

A logic of situated resource-bounded agents

Natasha Alechina and Brian Logan

University of Nottingham

School of Computer Science

Nottingham NG8 1BB, UK

{nza,bsl}@cs.nott.ac.uk

Abstract. We propose a framework for modelling situated resource-bounded agents. The framework is based on an objective ascription of intentional modalities and can be easily tailored to the system we want to model and the properties we wish to specify. As an elaboration of the framework, we introduce a logic, *OBA*, for describing the observations, beliefs, goals and actions of simple agents, and show that *OBA* is complete, decidable and has an efficient model checking procedure, allowing properties of agents specified in *OBA* to be verified using standard theorem proving or model checking techniques.

Keywords: agents, belief ascription, verification

1. Introduction

A major goal in intelligent agent research is the formal modelling of agent-environment systems. Such an account is key both in deepening our understanding of the notion of agency, e.g., the relationships between agent architectures, environments and behaviour, and for the principled design of agent systems. A common approach is to model the agent and its environment in some logic and prove theorems about the agent's behaviour in that logic. It is perhaps most natural to reason about the behaviour of the agent in an epistemic logic; epistemic notions such as knowledge and belief provide a compact and powerful way of reasoning about the structure and behaviour of agents (McCarthy, 1978), and there has been a considerable amount of work in this area, for example, (Levesque, 1984; Lakemeyer, 1986; Parikh, 1987; Fagin and Halpern, 1988; Rao and Georgeff, 1991; Fagin et al., 1995; Moore, 1995; van der Hoek et al., 1999; Singh, 1999; Wooldridge and Lomuscio, 2001; van Ditmarsch et al., 2007).

Much of this work has focused on postulating properties of agents in general, such as the relationships between beliefs, desires and intentions investigated by Rao and Georgeff (1991). While useful, such an approach can only provide very general guidance to the agent designer, since it abstracts away from the specifics of particular agent-environment systems. However, many interesting logical properties of agents depend on the agent's architecture and program and the environment in which it is situated. In addition, from a technical point of view, existing work often makes strong assumptions which can limit its applicability when considering feasible (i.e., implementable) agent



© 2008 Kluwer Academic Publishers. Printed in the Netherlands.

designs, e.g., the assumption that agents are logically omniscient (Hintikka, 1962). While this is a reasonable assumption in some situations, for example where the agent is capable of only a very restricted class of inferences (Alechina and Logan, 2002), there are many cases where the time taken to do deliberation is of critical importance. Practical agents take time to derive the consequences of their beliefs, and, in a dynamic environment, the time required by an agent to derive the consequences of its observations will determine whether such derivations can play an effective role in action selection.

In this paper, we propose a logical framework for modelling agent-environment systems. We adopt an explicitly design-oriented view in the sense that our framework makes only minimal assumptions about agents in general, and those assumptions we do make are motivated by consideration of feasible agent designs. For example, we assume that agents will have a finite state and will require time to perform inferences. We assume the agents execute a ‘sense-think-act’ cycle and consider beliefs and goals of the agent at various points in the cycle; this allows us to express fine-grained properties of the agent’s beliefs, such as ‘after sensing, the agent’s observational beliefs are always correct’.

The remainder of the paper is organised as follows. In sections 2 and 3 we present our modelling framework and develop a model of a simple agent-environment system within the framework. In section 4 we introduce a new logic based on these ideas, *OBA*, which can be used to model a resource-bounded agent’s observations, beliefs, goals and actions, and state some complexity results for *OBA*. In section 5 we illustrate our approach with a simple example based on the well known Tileworld domain (Pollack and Ringuette, 1990) and show how *OBA* can be used to specify properties of a Tileworld agent. In section 6 we discuss related work, and in section 7 we conclude with some ideas for future work.

2. The Agent-Environment Model

In this section we present a logical framework for modelling agent-environment systems based on state transition systems. We first describe how to ascribe beliefs to agents on the basis of the contents of their state and then outline how to specify properties of agent-environment systems on the basis of transitions between states.

The state of an agent-environment system, w , consists of two parts: the environment state $e(w)$ and the agent’s state $s(w)$.¹ The environment state is a description of the (physical or computational) environment in which the

¹ For simplicity, we consider only a single agent.

agent is situated. The agent's state contains all the internal representations which determine the behaviour of the agent. We assume that some parts of the agent's state can be interpreted as referring to (real or hypothetical) objects or events in the environment, e.g., that there is an obstacle dead ahead, or to properties of the agent itself, e.g., the level of the agent's battery. The combined agent-environment states are assumed to be finite (i.e. can be specified using finitely many variables) and we assume that there are finitely many of them.

We describe the properties of the environment in a language built from a set of propositional variables \mathcal{P} . Beliefs ascribable to the agent come from a finite set of literals (propositional variables or their negations) $\mathcal{L}_b = \{(\neg)p_1, \dots, (\neg)p_m\}$ (where $p_i \in \mathcal{P}$). Following, for example Rosenschein and Kaelbling (1995), we assume that the agent's state consists of finitely many 'memory locations' $\mathbf{l}_1, \dots, \mathbf{l}_n$, and that each location \mathbf{l}_i can contain (exactly) one of finitely many values, v_{i_1}, \dots, v_{i_k} . For example, we could have a location \mathbf{l}_t for the output of a temperature sensor which may take an integer value between -50 and 50. Based on those values, we can ascribe beliefs about the external world to the agent: for example, based on $\mathbf{l}_t = 20$ we ascribe to the agent a belief that the outside temperature is 20 C. Each $\phi \in \mathcal{L}_b$ corresponds to the fact that a given memory location \mathbf{l}_i (or set of memory locations) has a given a set of values, but 'translates' this into a statement about the world. We assume a mapping Bel assigning to each state $s(w)$ of the agent a set of propositional variables and their negations which form beliefs of the agent in state $s(w)$. Note that this 'translation' is fixed and does not depend on the truth or falsity of the formulas in the real world. In general, there is no requirement that $Bel(s(w))$ be consistent; if a propositional variable and its negation are associated with two different memory locations (e.g., in an agent which has two temperature sensors) then the agent may simultaneously believe that p and $\neg p$. $Bel(s(w))$ does not have to map a single value to a single belief, for example, all values of $\mathbf{l}_t > 20$ could be mapped to a single belief that it's "warm". Conversely, we do not assume that for every propositional variable $p \in \mathcal{P}$, either p or $\neg p$ belong to $Bel(s(w))$; if a location \mathbf{l}_i has no value (e.g., if a sensor fails) or has a value that does not correspond to any proposition, then the agent may have no beliefs about the outside world at all. Other intentional notions such as goals can be modelled analogously to beliefs, i.e., by introducing an explicit translation from the contents of the agent's state into the set of goals. We elaborate belief and goal ascription using the notion of a memory location rather than assuming that agents have an internal representation of beliefs or goals e.g., as a list of literals, for reasons of generality. The ascription mechanism described above is applicable to arbitrary agents, not only those with an explicit internal representation of beliefs and goals.

Our aim is to model the transitions of the agent-environment system as a kind of Kripke structure and express properties of the agent in some modal logic. We consider transition systems similar to the interpreted systems of Fagin et al. (1995), except that the agent’s beliefs are modelled as a local property of the agent’s state.² To be more precise, the structure consists of a set of agent-environment states W , each state $w \in W$ has an agent part $s(w)$ and an environment part $e(w)$. The state of the agent $s(w)$ comes equipped with a set of formulas $Bel(s(w))$ corresponding to the agent’s beliefs; the state of the environment $e(w)$ corresponds to a classical possible world, or a complete truth assignment to propositional variables in \mathcal{P} . The agent believes that p in state w , $w \models Bp$, if $p \in Bel(s(w))$. Note that this truth definition for B does not give rise to any interesting logical properties of B , e.g. to KD45 axioms. This is intentional: we do not want our agents to be logically omniscient and the logical properties of agent’s beliefs should be determined by the agent’s architecture and program.

Within this basic set up, we can express some properties of the agent in a suitable logic, such as CTL (Clarke and Emerson, 1981), and use a standard model checker such as SMV to check whether the properties are true of the agent-environment system. For example, we may want to check whether the agent always has true beliefs about a property p , i.e., whether $AG(Bp \rightarrow p)$. The set of states is generated by the agent’s program together with appropriate environment responses. Note that the check for beliefs is done locally in state w and consists of checking whether a given formula is in the set $Bel(s(w))$, so the problem can be easily solved by a standard model checking techniques (e.g., with beliefs encoded as propositional variables).

However such an approach is too coarse grained even for analysing simple agents. For example, if p is a property observable by the agent and observation of p is always reliable, then the agent will have true beliefs about p immediately after it has performed an observation. However, this does not guarantee that the agent will always have true beliefs about p in every state— if the agent doesn’t sense its environment continuously, it may have false beliefs about p in some states, for example, in states resulting from a change in the environment. We want to be able to express the fact that in all states resulting from the agent performing an observation, Bp is true only if p is true. The most natural way to express this is to label transitions of the system by ‘moves’ of the agent and the environment. We have chosen a variant of the Propositional Dynamic Logic PDL (Pratt, 1976) because it allows us to reason about labelled transition systems (another option would have been for example Multi-Modal CTL (Ågotnes et al., 2007)).

² Unlike in (Fagin et al., 1995), where beliefs are modelled using an equivalence relation on the set of agent’s states

In the remainder of this paper we work out this idea for a simple agent-environment system in which the agent senses the environment, updates its state (in particular, selects an action to perform), and sends the chosen action to the environment, which in turn changes according to one of the possible outcomes of the action (actions are assumed to be nondeterministic).

3. A Simple Agent-Environment System

In what follows we view the agent and its environment as a pair of interacting automata (cf (Rosenschein and Kaelbling, 1995)) performing a sense-think-act cycle (see Figure 1). At each cycle, the environment generates an input or ‘percept’ to the agent (updating several locations in the agent’s memory), followed by the agent generating an output or ‘action’ to the environment. We assume that the environment “computes” its response to the agent’s action instantaneously and that the agent produces its response to the input from the environment (chooses an action) in bounded time. We model perception and action as two non-deterministic transitions, *obs* and *act*. An *obs* transition takes the state of the environment and updates the agent’s state with a percept corresponding to the information returned by the agent’s sensors. We assume that perception either returns accurate information about the environment or one of finitely many outcomes of failed perception. The action transition *act* computes the new state of the environment given the current state of the environment and the action selected by the agent. An action either causes the environment to change in the desired way (the action succeeds), or results in one of finitely many outcomes for a failed action. If the response generated by the environment depends on both the agent’s action and the amount of time it took the agent to produce it, this approximates an asynchronous interaction between the agent and a dynamic environment. Intuitively, if the agent spends too much time selecting an action, then performing it does not produce the expected result.

The agent consists of some state and an internal transition *inf* which models the computation the agent uses to update its state (e.g. update its beliefs and select an action to perform). *inf* is assumed to depend on the agent’s percept at this cycle and its state from the previous cycle and to terminate in bounded time. This is a reasonable assumption for the types of agents which we consider in this paper (such as the simple Tileworld agent described in section 5).

As above, we describe the properties of the environment in a language built from a set of propositional variables \mathcal{P} , with a finite set \mathcal{L}_b of literals which can be ascribed to the agent as beliefs and goals. In addition to the mapping *Bel* introduced in the previous section, we define a mapping *Obs* which takes the agent’s state $s(w)$ and returns a set of formulas, as a restriction of




Figure 1. Simple agent-environment system

Bel to the locations holding percepts. Since observational beliefs are a subset of all beliefs, for every agent state $s(w)$ we have $Obs(s(w)) \subseteq Bel(s(w))$. Analogously we ascribe a set of goals $Goal(s(w))$ to the agent based on a designer-specified translation of the agent's state.

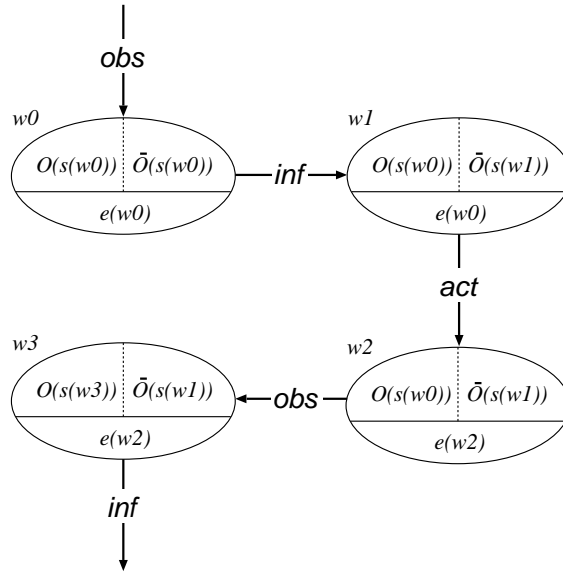


Figure 2. World Transitions, where $O(s(w)) = Obs(s(w))$ and $\bar{O}(s(w)) = -Obs(s(w))$, the set of non-observational beliefs. $Bel(s(w)) = Obs(s(w)) \cup -Obs(s(w))$.

The cycle of the combined agent-environment automaton consists of the three transitions: obs , inf and act (see Figure 2). At each cycle the environment updates the agent's state $s(w)$ with its percept for this cycle. The agent then derives the consequences of its new beliefs (if any) and performs an action on the basis of this new state. The agent's observational beliefs,

$Obs(s(w))$, only change after the *obs* transition and the rest of the agent's beliefs and goals only change after the *inf* transition. We assume that *inf* does not change the agent's observational beliefs $Obs(s(w))$, as these correspond to the agent's 'sensor readings', and, as such, should not be overwritten by inference.³ We further assume that the environment only changes in response to the agent's action.

Introducing three separate transitions rather than collapsing them into one single step (corresponding to a full *obs*, *inf* and *act* cycle) allows more precise modelling of changes in the agent's beliefs and goals. We want to be able to talk about the state of the agent after it has performed the observation but before it completed all the planning necessary to choose an action, or after it has performed an action but before it has sensed its effects. This is essential in analysing a resource-bounded agent which may be capable of selecting the "right" action to perform but is unable to do so before the environment changes so as to invalidate the action's preconditions.

4. A Logic for Agent-Environment Systems

In this section, we give a formal definition of the transition graphs of the agent-environment system described in the previous section, and a logical language to reason about the agent's observations, beliefs and goals.

We describe transition graphs in a logic which we call *OBA* (for 'Observation, Belief and Action'). *OBA* includes PDL (Propositional Dynamic Logic, see (Pratt, 1976)) with three atomic transition labels *o*, *i* and *a* corresponding to *obs*, *inf* and *act*, and in addition contains observation, belief and goal operators. We define a transition term x as $o \mid i \mid a \mid x; y \mid x^*$, where if x and y are transition terms then $x; y$ is their sequential composition (transition x followed by y) and x^* stands for 0 or finitely many iterations of x . (We could have added other constructs such as non-deterministic choice \cup but we do not need them for axiomatising the logic below.)

In addition to the set of propositional variables \mathcal{P} , a set of literals \mathcal{L}_b which are ascribable to the agent as beliefs and goals, and boolean connectives, the language of *OBA* contains:

- for each transition term x , a unary modality $\langle x \rangle$. $\langle x \rangle \phi$, where ϕ is any formula, stands for 'from here we can make an x transition after which ϕ

³ This limitation is not overly restrictive: the *obs* transition is non-deterministic (perception is not guaranteed to veridical) and the agent may still derive incorrect beliefs about the world on the basis of its percepts. Note however that those beliefs will be expressed by different literals, since observational and non-observational beliefs are disjoint. For example, the agent may have an observational belief 'Collision sensor indicates that there is no obstacle in front of me' and an inferential belief 'There is an obstacle in front of me' (which prevented the agent from moving forward).

holds'. A dual modality $[x]\phi$ is defined as $\neg\langle x\rangle\neg\phi$ and stands for 'after all x transitions we can make from here, ϕ holds'.

- unary modalities O , I and G , which can only be applied to elements of \mathcal{L}_b . $O\phi$ stands for 'the agent observes that ϕ ', $I\phi$ stands for 'the agent has a non-observational belief ϕ ', and $G\phi$ stands for 'the agent has a goal that ϕ ', $\phi \in \mathcal{L}_b$.

The belief modality B is defined as follows: $B\phi =_{df} O\phi \vee I\phi$.

In the models below we only allow legal states of the system as possible worlds, i.e., states which are possible given the characteristics of environment and the agent's program. We assume that beliefs are assigned to states as described in section 2

DEFINITION 1. *An OBA model M is a tuple $\langle W, V, Obs, -Obs, Goal, obs, inf, act \rangle$ where*

W is a set of worlds. Each world w is a pair $(s(w), e(w))$ where $s(w)$ is the state of the agent and $e(w)$ the state of the environment.

$Obs(s(w)) \subseteq \mathcal{L}_b$ is a finite set of observational beliefs associated with $s(w)$,

$-Obs(s(w)) \subseteq \mathcal{L}_b$ is a finite set of non-observational beliefs associated with $s(w)$; $Bel(s(w))$, the set of agent's beliefs at $s(w)$, is $Obs(s(w)) \cup -Obs(s(w))$.

$Goal(s(w)) \subseteq \mathcal{L}_b$ is a finite set of goals associated with $s(w)$;

V assigns true or false to pairs $(e(w), p)$ where $e(w)$ is a state of the environment and $p \in \mathcal{P}$.

obs, inf, act are binary relations on W . To say that a world v is reachable from a world w by an atomic transition x (e.g., obs) we will use the notation $w \xrightarrow{x} v$ (e.g., $w \xrightarrow{obs} v$).

The following restrictions on obs, inf and act hold:

Availability *For every w , exactly one of the following is true: $\exists v(w \xrightarrow{obs} v)$, $\exists v(w \xrightarrow{inf} v)$, $\exists v(w \xrightarrow{act} v)$*

Order *On all paths in the transition system, atomic transitions succeed each other in the order $obs; inf; act$. (We require this property since we want to model the sense-think-act cycle of the agent).*

Change(obs) *If $w \xrightarrow{obs} w'$ then $Goal(s(w)) = Goal(s(w'))$, $-Obs(s(w)) = -Obs(s(w'))$ and $e(w) = e(w')$. (Observation may change the percepts $Obs(s(w))$ but not the environment, non-observable beliefs or goals.)*

Change(*inf*) If $w \xrightarrow{inf} w'$ then $Obs(s(w)) = Obs(s(w'))$ and $e(w) = e(w')$. (Inference does not change the environment or percepts.)

Change(*act*) If $w \xrightarrow{act} w'$ then $Bel(s(w')) = Bel(s(w))$ and $Goal(s(w')) = Goal(s(w))$. (The agent's beliefs and goals do not change as a consequence of action.)⁴

Before we give the truth definition for OBA, we define the relations R_x corresponding to transition terms x inductively as follows:

$$R_o = obs, R_i = inf, R_a = act$$

$$R_{x;y} = \{(w, w') : \exists v (R_x(w, v) \wedge R_y(v, w'))\}$$

$$R_{x^*} = \{(w, w') : \exists n \geq 0 \exists v_0 \dots \exists v_n (w = v_0 \wedge R_x(v_0, v_1) \wedge \dots \wedge R_x(v_{n-1}, v_n) \wedge v_n = w')\}$$

The relation of a formula ϕ being true at a world w in M , $M, w \models_{OBA} \phi$ (or, for simplicity when no confusion can arise $M, w \models \phi$) is defined as follows:

$$M, w \models p \text{ iff } V(e(w), p) = true;$$

$$M, w \models \neg\psi \text{ iff } M, w \not\models \psi$$

$$M, w \models \psi \wedge \chi \text{ iff } M, w \models \psi \text{ and } M, w \models \chi$$

$$M, w \models \langle o \rangle \psi \text{ iff there exists a world } w' \text{ such that } w \xrightarrow{obs} w' \text{ and } M, w' \models \psi.$$

$$M, w \models \langle i \rangle \psi \text{ iff there exists a world } w' \text{ such that } w \xrightarrow{inf} w' \text{ and } M, w' \models \psi.$$

$$M, w \models \langle a \rangle \psi \text{ iff there exists a world } w' \text{ such that } w \xrightarrow{act} w' \text{ and } M, w' \models \psi.$$

$$M, w \models \langle x; y \rangle \psi \text{ iff there exists a world } w' \text{ such that } R_{x;y}(w, w') \text{ and } M, w' \models \psi.$$

$$M, w \models \langle x^* \rangle \psi \text{ iff there exists a world } w' \text{ such that } R_{x^*}(w, w') \text{ and } M, w' \models \psi.$$

$$M, w \models O\psi \text{ iff } \psi \in Obs(s(w));$$

$$M, w \models I\psi \text{ iff } \psi \in -Obs(s(w));$$

$$M, w \models G\psi \text{ iff } \psi \in Goal(s(w));$$

⁴ We could have added a condition that the outcome of an action depends only on the state of the environment (and not on the state of the agent) but in the current setting something like 'if $e(w) = e(w')$ then $\{v : w \xrightarrow{act} v\} = \{v : w' \xrightarrow{act} v\}$ ' is false because we abstract from the actual action performed during the *act* transition, and different actions would result in different sets of reachable states).

A formula ϕ is *OBA-satisfiable* if there is an *OBA* model M and a world w in M such that $M, w \models \phi$. A formula ϕ is *OBA-valid* ($\models_{OBA} \phi$) if its negation is not satisfiable.

THEOREM 1. *The axiom system below is sound and weakly complete for OBA, that is, for every formula ϕ ,*

$$\vdash_{OBA} \phi \text{ iff } \models_{OBA} \phi.$$

$\vdash_{OBA} \phi$ stands for ϕ is either one of the axioms below or is obtained from axioms by application of inference rules given below:

Classical propositional logic:

CL *Axioms of classical propositional logic;*

MP $\vdash \phi, \phi \rightarrow \psi \implies \vdash \psi$

Normal modal logic axioms for $[x]$, where x is any transition

K $[x](\phi \rightarrow \psi) \rightarrow ([x]\phi \rightarrow [x]\psi)$

N $\vdash \phi \implies \vdash [x]\phi$

Axioms for composition and iterations

C $[x; y]\phi \leftrightarrow [x][y]\phi$

It1 $\langle x^* \rangle \phi \leftrightarrow (\phi \vee \langle x \rangle \langle x^* \rangle \phi)$

It2 $[x^*](\phi \rightarrow [x]\phi) \rightarrow (\phi \rightarrow [x^*]\phi)$

Availability and order of transitions

T1 $(\langle o \rangle \top \wedge \neg \langle i \rangle \top \wedge \neg \langle a \rangle \top) \vee (\langle i \rangle \top \wedge \neg \langle o \rangle \top \wedge \neg \langle a \rangle \top) \vee (\langle a \rangle \top \wedge \neg \langle i \rangle \top \wedge \neg \langle o \rangle \top)$

T2 $[o]\langle i \rangle \top \wedge [i]\langle a \rangle \top \wedge [a]\langle o \rangle \top$

Changes after transitions

C(obs) $\phi \rightarrow [o]\phi$, where ϕ does not contain subformulas of the form $O\psi$

C(inf) $\phi \rightarrow [i]\phi$, where ϕ does not contain subformulas of the form $I\psi$ or $G\psi$

C(act) $X\phi \rightarrow [a]X\phi$, where $X \in \{O, I, G\}$.

Note that there are no axioms connecting beliefs and goals, in particular no requirement that the agent does not intend what it already believes (cf Rao and Georgeff (1991)). We will see in the next section that it is possible for a resource bounded agent to both believe and have a goal that ϕ , e.g., while it is updating its goals.

Proof. To prove soundness, we need to show that all axioms are valid and rules preserve validity. The argument for standard axioms and rules such as **CL**, **MP**, **K**, **N**, **C**, **It1**, **It2** is omitted. The axiom **T1** is valid because from every state w , exactly one of the three transitions is possible. The axiom **T2** is valid because if it is possible to make an x transition from w then from the resulting state it is possible to make a y transition, provided $(x, y) \in \{(o, i), (i, a), (a, o)\}$. Validity of **C(obs)** follows from **Change(obs)**. We can show by induction on subformulas χ of ϕ built from propositional variables and subformulas of the form $I\psi$ and $G\psi$ that $M, w \models \chi$ iff $M, w' \models \chi$ where w' is accessible from w by *obs*. Analogously for **C(inf)** and **C(act)**.

Now we prove completeness. Let ϕ be an *OBA*-consistent formula. We will construct a finite satisfying model for ϕ in a standard way (see, for example, Blackburn et al. (2001)). Let $Cl(\phi)$ be the Fisher-Ladner closure of the set of subformulas of ϕ together with three formulas $\langle o \rangle \top$, $\langle i \rangle \top$ and $\langle a \rangle \top$. The satisfying model for ϕ , M_ϕ , is defined as follows:

The set of worlds W_ϕ is the set of all maximal consistent subsets of $Cl(\phi)$

$$V_\phi(e(w), p) = \text{true iff } p \in w$$

$$Obs_\phi(s(w)) = \{\psi : O\psi \in w\}$$

$$\neg Obs_\phi(s(w)) = \{\psi : I\psi \in w\}$$

$$Goal_\phi(s(w)) = \{\psi : G\psi \in w\}$$

$w \xrightarrow{obs} v$ in M_ϕ iff the conjunction of formulas in w , which we denote by \hat{w} , is consistent with $\langle o \rangle \hat{v}$, similarly for *inf* and *act*.

For every regular expression π composed from *obs*, *inf* and *act* using sequential composition and finite iteration, we define the corresponding accessibility relation as $S_\pi(w, v)$ iff the formula $\hat{w} \wedge \langle \pi \rangle \hat{v}$ is *OBA*-consistent. The proof that relations defined this way do indeed correspond to sequential composition and finite iteration, as well as the existence lemma and the truth lemma, are standard. Note that the restriction on the language ensures that $Bel_\phi(s(w))$ and $Goal_\phi(s(w))$ only contain propositional variables and their negations.

We need to show that the special conditions on the *OBA* models hold: namely, that the transitions follow each other in the right way, and change the world in the way constrained by **Availability**, **Order**, **Change(obs)**, **Change(inf)** and **Change(act)**.

Let us first show that in each world, exactly one of the transitions *obs*, *inf*, *act* is possible. Note that the conjunction of formulas in each world is consistent with exactly one of $\langle o \rangle \top$, $\langle i \rangle \top$ and $\langle a \rangle \top$ (because the world is maximally consistent with respect to **T1** and contains some $\langle x \rangle \top$). This means that $\hat{w} \wedge \langle x \rangle \top$ for some x is consistent, and \top can be expanded into a maximally consistent conjunction of formulas \hat{v} from $Cl(\phi)$ by forcing choices; this gives us a world v accessible by x from w . For all other atomic transitions y , there is no such world v , because if $\hat{w} \wedge \langle y \rangle \hat{v}$ is consistent, then so is $\hat{w} \wedge \langle y \rangle \top$, but w is inconsistent with **T1**.

Next, let us show that the **Order** property holds. As an example, we show that *obs* cannot be followed by *act*: proofs for the other conditions are similar. Suppose some \hat{w} is consistent with $\langle o \rangle \hat{v}$ and \hat{v} is consistent with $\langle a \rangle \hat{u}$. Note that in this case \hat{v} contains $\langle a \rangle \top$. So it follows that $\langle o \rangle \langle a \rangle \top$ is consistent, but this contradicts **T2**.

To prove that **Change**(*obs*) holds, we need to show that whenever $\hat{w} \wedge \langle o \rangle \hat{v}$ is consistent, the state of the environment is the same in w and v , that is for every propositional variable p , $p \in w$ iff $p \in v$. Suppose $p \in w$ but $p \notin v$. Then it follows that $p \wedge \langle o \rangle \neg p$ is consistent, but it contradicts **C**(*obs*). Similarly for the case $p \notin w$ but $p \in v$. Non-observational beliefs are the same in w and v , that is $\neg Obs_\phi(s(w)) = \neg Obs_\phi(s(v))$: this follows from $I\psi \in w$ iff $I\psi \in v$, by **C**(*obs*). Similarly for the set of goals in w and v .

The case of **Change**(*inf*) is analogous.

Finally, **Change**(*act*) holds if $Obs_\phi(s(w))$, $\neg Obs_\phi(s(w))$ and $Goal(s(w))$ do not change as a result of an *act* transition. In other words, if $\hat{w} \wedge \langle a \rangle \hat{v}$ is consistent, then $X\phi \in w$ iff $X\phi \in v$, where $X \in \{O, I, G\}$. But this is exactly what the axiom **C**(*act*) guarantees. Note that the form of the axiom is as it is because *act* may change the truth value of arbitrary formulas, including non-modal formulas, but cannot change the agent's beliefs and goals. \square

THEOREM 2. *The satisfiability problem for OBA is decidable.*

Proof. Decidability follows from the bounded model property for OBA which can be established as a by-product of Theorem 1: every satisfiable formula ϕ has a model of size at most $2^{|\phi|}$. This gives the following NEXPTIME decision procedure: guess a model of size $2^{|\phi|}$ and check whether ϕ is true there. \square

In fact, the complexity of the decision procedure can be improved to EXPTIME analogously to PDL (see Blackburn et al. (2001)).

THEOREM 3. *Given a formula ϕ and a model, state pair M, w there is an $O(|M| \times |\phi|)$ algorithm for checking whether $M, w \models_{OBA} \phi$.*

Proof. This follows from the result on complexity for model checking for PDL (Emerson and Lei, 1986). In addition to PDL modalities we have modal operators O , I and G in the language, but checking whether $O\psi$, $I\psi$ or $G\psi$ is true in a world w can be done in constant time; subformulas of the form $O\psi$, $I\psi$ and $G\psi$ can be treated as propositional variables. \square

In *OBA*, we can express properties of a particular agent design and verify whether they hold as expected. There are two possible ways of doing this. We can axiomatise the agent's program and the interaction between the agent and its environment, and prove that the property follows from those extra axioms and the *OBA* axioms. Alternatively, given the agent's program and all possible responses to the agent's actions by the environment, we can generate the state transition graph and check whether the property is true in all states in the graph. Both approaches are feasible given the decidability and low cost of model checking of *OBA*.

In this paper, we are concerned with axiomatising phases in the agent's execution cycle rather than properties of a particular inference mechanism or of the agent's actions. This limitation can easily be relaxed. Namely, we can consider agents capable of performing a finite set of actions a_1, \dots, a_n , and define an action transition a as $a_1 \cup \dots \cup a_n$ (union of basic actions). In addition to the general conditions on the models above, we would add pre- and postconditions of all the basic actions, and the corresponding axioms to the logic. Similarly, we could add more conditions on the inference mechanism of the agent in the semantics and add corresponding axioms. For example, if an agent's program has a rule 'if p is in the set of beliefs, then adopt a goal q ' then the semantic condition on *inf* would include a corresponding condition on assignments before and after *inf* and the logic will have an axiom $Bp \rightarrow \langle i \rangle Gq$.

5. An example: the Tileworld

In this section we illustrate our approach with an example based on the well-known Tileworld domain (Pollack and Ringuette, 1990). The Tileworld is a testbed for the evaluation of agent architectures, and has been used to investigate agent commitment strategies (when an agent should abandon its current goal and replan) (Pollack et al., 1994) and in comparisons of reactive and deliberative agent architectures (Pollack and Ringuette, 1990). The Tileworld environment consists of a rectangular grid. Each square contains either a 'tile' or a 'hole' or is empty (i.e., contains neither a tile or a hole). The goal of a Tileworld agent is to fill the holes in the grid with tiles. In the original Tileworld, tiles and holes appear and disappear at random, and the agent's success is measured by the number of holes it is able to fill with

tiles in a given amount of (real or simulated) time. For reasons of brevity, our Tileworld is very simple. The environment consists of an unbounded strip of grid squares. Each grid square contains either a tile or a hole or is empty, and tiles and holes are distributed randomly along the strip of grid squares.⁵ The environment contains a single agent which can move forward on rails laid over the squares and has a sensor which allows it to see whether the square directly underneath it contains a tile or a hole or is empty. If it sees a tile it can pick it up and carry it. If it is holding a tile and is located above a hole it can drop it in the hole. Holes are one tile deep and dropping a tile into a hole ‘fills it in’, turning the hole into an empty square. For simplicity we have assumed that the agent’s sensors are veridical in the sense that they either return correct information about the environment or an ‘undefined’ value indicating the sensor returned no data at this cycle. We also assume that the actions of picking up or dropping a tile can fail, leaving the world unchanged; moving forward always succeeds and leaves the agent above the next square. We stress that these assumptions are not essential, and *OBA* can be used to reason about more complex environments and more realistic agents with unreliable perception and more complex actions failures.

The goal of the Tileworld agent is put a tile in a hole. Informally, the program of the agent consists of the following simple rules:

- if holding a tile and above a hole, repeatedly try to drop the tile until the tile is in the hole;
- if not holding a tile and above a tile, repeatedly try to pick it up until the square below the agent is empty;
- otherwise, move forward.

The agent’s state consists of four memory locations, **p**, **g**, **h**, and **a**:

p is for storing percepts, and can take one of four values: 0 for the agent is above a tile, 1 for the agent is above a hole, 2 for the agent is above an empty square and 3 for ‘undefined’.

h is for storing information about whether the agent is holding a tile: 1 for the agent is holding a tile, and 0 for the agent is not holding a tile. For the sake of the example, we assume that there is no sensor which can determine whether an agent is holding a tile, but instead this information has to be inferred by the agent.

g is for storing goals, and can take one of four values: 0 for looking for a tile, 1 for picking up a tile, 2 for looking for a hole, and 3 for dropping a tile.

⁵ Unlike the original Tileworld (Pollack and Ringuette, 1990), there are no obstacles and tiles and holes do not appear and disappear randomly, i.e., the environment is static.

\mathbf{a} is for storing selected actions, and can also take one of four values: 0 for pick up the tile beneath the agent, 1 for drop the tile the agent is holding, 2 for move forward one grid cell, and 3 for a special ‘no-op’ action which does not change the environment.

The execution cycle of the agent then becomes:

obs : sensing the environment to find out what is underneath the agent and updating the \mathbf{p} location;

inf : updating the values of \mathbf{h} , \mathbf{g} and \mathbf{a} according to the agent’s program; and

act : sending the value in \mathbf{a} to the environment.

To ascribe beliefs and goals to the Tileworld agent, we do not need a distinct propositional variable for every combination of location and value. The set $\mathcal{P}_{TL} = \{p_0, p_1, p_2, h, a_0, a_1, a_2, a_3\}$ is sufficient to capture those beliefs and goals which are relevant to the choice of the agent’s actions. To simplify fully automated belief ascription, we could use a propositional variable for every combination of location and value.

For the sake of brevity, we do not specify *Obs*, *Bel* and *Goal* completely, but illustrate them via examples. Let p_0 mean ‘the agent is above a tile’.

$p_0 \in Obs(s(w))$ iff in $s(w)$, $\mathbf{p} = 0$ (the agent observes that it is above a tile); similarly for p_1 and p_2 .

$\neg p_0 \in Obs(s(w))$ iff in $s(w)$, $\mathbf{p} = 1$ or $\mathbf{p} = 2$ (the agent observes that it is not above a tile); similarly for $\neg p_1$ and $\neg p_2$.

$p_0 \in Goal(s(w))$ iff in $s(w)$, $\mathbf{g} = 0$ (the agent has a goal to be above a tile); similarly for p_1 and $\mathbf{g} = 2$.

Note that it is possible for example for neither p_0 nor $\neg p_0$ to be in $Obs(s(w))$ if the agent’s sensors return no data at this cycle ($\mathbf{p} = 3$).

The agent’s program is given by the following rules:

- $r1. \mathbf{p} = 1, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} \neq 0 \implies \mathbf{p} = 1, \mathbf{h} = 1, \mathbf{g} = 3, \mathbf{a} = 1$
 $r2. \mathbf{p} = 2, \mathbf{h} = 1, \mathbf{g} = 3, \mathbf{a} = \{1, 3\} \implies \mathbf{p} = 2, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = 2$
 $r3. \mathbf{p} = 1, \mathbf{h} = 1, \mathbf{g} = 3, \mathbf{a} = \{1, 3\} \implies \mathbf{p} = 1, \mathbf{h} = 1, \mathbf{g} = 3, \mathbf{a} = 1$
 $r4. \mathbf{p} = 0, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = \{2, 3\} \implies \mathbf{p} = 0, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = 2$
 $r5. \mathbf{p} = 2, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = \{2, 3\} \implies \mathbf{p} = 2, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = 2$
 $r6. \mathbf{p} = 3, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = \{2, 3\} \implies \mathbf{p} = 3, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = 2$
 $r7. \mathbf{p} = 3, \mathbf{h} = 1, \mathbf{g} = 3, \mathbf{a} = \{1, 3\} \implies \mathbf{p} = 3, \mathbf{h} = 1, \mathbf{g} = 3, \mathbf{a} = 3$
 $r8. \mathbf{p} = 0, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} \neq 1 \implies \mathbf{p} = 0, \mathbf{h} = 0, \mathbf{g} = 1, \mathbf{a} = 0$
 $r9. \mathbf{p} = 2, \mathbf{h} = 0, \mathbf{g} = 1, \mathbf{a} = \{0, 3\} \implies \mathbf{p} = 2, \mathbf{h} = 1, \mathbf{g} = 2, \mathbf{a} = 2$
 $r10. \mathbf{p} = 0, \mathbf{h} = 0, \mathbf{g} = 1, \mathbf{a} = \{0, 3\} \implies \mathbf{p} = 0, \mathbf{h} = 0, \mathbf{g} = 1, \mathbf{a} = 0$
 $r11. \mathbf{p} = 1, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = \{2, 3\} \implies \mathbf{p} = 1, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = 2$
 $r12. \mathbf{p} = 2, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = \{2, 3\} \implies \mathbf{p} = 2, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = 2$
 $r13. \mathbf{p} = 3, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = \{2, 3\} \implies \mathbf{p} = 3, \mathbf{h} = 0, \mathbf{g} = 0, \mathbf{a} = 3$
 $r14. \mathbf{p} = 3, \mathbf{h} = 0, \mathbf{g} = 1, \mathbf{a} = \{0, 3\} \implies \mathbf{p} = 3, \mathbf{h} = 0, \mathbf{g} = 1, \mathbf{a} = 3$

where $\mathbf{a} = \{1, 3\}$ means the agent is in a state in which the value of location \mathbf{a} is either 1 or 3. The first seven rules cover the cases in which the agent holding a tile and is either searching for a hole to put it in, or has found a hole and is trying to drop the tile. Rules 8–14 cover the cases in which the agent is not holding a tile and is either searching for a tile, or trying to pick up a tile it has found. Rule 1 initiates a drop action when the agent is above a hole and holding a tile, rule 2 notices that the drop was successful and starts the agent looking for a new tile (rule 12), and rule 3 repeats the drop action if it was unsuccessful. Rules 4 and 5 cause the agent to continue searching (move forward) if it is holding a tile and not above a hole. Rules 8–12 are similar but handle the case where the agent is trying to pick up a tile. Rules 6 and 7 handle perception failures when the agent is looking for a hole and dropping a tile respectively, and rules 13 and 14 handle perception failures while looking for and picking up a tile. If the agent’s sensors return no data at this cycle (indicated by a value of 3 in location \mathbf{p}), the agent’s beliefs and goals don’t change and it selects no-op action, in the hope that sensing at the next cycle will return an informative value for \mathbf{p} .

OBA can be used to specify properties of the Tileworld agent that a designer may want to verify. For example, we can state that the agent will achieve a particular goal, e.g., finding a tile: $Gp_0 \rightarrow \langle (a; o; i)^* \rangle p_0$ (‘if the agent has a goal to be above a tile then after finitely many cycles it will be above a tile’). This statement holds in all states provided the sensors are veridical, sensor failures only last for finitely many cycles and it is always possible to reach a tile after finitely many moves forward. The veridicality of sensors can be expressed as $[o](O\phi \rightarrow \phi)$ (note that $O\phi \rightarrow \phi$ is only

guaranteed to hold after an *obs* transition).⁶ Another important property is commitment to goals: if the agent has a goal ϕ , it will not give it up until ϕ becomes true: $G\phi \rightarrow [a; o; i](\phi \vee G\phi)$. As we mentioned in the previous section, it is possible for the agent both to believe and intend the same formula: for example, if the agent's goal is to be above a tile and it has performed a move forward action which brought it above a tile, then after a successful *obs* transition it will believe that it is above a tile (Op_0 hence Bp_0), but will still have the goal of being above the tile (the goals will be updated after the *inf* transition); hence $B\phi \rightarrow \neg G\phi$ is not valid.

6. Related Work

Our approach has some similarities with the situated automata work of Rosenschein and Kaelbling (1995). Rosenschein and Kaelbling model an agent using two functions, a state update function, f , which maps the input (percepts) and the value of the agent's internal state at the last cycle into a new value of the internal state, and an output function, g , that maps the input and the value of the agent's internal state at the last cycle into the output (actions). They ascribe belief (or rather knowledge) to the agent by associating with every combination of location/value pair (l, v) the most informative proposition ϕ such that in all runs of the system when $l = v$ in the agent's state, ϕ holds. This forces the agent to have true beliefs and its beliefs are closed under logical consequence. The resulting notion corresponds to implicit knowledge as defined by Fagin et al. (1995), and satisfies the properties of an S5 modality. van der Hoek, Linder and Meyer (1999) formalise all relevant aspects of the agent's architecture and use PDL to reason about state transitions as we do. However they model agents as logically omniscient, consider only deterministic actions, and model knowledge and beliefs using an additional Kripke structure associated with every global state. In contrast, our approach combines the modelling of belief change and action performance in a single structure and is much simpler, giving a more efficient model-checking procedure.

A compact modular representation of agent-environment systems has been recently proposed by Jamroga and Ågotnes (2007); we believe our logic can be interpreted on such systems rather than on combined transition graphs.

Other related work on epistemic logic for agents includes, for example, Moore (1995), Singh (1999), Wooldridge and Lomuscio (2001) and recent work on dynamic epistemic logic (van Ditmarsch et al., 2007). However, as in (van der Hoek et al., 1999) they consider deterministic actions and a more traditional definition of knowledge using possible worlds semantics.

⁶ The condition that sensor failures only last for finitely many cycles cannot in be expressed in PDL in the non-deterministic case, but can be imposed as a condition on models.

Our work has more similarities with the approach of Konolige (Konolige, 1986) where beliefs are modelled as formulas, rather than as sets of possible worlds (propositions), and are not automatically closed under logical consequence. Other syntactic approaches to epistemic logic (syntactic belief and modelling steps of inference) include Step Logic (Elgot-Drapkin and Perlis, 1990). Step Logic provides a mechanism for reasoning about non-omniscient time-bounded reasoners; an agent's beliefs are indexed by time points or steps corresponding to stages in the agent's reasoning. Other recent approaches that avoid the problem of logical omniscience and model an agent's beliefs as syntactic objects are Duc's dynamic epistemic logic (Duc, 1995; Duc, 1997), Sierra's et al.'s (1996) work on logics for reflective architectures, and Ågotnes's logic of finite syntactic epistemic states (Ågotnes, 2004) and (Ågotnes and Alechina, 2007). The approach presented in this paper concentrates on the interaction of a resource-bounded agent with its environment, rather than on stepwise modelling of its inference process.

In other work, we have applied a syntactic approach to modelling agent's beliefs to verifying properties of an agent implemented in a simplified version of the 3APL agent programming language (Alechina et al., 2007). However, in that work we made a simplifying assumption that the agent's beliefs are veridical and actions successful, to avoid explicitly modelling the agent's environment. The current paper provides a way to overcome this limitation.

7. Summary and Future Work

In this paper, we propose a new approach to the formal modelling of agent-environment systems which focuses on particular agent designs. Properties of the agent-environment system depend on the agent's architecture and program and the characteristics of its environment rather than on a priori assumptions about agents in general, e.g., that agents are rational or logically omniscient. We show how to ascribe beliefs and other intentional modalities based on designer-stipulated correlations between the values of locations in the state of the agent and the state of the environment, and how to model operations within the agent and interactions between agent and environment as state transitions. Our approach to belief ascription is local in the sense that belief is not defined in terms of all possible runs of the agent-environment system (cf Fagin et al. (1995)) and the set of transitions can be tailored to the system we want to model and the properties we wish to specify. For example, to model the agent's deliberation more precisely, we can introduce a finer grain in the agent's state update transition.

Drawing on these ideas, we define a new logic, *OBA*, which can be used to model a resource bounded agent's observation, beliefs, goals and actions. We prove completeness and decidability results for *OBA* and show that it has

a reasonable model checking procedure ($O(|M| \times |\phi|)$, where $|M|$ is the size of the model and $|\phi|$ the size of the formula). *OBA* allows for both failed perception and failed actions, and the agent's beliefs are not required to be consistent or closed under logical consequence. We explicitly introduce the cycle of the agent in the logic and analyse the agent's beliefs at each point in the cycle, allowing us to distinguish the agent's beliefs, e.g., after a percept is received but before the rest of the state is updated. As an illustration, we model a simple Tileworld agent and express some properties of the resulting agent-environment system in *OBA*. Such properties can be efficiently checked using standard model checking or theorem proving techniques.

In future work we plan to develop the framework outlined above to analyse more complicated agent behaviours, for example, to model trial and error attempts to achieve a goal, or characterise robust behaviours.

Acknowledgements

The authors would like to thank Thomas Ågotnes and the anonymous referees for comments on an earlier version of the paper. The work was partially supported by EPSRC research grant number EP/E031226/1.

References

- Ågotnes, T.: 2004, 'A Logic of Finite Syntactic Epistemic States'. Ph.D. thesis, Department of Informatics, University of Bergen, Norway.
- Ågotnes, T. and N. Alechina: 2007, 'The Dynamics of Syntactic Knowledge'. *Journal of Logic and Computation* **17**(1), 83–116.
- Ågotnes, T., W. van der Hoek, and M. Wooldridge: 2007, 'Completeness and Complexity of Multi-Modal CTL'. In: C. Areces and S. Demri (eds.): *Proceedings of the 5th Workshop on Methods for Modalities (M4M5)*. pp. 75–90.
- Alechina, N., M. Dastani, B. Logan, and J.-J. C. Meyer: 2007, 'A Logic of Agent Programs'. In: *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence (AAAI 2007)*. pp. 795–800.
- Alechina, N. and B. Logan: 2002, 'Ascribing beliefs to resource bounded agents'. In: *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, Vol. 2. Bologna, pp. 881–888.
- Blackburn, P., M. de Rijke, and Y. Venema: 2001, *Modal Logic*, Vol. 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.
- Clarke, E. M. and E. A. Emerson: 1981, 'Design and Synthesis of Synchronization Skeletons Using Branching Time Temporal Logic'. In: D. Kozen (ed.): *IBM Logics of Programs Workshop*. pp. 52–71.
- Duc, H. N.: 1995, 'Logical Omniscience vs. Logical Ignorance on a Dilemma of Epistemic Logic'. In: C. A. Pinto-Ferreira and N. J. Mamede (eds.): *Progress in Artificial Intelligence, 7th Portuguese Conference on Artificial Intelligence, EPIA '95, Funchal, Madeira Island, Portugal, October 3-6, 1995, Proceedings*, Vol. 990 of *Lecture Notes in Computer Science*. pp. 237–248.

- Duc, H. N.: 1997, 'Reasoning about rational, but not logically omniscient, agents'. *Journal of Logic and Computation* **7**(5), 633–648.
- Elgot-Drapkin, J. J. and D. Perlis: 1990, 'Reasoning Situated in Time I: Basic Concepts'. *Journal of Experimental and Theoretical Artificial Intelligence* **2**, 75–98.
- Emerson, E. and C.-L. Lei: 1986, 'Efficient Model Checking in Fragments of the Propositional Mu-Calculus'. In: *Proceedings LICS'86*. pp. 267–278.
- Fagin, R. and J. Y. Halpern: 1988, 'Belief, awareness, and limited reasoning'. *Artificial Intelligence* **34**, 39–76.
- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi: 1995, *Reasoning about Knowledge*. Cambridge, Mass.: MIT Press.
- Hintikka, J.: 1962, *Knowledge and belief*. Ithaca, NY: Cornell University Press.
- Jamroga, W. and T. Ågotnes: 2007, 'Modular Interpreted Systems'. In: M. Huhns and O. Shehory (eds.): *Proceedings of the Sixth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*. pp. 892–899.
- Konolige, K.: 1986, *A Deduction Model of Belief*. San Francisco, Calif.: Morgan Kaufmann.
- Lakemeyer, G.: 1986, 'Steps towards a first-order logic of explicit and implicit belief'. In: J. Y. Halpern (ed.): *Theoretical Aspects of Reasoning About Knowledge: Proceedings of the 1986 Conference*. San Francisco, Calif., pp. 325–340.
- Levesque, H. J.: 1984, 'A logic of implicit and explicit belief'. In: *Proceedings of the Fourth National Conference on Artificial Intelligence, AAAI-84*. pp. 198–202.
- McCarthy, J.: 1978, 'Ascribing mental qualities to machines'. Technical report, Stanford AI Lab.
- Moore, R. C.: 1995, *Logic and Representation*, No. 39 in CSLI Lecture Notes. CSLI Publications.
- Parikh, R.: 1987, 'Knowledge and the Problem of Logical Omniscience'. In: *Methodologies for Intelligent Systems, Proceedings of the Second International Symposium*. pp. 432–439.
- Pollack, M. E., D. Joslin, A. Nunes, S. Ur, and E. Ephrati: 1994, 'Experimental Investigation of an Agent Commitment Strategy'. Technical Report TR 94–31, University of Pittsburgh, Pittsburgh, PA 15260.
- Pollack, M. E. and M. Ringuette: 1990, 'Introducing the Tileworld: Experimentally Evaluating Agent Architectures'. In: *Proceedings of the Ninth National Conference on Artificial Intelligence*. Boston, MA, pp. 183–189.
- Pratt, V. R.: 1976, 'Semantical considerations on Floyd-Hoare logic'. In: *Proceedings of the Seventeenth IEEE Symposium on Computer Science*. pp. 109–121.
- Rao, A. S. and M. P. Georgeff: 1991, 'Modeling rational agents within a BDI-architecture'. In: *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*. pp. 473–484.
- Rosenschein, S. J. and L. P. Kaelbling: 1995, 'A situated view of representation and control'. *Artificial Intelligence* **73**, 149–173.
- Sierra, C., L. Godo, R. L. de Mántaras, and M. Manzano: 1996, 'Descriptive dynamic logic and its application to reflective architectures'. *Future Gener. Comput. Syst.* **12**(2-3), 157–171.
- Singh, M. P.: 1999, 'Know-How'. In: M. Wooldridge and A. Rao (eds.): *Foundations of Rational Agency*. Dordrecht: Kluwer Academic, pp. 81–104.
- van der Hoek, W., B. van Linder, and J.-J. C. Meyer: 1999, 'An Integrated Modal Approach to Rational Agents'. In: M. Wooldridge and A. Rao (eds.): *Foundations of Rational Agency*. Dordrecht: Kluwer Academic, pp. 133–168.
- van Ditmarsch, H., W. van der Hoek, and B. Kooi: 2007, *Dynamic Epistemic Logic*, Vol. 337 of *Synthese Library*. Springer.
- Wooldridge, M. and A. Lomuscio: 2001, 'A Computationally Grounded Logic of Visibility, Perception, and Knowledge'. *Logic Journal of the IGPL* **9**(2), 273–288.