

Ideal agents sharing (some!) knowledge

Alessio Lomuscio¹ and Mark Ryan¹

Abstract. A well-known framework by Fagin, Halpern, Moses and Vardi models knowledge-based agents as “Interpreted Systems”. In this paper we analyse a particular class of interpreted systems, which we call hypercube systems, that share information among themselves. Hypercube systems arise by taking the full Cartesian product of the state space of interpreted systems. We analyse hypercube systems by taking their semantically equivalent Kripke frames and we present a sound and complete axiomatisation for them. The logic thus obtained, which we study in some detail, is a proper extension of the system $S5_n$, commonly studied for modelling knowledge for a community of ideal agents.

1 INTRODUCTION

The need for specifications of complex systems in AI, as in mainstream computer science, has brought forward the use of logic as formal tool for reasoning and proving properties about systems. In this respect, Multi-Agent Systems (MAS) constitute no exception and in the last thirty years many logics for modelling MAS have been proposed.

The design of a knowledge based agent is a central issue in agent theory, as knowledge is a key property of any intelligent system. Arguably the most successful approach is the modal logic $S5_n$, which has been used in Distributed Computing Theory by Halpern and Moses (see for example [8]), and others. $S5_n$ is the generalisation to a multi-agent scenario of the logic $S5$ which was originally proposed by Hintikka ([9]) in Philosophical Logic.

The logic $S5_n$ models a community of *ideal* knowledge agents. Ideal knowledge agents have, among others, the properties of veridical knowledge (everything they know is true), positive introspection (they know what they know) and negative introspection (they know what they do not know). The modal logic $S5_n$ (see for example [10] and [5]) can be axiomatised by taking all the propositional tautologies; the schemas of axioms $\Box_i(\phi \Rightarrow \psi) \Rightarrow \Box_i\phi \Rightarrow \Box_i\psi$, $\Box_i\phi \Rightarrow \phi$, $\Box_i\phi \Rightarrow \Box_i\Box_i\phi$, $\Diamond_i\phi \Rightarrow \Box_i\Diamond_i\phi$, where $i \in A$ represents an agent in the set of agents $A = \{1, \dots, n\}$; and the inference rules Modus Ponens and Necessitation.

The logic $S5_n$ has also been extended to deal with properties that arise when we investigate the state of knowledge of the group. Subtle concepts like common knowledge and distributed knowledge have been very well investigated ([2]). The logic $S5_n$ is a successful tool for the agent theorist also because, even in its extensions to common knowledge and distributed knowledge, it has important meta-properties like

closure under substitution, completeness and decidability (see for example [15]).

Notwithstanding this, there are some issues that epistemic theories based on $S5_n$ still do not address. One of the weakness of $S5_n$ is that all agents have equal status: the logic treats them in a uniform way. This is counterintuitive to the needs of AI, in which interesting scenarios arise when the agents are conceptually different. For example, consider a distributed system composed by a group of agents $A = \{1, \dots, n\}$ and the two following situations:

One agent knowing everything the others know. An agent j is the central processing unit of a distributed system of agents whose non-specialised entities transmit their knowledge to a central unit j .

Linear order in agents’ private knowledge. Several processes with the same information at disposal but different computational power are running the same program. Under certain assumptions, it is reasonable to assume that the knowledge of the agents, seen as knowledge bases, increases with the order of computational power at disposal.

These and similar scenarios can still be modelled by extensions of $S5_n$ in which *interaction axioms* are imposed. The first example of the description above can be modelled by the logic $S5_n$ plus the axiom schema

$$\Box_i\phi \Rightarrow \Box_j\phi, \forall i \in A.$$

The second scenario can be described by assuming an order on the set of agents $A = \{1, \dots, n\}$, reflecting their increasing computational power, and by taking $S5_n$ plus the axiom:

$$\Box_i\phi \Rightarrow \Box_j\phi, \text{ with } i \leq j.$$

These are just two isolated examples but there is actually a *broad spectrum* of possible specifications on how private states of knowledge are affected by other agents’ knowledge. At one end of the spectrum we have the system $S5$ in which all the knowledge bases are effectively equivalent. This can be modelled by taking an extension of $S5_n$ in which the axiom

$$\Box_i\phi \Leftrightarrow \Box_j\phi, \forall i, j \in A \quad (1)$$

holds. This is a very strong constraint. At the other end of the spectrum is simply $S5_n$, in which the only interesting interaction appears to be:

$$\vdash_{S5_n} \Box_i\phi \Rightarrow \Diamond_j\phi. \quad (2)$$

¹ School of Computer Science, University of Birmingham, UK

The formula above states that agents cannot rule out the possibility of a fact known by another agent. This is not surprising once we remember that all the known facts must be true at the real world. Quite surprisingly little work has been done to analyse systematically such interaction schemas - the only exceptions to this that we are aware of are [3], and [1] in which a limited class of interactions between the agents is proposed.

In this paper we isolate and study a special class of interpreted systems that can be modelled by an extension of $S5_n$ that falls into the above described spectrum. The systems we investigate, that we call hypercube systems or simply hypercubes, are defined by taking not an arbitrary subset (as interpreted systems are defined) but the *full* Cartesian product of the local states for the agents. We show that hypercubes not only satisfy the usual properties of veridical knowledge and complete introspection, but also an interesting interaction property: these agents do not rule out conjunctions of facts that are known by some other agent or considered possible by some agent to be known by some other agent. In this way, they share some knowledge among other agents in the community.

The paper is organised as follows: In Section 2 we recall the two standard semantics for MAS (Kripke models and interpreted systems), and we discuss hypercube systems. In Section 3 we present a sound and complete axiomatisation for the hypercubes. In Section 4 we explore this resulting logic and we provide an alternate axiomatisation which is more intuitive in terms of MAS. In Section 5 we draw our conclusions and we suggest further work.

This is technical paper - the reader is assumed to be familiar with the notation and the standard techniques of modal logic. Excellent introductions to the subject are [10] and [5]. The reader is also referred to [13] for the proofs of the results presented in this paper.

2 HYPERCUBE SYSTEMS

As far as logics for knowledge are concerned, two semantic treatments are available: *interpreted systems* and *Kripke models*. The two approaches have different advantages and disadvantages. On the one hand, interpreted systems are more intuitive to model real MAS, on the other hand Kripke models come with an heritage of fundamental techniques that help the user prove properties about her specification.

We briefly remind the key definitions of Kripke frames, interpreted systems and we define hypercube systems.

2.1 Kripke models

Kripke models ([12]) have been first formally proposed in Philosophical Logic and later used in Logic for AI as semantic structures for logics for belief, logics for knowledge, temporal logics, logics for actions, etc., all of which are modal logics. Over the last thirty years, many formal techniques have been developed for the study of modal logics grounded on Kripke semantics, such as completeness proofs via canonical models, decidability via the finite model property [10], and more recently, techniques for combining logics [11, 4].

The epistemic logic $S5_n$ is based on the notion of equivalence frame.

Definition 2.1 (Frames) A frame is a tuple $F = (W, R_1, \dots, R_n)$, where W is a non-empty set and for every i in A , R_i is a relation in $W \times W$. Elements of W are called worlds and are denoted as: w_1, w_2, \dots . If all relations are equivalence relations, the frame is an equivalence frame and we write \sim_i for R_i . \mathcal{F} denotes the class of equivalence frames.

(In the following we will denote by id_W the identity relation on W .)

2.2 Interpreted systems

Interpreted systems have first been proposed by Fagin, Halpern, Moses and Vardi [7] to model distributed systems. The growing interest in MAS and their specifications has brought forward the concept of interpreted system as useful formal tool to model key characteristics of the agents, such as the evolution of their knowledge, communication, etc. This work has culminated in the publication of [2] and [6] in which the authors use the notion of interpreted system to explore systematically fundamental classes of MAS (such as synchronous, asynchronous, with perfect recall ability, etc.) by the use of interpreted systems.

Interpreted systems can be defined as follows ([2]). Consider n sets of local states, one for every agent of the MAS, and a set of states for the environment.

Definition 2.2 (Global states of interpreted systems)

A set of global states for an interpreted system is a subset of the Cartesian product $S \subseteq L_e \times L_1 \times \dots \times L_n$, where L_e, L_1, \dots, L_n are non-empty sets. The set L_i represents the local states possible for agent i and L_e represents the possible states of the environment.

A global state represents the situation of all the agents and of the environment at a particular instant of time. The idea behind considering a subset is that some of the tuples that originate from the Cartesian product might not be possible because of explicit constraints present in the MAS. By considering functions (runs) from the natural numbers to the set of global states, it is possible to represent the temporal evolution of the system. An interpreted system is a set of functions on the global states with a valuation for the atoms of the language. Since here we carry out an analysis of the static properties of knowledge, we will not consider runs explicitly.

As it is shown in [2], interpreted systems can represent the knowledge of the MAS by considering two global states to be indistinguishable for an agent if its local state is the same in the two global states. Thus, a set of global states S denotes the Kripke frame $F = (W, \sim_1, \dots, \sim_n)$, if $W = S$, and $(l_1, \dots, l_n) \sim_i (l'_1, \dots, l'_n)$, if $l_i = l'_i, i \in A$.

2.3 Hypercube systems

Given n sets of local states for the agents of the MAS, the interpreted systems we analyse in this paper and that we call hypercube systems or hypercubes, result by considering the admissible state space of the MAS to be described by the full Cartesian product of its sets of local states. This means that every global state is in principle possible, i.e. there are no mutually exclusive configurations between local states. In

these cases the state space of the system is the whole full Cartesian product of the sets of local states for the agents².

With hypercubes we are imposing a further simplification on the notion presented in Definition 2.2: in the tuples representing the configuration of the system we do not consider a slot for the environment. While we will later observe that this restriction is not relevant in terms of the resulting interaction of private knowledge within the community, the following is worth pointing out. The presence of the environment in the notion of Fagin et al. is motivated in order to keep track of the changes in the system and in general to represent everything that cannot be captured by the local states of the single agents (most importantly messages in transit, etc.). By neglecting the dimension of the environment or, which comes to be the same thing, by treating it as a constant, we are projecting the notion of Fagin et al. of a time-dependent interpreted system to the product of its local states. Since we are focusing on a static case, in a way we can see this restriction as fixing the environment at the time in analysis, and investigate the possible configurations of the states of the agents. We now formally define hypercube systems.

Definition 2.3 (Hypercube systems) *A hypercube system, or hypercube, is a Cartesian product $H = L_1 \times \dots \times L_n$, where L_1, \dots, L_n are non-empty sets. The set L_i represents the local states possible for agent i . Elements of a local state L will be indicated with l_1, l_2, \dots . The class of hypercube systems is denoted by \mathcal{H} .*

Semantic structures which can be described by full Cartesian product can be proven to be semantically equivalent to particular Kripke frames, as it is shown in [14]. This particular class of frames is the following:

Definition 2.4 *Let \mathcal{G} be the class of equivalence frames that satisfy properties:*

1. $\bigcap_{i \in A} \sim_i = id_W$;
2. For any w_1, \dots, w_n in W there exists a \bar{w} such that $\bar{w} \sim_i w_i$, $i = 1, \dots, n$

We have the following result ([14]):

Theorem 1 *For all ϕ , $\mathcal{H} \models \phi$ if and only if $\mathcal{G} \models \phi$.*

Theorem 1 shows that hypercube systems are semantically equivalent to equivalence frames with properties as in Definition 2.4, and therefore an axiomatisation of the latter is also an axiomatisation of the former. In the next section we provide such an axiomatisation.

3 AXIOMATISATION OF HYPERCUBE SYSTEMS

In this section we aim to axiomatise the class of hypercube systems by axiomatising the class \mathcal{G} of frames.

We start by looking at the first property of Definition 2.4. By denoting with D_A the distributed knowledge of the group A of agents as it is defined in [2], we can prove the following:

² It has also been proved by Ron van der Meyden that there is a connection between Hypercubes and the states of knowledge in certain classes of broadcast systems (personal communication).

Lemma 3.1 *Consider an equivalence frame $F = (W, \sim_1, \dots, \sim_n)$.*

$\bigcap_{i \in A} \sim_i = id_W$ if and only if $F \models \phi \Leftrightarrow D_A \phi$.

The previous Lemma is a first surprising result: in the hypercubes the notion of distributed knowledge collapses to the truth of the formula. This follows from our simplification to consider global states that do not represent the environment. Notwithstanding this, the same result would have been achieved had we considered a MAS whose environment is constant, i.e. set of states for the environment composed by a singleton.

From now on we shall focus on a language that does not include a modal operator for distributed knowledge. Since the operator of distributed knowledge cannot be rewritten in terms of standard knowledge operators, in the following we will not be able to distinguish between extensions of $S5_n$ that satisfy property 1 and the ones which do not.

Property 2 of Definition 2.4 is more difficult to analyse. We begin by noting the following:

Lemma 3.2 *If F is a frame such that, $\forall w_1, \dots, w_n, \exists \bar{w}$ such that $\bar{w} R_i w_i$, $i = 1, \dots, n$, then $F \models \Diamond_i \Box_j \phi \Rightarrow \Box_j \Diamond_i \phi$, where $i \neq j$.*

Lemma 3.2 says that the agents described by hypercubes have the property that if agent i considers possible that agent j knows ϕ , then agent j knows that agent i considers ϕ to be possible. This is a constraint on the agents' knowledge because it implies that two agents i and j cannot be in a situation in which i considers that j might know a fact and j considers that i might know the negation of the same fact. We also notice that the formula expressed in Lemma 3.2 is not provable in $S5_n$. To see this, consider the following $S5_n$ counter-model (at w_1) for $\Diamond_1 \Box_2 p \Rightarrow \Box_2 \Diamond_1 p$, in which we do not indicate the reflexive relations:

$$M = (\{w_1, w_2, w_3\}, \{w_1 \sim_1 w_2, w_1 \sim_2 w_3\}, \pi(p) = \{w_1, w_2\}).$$

So, intuitively, the axiomatisation of the hypercube systems will have to be an extension of $S5_n$.

To understand better the property and its impact on the logic we need to proceed formally.

We call n-directedness (nD) property 2 of Definition 2.4.

Definition 3.1 (nD) *A frame $F = (W, R_i)$ is n-Directed (nD) if $\forall w_1, \dots, w_n, \exists \bar{w}$ such that $\bar{w} R_i w_i$, $i = 1, \dots, n$.*

Since property 1 of Definition 2.4 cannot be captured by the standard modal boxes, our aim is just to axiomatise nD equivalence frames. Using the standard $S5_n$ language, an axiomatisation for nD equivalence frames is also complete with respect to nD equivalence frames that satisfy property 1.

We know equivalence frames can be axiomatised by $S5_n$. The usual way to axiomatise a class of frames is to work out the correspondent axioms and try to prove completeness for that axiomatisation.

Unfortunately we note the following:

Lemma 3.3 *No modal formula corresponds to n-directedness.*

To find a correspondence result, we need to look at a weaker property. In the following P_n is the set of all the permutations of $\{1, \dots, n\}$, without fixed-points, i.e. if $(x_1, \dots, x_n) \in P_n$, then $x_i \neq i$.

Definition 3.2 (*nWD*) A frame $F = (W, R_i)$ is *n-weakly-directed* (*nWD*) if $\forall w, w_1, \dots, w_n \in W$, such that $wR_i w_i, i = 1, \dots, n$, and $\forall (x_1, \dots, x_n) \in P_n, \exists \bar{w}$ such that $w_i R_{x_i} \bar{w}$, for all $i = 1, \dots, n$.

When n is clear from the context we just use we refer to *nWD* just as *WD*. The property *nWD* for $n = 2$ was discussed in [16] and [1]; *nWD* is a generalisation of it.

It is immediate to note that:

Lemma 3.4 *If a frame is directed then it is weakly-directed.*

We analyse extensions of $S5_n$ with respect to the axiom:

$$\bigwedge_{(x_1, \dots, x_n) \in P_n} (\diamond_1 \square_{x_1} p_1 \wedge \dots \wedge \diamond_{n-1} \square_{x_{n-1}} p_{n-1}) \Rightarrow \square_n \diamond_{x_n} (\bigwedge_{i=1}^{n-1} p_i) \quad \mathbf{WD}$$

We have the correspondence result:

Lemma 3.5 $F \models \mathbf{WD}$ if and only if F is weakly-directed.

It is now possible to prove completeness:

Theorem 2 *The logic $S5_nWD$ is sound and complete with respect to the class of reflexive, symmetric, transitive and WD frames.*

As it is shown in [13], by means of a few laborious lemmas it is possible to strengthen Theorem 2 and prove the following:

Theorem 3 *The logic $S5_nWD$ is sound and complete with respect to the class of reflexive, symmetric, transitive and directed frames.*

Theorem 3 and Theorem 1 provide the axiomatisation of the hypercube systems in a multi-modal language.

4 THE LOGIC $S5_nWD$

In this section we try to explore the logic $S5_nWD$. In particular we present an equivalent formulation that can be interpreted more easily in terms of knowledge agents.

Let us start by analysing the type of constraint imposed by **WD** on the community of agents in the case $n = 2$:

$$\diamond_1 \square_2 p \Rightarrow \square_2 \diamond_1 p \quad \mathbf{2WD}$$

We do not need to make the other conjunct explicit, as this can be obtained by taking the contrapositive of **2WD**.

2WD can be read as “If it is conceivable for agent 1 that agent 2 knows the fact p , then agent 2 knows that p is conceivable for agent 1”. In other words, we are ruling out a situation in which agent 1 considers possible that agent 2 knows p , while agent 2 considers possible that agent 1 knows not p . We can say that **2WD** imposes a sort of homogeneity on the knowledge considered possible by other agents.

It is interesting to note that this constraint is logically equivalent to instantaneous message passing of possible knowledge of other agents. To see this, suppose that every time an agent considers possible that the other agent knows p , it broadcasts this information state, and this message is always safely received by the other agent. We also consider the communication to be always truthful. We have the axiom:

$$(\diamond_1 \square_2 p \Rightarrow \square_2 \diamond_1 \square_2 p) \wedge (\diamond_2 \square_1 q \Rightarrow \square_1 \diamond_2 \square_1 q) \quad \mathbf{2WD'}$$

In $S5_n$, we have the following:

Lemma 4.1 $\vdash_{S5_n} \mathbf{2WD} \Leftrightarrow \mathbf{2WD'}$

So, in the case of $n = 2$, the logic $S5_nWD$ specifies ideal agents of knowledge with an interaction between the two agents’ knowledge that can be simulated by truthful communication of possible knowledge between the agents.

Let us now analyse the case $n = 3$. We have:

$$\begin{aligned} (\diamond_1 \square_2 p_1 \wedge \diamond_2 \square_3 p_2 &\Rightarrow \square_3 \diamond_1 (p_1 \wedge p_2)) \wedge \\ (\diamond_1 \square_3 q_1 \wedge \diamond_2 \square_1 q_2 &\Rightarrow \square_3 \diamond_2 (q_1 \wedge q_2)) \end{aligned} \quad \mathbf{3WD}$$

The reading of the first conjunct of **3WD** is “If it is conceivable for agent 1 that agent 2 knows p_1 and it is conceivable for agent 2 that agent 3 knows p_2 , then agent 3 knows that p_1 and p_2 is conceivable for agent 1”.

Considering the first conjunct with the special cases of $p_2 = \top, p_1 = \top, p_2 = \neg p_1$, it can be checked that **3WD** implies the formulas: $\diamond_1 \square_2 p \Rightarrow \square_3 \diamond_1 p$, $\diamond_2 \square_3 p \Rightarrow \square_3 \diamond_1 p$, and $\diamond_1 \square_2 p \Rightarrow \square_2 \diamond_3 p$. More generally, it is easy to see that **3WD** implies $\diamond_i \square_j p \Rightarrow \square_k \diamond_l p$ ($i, j, k, l \in \{1, 2, 3\}$), that may be rewritten in the shape of **2WD'**.

Notwithstanding this, the intuition is that **3WD** is stronger than the simple conjunction of all these formulas. In fact, the semantic condition on the frames that corresponds to $\diamond_i \square_j p \Rightarrow \square_k \diamond_l p$ is $\forall w, w_1, w_2$ such that $wR_i w_1$ and $wR_k w_2$, $\exists \bar{w}$ such that $w_1 R_j \bar{w}$ and $w_2 R_l \bar{w}$. This condition appears to be weaker than $3WD$, that corresponds to **3WD**, and so it is reasonable to expect $\not\vdash_{S5_n} (\bigwedge_{j,k,l \in \{1,2,3\}} \diamond_i \square_j p \Rightarrow \square_k \diamond_l p) \Rightarrow \mathbf{3WD}$. In other words, it is unlikely that we can see the constraint imposed by **3WD** on the private knowledge of the agents as the result of a relatively simple message passing activity among the agents.

If we consider the case for arbitrary n , it is indeed quite hard to have a clear picture of the constraint imposed by the axiom **WD**; still this does impose a constraint on the information they possess. To understand better the implications of the logic $S5_nWD$ we axiomatise it in a slightly different way. Here we only sketch how this can be done and we refer to [13] for a technical presentation of the matter.

In Definition 3.2 we considered P_n to be the set of all the permutations (x_1, \dots, x_n) without fixed-points $x_i = i$. But it is possible to extend the results of the previous section and prove correspondence and completeness for the case of arbitrary permutations. It can then be observed that any permutation can be obtained by a sequence of swaps between two elements. By exploiting this observation it is possible to define a logic which is still equivalent to $S5_nWD$, and therefore complete with respect to the same class of frames, but whose interaction axiom can more easily be understood in terms of knowledge.

The logic in question, that we call $S5_nWD^*$, is $S5_n$ enriched by the following axiom:

$$\bigwedge_{1 \leq j, k, l \leq n} \bigwedge_{\substack{i=1 \\ i \neq j, i \neq k, i \neq l}}^{n-1} \left(\square_i p_i \wedge \diamond_j \square_k p_j \wedge \diamond_k \square_j p_k \Rightarrow \diamond_l (\bigwedge_{i=1}^{n-1} p_i) \right) \quad \mathbf{WD^*}$$

Theorem 4 *The logic $S5_nWD^*$ is equivalent to $S5_nWD$ and therefore sound and complete with respect to the class of reflexive, symmetric, transitive and WD frames.*

Theorem 4 permits us to reason about hypercubes by considering the logic $S5_nWD^*$. The advantage of WD^* over WD is that its meaning is more immediate in terms of the knowledge of the agents. In fact, axiom WD^* specifies a class of knowledge agents with the following property:

Observation 4.1 (Knowledge sharing in hypercubes)

Let $A = \{1, \dots, n\}$ be a community of agents modelled by an hypercube and consider any three agents, j, k, l . Then agent l thinks that the conjunction of

- anything that j thinks may be known by k ,
- anything that k thinks may be known by j ,
- anything known by any other agent,

is possible.

MAS modelled by hypercube systems share knowledge among themselves following Observation 4.1. They have full introspection capabilities and veridical knowledge, but they also consider possible facts known by some agents, or regarded by some agent to be possibly known by some other agent. Observation 4.1 specifies how private knowledge is shared in such community of agents. Note how the symmetry of the act of broadcasting is reflected in the symmetry of the interaction axiom.

5 CONCLUSIONS AND FURTHER WORK

In this paper we have formally investigated hypercube systems, a particular class of interpreted systems that arises by considering the full Cartesian product of the local states. We have argued that hypercubes model interpreted systems with no a priori constraints on their possible configurations. Thanks to a result that relates interpreted systems to Kripke frames, in order to axiomatise hypercubes we have analysed their images in a particular class of equivalence Kripke frames.

On the conceptual level the paper tries to offer two different contributions.

First, we have motivated that epistemic logics for MAS need to address the issue of interaction among agents' knowledge. We hope that in the near future $S5_nWD$ will be just one of the formal options available in the literature and that, thanks to more examples, a systematic study of types of interactions can be carried out.

Second, we have analysed the interesting class of nD equivalence frames. The variation of the canonical model technique we used to prove completeness heavily relies on the frames being reflexive, symmetric and transitive, properties guaranteed by the fact that we were analysing extensions of $S5_n$. The question is whether it is possible to prove similar results for weaker logics, such as $S4_n$, whose agents differently from $S5_n$ do not have negative introspection capabilities. In the future we will also analyse completeness for $S5_nWD\{\phi \Leftrightarrow D_A\phi\}$ with respect to equivalence frames with properties 1 and 2 of Definition 2.4, and we will try to prove decidability of $S5_nWD$ through finite model property.

Hypercube systems seem to be related to a special class of agents broadcasting information on a channel and this issue will be investigated shortly. In such a case the symmetry of the act of broadcasting would be reflected in the symmetry

of the interaction axiom WD . Future work also include an analysis of common knowledge in hypercube systems.

ACKNOWLEDGEMENTS

We wish to thank Ron van der Meyden for extensively commenting previous drafts of this paper. We also would like to thank Robert Goldblatt, Joe Halpern, Mathias Kegelmann and J.-J. Meyer for useful discussions on some of the issues addressed in this note.

We acknowledge partial financial support from the EU via Esprit WGs Fireworks and Aspire.

REFERENCES

- [1] L. Catach, 'Normal multimodal logics', in *Proceedings of the 7th National Conference on Artificial Intelligence*, ed., Tom M. Smith, Reid G.; Mitchell, pp. 491–495, St. Paul, MN, (August 1988). Morgan Kaufmann.
- [2] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning about Knowledge*, MIT Press, Cambridge, 1995.
- [3] Ronald Fagin, Joseph Y. Halpern, and Moshe Y. Vardi, 'What can machines know? on the properties of knowledge in distributed systems', *Journal of the ACM*, **39**(2), 328–376, (April 1992).
- [4] D. M. Gabbay, 'Fibred semantics and the weaving of logics. Part 1. Modal and intuitionistic logics', *Journal of Symbolic Logic*, **61**(4), 1057–1120, (December 1996).
- [5] R. Goldblatt, *Logics of Time and Computation, Second Edition, Revised and Expanded*, volume 7 of *CSLI Lecture Notes*, CSLI, Stanford, 1992 (first edition 1987). Distributed by University of Chicago Press.
- [6] J. Halpern, R. van der Meyden, and M. Y. Vardi, 'Complete axiomatisations for reasoning about knowledge and time'. submitted, 1997.
- [7] J. Y. Halpern and R. Fagin, 'A formal model of knowledge, action, and communication in distributed systems: Preliminary report', in *Proceedings of the Fourth ACM Symposium on Principles of Distributed Computing*, pp. 224–236, (1985).
- [8] Joseph Halpern and Yoram Moses, 'Knowledge and common knowledge in a distributed environment', *Journal of the ACM*, **37**(3), 549–587, (1990). A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.
- [9] J. Hintikka, *Knowledge and Belief, an introduction to the logic of the two notions*, Cornell University Press, Ithaca (NY) and London, 1962.
- [10] G. E. Hughes and M. J. Cresswell, *A new introduction to modal logic*, Routledge, New York, 1968.
- [11] M. Kracht and F. Wolter, 'Properties of independently axiomatizable bimodal logics', *The Journal of Symbolic Logic*, **56**(4), (1991).
- [12] S. A. Kripke, 'Semantic analysis of modal logic (abstract)', *Journal of Symbolic Logic*, **24**, 323–324, (1959).
- [13] A. Lomuscio and M. Ryan, 'Ideal agents sharing (some!) information', Technical Report CSR-98-1, School of Computer Science, University of Birmingham, UK, (January 1998). <ftp://ftp.cs.bham.ac.uk/pub/tech-reports/1998/CSR-98-01.ps.gz>.
- [14] A. Lomuscio and M. Ryan, 'On the relation between interpreted systems and Kripke models', in *Agent and Multi-Agent Systems - Proceedings of the AI97 Workshop on the theoretical and practical foundations of intelligent agents and agent-oriented systems*, eds., M. Pagnucco, W. R. Wobcke, and C. Zhang, volume 1441 of *Lecture Notes in Artificial Intelligence*. Springer Verlag, Berlin, (May 1998).
- [15] J.-J. Ch. Meyer and W. van der Hoek, *Epistemic Logic for AI and Computer Science*, volume 41, Cambridge University Press, 1995.
- [16] S. Popkorn, *First Steps in Modal Logic*, Cambridge University Press: Cambridge, England, 1994.