completeness of QKT_m^2 with respect to $\mathcal{QIS}_m^{pr,uis}$ follows by Remark 25.

The Classes $\mathcal{QIS}_m^{sync,pr}$, $\mathcal{QIS}_m^{sync,pr,uis}$

The completeness of QKT_m^1 with respect to $\mathcal{QIS}^{sync,pr}$ is proved similarly to the previous case by using the next lemma instead of Lemma 26.

Lemma 29 *For ι-points $\mathfrak{P}_1$, $\mathfrak{P}_2$ and $\iota\sharp i$-point $\mathfrak{P}_2'$, if $\mathfrak{P}_1 \Rightarrow \mathfrak{P}_2$ and $\mathfrak{P}_2 \approx_i \mathfrak{P}_2'$ then there is a $\iota\sharp i$-point $\mathfrak{P}_1'$ such that $\mathfrak{P}_1 \approx_i \mathfrak{P}_1'$ and $\mathfrak{P}_1' \Rightarrow \mathfrak{P}_2'$.*

Completeness of QKT_m^1 with respect to $\mathcal{QIS}_m^{pr,sync,uis}$ follows again by Remark 25.

The Class $\mathcal{QIS}_m^{nl}$

First, we give the following definition, which will be useful in the completeness proof.

Definition 30 *Two sequences of types Σ and Σ' are $\approx_i$-concordant if there is some $n \in \mathbb{N}$ (n may be ∞) and non-empty consecutive intervals $\Sigma_1, \ldots, \Sigma_n$ of Σ and $\Sigma_1', \ldots, \Sigma_n'$ of Σ' such that for all $\mathfrak{s} \in \Sigma_j$ and $\mathfrak{s}' \in \Sigma_j'$ we have $\mathfrak{s} \approx_i \mathfrak{s}'$ for $j \leq n$.*

To prove the completeness of QKT_m^3 for $\mathcal{QIS}_m^{nl}$ we need the following lemma, which is dual to Lemma 26.

Lemma 31 *For ι-points $\mathfrak{P}_1 = \langle \mathfrak{C}_1, \mathfrak{t}_1 \rangle$, $\mathfrak{P}_2 = \langle \mathfrak{C}_2, \mathfrak{t}_2 \rangle$ and $\iota\sharp i$-type $\mathfrak{t}_1'$, if $\mathfrak{P}_1 \Rightarrow \mathfrak{P}_2$ and $\mathfrak{t}_1 \approx_i \mathfrak{t}_1'$ then there is a $\iota\sharp i$-point $\mathfrak{P}_1' = \langle \mathfrak{C}_1', \mathfrak{t}_1' \rangle$ and a $\Rightarrow$-sequence $\mathfrak{P}_1' = \mathfrak{S}_0 \Rightarrow \ldots \Rightarrow \mathfrak{S}_n$ of $\iota\sharp i$-points such that $\mathfrak{S}_k = \langle \mathfrak{D}_k, \mathfrak{s}_k \rangle$ and $\mathfrak{s}_k \approx_i \mathfrak{t}_1$ for $k < n$, and $\mathfrak{t}_2 \approx_i \mathfrak{s}_n$.*

As pointed out in (Halpern, Meyden, and Vardi 2003) Lemma 31 is not sufficient to construct a quasimodel$^+$ satisfying the no learning condition. In fact, given a $\Rightarrow$-sequence $\Sigma = \mathfrak{C}_0, \mathfrak{C}_1, \ldots$ of ι-state candidates and a $\iota\sharp i$-point $\mathfrak{t}_0'$ such that $\mathfrak{t}_0 \approx_i \mathfrak{t}_0'$ for $\mathfrak{t}_0 \in \mathfrak{C}_0$ by Lemma 31 we can find a $\Rightarrow$-sequence $\Sigma' = \mathfrak{C}_0', \mathfrak{C}_1', \ldots$ such that $\mathfrak{t}_0' \in \mathfrak{C}_0'$ and satisfying the no learning condition. However, it does not follow from the acceptability of Σ that Σ' is also acceptable. So, as in the propositional case, we have to work with trees of state candidates. Hereafter we extend to the first order the definitions given in (Halpern, Meyden, and Vardi 2003) for the propositional case.

Definition 32 *A k-tree for ϕ (for $k \leq ad(\phi)$) is a set Π of ι-state candidates for ϕ with $|\iota| \leq k$ that contains a unique ϵ-state candidate, i.e., the root, and for every ι-point $\mathfrak{t}$ in some $\mathfrak{C} \in \Pi$,*

- *if $\mathfrak{t}'$ is a $\iota\sharp i$-type such that $\mathfrak{t} \approx_i \mathfrak{t}'$ and $|\iota\sharp i| \leq k$ then there is some $\mathfrak{C}' \in \Pi$ such that $\mathfrak{t}' \in \mathfrak{C}'$;*
- *if $\iota = \iota'\sharp i$ then there is a ι'-state candidate $\mathfrak{C}' \in \Pi$ and a ι'-type $\mathfrak{t}' \in \mathfrak{C}'$ such that $\mathfrak{t} \approx_i \mathfrak{t}'$.*

Intuitively, a k-tree is a view of the epistemic state of a quasimodel from a particular type $\mathfrak{t}$, up to k steps from $\mathfrak{t}$. We now extend the relation $\Rightarrow$ to k-trees.

Definition 33 *If Π and Π' are k-trees for ϕ then $\Pi \Rightarrow_f \Pi'$ if f is a function associating with each ι-type $\mathfrak{t} \in \mathfrak{C}$, for $\mathfrak{C} \in \Pi$, a finite $\Rightarrow$-sequence of ι-types in $\Pi \cup \Pi'$ such that:*

1. *if $f(\mathfrak{t}) = \mathfrak{t}_0 \Rightarrow \ldots \Rightarrow \mathfrak{t}_k$ then (a) $\mathfrak{t} = \mathfrak{t}_0$; (b) $\mathfrak{t}_j \in \mathfrak{C}_j$ for some $\mathfrak{C}_j \in \Pi$ for $j < k$ and $\mathfrak{t}_k \in \mathfrak{C}_k$ for some $\mathfrak{C}_k \in \Pi'$;*
2. *if $\mathfrak{t} \approx_i \mathfrak{t}'$ then $f(\mathfrak{t})$ and $f(\mathfrak{t}')$ are $\approx_i$-concordant;*
3. *for at least one $\mathfrak{t}$ the sequence $f(\mathfrak{t})$ has length at least 2.*

We now show how to obtain acceptable sequences of state candidates from sequences of trees. Given two sequences of ι-state candidates $\lambda = \mathfrak{C}_0, \ldots, \mathfrak{C}_k$ and $\mu = \mathfrak{C}_0', \ldots$, where λ is finite, the *fusion* $\lambda \cdot \mu$ is defined as $\mathfrak{C}_0, \ldots, \mathfrak{C}_{k-1}, \mathfrak{C}_0', \ldots$ only if $\mathfrak{C}_k = \mathfrak{C}_0'$. Furthermore, given an infinite sequence $\Theta = \Pi_0 \Rightarrow_{f_0} \Pi_1 \Rightarrow_{f_1} \ldots$ of k-trees, we say that a sequence λ of ι-state candidates is *compatible with* Θ if there exists some $h \in \mathbb{N}$ and ι-state candidates $\mathfrak{C}_h, \mathfrak{C}_{h+1}, \ldots$, with $\mathfrak{C}_j \in \Pi_j$ for $j \geq h$, such that $\lambda = f_h(\mathfrak{C}_h) \cdot f_{h+1}(\mathfrak{C}_{h+1}) \cdot \ldots$. The sequence Θ is *acceptable* if every $\Rightarrow$-sequence compatible with Θ is infinite and acceptable. The basic idea of the completeness proof is to define the quasimodel$^+$ starting from an acceptable sequence Θ.

Lemma 34 *If $\phi \in \mathcal{L}_m^1$ is consistent with QKT_m^3 then there is an acceptable sequence Θ of $ad(\phi)$-trees such that ϕ belongs to the root of the first tree.*

The proof of this lemma relies on Lemma 31. Now let $\mathcal{R}$ consist of all acceptable $\Rightarrow$-sequences compatible with Θ, while the function $\mathfrak{f}$ is defined as for perfect recall. Furthermore, $\mathcal{O}$ is the set of all object$^+$ on the frame $\mathcal{F} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\mathcal{F}$ satisfies no learning. We can now state the following lemma.

Lemma 35 *The tuple $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a quasimodel$^+$ for ϕ.*

This completes the proof for QKT^3 with respect to $\mathcal{QIS}_m^{nl}$.

The Class $\mathcal{QIS}_m^{nl,sync}$

To show that QKT_m^4 is a complete axiomatisation for $\mathcal{QIS}_m^{nl,sync}$ we need the following analogue of Lemma 31.

Lemma 36 *For ι-points $\mathfrak{P}_1$, $\mathfrak{P}_2$ and $\iota\sharp i$-point $\mathfrak{P}_1'$, if $\mathfrak{P}_1 \Rightarrow \mathfrak{P}_2$ and $\mathfrak{P}_1 \approx_i \mathfrak{P}_1'$ then there is a $\iota\sharp i$-point $\mathfrak{P}_2'$ such that $\mathfrak{P}_1' \Rightarrow \mathfrak{P}_2'$ and $\mathfrak{P}_2' \approx_i \mathfrak{P}_2$.*

Further, if Π and Π' are k-trees then $\Pi \Rightarrow_f^{sync} \Pi'$ only if $\Pi \Rightarrow_f \Pi'$ and for all $\mathfrak{t} \in \Pi$, $f(\mathfrak{t})$ has exactly length 2. A *sync*-acceptable sequence of trees is defined as an acceptable sequence where the relation $\Rightarrow$ is substituted by the relation $\Rightarrow^{sync}$. The following analogue of Lemma 34 holds.

Lemma 37 *If $\phi \in \mathcal{L}_m^1$ is consistent with QKT_m^4 then there is a* sync-*acceptable sequence Θ of $ad(\phi)$-trees such that ϕ belongs to the root of the first tree.*

Let $\mathcal{R}$ consist of all acceptable $\Rightarrow$-sequences compatible with Θ. The function $\mathfrak{f}$ is defined as in the previous section, and $\mathcal{O}$ is the set of all object$^+$ on the frame $\mathcal{F} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\mathcal{F}$ satisfies synchronicity and no learning. As in the previous sections the tuple $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a quasimodel$^+$ for ϕ. This completes the proof for QKT_m^4 with respect to $\mathcal{QIS}_m^{nl,sync}$.

The Classes $\mathcal{QIS}_m^{nl,pr}$ and $\mathcal{QIS}_1^{nl,pr,uis}$

If $\phi \in \mathcal{L}_m^1$ is consistent with $\text{QKT}_m^{2,3}$ then by Lemma 34 there exists an acceptable sequence Θ of $ad(\phi)$-trees such that the consistent formula ϕ belongs to the root of the first tree. Let $\mathcal{R}$ be the set of all acceptable $\Rightarrow$-sequences *that have a suffix* that is compatible with Θ, while the function $\mathfrak{f}$ is defined as in the previous section. Finally, $\mathcal{O}$ is the set of all object$^+$ on the frame $\mathcal{F} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\mathcal{F}$ satisfies perfect recall and no learning. We can prove that the tuple $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a quasimodel$^+$ for ϕ. This establishes the completeness of $\text{QKT}^{2,3}$ with respect to $\mathcal{QIS}_m^{nl,pr}$. Completeness with respect to $\mathcal{QIS}_1^{nl,pr,uis}$ follows from the next remark, whose proof is the same as in the propositional case.

Remark 38 *A formula $\phi \in \mathcal{L}_1^1$ is satisfiable in $\mathcal{QIS}_1^{nl,pr}$ (resp. $\mathcal{QIS}_1^{nl,pr,sync}$) iff it is satisfiable in $\mathcal{QIS}_1^{nl,pr,uis}$ (resp. $\mathcal{QIS}_1^{nl,pr,sync,uis}$).*

The Class $\mathcal{QIS}_m^{nl,pr,sync}$

To prove the completeness of $\text{QKT}_m^{1,4}$ with respect to $\mathcal{QIS}_m^{nl,pr,sync}$ we combine the results of the previous two sections. If $\phi \in \mathcal{L}_m^1$ is consistent with $\text{QKT}_m^{1,4}$ then by Lemma 37 there is a *sync*-acceptable sequence Θ of $ad(\phi)$-trees such that ϕ belongs to the root of the first tree. Let $\mathcal{R}$ be the set of all acceptable $\Rightarrow$-sequences with suffixes that are compatible with Θ; the function $\mathfrak{f}$ is defined as above. Finally, $\mathcal{O}$ is the set of all object$^+$ on the frame $\mathcal{F} = \langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ such that $\mathcal{F}$ satisfies perfect recall, no learning and synchronicity. Also in this case we can show that the tuple $\langle \mathcal{R}, \{\sim_{i,\rho}\}_{i \in A, \rho \in \mathcal{O}}, \mathcal{O}, \mathfrak{f} \rangle$ is a quasimodel$^+$ for ϕ. This completes the proof.

The Classes $\mathcal{QIS}_m^{nl,sync,uis}$ and $\mathcal{QIS}_m^{nl,pr,sync,uis}$

We now show that the system $\text{QKT}_m^{1,4,5}$ is complete with respect to the classes $\mathcal{QIS}_m^{nl,sync,uis}$ and $\mathcal{QIS}_m^{nl,pr,sync,uis}$. First, we remark that these two classes share the same set of validities. By this remark and axiom KT5 it is sufficient to prove the completeness of $\text{QKT}_1^{1,4}$ with respect to $\mathcal{QIS}_1^{nl,pr,sync,uis}$. From the previous section $\text{QKT}_1^{1,4}$ is complete with respect to $\mathcal{QIS}_1^{nl,pr,sync}$ and the result follows by Remark 38.

Security Protocols as Message Passing Systems

In this section we present a demonstration of the formal machinery developed thus far. Specifically, we model a security protocol as a message passing system (Fagin et al. 1995; Lamport 1978) in the framework of QIS. First of all, a message passing system (MPS) is a MAS in which the only actions for the agents are sending and receiving messages. This setting is common to a variety of distributed systems, well beyond the realms of MAS and AI.

To define message passing QIS we introduce a set *Msg* of messages $\mu_1, \mu_2, \ldots$, and define the local state l_i for agent i as a *history* over *Msg*, that is, a sequence of events of the form $send(i, j, \mu)$ and $rec(j, \mu)$, for $i, j \in A$, $\mu \in Msg$. Intuitively, $send(i, j, \mu)$ represents the event where *agent i sends agent j message μ*, while the meaning of $rec(j, \mu)$ is that *agent j receives message μ*. A global state $s \in \mathcal{S}$ is a tuple $\langle l_e, l_1, \ldots, l_n \rangle$ where $l_1, \ldots, l_n$ are local states as above and l_e contains all the events in $l_1, \ldots, l_n$.

We define the protocol for message passing systems as follows:

- $P(l_i) = \{\lambda, send(j, \mu) \mid j \in A, \mu \in Msg\}$
- $P(l_e) = \{\lambda, del(j, \mu) \mid j \in A, \mu \in Msg\}$

In each local state agent i can either perform the null action λ or send a message. The environment can either do nothing or deliver a message. Further, we define the transition function:

- $\tau(a_e, a_1, \ldots, a_m)(s_e, s_1, \ldots, s_m) = (s_e', s_1', \ldots, s_m')$ if $a_e = del(j, \mu)$ implies $s_j' = (s_j, rec(j, \mu))$ and $a_i = send(j, \mu)$ implies $s_i' = (s_i, send(i, j, \mu))$.

A run r is a function from the naturals $\mathbb{N}$ to $\mathcal{S}$ that respects the transition function τ. By the definitions of local states, protocols and transition function it is easy to check that the following specifications MP1 and MP3 are satisfied.

MP1 for every $n \in \mathbb{N}$, $r_i(n)$ is a history over *Msg*;

MP2 for $i \in A$, $r_i(0)$ is the empty sequence

MP3 for $i \in A$, $r_i(n+1)$ is either identical to $r_i(n)$ or results from appending an event to $r_i(n)$.

By MP1 the local state of each agent records the messages she has sent or received, so the system satisfies *perfect recall*. We assume also MP2, which enforces a unique initial state in the system.

We define message passing QIS (MPQIS) as the class of quantified interpreted systems $\mathcal{P} = \langle \mathcal{R}, \mathcal{D}, I \rangle$ where $\mathcal{R}$ is a non-empty set of runs defined as above, $\mathcal{D}$ contains the agents in A and the messages in *Msg*, and I is an interpretation for $\mathcal{L}_m$. By MP1-3 message passing QIS belong to

the class $\mathcal{QIS}^{pr,uis}$ of QIS with perfect recall and a unique initial state. In what follows we use the same notation for objects in the model and syntactic elements.

For the specification of MPQIS we introduce a predicative constant $Send$ such that $(\mathcal{P}^{\sigma}, r, n) \models Send(i, j, \mu)$ if event $send(i, j, \mu)$ occurs to agent i at time n in run r, i.e., $r_i(n)$ is the result of appending $send(i, j, \mu)$ to $r_i(n-1)$. Also, we introduce the predicate $Sent$ such that $(\mathcal{P}^{\sigma}, r, n) \models Sent(i, j, \mu)$ if event $send(i, j, \mu)$ occurs to agent i before time n in run r, i.e., $send(i, j, \mu)$ appears in $r_i(n)$. The predicates $Rec(j, \mu)$ and $Rec'ed(j, \mu)$ are similarly defined for event $rec(j, \mu)$. Finally, $Rec(i, j, \mu)$ is a shorthand for $Rec(j, \mu) \wedge Sent(i, j, \mu)$.

We briefly explore the range of specifications that can be expressed in this formalism. A property often required in MPQIS is that there are no "ghost" messages: if agent j receives a message μ, then j knows that μ must actually have been sent by some agent i. We can express this requirement as a monodic formula:

$$\forall\mu(Rec(j, \mu) \rightarrow K_j \exists i Sent(i, j, \mu)) \tag{1}$$

This specification is not satisfied by the present definition of MPQIS, but we can modify the protocol for the environment as follows in order to enforce the validity of formula (1) on MPQIS.

- $P(l_e) = \{\lambda, \mathrm{del}(j, \mu) \mid \mathrm{send}(j, \mu) \in l_e, j \in A, \mu \in Msg\}$

We compare (1) with a further relevant property of MPQIS, i.e., *authentication*: if agent i has received a message μ from agent j, then i knows that μ had actually been sent by j:

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

Note that, differently from (1), (2) is not monodic.

We now introduce the basic constructors to specify cryptographic protocols within the framework of MPQIS. Specifically, we model a security protocol as a MPQIS that exchanges encrypted messages. We assume atomic messages $m_1, m_2, \ldots$, nonces $N_i, N_i', \ldots$, and symmetric encryption keys $k_{ij}, \ldots$ for principals $i, j \in A$. The encrypted messages in the domain $\mathcal{D}$ are inductively defined as follows:

Definition 39 ($Term$)

$$\mu ::= m \mid N_i \mid k_{ij} \mid \mu, \mu' \mid fst(\mu) \mid snd(\mu) \mid \{\mu\}_{k_{ij}} \mid \{\mu\}_{k_{ij}^{-1}}$$

We have that μ, μ' is the concatenation of messages μ and μ', $fst(\mu)$ and $snd(\mu)$ are the first and second projection of μ respectively, while $\{\mu\}_{k_{ij}}$ is the encryption of message μ with the key k_{ij}. Similarly, the decryption function k_{ij}^{-1} applies to an encrypted message μ to return a decrypted message $\{\mu\}_{k_{ij}^{-1}}$.

We now introduce equational cryptographic theories to reason about the meaning of encrypted messages.

Definition 40 *An equational cryptographic theory (e.c.t.) is a couple $\mathcal{E} = \langle X, \equiv_X \rangle$ where (i) $X \subseteq Term$, and (ii) $\equiv_X$ is an equivalence relation on X such that for $\mu, \mu' \in$*

X *(i)* $fst(\mu, \mu') \equiv_X \mu$; *(ii)* $snd(\mu, \mu') \equiv_X \mu'$; *(iii)* $\{\{\mu\}_{k_{ij}}\}_{k_{ij}^{-1}} \equiv_X \mu$. An e.c.t. $\mathcal{E}$ is clear for $i \in A$ if for all $\mu \in \mathcal{E}$, $\{\mu\}_{kij} \equiv_X \mu$.

We extend the definition of local state for a principal $i \in A$ by adding an e.c.t. $\mathcal{E}_i = \langle X_i, \equiv_i \rangle$ to every l_i. Further, for $s = \langle l_e, l_1, \ldots, l_m \rangle$ let $\mathcal{E}_s = \langle X_s, \equiv_s \rangle$ for $X_s = \bigcup_{i \in A} X_i$ and $\equiv_s = \bigcup_{i \in A} \equiv_i$ be the e.c.t. for the state s. Thus, we have

$$(\mathcal{P}^{\sigma}, r, n) \models \mu \equiv \mu' \quad \text{if} \quad I^{\sigma}(\mu), I^{\sigma}(\mu') \in X_{(r,n)}$$
$$\text{and} \quad I^{\sigma}(\mu) \equiv_{(r,n)} I^{\sigma}(\mu')$$

Notice that the formula

$$\forall\mu, \mu'(\mu \equiv \mu' \rightarrow K_i(\mu \equiv \mu'))$$

does not hold in general. This is a desirable property of MPQIS, since it expresses the limits of one agent's knowledge as to the meaning of encrypted messages.

The language $\mathcal{L}_m$ is suitable for specifying a wealth of properties of security protocols. Since our language $\mathcal{L}_m$ does not contain functors, we define some shorthands in order to simplify the notation of specifications in the next section. First, we introduce a predicative constant $Conc$ such that

$$(\mathcal{P}^{\sigma}, r, n) \models Conc(\mu, \mu', \mu'') \quad \text{if} \quad \mu'' \equiv \mu, \mu'$$

By using $Conc$ we can define what it means for a message μ' to be the first (resp. second) projection of a term μ:

$$fst(\mu) = \mu' \quad ::= \quad \exists\mu'' Conc(\mu', \mu'', \mu)$$
$$snd(\mu) = \mu' \quad ::= \quad \exists\mu'' Conc(\mu'', \mu', \mu)$$

Further, we introduce a predicative constant Enc to express message encryption:

$$(\mathcal{P}^{\sigma}, r, n) \models Enc(\mu, \mu', k) \quad \text{if} \quad \mu \equiv \{\mu'\}_k$$

We normally write $\mu \equiv \{\mu'\}_k$ instead of $Enc(\mu, \mu', k)$.

Now we can define what it means that a principal knows a cryptographic key, i.e., *a principal j knows a key k iff she knows the identites of every message encrypted with k*:

$$K_j Key(k) ::= \forall\mu, \mu'(\mu \equiv \{\mu'\}_k \rightarrow K_j(\mu \equiv \{\mu'\}_k)) \tag{3}$$

The concepts here introduced will be useful in the analysis in the next section.

The Otway-Rees protocol

In this section we apply the formal machinery developed thus far to the analysis of the Otway-Rees protocol (Otway and Rees 1987). This is a shared-key authentication protocol, in which two principals A and B use a trusted server S to generate a session key k_{AB}. Further, k_{AS} is the key shared between A and S, k_{BS} is shared between B and S, N_A and N_B are nonces, and μ_{AB} is the primitive message whose intuitive meaning is "A *wants to communicate with* B". We represent the protocol in the Alice-Bob notation as follows:

$$
\begin{aligned}
A \rightarrow B \quad &: \quad \mu_{AB}, \{N_A, \mu_{AB}\}_{k_{AS}} \\
B \rightarrow S \quad &: \quad \{N_A, \mu_{AB}\}_{k_{AS}}, \{N_B, \mu_{AB}\}_{k_{BS}} \\
S \rightarrow B \quad &: \quad \{N_A, k_{AB}\}_{k_{AS}}, \{N_B, k_{AB}\}_{k_{BS}} \\
B \rightarrow A \quad &: \quad \{N_A, k_{AB}\}_{k_{AS}}
\end{aligned}
$$

Principal A sends B the encrypted message $\{N_A, \mu_{AB}\}_{k_{AS}}$ together with enough information for B (i.e. the message μ_{AB}) to send a similar encrypted message to S. Principal B forwards $\{N_A, \mu_{AB}\}_{AS}$ to S together with the encrypted message $\{N_B, \mu_{AB}\}_{BS}$. When S receives the message he checks whether the components μ_{AB}, N_A, N_B, A and B match in the encrypted messages. If this is the case, S generates a new session key k_{AB}, encrypts it with k_{AS} and k_{BS}, then sends both messages to B, who forwards the appropriate part to A. Finally, A and B decript the messages, check the nonces and use k_{AB} as the new session key.

We can represent the Otway-Reese protocol as a MPQIS in which the protocol is defined as follows:

- $P_A(\epsilon) = \text{send}(B, (\mu_{AB}, \{N_A, \mu_{AB}\}_{k_{AS}}))$
- $P_B(\langle rec(B, (\mu_{AB}, \mu'))\rangle) = \text{send}(S, (\mu', \{N_B, \mu_{AB}\}_{k_{BS}}))$
- $P_S(\langle rec(S, (\{N_A, \mu_{AB}\}_{k_{AS}}, \{N_B, \mu_{AB}\}_{k_{BS}}))\rangle) = \text{send}(B, (\{N_A, k_{AB}\}_{k_{AS}}, \{N_B, k_{AB}\}_{k_{BS}}))$
- $P_B(\langle rec(B, (\mu, \{N_B, k_{AB}\}_{k_{BS}}))\rangle) = \text{send}(A, \mu)$

Let $\mathcal{P}$ include all runs consistent with the definitions above, and let r^* be the run such that for all $n \in \mathbb{N}$ the e.c.t. in $r_i^*(n)$ is clear for $i \in A$. Intuitively, r^* is the run in which each agent knows the meaning of the messages encrypted with her keys. Also, let the environment e model a Dolev-Yao intruder I, which can eavesdrop all the communications between A, B and S.

In the first step of the Otway-Rees protocol principal A sends the message $\mu_{AB}, \{N_A, \mu_{AB}\}_{k_{AS}}$ to B. Let $Init$ be a propositional constant such that $(\mathcal{P}, r, n) \models Init$ iff $r_j(n)$ is the empty list for all $j \in A \cup \{e\}$. We represent the first step by means of the following specification:

$$Init \rightarrow Send(A, B, (\mu_{AB}, \{N_A, \mu_{AB}\}_{k_{AS}}))) \qquad (4)$$

We can check that formula (4) holds in the QIS $\mathcal{P}$ representing the Otway-Rees protocol.

The message $\{N_A, \mu_{AB}\}_{k_{AS}}$ is the encryption of (N_A, μ_{AB}) with key k_{AS}. Therefore, $\{N_A, \mu_{AB}\}_{k_{AS}} \equiv (N_A, \mu_{AB})$ holds in the e.c.t. at $r_A^*(1)$ as it is clear for A. Principal A knows the key k_{AS}, hence by (3) A knows the meaning of the encrypted message $\{N_A, \mu_{AB}\}_{k_{AS}}$:

$$(\mathcal{P}, r, 1) \models K_A(\{N_A, \mu_{AB}\}_{k_{AS}} \equiv (N_A, \mu_{AB}))$$

In the second step of the protocol principal B receives a message $(\mu_{AB}, \{N_A, \mu_{AB}\}_{k_{AS}})$ and forwards it to the server S after appending the encrypted message $\{N_B, \mu_{AB}\}_{k_{BS}}$. We represent this step by means of the following specification:

$$\forall \nu \; (Rec(B, \nu) \wedge fst(\nu) \equiv \mu_{AB} \rightarrow$$
$$\rightarrow \bigcirc Send(B, S, (\nu, \{N_B, \mu_{AB}\}_{k_{BS}})))$$

We assumed that the MPQIS representing the Otway-Rees protocol does not validate version (1) of authentication, hence B does not know the identity of the sender. This means that he might consider a point (r', n) where the local state of B is the same as in $r^*(2)$, i.e., $r_B'(n) = r_B^*(2)$, while the intruder I has delivered the message to B pretending to be A. Thus, the following specification is not satisfied:

$$\forall \nu (Rec(B, (\mu_{AB}, \nu)) \rightarrow K_B Sent(A, B, (\mu_{AB}, \nu)))$$

Since B does not know the key shared between A and S, he cannot decrypt the message sent by A. Thus, at $(r^*, 2)$ B does not know the meaning of $\{N_A, \mu_{AB}\}_{k_{AS}}$:

$$(\mathcal{P}, r^*, 2) \not\models K_B(\{N_A, \mu_{AB}\}_{k_{AS}} \equiv (N_A, \mu_{AB}))$$

In the third step, when the server S receives the message from B, he checks whether the components μ_{AB}, N_A, N_B, A and B match in the encrypted messages, then sends the encrypted keys to B.

$$\forall \nu ((Rec(S, \nu) \wedge K_S(fst(\nu) \equiv (N_A, \mu_{AB})) \wedge$$
$$\wedge K_S(snd(\nu) \equiv (N_B, \mu_{AB}))) \rightarrow$$
$$\rightarrow \bigcirc Send(S, B, (\{N_A, k_{AB}\}_{k_{AS}}, \{N_B, k_{AB}\}_{k_{BS}}))) (5)$$

Also, S knows that the messages were actually sent by A and B:

$$\forall \nu ((Rec(S, \nu) \wedge K_S(fst(\nu) \equiv (N_A, \mu_{AB})) \wedge$$
$$\wedge K_S(snd(\nu) \equiv (N_B, \mu_{AB}))) \rightarrow$$
$$\rightarrow K_S Sent(A, B, (N_A, \mu_{AB})) \wedge$$
$$\wedge K_S Sent(B, S, (N_B, \mu_{AB})))$$

If S knows that the messages were actually sent by A and B, then he generates a new session key k_{AB}, encrypts it with k_{AS} and k_{BS}, and sends both messages to B. We represent this by the following specification:

$$(K_S Sent(A, B, (N_A, \mu_{AB})) \wedge$$
$$\wedge K_S Sent(B, S, (N_B, \mu_{AB})) \rightarrow$$
$$\rightarrow \bigcirc Send(S, B, (\{N_A, k_{AB}\}_{k_{AS}}, \{N_B, k_{AB}\}_{BS})))$$

On the other hand, if the intruder I had eavesdropped the communication between A and B, he could have sent a message $(\{N_A, \mu_{AB}\}_{k_{AS}}, \nu)$ to S pretending to be B. This situation is represented by a state $r''(n'')$ such that $r_A''(n'') = r_A^*(3)$, $r_B''(n'')$ is empty and the local state of the server is defined as follows:

$$r_S''(n'') = \langle rec(S, \{N_A, \mu_{AB}\}_{k_{AS}}, \nu)\rangle$$

By checking nonces and keys as specified in (5) the server S understands that she received a message from an individual (i.e. I) different from the one principal A wants to communicate with (i.e. B). Since the precondition $K_S(\nu \equiv (N_B, \mu_{AB}))$ in (5) is not satisfied, the server does not distribute keys to the principals.

Assuming no intruder eavesdropped messages and the protocol went on smoothly, in the fourth step B receives the key from S, checks the message, and forwards the appropriate part to A:

$$\forall \nu \; (Rec(B, \nu) \wedge K_B(snd(\nu) \equiv (N_B, k_{AB}))) \rightarrow$$
$$\rightarrow \bigcirc Send(B, A, fst(\nu)))$$

In fact, $Rec(B, (\{N_A, k_{AB}\}_{k_{AS}}, \{N_B, k_{AB}\}_{k_{BS}}))$ holds at $(r^*, 4)$ and by (3) B knows the encrypted message sent by S:

$$(\mathcal{P}, r^*, 4) \models K_B(\{N_B, k_{AB}\}_{k_{BS}} \equiv (N_B, k_{AB}))$$

Finally, in the fifth step A receives the message from B, check the nonce, and knows that k_{AB} is the new session key:

$$Rec(A, \{N_A, k_{AB}\}_{k_{AS}}) \rightarrow K_A(\{N_A, k_{AB}\}_{k_{AS}} \equiv (N_A, k_{AB}))$$

This completes our analysis of the Otway-Rees protocol.

Conclusions and Further Work

In this paper we presented a number of classes of quantified interpreted systems satisfying conditions such as synchronicity, no learning, perfect recall, and unique initial state, which are relevant for applications in real scenarios. In Theorem 19 we proved that the sets of monodic validities in these classes of QIS are axiomatisable. These results extend previous contributions on pure first-order epistemic and temporal logic (Sturm, Wolter, and Zakharyaschev 2000; Wolter and Zakharyaschev 2002) and propositional temporal epistemic logic (Halpern, Meyden, and Vardi 2003).

Still, further work is required in this line of research. On the temporal dimension, it would be of interest to pursue an analysis of CTL modalities interpreted on quantified interpreted systems. In this area there are contributions on pure branching time logic (Gabbay et al. 2003; Hodkinson, Wolter, and Zakharyaschev 2002). On the epistemic dimension, it would be of interest for applications to add epistemic operators for group knowledge. We anticipate that common knowledge is likely to lead to increased complexity, as this already happens at the propositional level (Halpern and Vardi 1989). Both dimensions need to be explored in more detail. Finally, it seems worthwhile to explore the issues pertaining to the decidability of the logics here discussed. An obvious starting point here are the results in (Hodkinson, Wolter, and Zakharyaschev 2000; Wolter and Zakharyaschev 2001).

References

Belardinelli, F., and Lomuscio, A. 2008. A complete first-order logic of knowledge and time. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Eleventh International Conference*, 705–714. AAAI Press.

Belardinelli, F., and Lomuscio, A. 2009a. First-order linear-time epistemic logic with group knowledge: An axiomatisation of the monodic fragment. In *Logic, Language, Information and Computation, 16th International Workshop, WoLLIC 2009*, 140–154. Springer.

Belardinelli, F., and Lomuscio, A. 2009b. Quantified epistemic logics for reasoning about knowledge in multi-agent systems. *Artificial Intelligence* 173(9-10):982–1013.

Cohen, P., and Levesque, H. 1995. Communicative actions for artificial agents. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, 65–72. AAAI Press.

Deutsch, A.; Sui, L.; and Vianu, V. 2004. Specification and verification of data-driven web services. In *Proceedings of the Twenty-third ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 71–82. ACM.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning about Knowledge*. Cambridge: MIT Press.

Gabbay, D.; Kurucz, A.; Wolter, F.; and Zakharyaschev, M. 2003. *Many-Dimensional Modal Logics: Theory and Applications*, volume 148 of *Studies in Logic*. Elsevier.

Hallé, S., and Villemaire, R. 2009. Browser-based enforcement of interface contracts in web applications with beep-beep. In *Computer Aided Verification, 21st International Conference, CAV 2009*, 648–653. Springer.

Halpern, J., and Moses, Y. 1992. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence* 54:319–379.

Halpern, J. Y., and Vardi, M. Y. 1989. The complexity of reasoning about knowledge and time 1: lower bounds. *Journal of Computer and System Sciences* 38(1):195–237.

Halpern, J.; Meyden, R.; and Vardi, M. Y. 2003. Complete axiomatisations for reasoning about knowledge and time. *SIAM Journal on Computing* 33(3):674–703.

Hodkinson, I. M.; Kontchakov, R.; Kurucz, A.; Wolter, F.; and Zakharyaschev, M. 2003. On the computational complexity of decidable fragments of first-order linear temporal logics. In *10th International Symposium on Temporal Representation and Reasoning (TIME-ICTL 2003)*, 91–98. IEEE Computer Society.

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

Hodkinson, I. M.; Wolter, F.; and Zakharyaschev, M. 2002. Decidable and undecidable fragments of first-order branching temporal logics. In *17th IEEE Symposium on Logic in Computer Science*, 393–402. IEEE Computer Society.

Hodkinson, I. 2006. Complexity of monodic guarded fragments over linear and real time. *Annals of Pure and Applied Logic* 138:94–125.

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

Otway, D., and Rees, O. 1987. Efficient and timely mutual authentication. *Operating Systems Review* 21(1):8–10.

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

Rao, A., and Georgeff, M. 1991. Deliberation and its role in the formation of intentions. In *Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence*, 300–307. Morgan Kaufmann Publishers.

Sturm, H.; Wolter, F.; and Zakharyaschev, M. 2000. Monodic epistemic predicate logic. In *Logics in Artificial Intelligence, European Workshop*, 329–344. Springer.

Wolter, F., and Zakharyaschev, M. 2001. Decidable fragments of first-order modal logics. *Journal of Symbolic Logic* 66(3):1415–1438.

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

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