Knowledge in Multi-Agent Systems: Initial Configurations and Broadcast

Alessio Lomuscio
Department of Computing, Imperial College, London, United Kingdom
Ron van der Meyden
School of Computer Science and Engineering, University of New South Wales, Australia
and
Mark Ryan
School of Computer Science, University of Birmingham, Birmingham, United Kingdom

The semantic framework for the modal logic of knowledge due to Halpern and Moses provides a way to ascribe knowledge to agents in distributed and multi-agent systems. In this paper we study two special cases of this framework: *full systems* and *hypercubes*. Both model static situations in which no agent has any information about another agent's state. Full systems and hypercubes are an appropriate model for the initial configurations of many systems of interest. We establish a correspondence between full systems and hypercube systems and certain classes of Kripke frames. We show that these classes of systems correspond to the same logic. Moreover, this logic is also the same as that generated by the larger class of *weakly-directed frames*. We provide a sound and complete axiomatisation, $\mathrm{S5WD}_n$, of this logic. Finally, we show that under certain natural assumptions, in a model where knowledge evolves over time, $\mathrm{S5WD}_n$ characterises the properties of knowledge not just at the initial configuration, but also at all later configurations. In particular, this holds for *homogeneous broadcast systems*, which capture settings in which agents are initially ignorant of each others local states, operate synchronously, have perfect recall and can communicate only by broadcasting.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic and Formal Languages**]: Mathematical Logic—*Modal Logic*; I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Methods—*Modal Logic*; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms: Knowledge Representation in Distributed Systems, Modal Logic

Additional Key Words and Phrases: Completeness, Computational Complexity

A. Lomuscio acknowledges funding from the University of Birmingham, and from EU Framework V research project ALFEBIITE IST-1999-10298. Work of Ron van der Meyden supported by an Australian Research Council Large Grant.
Authors' addresses. A. Lomuscio, Department of Computing, Imperial College of Science, Technology and Medicine, 180 Queen's Gate, London SW7 2BZ, United Kingdom, e-mail: A.Lomuscio@doc.ic.ac.uk. R. van der Meyden, School of Computer Science and Engineering, University of New South Wales, Sydney 2052, Australia, e-mail: Meyden@cse.unsw.edu.au. M. Ryan, School of Computer Science, University of Birmingham, Edgbaston, Birmingham B15 2TT, United Kingdom, e-mail: M.D.Ryan@cs.bham.ac.uk.

1. INTRODUCTION

Modal logics of knowledge have been proposed as a formal tool for specifying and reasoning about multi-agent systems in a number of disciplines, including Distributed Computing [Halpern and Moses 1990], Artificial Intelligence [McCarthy 1989] and Economics [Aumann 1976; Rubenstein and Wolinsky 1990]. The logic most commonly applied in this area is the logic $S5_n$, a generalisation to a multi-agent setting of the logic S5 (see, e.g., [Hughes and Cresswell 1996] and [Goldblatt 1992]). The logic $S5_n$ includes for each $i = 1 \ldots n$ an operator $\Box_i$. The intended interpretation is that each $i = 1, \ldots, n$ represents an agent, and $\Box_i \phi$ expresses "agent i knows that ϕ." The logic $S5_n$ can be axiomatised by taking all the propositional tautologies; the axiom schemas $\Box_i(p \Rightarrow q) \Rightarrow \Box_i p \Rightarrow \Box_i q$, and $\Box_i p \Rightarrow p$, and $\Box_i p \Rightarrow \Box_i \Box_i p$, and $\neg \Box_i \neg p \Rightarrow \Box_i \neg \Box_i \neg p$, and the inference rules Modus Ponens, Necessitation and Uniform Substitution. The logic $S5_n$ has also been extended to deal with issues that arise when we investigate the state of knowledge of groups of agents. Subtle concepts like common knowledge and distributed knowledge have been investigated, as has the combination of the logic of knowledge and time (see [Fagin et al. 1995; Meyer and Hoek 1995] for extensive treatments of this literature.)

In $S5_n$, each agent's knowledge is independent of the knowledge of the other agents. In this paper, we investigate extensions of $S5_n$ which allow us to describe various kinds of dependencies. For example, consider a distributed system composed of a group of agents $A = \{1, \ldots, n\}$ and the following situations:

One agent knowing everything the others know. An agent j is the central librarian of a distributed system of agents that rely on j to maintain all their knowledge.

Linear order in agents' private knowledge. The agents operate within a chain of command subject to security restrictions. Each agent in the chain has a higher security clearance than the previous agent, and has access to a larger set of information sources.

These and similar scenarios can be modelled by extensions of $S5_n$ in which *interaction axioms* are imposed. Write $S_{i,j}$ for the axiom schema $\Box_i \phi \Rightarrow \Box_j \phi$. Then the first example above can be modelled by the logic $S5_n$ plus $S_{i,j}$ for all $i \in A$. The second scenario can be described by assuming an order on the set of agents reflecting their increasing information, and by taking $S5_n$ plus $S_{i,j}$ for all $i \leq j$. These are just two isolated examples; one may imagine *broad spectrum* of possible specifications on how private states of knowledge are affected by other agents' knowledge (see [Lomuscio and Ryan 2000] for a detailed exposition). At one end of the spectrum we have the system S5 in which all agents have the same knowledge. This can be modelled by taking an extension of $S5_n$ in which the axiom $S_{i,j}$ holds for *all* $i, j \in A$, making all the modalities collapse onto each other. This is a very strong constraint. At the other end of the spectrum is simply $S5_n$. Among others, Catach [Catach 1988] has studied a limited class of such interactions between knowledge of the agents.

In this paper, we work with the notion of *interpreted system* of Halpern and

Moses [Halpern and Moses 1990; Fagin et al. 1995]. We introduce and study two special cases of the interpreted systems model, which we call *full systems* and *hypercubes*. Both are systems in which *every* possible combination of individual agents' local states occurs in some global state in the system. In hypercubes we require additionally that every combination of state of the environment and the agents' local states occurs. Full systems and hypercubes are appropriate classes of systems for modelling the *initial configurations* of many systems of interest, in which no agent has any information concerning any other agent's state.

Full systems and hypercubes may be shown to satisfy an axiom that does not follow from $S5_n$. This axiom states in a quite intuitive fashion the property that every combination of the individual agent's local states occurs in some global state of the system. By characterising full systems and hypercubes in terms of certain classes of Kripke frames, we establish a sound and complete axiomatisation of the logic of knowledge in these classes of systems. Interestingly, the two classes correspond to the *same* logic, which we call $S5WD_n$. The nomenclature arises from the fact that we show that a further class of frames, the *weakly-directed* frames, corresponds to the same logic. We also show that $S5WD_n$ is decidable, and we investigate its computational complexity.

The definition of full systems and hypercubes takes a static viewpoint of multi-agent systems that does not use the full power of the interpreted systems model, which is also capable of modelling the evolution of knowledge over time. As noted above, these definitions provide an appropriate characterisation of the agents' knowledge in the initial configurations of many distributed systems. However, we show that the logic $S5WD_n$ has broader applicability than simply reasoning about such initial configurations. We also study in this paper the dynamic behaviour of knowledge in *homogeneous broadcast environments*. These model a particular communication architecture, in which agents operate synchronously and can communicate only by broadcasting information to *all* agents. We assume that agents have perfect recall, and that their initial configuration is characterised by a hypercube. We show that not just the initial configuration, but *all* configurations arising in such a system can be characterised using a full system. It follows that $S5WD_n$ exactly captures the properties of knowledge in homogeneous broadcast systems. Since $S5WD_n$ extends $S5_n$, this provides another example of a natural situation in which $S5_n$ is an incomplete characterisation of the logic of knowledge, analogous to the results of Fagin et al. [Fagin et al. 1992; Fagin and Vardi 1986].

The paper is organised as follows. In Section 2 we recall the two standard semantics for knowledge in multi-agent systems (Kripke models and interpreted systems), and we introduce full systems and hypercubes. In Section 3 we formally relate full systems and hypercubes to Kripke models by identifying corresponding classes of Kripke frames. We also show that with respect to the logic of knowledge we consider in this paper these classes of frames generate the same logic. In Section 4 we present a sound and complete axiomatisation $S5WD_n$ for this logic. We prove the logic decidable in Section 5, and, by relating it to an existing area of study, we identify its computational complexity. These sections all deal with a static framework. In Section 6 we go on to consider a dynamic framework that models how agents' knowledge changes over time. We define homogeneous broadcast systems, and show that agents' states of knowledge in such systems can be characterised by a hypercube at each point of time, thereby showing that $S5WD_n$ is also a sound and complete axiomatisation of the logic of knowledge in homogeneous broadcast systems. We illustrate the theory with an example of a two-person game. Finally, in Section 7 we draw our conclusions and we suggest further work.

2. PRELIMINARY DEFINITIONS

Amongst the approaches that have been proposed to the semantics of logics for knowledge are *interpreted systems* and *Kripke models*. The two approaches have different advantages and disadvantages. On the one hand, interpreted systems provide a more concrete and intuitive way to model real systems, but on the other hand Kripke models come with an heritage of fundamental techniques that may be used to prove properties of the logic.

In order to fix the notation used in this paper, in this section we briefly recall the key definitions of Kripke frames and interpreted systems. (For more technical details and motivation, the reader is referred to an introduction to modal logic, such as [Hughes and Cresswell 1996] or [Goldblatt 1992] or [Hughes and Cresswell 1984], and to [Fagin et al. 1995] for an introduction of its use in Computer Science.) We then we define hypercube systems and full systems, the particular classes of systems that are the focus of this paper.

We use the following mathematical notations throughout. If W is a set, we write $|W|$ for its cardinality. If $\sim$ is an equivalence relation on W and $w \in W$, then we write $W/\sim$ for the set of equivalence classes of $\sim$, and write $[w]_\sim$ for the equivalence class containing w. We write id_W for the identity relation on W, i.e., the relation $\{(w,w) \mid w \in W\}$.

We assume a set $Atoms = \{p, \dots\}$ of *propositional atoms*, and a finite $A = \{1, \dots, n\}$ of *agents*. We will deal primarily with a formal language given by the following grammar:

$$\phi \ ::= \ p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \Box_i \phi$$

where $p \in Atoms$ and $i \in A$. We write $\mathcal{L}_n$ for the set of formulae generated by this grammar when $A = \{1, \dots, n\}$. Other connectives such as $\vee$, $\rightarrow$, and $\Diamond_i$ can be defined in the usual way.

2.1 Kripke models

Definition 2.1 (Kripke frames and Kripke models). *A* frame *F is a tuple $F = (W, R_1, \dots, R_n)$, where W is a non-empty set (called set of worlds) and for each $i \in A$ the component R_i is a binary relation on W. If all relations are equivalence relations, the frame is an* equivalence frame *and we write $\sim_i$ for R_i.*

A model *M is a frame together with an interpretation $\pi : W \rightarrow 2^{Atoms}$ for the atoms. An* equivalence model *is a model whose underlying frame is an equivalence frame.*

We will call equivalence frames *E-frames* and equivalence models *E-models*. The class of equivalence frames will be denoted by $\mathcal{F}_E$. The class of all equivalence models is frequently taken to be the appropriate class of structures for the logic of knowledge. The logic corresponding to this class of structures is the logic $S5_n$ [Halpern and Moses 1990].

We assume the standard notions of satisfaction and validity, but we recall the definition of p-morphism for the multi-modal case:

Definition 2.2 (p-morphism). *Let* $F = (W, R_1, \ldots, R_n)$ *and* $F' = (W', R'_1, \ldots, R'_n)$ *be frames. A* frame p-morphism *from* F *to* F' *is a mapping* $p : W \to W'$ *that satisfies*

(1) the function p *is surjective, and*

(2) for all $u, v \in W$ *and each* $i = 1 \ldots n$*, if* uR_iv *then* $p(u)R_ip(v)$*, and*

(3) for each $i = 1 \ldots n$ *and* $u \in W$ *and* $v' \in W'$*, if* $p(u)R'_iv'$ *then there exists* $v \in W$ *such that* uR_iv *and* $p(v) = v'$*.*

If $M = (W, R_1, \ldots, R_n, \pi)$ *and* $M' = (W', R'_1, \ldots, R'_n, \pi')$ *are Kripke structures, then a* model p-morphism *from* M *to* M' *is a mapping* $p : W \to W'$ *that is a frame p-morphism from* $(W, R_1, \ldots, R_n)$ *to* $(W', R'_1, \ldots, R'_n)$ *and satisfies* $q \in \pi'(p(w))$ *if and only if* $q \in \pi(w)$ *for all propositions* $q \in Atoms$ *and all worlds* $w \in W$*.*

One further property of the language $\mathcal{L}_n$ that will be of use to us is the fact that satisfaction of a formula at a world depends only on worlds connected to that world. Say that two worlds w, w' of a frame $F = (W, \sim_1, \ldots, \sim_n)$ are *connected* if there exists a finite sequence $w = w_0, \ldots, w_k = w'$ of worlds in W such that for $j = 0 \ldots k-1$ we have $w_j \sim_i w_{j+1}$ for some i. Say that F is *connected* if all pairs of worlds $w, w' \in W$ are connected. The *connected component* of F containing a world w is the frame $F_w = (W_w, \sim'_1, \ldots, \sim'_n)$ where W_w is the set of worlds of F connected to w, and each $\sim'_i$ is the restriction of $\sim_i$ to W_w. Similarly, the connected component of a model $M = (F, \pi)$ containing a world w is the model $M_w = (F_w, \pi')$ where π' is the restriction of π to W_w. The model M_w is also called the model generated by w from M. The following result (see, e.g., [Hughes and Cresswell 1984] page 80) makes precise the claim that satisfaction of a formula of $\mathcal{L}_n$ at a world depends only on connected worlds.

Lemma 2.1. *For all worlds* w *of a model* M *and for all formulae* $\psi \in \mathcal{L}_n$ *we have* $M \models_w \phi$ *if and only if* $M_w \models_w \phi$*.*

2.2 Interpreted systems

Interpreted systems are a model for distributed and multi-agent systems proposed by Fagin, Halpern, Moses and Vardi [Fagin et al. 1995], based on an earlier model of Halpern and Moses [Halpern and Moses 1990]. They provide a general theoretical framework within which it is possible to model a variety of modes of communication, failure properties of communication channels, and assumptions about coordination such as synchrony and asynchrony. Its specific focus is to enable states of knowledge to be ascribed to the agents in the system, and to study the evolution of this knowledge as agents communicate. For discussion of axiomatic properties of this model see [Fagin et al. 1995] and [Halpern et al. 1997].

The key aspect of interpreted systems that allows knowledge to be ascribed to agents is the notion of local state. Intuitively, the local state of an agent captures the complete scope of the information about the system that is accessible to the agent. This may include the values of its personal variables and data structures, its record of prior communications, etc. The agents' local states, together with a state of the environment within which they operate, determines the global state of the system at any given time.

Consider n sets of local states, one for every agent of the system, and a set of states for the environment. We denote by L_i the non-empty sets of local states possible for agent i, and by L_e the non-empty set of possible states for the environment. Elements of L_i will be denoted by $l_1, l_2, \ldots$. Elements of L_e will be denoted by $l_e, \ldots$.

DEFINITION 2.3 (SYSTEM OF GLOBAL STATES). A system of global states for n agents *is a subset of a Cartesian product* $L_e \times L_1 \times \cdots \times L_n$. *An* interpreted *system of global states is a pair* (S, π) *where* S *is a system of global states and* $\pi : S \rightarrow 2^{Atoms}$ *is an interpretation function for the atoms.*

The reason for considering a subset is that some of the tuples in the Cartesian product might not be possible because of explicit constraints present in the multi-agent system. The framework of Fagin et al. [Fagin et al. 1995] models the temporal evolution of a system by means of *runs*, which are functions from the natural numbers to the set of global states. An *interpreted system*, in their terminology, is a set of runs over global states together with a valuation for the atoms of the language on points of these runs. We simplify this notion here, since we will deal initially with an atemporal setting. However, we will consider a run-like construct in Section 6.

As shown in [Fagin et al. 1995], interpreted systems can be used to ascribe knowledge to the agents by considering two global states to be indistinguishable for an agent if its local state is the same in the two global states. We formulate this here as a mapping from systems of global states to Kripke frames.

DEFINITION 2.4. *The function* F *mapping systems of global states to Kripke frames is defined as follows: if* $S \subseteq L_e \times L_1 \times \ldots \times L_n$ *is a set of global states for* n *agents then* $F(S)$ *is the Kripke frame* $(W, \sim_1, \ldots, \sim_n)$, *with* $W = S$, *and for each* $i = 1 \ldots n$ *the relation* $\sim_i$ *defined by* $\langle l_1, \ldots, l_n \rangle \sim_i \langle l'_1, \ldots, l'_n \rangle$ *if* $l_i = l'_i$. *The function* F *is naturally extended to map interpreted systems of global states to Kripke models as follows: if* $F(S) = (W, \sim_1, \ldots, \sim_n)$ *then* $F(S, \pi) = (W, \sim_1, \ldots, \sim_n, \pi)$.

Note that for all systems of global states S, the frame $F(S)$ is an equivalence frame. Combined with the semantic interpretation of the language $\mathcal{L}_n$ on Kripke models, this mapping provides a way to interpret $\mathcal{L}_n$ on interpreted systems of global states (for n agents). We say that $\phi \in \mathcal{L}_n$ is valid on an interpreted system of global states (S, π) if ϕ is valid on the model $F(S, \pi)$. Similarly, ϕ is valid on the system S of global states if ϕ is valid on $F(S, \pi)$ for all interpretations π.

2.3 Hypercube systems and full systems

We now define two classes of systems of global states, *full systems* and *hypercube systems*, that provide an intuitive model for the initial situation in many systems of interest. These classes both capture situations in which the agents do not have information about each others' local states (common knowledge might be present, but we are not concerned with this concept in this paper). In hypercubes the environment is assumed trivial, so there is no interesting correlation between the agents' states and the environment.

DEFINITION 2.5 (HYPERCUBE SYSTEM). *A* hypercube system, *or* hypercube, *is a Cartesian product* $H = L_e \times L_1 \times \cdots \times L_n$, *where* L_e *is a singleton and* $L_1, \ldots, L_n$ *are non-empty sets. The class of hypercube systems is denoted by* $\mathcal{H}$.

In full systems, the agents may, however, have some information about how their local state correlates with the state of the environment.

DEFINITION 2.6 (FULL SYSTEM). *A system $S \subseteq L_e \times L_1 \times \cdots \times L_n$ is* full *if for every tuple $\langle l_1, \ldots, l_n\rangle \in L_1 \times \cdots \times L_n$ there exists $s \in L_e$ such that $\langle s, l_1, \ldots, l_n\rangle \in S$. The class of full systems is denoted by $\mathcal{FS}$.*

Clearly, every hypercube is full. The converse is not true. The following example illustrates these definitions.

EXAMPLE 2.1. *Distributed consensus protocols [Lynch 1996; Dwork and Moses 1990] are a class of protocols in which each of n processes is assumed to start the protocol with some choice of a value v from a set V. During the protocol, some subset of the n processes may fail, and this failure may be of a variety of types, e.g. crash failure, message omission, message loss or arbitrary (Byzantine) behaviour. The goal of the protocol is for the non-failed processes to reach agreement upon some output value in V (which is required to be related to the initial values in some way.) We may model the initial state of knowledge of the processes participating in such a protocol by the system $L_e \times L_1 \times \ldots \times L_n$, where L_e is a singleton and, for each $i = 1 \ldots n$, the local states $L_i = V$ correspond to the initial values. Clearly, this system is a hypercube.*

It is common to model fault-tolerant protocols by representing the environment's future failure pattern in the initial state [Dwork and Moses 1990]. This can be done by taking L_e to be the set of possible future failure patterns. For example, when dealing with crash failures, L_e might be taken to be the set of all subsets of $\{1, \ldots, n\}$ containing k or fewer elements, representing the processes that may crash during the run of the protocol. In this case $L_e \times L_1 \times \ldots \times L_n$ is a full system, but not a hypercube.

For an example of a full system that is not a Cartesian product, consider a parallel computation which in its first step initialises n processors in such a way that the local memory of process i stores the value of the i-th component of a vector $(v_1, \ldots, v_n) \in V^n$ held in a shared memory. This initial state of the computation may therefore be represented as the full system $\{\langle (v_1, \ldots, v_n), v_1, \ldots v_n\rangle \mid v_i \in V,\ i = 1 \ldots n\}$, which is a subset of the system $L_e \times L_1 \times \ldots L_n$, where $L_e = V^n$ and for $i = 1 \ldots n$ we have $L_i = V$. Here L_e represents the state of shared memory and, for each $i = 1, \ldots n$, L_i represents the state of the local memory of process i.

We will show in Section 6 that the applicability of the class of full systems and hypercubes goes beyond that of modelling the initial configurations of naturally occurring systems. We will define a dynamic framework that shows how agents' knowledge changes over time, and show that under certain conditions, the agents' states if knowledge can be characterised by a hypercube at all times, not just at the initial state. Our first aim, however, will be to axiomatise the class of hypercubes and full systems. In order to use the tools of modal logics for this aim we formally relate these classes of systems to several classes of Kripke frames.

3. CLASSES OF FRAMES CORRESPONDING TO HYPERCUBES AND FULL SYSTEMS

In this section we identify a number of properties of frames that can be used to characterise the frames corresponding to full systems and hypercubes up to isomorphism. We also show that, somewhat surprisingly, full systems and hypercubes generate precisely the same set of valid formulae of the language $\mathcal{L}_n$. We obtain this result by establishing the existence of p-morphisms between frames in the classes of frames corresponding to these classes of systems.

3.1 Directed frames

We have seen above that every system of global states generates a frame. Our aim in this section is to characterise the frames generated by hypercubes and by full systems. The following result identifies some properties of the resulting frames based on the properties of the system of global states.

LEMMA 3.1. *Let S be a system of global states, and let $F(S) = (W, \sim_1, \dots, \sim_n)$ be the frame defined from it by Definition 2.4.*

(1) If S is a hypercube then $F(S)$ is such that $\bigcap_{i \in A} \sim_i = id_W$;
(2) If S is full and $n \geq 2$ then $F(S)$ is connected.
(3) If S is full then for any $w_1, \dots, w_n \in W$ there exists a $\overline{w} \in W$ such that $w_i \sim_i \overline{w}$ for each $i = 1, \dots, n$.

PROOF. For (1), suppose that S is a hypercube and consider any two elements $w = (l_e, l_1, \dots, l_n)$, $w' = (l_e, l'_1, \dots, l'_n)$ in W such that $w(\bigcap_{i \in A} \sim_i)w'$. (Note that the first component of these tuples must be the same if S is a hypercube.) Then for all i in A, $(l_e, l_1, \dots, l_n) \sim_i (l_e, l'_1, \dots, l'_n)$. Therefore, by definition of the relations $\sim_i$, for all i in A we have $l_i = l'_i$, that is $w = w'$.

For (2), suppose that S is full and let $w = (l_e, l_1, \dots, l_n)$ and $w' = (l'_e, l'_1, \dots, l'_n)$ be points in S. Since S is full there exists l''_e such that $w'' = (l''_e, l_1, l'_2 \dots, l'_n)$ is in W. Clearly $w \sim_1 w'' \sim_2 w'$. Thus, there is a path from w to w' of length two.

For (3), suppose that S is full and consider any $w_1 = (l^1_e, l^1_1, \dots, l^1_n), \dots, w_n = (l^n_e, l^n_1, \dots, l^n_n)$. Since S is full there exists l_e such that $\overline{w} = (l_e, l^1_1, \dots, l^n_n) \in S$. By Definition 2.4, the world $\overline{w}$ is in W and by construction for each $i = 1, \dots, n$, we have $w_i \sim_i \overline{w}$. □

This shows that Kripke frames that we build from the hypercubes and full systems by means of the standard technique ([Fagin et al. 1995]) constitute a proper subclass of the class of equivalence frames. We will show that the properties of Lemma 3.1 can be used to characterise the images of the hypercubes and full systems.

We will say that a frame $(W, \sim_1, \dots, \sim_n)$ has the *identity intersection property*, or is an *I frame*, if $\bigcap_{i \in A} \sim_i = id_W$. Similarly, we say that a frame is *directed*, or is a *D frame*, if for any $w_1, \dots, w_n \in W$ there exists a $\overline{w} \in W$ such that $w_i \sim_i \overline{w}$ for each $i = 1, \dots, n$. We will also use combinations of these letters to refer to frames satisfying several of these properties. Thus, directed equivalence frames with the identity intersection property will be called *EDI frames*. Similarly, we subscript $\mathcal{F}$ by these letters to indicate the class of frames have the corresponding properties; thus $\mathcal{F}_{EDI}$ denotes the class of EDI frames. Lemma 3.1 states that the image of a full system under F is an ED frame and since every hypercube is full, the image of a hypercube is an EDI frame.

The converse of these properties is not true, e.g., it is not the case that every ED frame is the image of a hypercube. However, something very close to this is the case:

LEMMA 3.2. *(1) For every ED frame F there exists a full system S such that $F(S) \equiv F$.*

(2) For every EDI frame F there exists a hypercube S such that $F(S) \equiv F$.

PROOF. We first show part (1). Let $F = (W, \sim_1, \dots, \sim_n)$ be an ED frame. Take $S \subseteq W \times W/{\sim_1} \times \dots \times W/{\sim_n}$ to be the set of tuples $\langle w, [w]_{\sim_1}, \dots, [w]_{\sim_n} \rangle$ where $w \in W$. We show that S is a full system and such that $F(S) \equiv F$.

To see that S is full, let $w_1, \dots w_n \in W$. We show that there exists $w \in W$ such that $\langle w, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle \in S$. Since F is a D frame, there exists a world w such that $w \sim_i w_i$ for each $i = 1 \dots n$. Thus $\langle w, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle = \langle w, [w]_{\sim_1}, \dots, [w]_{\sim_n} \rangle \in S$. This shows that S is full.

Write $F(S) = (S, \sim'_1, \dots, \sim'_n)$. To show that $F(S) \equiv F$, define the mapping $h : S \to W$ by $h(\langle w, [w]_{\sim_1}, \dots, [w]_{\sim_n} \rangle) = w$. It is clear that h is a bijection. Moreover,

$$\begin{aligned}
\langle w_1, [w_1]_{\sim_1}, \dots, [w_1]_{\sim_n} \rangle \sim'_i\ & \langle w_2, [w_2]_{\sim_1}, \dots, [w_2]_{\sim_n} \rangle \\
& \text{iff } [w_1]_{\sim_i} = [w_2]_{\sim_i} \\
& \text{iff } w_1 \sim_i w_2 \\
& \text{iff } h(\langle w_1, [w_1]_{\sim_1}, \dots, [w_1]_{\sim_n} \rangle) \sim_i h(\langle w_2, [w_2]_{\sim_1}, \dots, [w_2]_{\sim_n} \rangle).
\end{aligned}$$

Thus, h is a frame isomorphism, establishing $F(S) \equiv F$. This completes the proof of part (1).

For part (2), let $F = (W, \sim_1, \dots, \sim_n)$ be an EDI frame. Define $S = \{1\} \times W/{\sim_1} \times \dots \times W/{\sim_n}$. Clearly S is a hypercube. Write $F(S) = (S, \sim'_1, \dots, \sim'_n)$. We show that $F(S) \equiv F$.

Consider an element $\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle$ of S. Since F is a D frame, there exists $w \in W$ such that $w \in [w_i]_{\sim_i}$ for each $i = 1, \dots, n$. Moreover, because F is an I frame this w is unique. Define the mapping $h : S \to W$ by taking $h(\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle)$ to be the unique w such that $w \in [w_i]_{\sim_i}$ for each $i = 1, \dots, n$. The mapping h is surjective because for each $w \in W$ we have $h(\langle 1, [w]_{\sim_1}, \dots, [w]_{\sim_n} \rangle) = w$. Moreover h is injective because if $h(\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle) = w = h(\langle 1, [w'_1]_{\sim_1}, \dots, [w'_n]_{\sim_n} \rangle)$ then for each $i = 1 \dots n$ we have that w is in both $[w_i]_{\sim_i}$ and $[w'_i]_{\sim_i}$. Thus, these equivalence classes must be the identical, and hence the tuples $\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle$ and $\langle 1, [w'_1]_{\sim_1}, \dots, [w'_n]_{\sim_n} \rangle$ are identical.

It remains to show that h has the homomorphism property. For this, note that by construction, for each $i = 1, \dots, n$ we have $h(\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle) \sim_i w_i$. Thus if $\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle \sim'_i \langle 1, [w'_1]_{\sim_1}, \dots, [w'_n]_{\sim_n} \rangle$ then $[w_i]_{\sim_i} = [w'_i]_{\sim_i}$, hence $h(\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle) \sim_i w_i \sim_i w'_i \sim_i h(\langle 1, [w'_1]_{\sim_1}, \dots, [w'_n]_{\sim_n} \rangle)$. Conversely, suppose $u = h(\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle)$ and $v = h(\langle 1, [w'_1]_{\sim_1}, \dots, [w'_n]_{\sim_n} \rangle)$ and $u \sim_i v$. By definition of h, $u \in [w_i]_{\sim_i}$ and $v \in [w'_i]_{\sim_i}$. Since $u \sim_i v$ it follows that $[w_i]_{\sim_i} = [w'_i]_{\sim_i}$. Thus, $\langle 1, [w_1]_{\sim_1}, \dots, [w_n]_{\sim_n} \rangle \sim'_i \langle 1, [w'_1]_{\sim_1}, \dots, [w'_n]_{\sim_n} \rangle$. Thus, h is a frame isomorphism, establishing $F(S) \equiv F$. This completes the proof of part (2). $\square$

Using the fact that isomorphic frames satisfy the same formula schemes, it follows from Lemma 3.1 and Lemma 3.2 that from the point of view of the language $\mathcal{L}_n$, full systems and ED frames are equivalent, as are hypercubes and EDI frames. Stated more precisely, we have the following.

THEOREM 3.1. *$F(\mathcal{H}) \equiv_{\mathcal{L}_n} \mathcal{F}_{EDI}$ and $F(\mathcal{FS}) \equiv_{\mathcal{L}_n} \mathcal{F}_{ED}$.*

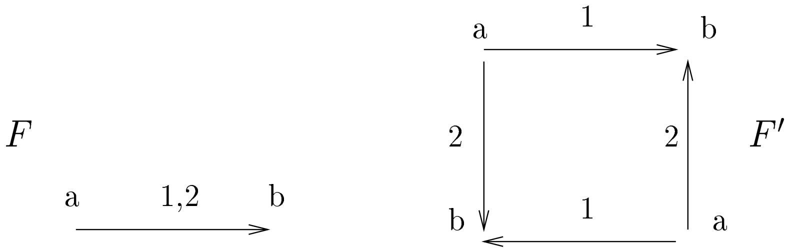

Fig. 1. Two p-morphic frames used in the proof of Lemma 3.3

Our strategy for axiomatising these classes of systems will be to focus on the corresponding classes of frames instead. One approach to this would be to seek axioms that correspond to the properties D and I. This turns out not to be possible.

LEMMA 3.3. *If A is not a singleton, then no modal formula corresponds to property I even in the case of equivalence frames.*

PROOF. Suppose the opposite and assume there is a formula ϕ that corresponds to property I. Consider the frame F' in Figure 1. The frame F' is an I frame, so $F' \models \phi$. Consider now the frame F and a function $p : F' \to F$ such that p maps points in F according to the names in the Figure. It is easy to see that p is a p-morphism from F' to F. Since p-morphisms preserve validity on frames (see, e.g., [Hughes and Cresswell 1984] page 73) we have that $F \models \phi$. But F is not an I frame and we have a contradiction.

In the case of equivalence frames the argument is analogous, but one should consider the symmetric closure of the frames pictured in Figure 1. □

Similar reasoning shows that the above holds even by restricting to equivalence frames.

As an aside, we note that this result is very sensitive to the language under consideration. There are extensions of the language under which it fails. For example, consider a language containing an operator for distributed knowledge [Fagin et al. 1992]. This operator is used to express the knowledge that the group of all agents would have if they pooled their information. Formally, if ϕ is a formula, then so is $D_A\phi$. The formula $D_A\phi$ is interpreted by associating the relation $\sim = \bigcap_{i\in A} \sim_i$ to the operator D_A in the standard Kripke-style interpretation, i.e., we define $M \models_w D_A\psi$ if $M \models_{w'} \psi$ for all $w' \sim w$. Using this operator, we can prove a correspondence result for the intersection property.

LEMMA 3.4. *An equivalence frame F is an I frame if and only if $F \models \phi \Leftrightarrow D_A\phi$.*

PROOF. Left to right. Let M be a model based on F such that $M \models_w \phi$. Since $\bigcap_{i\in A} \sim_i = id_W$, then $M \models_w D_A\phi$. Analogously, suppose $M \models_w D_A\phi$. Since $w(\bigcap_{i\in A} \sim_i)w$ we have $M \models_w \phi$.
Right to left. Suppose $F \models \phi \Leftrightarrow D_A\phi$ and for all i we have $w_1 \sim_i w_2$. Take a valuation π such that $p \in \pi(w)$ if and only if $w = w_1$. Since $F, \pi \models_{w_1} p \Leftrightarrow D_A p$ and $(F,\pi) \models_{w_1} p$, we have $(F,\pi) \models_{w_1} D_A p$ and so $(F,\pi) \models_{w_2} p$. But since $\pi(p) = \{w_1\}$, it must be that $w_1 = w_2$. □

This result suggests that the axiom $\phi \Leftrightarrow D_A\phi$ could be used as part of an axiomatisation of the class of EDI frames when the language includes the distributed

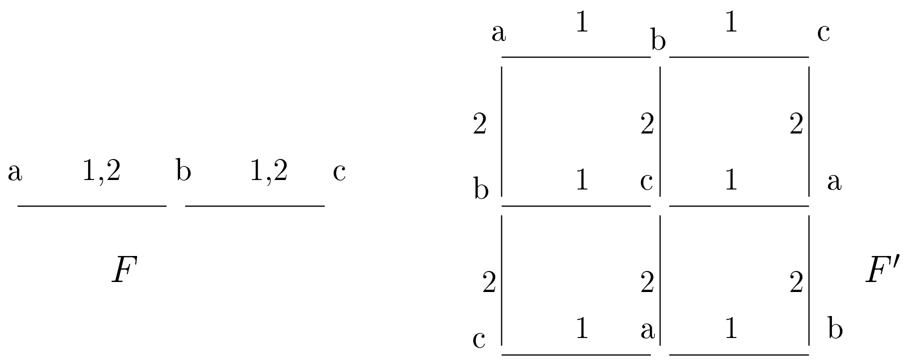

Fig. 2. An EDI frame mapping an ED frame via a p-morphism

knowledge operator. We are concerned in this paper, however, with a weaker language. Lemma 3.3 suggests that it may be inappropriate to focus on the identity intersection property in seeking to obtain the axiomatisation. Indeed, it turns out that this property has no impact on the set of valid formulae of $\mathcal{L}_n$ in the context of interest to us. More precisely, we have the following result.

THEOREM 3.2. $\mathcal{F}_{EDI} \equiv_{\mathcal{L}_n} \mathcal{F}_{ED}$

To establish Theorem 3.2, we prove that any ED frame can be seen as the target of a p-morphism from an EDI frame; the result will then follow using the fact that p-morphisms between frames preserve validity and that the class of DI frames is a subclass of the class of D frames. (Note that the identity map on a frame is a frame isomorphism, hence a p-morphism.)

Consider an ED frame $F = (W, \sim_1, \ldots, \sim_n)$. Write $\sim$ for the relation $\bigcap_{i=1\ldots n} \sim_i$; since each of the $\sim_i$ is an equivalence relation, so is $\sim$. The frame F can then be viewed as the union of equivalence classes of the relation $\sim$, which we call *clusters*. Clusters containing more than a single point are sub-frames in which property I clearly does not hold; in general a cluster may be infinite in size.

If we want to construct an EDI frame that maps to a particular ED frame by a p-morphism, one way is to replace every cluster of the ED frame with a sub-frame that is EDI but that can still be mapped into the cluster. Figure 2 depicts the relatively simple case of an equivalence frame F composed by three points a, b, c connected by all the relations: $\sim_1, \sim_2$, in this case; F clearly is ED but not EDI[1]. The frame F' on the right of the figure is an EDI frame; the names of its points represent the targets of the p-morphism from F' onto F. So, for example the top left point of F' is mapped onto a of F; the relations are mapped in the intuitive way. It is an easy exercise to show that F is indeed a p-morphic image of F' and will therefore validate every formula which is valid on F'.

The aim of the following is to define precisely how to build, given any ED frame, a new EDI frame in which every cluster is "unpacked" into an appropriate similar structure and to define the relations appropriately.

[1]The relations are supposed to be the reflective transitive closure of the the ones depicted in the figure.

In order to achieve the above, we present two set theoretic results. In Lemma 3.5 we show that every infinite set X can be seen as the image of a product X^n under a function p. Intuitively, this lemma will be used by taking the set X as one of the clusters of an EDI frame F, the function p as the p-morphism and the product X^n (where n is the number of relations on the frame) as the sub-frame that will replace the cluster in the new frame F'. Lemma 3.6 extends the result of Lemma 3.5 to guarantee that even if the clusters differ in size it is always possible to find a single sub-frame that can replace each of them.

We assume n to be a natural number, such that $n \geq 2$.

LEMMA 3.5. *Given any infinite set X, there exists a function $p : X^n \to X$ such that the following holds.*

Let $i \in \{1, \ldots, n\}$. For all $u, x_i \in X$, there are $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n, \in X$, such that $p(x_1, \ldots, x_n) = u$.

PROOF. Consider the set $T = \{\tau_{x,y} \mid x, y \in X\}$ of the transpositions of X, i.e. functions $\tau_{x,y} : X \to X$; where $x, y \in X$, and such that $\tau_{x,y}(z) = y$ if $z = x$; $\tau_{x,y}(z) = x$ if $z = x$; $\tau_{x,y}(z) = z$ otherwise. We have $|X| \leq |T| \leq |X \times X|$. But by set theory ([Lang 1984] page 701 for example) $|X| = |X \times X|$, and so $|X| = |T|$. So, by induction, we have $|X^{n-1}| = |X| = |T|$. Call f the bijection $f : X^{n-1} \to T$, and define $p(x_1, \ldots, x_n) = f(x_1, \ldots, x_{n-1})(x_n)$. To prove the lemma holds we consider two cases: $i \neq n$ and $i = n$.

For $i \neq n$, assume any $u \in X$, and any $x_i \in X$. Take any $x_j, j \in \{1 \ldots n-1\} \setminus \{i\}$; $f(x_1, \ldots, x_{n-1})$ is a transposition of X. So, there exists an $x_n \in X$ such that $f(x_1, \ldots, x_{n-1})[0](x_n) = u$. So $p(x_1, \ldots, x_n) = u$.

For $i = n$, assume again any $u \in X$, and any $x_n \in X$. Consider the transposition $\tau_{x_n,u}$; we have $\tau_{x_n,u}(x_n) = u$. But $\tau_{x_n,u} = f(x_1, \ldots, x_{n-1})$ for some $x_1, \ldots, x_{n-1} \in X$. So $p(x_1, \ldots, x_n) = u$. □

Lemma 3.5 induces a similar result for sets whose cardinality is smaller than X.

LEMMA 3.6. *Given any infinite set X, and a set $C \neq \emptyset$, such that $|C| \leq |X|$, there exists a function $p : X^n \to C$ such that the following holds.*
Let $i \in \{1, \ldots, n\}$. For all $x_i \in X, u \in C$, there are $x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n, \in X$, such that $p(x_1, \ldots, x_n) = u$.

PROOF. Consider a set T such that $C \cup T$ and X have the same cardinality, and let g be a bijection from X to $(C \cup T)$. Then there is a function $p' : (C \cup T)^n \to (C \cup T)$, satisfying the property expressed by Lemma 3.5. Define now a function $p'' : (C \cup T) \to C$, such that $p''(x) = x$ if $x \in C$, otherwise $p''(x) = c$, where c is any element in C. Define the function $p : X^n \to C$ by $p(x_1, \ldots, x_n) = p''(p'(g(x_1), \ldots, g(x_n)))$. The following shows that p has the property required. For, let $i \in \{1, \ldots, n\}$ and take any $x_i \in X$ and $u \in C$. Then $g(x_i) \in (C \cup T)$, and so by Lemma 3.5 there exist $c_1, \ldots, c_{i-1}, c_{i+1}, \ldots, c_n \in C \cup T$, such that $p'(c_1, \ldots, c_{i-1}, g(x_i), c_{i+1}, \ldots, c_n) = u$. Define $x_j = g^{-1}(c_j)$ for $j \in \{1, \ldots, n\} \setminus \{i\}$. We then have $p(x_1, \ldots, x_n) = p''(p'(c_1, \ldots, c_{i-1}, g(x_i), c_{i+1}, \ldots, c_n)) = p''(u) = u$ since $u \in C$. □

We rely on the two results above to define a function p that maps tuples $\langle c, x_1, \ldots, x_n \rangle$ into c, where c is a cluster and $x_i \in X$, for some appropriate set X. The function p is defined as in Lemma 3.6 but it has an extra component for the cluster.

COROLLARY 3.1. *Let $\mathcal{C}$ be a set of nonempty subsets of a set W. Then there exists a set X and a function $p : \mathcal{C} \times X^n \to W$ such that*

(1) for all tuples $\langle c, x_1, \dots, x_n \rangle$ we have $p(\langle c, x_1, \dots, x_n \rangle) \in c$, and
(2) for all $c \in \mathcal{C}$, for all $u \in c$, for all $i = 1 \dots n$, and for all $x_i \in X$, for each $j \in \{1 \dots n\} \setminus \{i\}$ there exists $x_j \in X$, such that $p(\langle c, x_1 \dots x_n \rangle) = u$.

PROOF. Let X be an infinite set with cardinality at least as great as the cardinality of any $c \in \mathcal{C}$. This can be constructed by taking the union of these sets $c \in \mathcal{C}$ or by considering the set of the natural numbers $X = \mathbb{N}$ if all the sets $c \in \mathcal{C}$ are finite. For each $c \in \mathcal{C}$, let $p_c : X^n \to c$ be the function promised by Lemma 3.6. Define $p : \mathcal{C} \times X^n \to W$ by $p(c, x_1, \dots, x_n) = p_c(x_1, \dots, x_n)$. It is immediate that this function has the required property. □

THEOREM 3.3. *Given any equivalence D frame F, there exists an equivalence DI frame F', and a p-morphism p, such that $p(F') = F$.*

PROOF. Let $F = (W, \sim_1, \dots, \sim_n)$ be a frame with n relations on its support set W. Write $\sim$ for the relation $\bigcap_{i=1 \dots n} \sim_i$. Since each of the $\sim_i$ is an equivalence relation, so is $\sim$. Since the set of worlds W of the frame F is non-empty, it can be viewed as the union of the equivalence classes of the relation $\sim$, which we call *clusters*. Call $\mathcal{C}$ the set of clusters of F. Consider the infinite set X and a function p as described in Corollary 3.1, and define the frame $F' = (W', \sim'_1, \dots, \sim'_n)$ as follows:

—$W' = \mathcal{C} \times X^n$,
—$\langle c, x_1, \dots, x_n \rangle \sim'_i \langle d, y_1, \dots, y_n \rangle$ if $x_i = y_i$ and there exists worlds $u \in c$ and $v \in d$ such that $u \sim_i v$.

We can prove that:

(1) The frame F' is an equivalence DI frame.

PROOF. a) F' is clearly an equivalence frame.
b) We prove F' satisfies property I. Write $\sim'$ for $\bigcap_{i=1 \dots n} \sim'_i$. Suppose $\langle c, x_1, \dots, x_n \rangle \sim' \langle d, y_1, \dots, y_n \rangle$. Then for all $i = 1 \dots n$ we have that $x_i = y_i$, and there exist $u_i \in c$ and $v_i \in d$ such that $u_i \sim_i v_i$. Since c and d are equivalence classes of $\sim$, it follows from the latter that $u_1 \sim v_1$, and consequently that $c = d$. Thus, $\langle c, x_1, \dots, x_n \rangle = \langle d, y_1, \dots, y_n \rangle$.
c) We prove F' satisfies property D. Consider n tuples $\langle c_1, x_1^1, \dots, x_n^1 \rangle, \dots, \langle c_n, x_1^n, \dots, [0]x_n^n \rangle$ in W'. For each $i = i \dots n$ let u_i be a world in cluster c_i. Since F has property D, there exists a world w such that $w \sim_i u_i$ for each $i = 1 \dots n$. Let c be the cluster containing w. Then, by construction, for each $i = 1 \dots n$ we have $\langle c, x_1^1, \dots, x_n^n \rangle \sim'_i \langle c_i, x_1^i, \dots, x_n^i \rangle$. □

(2) The function p is a p-morphism from F' to F.

PROOF. That the function p is surjective follows from property (2) of Corollary 3.1.
Next, we show that p is a frame homomorphism. Consider two tuples $\langle c, x_1, \dots, x_n \rangle$, $\langle d, y_1, \dots, y_n \rangle$ in W' such that $\langle c, x_1, \dots, x_n \rangle \sim'_i \langle d, y_1, \dots, y_n \rangle$. Then there exists $u \in c$ and $v \in d$ such that $u \sim_i v$. By property (1) of Corollary 3.1, we have

$p(\langle c, x_1, \ldots, x_n \rangle) \sim_i u$ and $p(\langle d, y_1, \ldots, y_n \rangle) \sim_i v$. Since $\sim_i$ is an equivalence relation, it follows that $p(\langle c, x_1, \ldots, x_n \rangle) \sim_i p(\langle d, y_1, \ldots, [0]y_n \rangle)$.
To show the backward simulation property, consider a tuple $\mathbf{x} = \langle c, x_1, \ldots, x_n \rangle$, and assume $p(\mathbf{x}) \sim_i w$ for some world w of F. Let d be the cluster containing w. By Corollary 3.1(2), there exist y_j for $j \neq i$ such that if $\mathbf{y} = \langle d, y_1, \ldots y_{i-1}, x_i, y_{i+1}, \ldots y_n \rangle$, then $p(\mathbf{y}) = w$. Since $p(\mathbf{x}) \in c$ by Corollary 3.1(1), it is immediate that $\mathbf{x} \sim_i' \mathbf{y}$. □

This completes the proof of Theorem 3.2. Since we confine our attention in this paper to the language $\mathcal{L}_n$, this result, together with Theorem 3.1, shows that the set of valid formulae for the class of full frames is the same as that for the class of hypercubes. Both sets of valid formulae are equal to the set of formulae valid on ED frames. We now set about attempting to axiomatise the latter. It turns out to be necessary to introduce one more class of frames in order to achieve this.

3.2 Weakly-directed frames

In order to axiomatise the class of ED frames, we need to introduce one more class of frames. The reason for this is that the directedness property does not naturally correspond to any formula of $\mathcal{L}_n$.

LEMMA 3.7. *No modal formula corresponds to n-directedness.*

PROOF. Suppose the opposite and assume there is a formula ϕ that corresponds to n-directedness. Consider two disjoint frames, $F = (W, \sim_1, \ldots, \sim_n)$ and $F' = (W', \sim_1', \ldots, \sim_n')$, where $W \cap W' = \emptyset$, such that both F and F' are n-directed. Consider the frame $F \cup F' = (W \cup W', \sim_1 \cup \sim_1', \ldots, \sim_n \cup \sim_n')$. Since by assumption $F \models \phi$ and $F' \models \phi$, it follows from Lemma 2.1 that $F \cup F' \models \phi$. But, then ϕ is valid on a frame which, in general, is not n-directed. This is the opposite of what we assumed at the beginning. □

The problem here is rather superficial however. Any class of frames corresponding to a modal formula should be closed under disjoint unions. To address this problem, we define a slight weakening of the notion of directedness. We will show that the class of frames satisfying this weaker notion validates the same class of formulae.

DEFINITION 3.1 (WEAK-DIRECTEDNESS). *A frame $F = (W, \sim_1, \ldots, \sim_n)$ is* weakly-directed *when for all worlds $w_0, w_1, \ldots, w_n \in W$, if for each $i = 1, \ldots, n$ there exists $j \in \{1, \ldots, n\}$ such that $w_0 \sim_j w_i$, then there exists a world w such that $w_i \sim_i w$ for each $i = 1, \ldots, n$.*

That is, weak-directness is like directedness in requiring the existence of a world w such that $w_i \sim_i w$ for each $i = 1, \ldots, n$, but it does so only under the condition that the worlds w_i are each connected to some world through a single step through one of the relations $\sim_j$. We use the notation "WD" to refer to the property of weak-directedness. Thus, we write, e.g., $\mathcal{F}_{EWD}$ for the class of weakly-directed equivalence frames. Clearly, every directed frame is weakly-directed. Moreover, the class of weakly-directed frames is easily seen to be closed under disjoint unions. Indeed, this class of frames turns out to be the smallest class of frames containing the directed frames that is closed under disjoint unions. We first note the following.

LEMMA 3.8. *Every weakly-directed and connected equivalence frame is directed.*

PROOF. Suppose that $F = (W, \sim_1, \ldots, \sim_n)$ is weakly-directed and connected. Let $w_1, \ldots, w_n$ be any n worlds in W. We show that there exists a world w such that $w_i \sim_i w$ for each $i = 1, \ldots, n$. Since F is connected, the worlds $w_1, \ldots, w_n$ are in the same connected component, and hence all connected to some world w'. Since F is an equivalence frame the relations $\sim_i$ are symmetric, so we may assume that for each i there exists a path directed from w' to w_i. We now claim that none of these paths need to be any longer than one step, for if so, we can reduce their length. For, suppose without loss of generality that the path from w' to w_1 involves more than one step. Write this path as $w' \sim_{i_1} u \sim_{j_1} v \ldots w_1$ and write the remaining paths as $w' \sim_{i_2} w'_2 \ldots w_2$ to $w' \sim_{i_n} w'_n \ldots w_n$. Using weak-directedness (and an ordering of the worlds $u, w'_2 \ldots w'_n$ such that u occurs in position j_1), we obtain a world w'' such that $u \sim_{j_1} w''$ and for each $k \neq 1$ we have $w'_k \sim_{j_k} w''$ for some j_k. By symmetry of the relations, we obtain paths from w'' to the worlds w_i. For $k \neq 1$ these paths are of the form $w'' \sim_{j_k} w'_k \ldots w_k$ and have the same length as the path connecting w' to w_k. For $k = 1$ we have the path $w'' \sim_{j_1} u \sim_{j_1} v \ldots w_1$, which can be shortened to $w'' \sim_{j_1} v \ldots w_1$ by transitivity of $\sim_{j_1}$. This argument establishes that there exists a world w' such that for each $i = 1, \ldots, n$ we have $w' \sim_j w_i$ for some j. Since F is weakly-directed, it follows that there exists a world w such that $w_i \sim_i w$ for each $i = 1, \ldots, n$. □

We obtain two consequences of this result. First, the characterisation of the weakly-directed frames claimed above.

COROLLARY 3.2. *The class of weakly-directed equivalence frames is the smallest class of equivalence directed frames that is closed under arbitrary disjoint unions and isomorphism.*

PROOF. It is immediate from the definition that the class of weakly-directed equivalence frames contains the class of ED frames and is closed under disjoint unions and isomorphism. To show that it is the smallest such class, we show that any weakly-directed equivalence frame is isomorphic to a disjoint union of directed equivalence frames. For let F be weakly-directed, and let W' be a subset of the set of worlds of F containing exactly one world from each connected component of F. For each $w \in W$ let F_w denote the connected component of F containing w. By Lemma 3.8, each F_w is directed. It is then possible to show that F is isomorphic to the disjoint union of the frames F_w as w ranges over W'. □

The second consequence of Lemma 3.8 is the fact that the formulae of $\mathcal{L}_n$ validated by the equivalence weakly-directed frames is the same as the set validated by the ED frames.

COROLLARY 3.3. $\mathcal{F}_{EWD} \equiv_{\mathcal{L}_n} \mathcal{F}_{ED}$

PROOF. Since every ED frame is EWD, every formula valid on the EWD frames is valid on the ED frames. Conversely, suppose that $\phi \in \mathcal{L}_n$ is not valid on some EWD frame F. Then there exists a valuation π and a world w such that $M \models_w \neg\phi$, where $M = (F, \pi)$. Let M_w be the connected component of M containing w and F_w the corresponding frame. Then M_w is a directed equivalence model, and by Lemma 2.1 we have $M_w \models_w \neg\phi$. Consequently, ϕ is not valid on the ED frame F_w. □

This result, together with the results of the preceding sections, enables us to focus, in our quest for an axiomatisation of the full systems and hypercubes, on the class of weakly-directed equivalence frames.

4. AXIOMATISATION

We are now ready to present an axiomatisation of the full systems and hypercubes with respect to $\mathcal{L}_n$. The basis for the axiomatisation will be the property of weak-directedness identified in the previous section.

For convenience, we first introduce some notation and terminology. We will write $S\phi$ for the formula $\bigvee_{i=1,\ldots,n} \Diamond_i \phi$. Note that $M \models_w S\phi$ if there exists a world w' such that $M \models_{w'} \phi$ and $w \sim_i w'$ for some i. Intuitively, $S\phi$ asserts that at least one of the agents $1, \ldots, n$ considers ϕ possible.

For each $i = 1, \ldots, n$, we also define a formula to be i-local if it is a boolean combination of formulae of the form $\Box_i \phi$. Intuitively, an i-local formula expresses a property of agent i's state of knowledge. More precisely, we have the following fact, which may be proved by a straightforward induction.

Lemma 4.1. *Let ϕ be an i-local formula in $\mathcal{L}_n$, let M be an equivalence model on n agents, and let w and w' be two worlds of M with $w \sim_i w'$. Then $M \models_w \phi$ if and only if $M \models_{w'} \phi$.*

We analyse extensions of $\mathrm{S5}_n$ with respect to the axiom schema:

$$\left(\bigwedge_{i=1,\ldots,n} S\phi_i \right) \Rightarrow SS \left(\bigwedge_{i=1,\ldots,n} \phi_i \right) \qquad \mathbf{WD}$$

where each ϕ_i is required to be an i-local formula. There is a close relationship between this axiom, the property of weak directedness and the property defining full systems. Intuitively, the axiom states that if there are n worlds (each reachable in a single step from the present world), such that the i-th world is one in which agent i is in a state of knowledge described by ϕ_i, then there exist a single world (reachable in two steps from the present) that realizes these n states of knowledge. This intuitive relationship may be made precise by the following correspondence result:

Lemma 4.2. *For equivalence frames F, we have $F \models$ **WD** if and only if F is weakly-directed.*

Proof. We first show that if F is a WD frame then $F \models \mathbf{WD}$. For, suppose that π is an interpretation of F and w_0 a world of F such that $(F, \pi) \models_{w_0} (\bigwedge_{i=1,\ldots,n} S\phi_i)$. Then for each $i = 1, \ldots, n$ there exists a world w_i such that $w_0 \sim_{j_i} w_i$ for some j_i and $(F, \pi) \models_{w_i} \phi_i$. Since F is weakly-directed there exists a world w such that $w_i \sim_i w$ for each $i = 1, \ldots, n$. By Lemma 4.1, we have $(F, \pi) \models_w \bigwedge_{i=1,\ldots,n} \phi_i$. Since $w_0 \sim_{j_1} w_1 \sim_1 w$, it follows that $(F, \pi) \models_{w_0} SS \bigwedge_{i=1,\ldots,n} \phi_i$. This establishes $F \models \mathbf{WD}$.

Conversely, suppose $F \models \mathbf{WD}$. We show that F is weakly-directed. Let $w_0, w_1, \ldots, w_n$ be worlds of F such that for each $i = 1, \ldots, n$ there exists j_i such that $w_0 \sim_{j_i} w_i$. We need to show that there exists a world w such that $w_i \sim_i w$ for each $i = 1, \ldots, n$. To achieve this, let $p_1, \ldots, p_n$ be n distinct propositions and define the interpretation π by $p_i \in \pi(w)$ if and only if $w \sim_i w_i$, for each $i = 1, \ldots, n$.

(The interpretation π may be defined arbitrarily on all other propositions.) Note that we have $(F,\pi) \models_{w_0} \bigwedge_{i=1,\dots,n} S\Box_i p_i$. Since $F \models \mathbf{WD}$, and each formula $\Box_i p_i$ is i-local, it follows that $(F,\pi) \models_{w_0} SS \bigwedge_{i=1,\dots,n} \Box_i p_i$. In particular, there exists a world w such that $(F,\pi) \models_w \bigwedge_{i=1,\dots,n} \Box_i p_i$, hence $(F,\pi) \models_w \bigwedge_{i=1,\dots,n} p_i$. But, by definition of π, this means that $w_i \sim_i w$ for each $i = 1,\dots,n$, as required. $\Box$

The correspondence result above strongly indicates that the axiom **WD** can serve as a basis for an axiomatisation of the weakly-directed equivalence frames. We now establish that this is indeed the case. The proof will be by means of a standard technique for completeness proofs in modal logic, namely the construction of a canonical model. We now briefly review this technique to fix the notation, but refer the reader to [Chellas 1980; Hughes and Cresswell 1984] for details.

A logic L consists of a derivability relation $\vdash_L$ typically defined inductively using a basis of a set of axioms and closing under a set of inference rules. Given a logic L, a set of formulae Γ is *L-inconsistent* if there are formulae $\alpha_1,\dots,\alpha_m \in \Gamma$, such that $\vdash_L \neg(\alpha_1,\dots,\alpha_m)$, and *$L$-consistent* otherwise. A set of formulae Γ is *maximal* if for every α of the language either $\alpha \in \Gamma$ or $\neg\alpha \in \Gamma$. Under appropriate conditions, it is possible to prove that every L-consistent set admits a maximal L-consistent extension.

Given a multi-modal logic L, the canonical model $M_C^L = (W, R_1,\dots,R_n,\pi)$ is a model for the logic L, built as follows. The set W is made of all the maximal L-consistent sets of formulae, R_i is a family of relations on W^2 defined by wR_iw' if $\forall\alpha$ ($\Box_i\alpha \in w$ implies $\alpha \in w'$). The interpretation π for the atoms is defined as $p \in \pi(w)$ if $p \in w$. For normal modal logics L that are "compact" in the sense that all rules of inference have a finite number of antecedents, the canonical model has the property that $M_C^L \models \phi$ if and only if $\vdash_L \phi$.

A logic L is *sound* with respect to a class of frames $\mathcal{F}$ if $\vdash_L \phi$ implies $\mathcal{F} \models \phi$. A logic L is *complete* with respect to a class of frames $\mathcal{F}$ if $\mathcal{F} \models \phi$ implies $\vdash_L \phi$. Some logics are not only described by the canonical model but also by the frame of the canonical model, called the canonical frame. It can be proved that completeness of a logic L with respect to a class of frames $\mathcal{F}$ holds if the frame of the canonical model is in $\mathcal{F}$. Define the logic S5WD$_n$ to be the logic obtained from S5$_n$ by adding the axiom **WD**. It is possible to prove its completeness with respect to EWD frames.

THEOREM 4.1. *The logic S5WD$_n$ is sound and complete for $\mathcal{L}_n$ with respect to the class of EWD frames.*

PROOF. Soundness follows from what was proved in the first part of Lemma 4.2 and the fact that all axioms and rules of S5$_n$ are sound for equivalence frames [Halpern and Moses 1990]. To prove completeness we use the canonical model technique. It is easy to show that the frame $F_C^{\mathrm{S5WD}_n} = (W, R_1,\dots,R_n)$ of the canonical model for S5WD$_n$ is reflexive, symmetric, and transitive with respect to the n relations. We prove it is also WD.

Suppose that $w_0, w_1,\dots,w_n$ are worlds of $F_C^{\mathrm{S5WD}_n}$ such that $w_0R_{j_i}w_i$, for $i = 1,\dots,n$. Consider the set

$$\Gamma = \bigcup_{i=1}^{n} \{\phi : \Box_i\phi \in w_i\}.$$

We show that Γ is S5WD$_n$-consistent. It then follows by the maximal extension theorem that there is a maximal S5WD$_n$-consistent extension w, which satisfies $w_i R_i w$ by construction. This will establish that the frame is WD. To show Γ is S5WD$_n$-consistent, we assume it is not S5WD$_n$-consistent and obtain a contradiction. It follows from the assumption that for each $i = 1, \ldots, n$ there are formulae $\alpha_1^i, \ldots, \alpha_{m_i}^i$, with $\Box_i \alpha_1^j \in w_i$ for each $j = 1 \ldots m_i$, such that

$$\vdash_{\mathrm{S5WD}_n} \neg(\alpha_1^1 \wedge \cdots \wedge \alpha_{m_1}^1 \wedge \cdots \wedge \alpha_1^n \wedge \cdots \wedge \alpha_{m_n}^n)$$

Let us now call $\alpha_i = \wedge_{j=1}^{m_i} \alpha_j^i$. Note that by S5$_n$ reasoning, we have $\Box_i \alpha_i \in w_i$. It follows that $\Diamond_{j_i} \Box_i \alpha_i \in w_0$. (For else, $\Box_{j_i} \neg \Box_i \alpha_i \in w_0$, hence $\neg \Box_i \alpha_i \in w_i$, contradicting consistency of w_i.) By propositional logic we obtain $S \Box_i \alpha_i \in w_0$. Thus, $\bigwedge_{i=1,\ldots,n} S \Box_i \alpha_i \in w_0$. Now the formulae $\Box_i \alpha_i$ are i-local, so using **WD** it follows that $SS(\bigwedge_{i=1,\ldots,n} \Box_i \alpha_i) \in w_0$. By S5$_n$ reasoning we get $SS(\bigwedge_{i=1,\ldots,n} \alpha_i) \in w_0$. But by S5$_n$ reasoning and the fact that $\vdash_{\mathrm{S5WD}_n} \neg \bigwedge_{i=1,\ldots,n} \alpha_i$ this leads to the conclusion that w_0 is inconsistent. This is the contradiction promised. □

Applying the equivalences with respect to $\mathcal{L}_n$ established previously, we also obtain soundness and completeness with respect to several other semantics.

Corollary 4.1. *The logic S5WD$_n$ is sound and complete for $\mathcal{L}_n$ with respect to*

(1) the class of full systems;
(2) the class of hypercube systems;
(3) the class of EDI frames;
(4) the class of ED frames.

With Corollary 4.1 we have the axiomatisation of full systems and hypercubes systems that we aimed for. We remark that it can be shown that several other axioms could be used for this result instead of **WD**. For example, an analysis of the proofs of both the correspondence and completeness results reveals that the axiom

$$\left(\bigwedge_{i=1,\ldots,n} S \Box_i \phi_i \right) \Rightarrow SS \left(\bigwedge_{i=1,\ldots,n} \Box_i \phi_i \right)$$

(where the ϕ_i are not required to be i-local) also suffices. This is not surprising, since for any i-local formula ϕ it can be shown that $\phi \Leftrightarrow \Box_i \phi$ is S5$_n$-valid.

While the axiom **WD** is compact when expressed using the operator S, it is quite lengthy when expanded and involves considerable use of disjunction. It is possible to show that **WD** may be replaced by certain other axioms which are less symmetrical, but which state interactions between the agents' knowledge of a syntactically simpler form than the expansion of **WD**. For a discussion of a number of alternative axioms that can be shown to be equivalent to **WD**, we refer the reader to the thesis of Lomuscio [Lomuscio 1999]. One such alternative, for the case $n = 2$, has appeared in the literature before, as the axiom

$$\Diamond_1 \Box_2 p \Rightarrow \Box_2 \Diamond_1 p$$

due to Catach [Catach 1988], also discussed in [Popkorn 1994].

Theorem 4.2. **WD** *in the case $n = 2$ is S5$_2$-equivalent to Catach's axiom.*

PROOF. If $n = 2$ then **WD** is

$$(\Diamond_1\phi_1 \vee \Diamond_2\phi_1) \wedge (\Diamond_1\phi_2 \vee \Diamond_2\phi_2) \Rightarrow \bigvee_{i,j\in\{1,2\}} \Diamond_i\Diamond_j(\phi_1 \wedge \phi_2)$$

(with ϕ_i i-local).

WD to Catach: Put $\phi_1 = \Box_1\neg p$ and $\phi_2 = \Box_2 p$ (note that these are 1-local and 2-local respectively). **WD** now becomes

$$(\Diamond_1\Box_1\neg p \vee \Diamond_2\Box_1\neg p) \wedge (\Diamond_1\Box_2 p \vee \Diamond_2\Box_2 p) \Rightarrow \bot.$$

Now we drop the disjuncts $\Diamond_1\Box_1\neg p$ and $\Diamond_2\Box_2 p$ (this strengthens the antecedent and hence weakens the whole formula) to obtain as a consequence

$$\Diamond_2\Box_1\neg p \wedge \Diamond_1\Box_2 p \Rightarrow \bot,$$

which can be simply rearranged to obtain $\Diamond_1\Box_2 p \Rightarrow \Box_2\Diamond_1 p$ as required.

Catach to **WD**: From $\Diamond_1\Box_2 p \Rightarrow \Box_2\Diamond_1 p$ we want to obtain

$$(\Diamond_1\phi_1 \vee \Diamond_2\phi_1) \wedge (\Diamond_1\phi_2 \vee \Diamond_2\phi_2) \Rightarrow \bigvee_{i,j\in\{1,2\}} \Diamond_i\Diamond_j(\phi_1 \wedge \phi_2)$$

in the case that the ϕ_i are i-local. Since the ϕ_i are i-local, we have $\Box_i\phi_i \Leftrightarrow \phi_i$. Assume

$$(\Diamond_1\Box_1\phi_1 \vee \Diamond_2\Box_1\phi_1) \wedge (\Diamond_1\Box_2\phi_2 \vee \Diamond_2\Box_2\phi_2)$$

which, on distribution, is

$$(\Diamond_1\Box_1\phi_1 \wedge \Diamond_1\Box_2\phi_2) \vee (\Diamond_1\Box_1\phi_1 \wedge \Diamond_2\Box_2\phi_2) \vee (\Diamond_2\Box_1\phi_1 \wedge \Diamond_1\Box_2\phi_2) \vee (\Diamond_2\Box_1\phi_1 \wedge \Diamond_2\Box_2\phi_2)$$

We will derive from each of these disjuncts, either $\Diamond_1\Diamond_2(\phi_1 \wedge \phi_2)$ or $\Diamond_2\Diamond_1(\phi_1 \wedge \phi_2)$, thus proving WD. The derivations are as follows:

(1) From $(\Diamond_1\Box_1\phi_1 \wedge \Diamond_1\Box_2\phi_2)$, apply Catach's axiom together with uniform substitution to the second term to obtain $(\Diamond_1\Box_1\phi_1 \wedge \Box_2\Diamond_1\phi_2)$. Use the $S5_n$ axioms $\Diamond_1\Box_1\psi \Leftrightarrow \Box_1\psi$ and $\Box_2\Diamond_1\psi \Rightarrow \Diamond_1\psi$ to obtain $\Box_1\phi_1 \wedge \Diamond_1\phi_2$. From this we deduce $\Diamond_1(\phi_1 \wedge \phi_2)$ and from the axiom T: $p \Rightarrow \Diamond_2 p$ and substitution we obtain $\Diamond_2\Diamond_1(\phi_1 \wedge \phi_2)$.

(2) From $(\Diamond_1\Box_1\phi_1 \wedge \Diamond_2\Box_2\phi_2)$: the first conjunct gives $\Box_1\phi_1$, then ϕ_1, then $\Diamond_2\phi_1$ by $S5_n$ axioms. The second conjunct gives $\Box_2\phi_2$, so putting them together we have $\Diamond_2\phi_1 \wedge \Box_2\phi_2$, from which we obtain $\Diamond_2(\phi_1 \wedge \phi_2)$ as a consequence, and hence $\Diamond_1\Diamond_2(\phi_1 \wedge \phi_2)$.

(3) From $(\Diamond_2\Box_1\phi_1 \wedge \Diamond_1\Box_2\phi_2)$, we obtain $\Box_1\Diamond_2\phi_1 \wedge \Diamond_1\Box_2\phi_2$ by applying Catach to the first term. This now implies $\Diamond_1(\Diamond_2\phi_1 \wedge \Box_2\phi_2)$, which in turn implies $\Diamond_1\Diamond_2(\phi_1 \wedge \phi_2)$.

(4) From $(\Diamond_2\Box_1\phi_1 \wedge \Diamond_2\Box_2\phi_2)$: this case is similar to the first one.

□

5. DECIDABILITY AND COMPLEXITY

In this section we investigate two important meta-logical properties of the logics S5WD$_n$: decidability and computational complexity.

5.1 Decidability

In order to prove decidability of $\mathrm{S5WD}_n$, we show that the logic has the finite model property.

DEFINITION 5.1. *A logic L is said to have the finite model property (or fmp in short) if for any formula ϕ, $\not\vdash_L \phi$ implies that there is a finite model M for L such that $M \not\models \phi$.*

A logic can be proved to have the fmp in a number of different ways: algebraically as in [McKinsey 1941; Bergmann 1949], by the use of a "mini-canonical" model as in [Hughes and Cresswell 1996], etc. Here we use the another standard technique which is better suited for this case: *filtrations* (first presented in [Lemmon 1977]).

The idea of filtrations is the following. If a logic is complete, we know that if a formula ϕ is a non-theorem of L (i.e. if $\neg\phi$ is L-consistent), then ϕ is invalid on some model M for L. The model M might be infinite. Filtrations enable us to produce a model M' from M, such that M' is finite. If we can further prove that M' is also a model for L, then we have proved that the logic L has the finite model property.

We formally proceed as follows. Given a formula ϕ, define the set Φ_ϕ to be the set of formulae α that are either a sub-formula of ϕ or the negation of a sub-formula of ϕ. The set Φ_ϕ is obviously finite for any formula ϕ.

DEFINITION 5.2. *Let M be a model. Two worlds w, w' of M are equivalent with respect to Φ_ϕ (denoted $w \equiv_{\Phi_\phi} w'$, or simply $w \equiv w'$ if it is not ambiguous), if for every $\alpha \in \Phi_\phi$, we have $M \models_w \alpha$ if and only if $M \models_{w'} \alpha$.*

We can now define *filtrations* as follows.

DEFINITION 5.3. *Given a formula ϕ and a model $M = (W, R_1, \ldots, R_n, \pi)$, a* filtration *through Φ_ϕ is a model $M' = (W', R'_1, \ldots, R'_n, \pi')$ satisfying the following three properties:*

—*$W' = W/\equiv_{\Phi_\phi}$, where $\equiv_{\Phi_\phi}$ is the equivalence relation defined as in 5.2.*
—*For each $i \in A$, the relation R'_i is* suitable*, i.e. it satisfies the two properties:*
  (1) *For all $[w_1], [w_2] \in W'$, if there exists $u \in W$ such that $w_1 R_i u$ and $u \equiv w_2$, then $[w_1] R'_i [w_2]$.*
  (2) *For all $[w_1], [w_2] \in W'$, if $[w_1] R'_i [w_2]$ then for all formulae α such that $\Box_i \alpha \in \Phi_\phi$, if $M \models_{w_1} \Box_i \alpha$ then $M \models_{w_2} \alpha$.*
—*For any $p \in Atoms$, $p \in \pi'([w])$ if and only if $p \in \pi(w)$.*

Note that a model M' satisfying these conditions must be finite since Φ_ϕ is finite, so the number of equivalence classes under $\equiv_{\Phi_\phi}$ is finite. Indeed, the number of worlds in M' is at most $2^{|\phi|}$,

It can be proved by induction (see for example [Hughes and Cresswell 1984] page 139) that suitability of the relations R'_i guarantees the validity of the following:

THEOREM 5.1. *Given a model M, and any formula ϕ, a filtration M' of M through Φ_ϕ has the property that for any point $w \in W$ and and for any formula $\alpha \in \Phi$, we have $M' \models_{[w]} \alpha$ if and only if $M \models_w \alpha$*

We now proceed to the case of interest here: the logic $\mathrm{S5WD}_n$. Consider the canonical model M for $\mathrm{S5WD}_n$. We know (see Theorem 4.1) that M is a weakly-directed equivalence model. By Lemma 3.8, the model generated by any point of M is directed. Consider any formula ϕ. We consider the model M' defined as follows:

DEFINITION 5.4. *Given a model M and a formula ϕ define the model $M' = (W', \sim'_1, \ldots, \sim'_n, \pi')$ by*

—$W' = W/\equiv_{\Phi_\phi}$, *where $\equiv_{\Phi_\phi}$ is the equivalence relation defined by Definition 5.2.*

—$[w_1] \sim'_i [w_2]$ *if for all formulae α such that $\Box_i\alpha \in \Phi_\phi$, we have $M \models_{w_1} \Box_i\alpha$ if and only if $M \models_{w_2} \Box_i\alpha$.*

—*For any $p \in Atoms$, we have $p \in \pi'([w])$ if and only if $p \in \pi(w)$.*

Indeed the model M' defined by Definition 5.4 is a filtration as the following shows (stated in [Hughes and Cresswell 1984] page 145 for the mono-modal case).

LEMMA 5.1. *Given an equivalence model M and a formula ϕ, the model M' as described in Definition 5.4 is a filtration of M through Φ_ϕ.*

PROOF. All we need to prove is that the relations $\sim'_i$ are suitable.

Property 1. Consider worlds $[w_1], [w_2] \in W'$ and world $u \in W$ such that $w_1 \sim_i u$ and $u \equiv w_2$. We need to prove that $[w_1] \sim'_i [w_2]$, i.e. that for all formulae α such that $\Box_i\alpha \in \Phi_\phi$we have $M \models_{w_1} \Box_i\alpha$ if and only if $M \models_{w_2} \Box_i\alpha$. We prove it from left to right; the other direction is similar. Note that $M \models_{w_1} \Box_i\alpha$ if and only if $M \models_{w_1} \Box_i\Box_i\alpha$ because M is an equivalence model; but $w_1 \sim_i u$ and so $M \models_u \Box_i\alpha$. But $\Box_i\alpha \in \Phi$ and $w_2 \equiv u$, so $M \models_{w_2} \Box_i\alpha$, which is what we wanted to prove.

Property 2. Consider worlds $[w_1], [w_2] \in W'$ such that $[w_1] \sim'_i [w_2]$. This means that for all $\Box_i\alpha \in \Phi$, we have $M \models_{w_1} \Box_i\alpha$ if and only if $M \models_{w_2} \Box_i\alpha$. Since M is an equivalence model it follows that $M \models_{w_2} \alpha$. $\Box$

We now prove that the filtration defined above produces models for S5WD$_n$. We first consider the effect of the filtration on directed models.

LEMMA 5.2. *If M is an equivalence directed model, then the model M' defined in Definition 5.4, is also an equivalence directed model.*

PROOF. We prove that $F' = (W', \sim_1, \ldots, \sim'_n)$ is an ED frame. The relations $\sim'_i$ are clearly equivalence relations. All it remains to show is that F' is directed. To do that, consider any $[w_1], \ldots, [w_n] \in W'$. Since M is directed, there exists $w \in W$ such that $w_i \sim_i w$ for $i = 1, \ldots, n$. But each $\sim'_i$ is suitable and so, by a consequence of property 1 of suitability we have that $[w_i] \sim'_i [w]$, for $i = 1, \ldots, n$. Therefore the frame F' is directed. $\Box$

We are finally in the position to prove fmp.

THEOREM 5.2. *The logic* S5WD$_n$ *has the finite model property. Indeed, every formula ϕ with a counter-model has a counter-model with at most $2^{|\phi|}$ worlds.*

PROOF. Suppose $\not\models \phi$. Since by the proof of Theorem 4.1 the logic S5WD$_n$ is canonical, the canonical model $M = (W, \sim_1, \ldots, \sim_n, \pi)$ for S5WD$_n$ is an equivalence model, it is weakly-directed and there is a point $w \in W$, such that $M \models_w \neg\phi$. Consider the model M_w generated by w. By Lemma 2.1, we have $M_w \models_w \neg\phi$. The model M_w is clearly an equivalence model and, since it is connected it is also directed, by Lemma 3.8. Consider now the filtration M' of M_w through Φ_ϕ according to Definition 5.4; by Lemma 5.2, M' is an equivalence directed model and it is finite by construction because Φ_ϕ is a finite set. But M' is a filtration, and by Theorem 5.1, $M' \models_{[w]} \neg\phi$, which is what we needed to prove. The bound on the size of M' follows from the observation above. $\Box$

It is immediate from this result that the logic S5WD$_n$ is decidable.

5.2 Computational complexity

We now turn our attention to the computational complexity of the satisfaction problem of the logic $\mathrm{S5WD}_n$. In order to do this, we explore the interesting relation between hypercubes and products of modal logic, a topic of recent interest in modal logic[2].

"Combining logics" is a very active area of research in pure logic [Gabbay 1998], and various methodologies to build more complex logics from basic components have been analysed [Finger and Gabbay 1992; Kracht and Wolter 1991; Fine and Schurz 1996]. "Products of modal logics" [Gabbay and Shehtman 1998] constitute one of these techniques; they are semantically defined as follows. Consider two frames $F = (W, R), F' = (W', R')$; the product $F \times F' = (W \times W', S, T)$ is defined by:

$$\langle w, w' \rangle S \langle x, x' \rangle \text{ if } wRx \text{ and } w' = x';$$

$$\langle w, w' \rangle T \langle x, x' \rangle \text{ if } w'R'x' \text{ and } w = x.$$

If $\mathcal{F}, \mathcal{G}$ are two classes of frames, the product of the classes is defined as $\mathcal{F} \times \mathcal{G} = \{\mathcal{F} \times \mathcal{G} \mid \mathcal{F} \in \mathcal{F}, \mathcal{G} \in \mathcal{G}\}$. One of the questions that naturally arise in this setting is to axiomatise the product of two classes of frames. It is shown in [Gabbay and Shehtman 1998] that in certain cases the logic that arises by considering the product of two classes of frames $\mathcal{F}_1, \mathcal{F}_2$ is equal to the fusion [Kracht and Wolter 1991] (i.e. the union with an opportune renaming of the inference rules) of the corresponding two logics L_1, L_2 to which the following two interaction axioms are added:

$$\Diamond_1 \Box_2 p \Rightarrow \Box_2 \Diamond_1 p,$$

$$\Diamond_1 \Diamond_2 p \Rightarrow \Diamond_2 \Diamond_1 p.$$

Indeed, as is proven in [Gabbay and Shehtman 1998], this is the case for equivalence frames. So, the logic of the binary product of two classes of equivalence frames is $\mathrm{S5}_2$ to which one of the two interaction axioms above is added, as it is not hard to show that on equivalence frames the two axioms above correspond to equivalent first-order conditions[3].

The computational complexity of the logic above is shown in [Marx 1999].

THEOREM 5.3. *[Marx 1999] The satisfaction problem of S5 $\times$ S5 is* NEXPTIME-*complete.*

---

[2]We are grateful to an anonymous referee for encouraging us to investigate this issue. We also wish to thank Maarten Marx and Agnes Kurucz for valuable discussions on this topic.

[3]We further note that for the case $n = 2$, the frames generated by hypercubes are a special case of products: the ones in which the relations on the two components are total. The definition of binary products of frames reported above can be extended to the general case of n frames, by considering two tuples being related by relation i whenever all but the i-th components are equal, and the i-th components are related on the original i-th frame (the motivation for this comes from modelling multi-dimensional spaces). Note that this is "complementary" to the definition of hypercubes, and it intuitively generates quite a different class of frames. This is confirmed by the different meta-logical properties that products have. For instance, many products are known to be not finitely axiomatisable; indeed this is the case for K^n [Kurucz 2000], and $\mathrm{S5}^n$ [Johnson 1969]. We conclude that except in the case $n = 2$, the logics $\mathrm{S5WD}_n$ and $\mathrm{S5}^n$ are different. Products of modal logic of this type may also be related to cylindrical modal logic [Venema 1995]; we do not discuss this relation here.

By using Theorem 4.2, Theorem 3.1, and Theorem 4.1, we then have that for the case $n = 2$ the logic of hypercubes is the same as the logic S5 $\times$ S5.

THEOREM 5.4. *For $n \geq 2$ the satisfaction problem of $S5WD_n$ is* NEXPTIME-*complete.*

PROOF. The case of $n = 2$ is immediate from the considerations above. For the general case, we argue as follows.

By Theorem 5.2, to check that ϕ is satisfiable, it suffices suffices to check that ϕ has a model with at most $2^{|\phi|}$ worlds. This can be done by non-deterministically guessing a model of at most this size and checking whether it satisfies ϕ. This yields a NEXPTIME upper bound.

For the lower bound, we use a reduction from the case for $n = 2$. We claim that if ϕ is a formula in $\mathcal{L}_2$ and $n \geq 2$ then ϕ is satisfiable with respect to hypercubes for two agents iff it is satisfiable with respect to hypercubes for n agents. We first show that if ϕ is satisfiable in hypercubes for two agents then it is satisfiable in hypercubes for n agents. Suppose $(L_e \times L_1 \times L_2, \pi) \models_v \phi$, where L_e is a singleton and $v = (l_e, l_1, l_2)$. Take $L_3 = \ldots = L_n = \{l\}$ to be the singleton sets. Then it is straightforward to show that $(L_e \times L_1 \times L_2 \times L_3 \times \ldots \times L_n, \pi') \models_{v'} \phi$, where $v' = (l_e, l_1, l_2, l, \ldots, l)$ and π' is defined by $\pi'(l_e, x_1, x_2, l, \ldots, l) = \pi(l_e, x_1, x_2)$.

Conversely, we show that if ϕ is satisfied in a hypercube for n agents, then ϕ is satisfied in a hypercube for 2 agents. For this, suppose that $(L_e \times L_1 \times L_2 \times L_3 \times \ldots \times L_n, \pi) \models_v \phi$. Defining $L'_e = L_e \times L_3 \times \ldots \times L_n$, we may view the hypercube $(L_e \times L_1 \times L_2 \times L_3 \times \ldots \times L_n, \pi)$ for n agents as isomorphic to a full system $(L'_e \times L_1 \times L_2, \pi')$, so ϕ is satisfied in a full system for two agents. Thus, $\neg\phi$ is not valid with respect to full systems for two agents. By the results of Theorem 3.1, and Theorem 3.2, a formula is valid with respect to hypercubes for 2 agents iff it is valid in full systems for two agents (note that the constructions used in the proofs preserve the number of agents.) It follows that $\neg\phi$ is not valid with respect to hypercubes for two agents. Thus, there exists an interpreted hypercube for two agents on which ϕ is satisfied. $\square$

We remark that satisfiability of S5 is NP-complete [Ladner and Reif 1977], and the logic $\mathrm{S5}_n$ is known to be PSPACE-complete [Halpern and Moses 1992]. In the latter case, the bound is obtained by a tableau construction that improves upon the bound obtained from the finite model property. The fact that $\mathrm{S5WD}_n$ is NEXPTIME hard suggests no similar construction is possible for this logic.

6. HOMOGENEOUS BROADCAST SYSTEMS

Hypercubes were motivated above as an appropriate model for the initial configuration of a multi-agent system, in which all agents are ignorant of each other's local state. In this section we will show that for a particular class of systems, homogeneous broadcast systems with perfect recall, hypercubes are also an appropriate model of the states of knowledge of agents that acquire information over time. In this class of systems, all communication is by synchronous broadcast, agents have perfect recall, and the agents' knowledge in the initial configuration is characterised by a hypercube system. We establish that in such systems, the agents' knowledge can be characterised by a hypercube system not just at the initial time, but also at all subsequent times. It follows from this result that the logic of knowledge in homogeneous broadcast systems can be axiomatised by the logic $\mathrm{S5WD}_n$ studied in the previous section. Thus, the applicability of hypercubes as a model of agents' knowledge extends beyond initial configurations.

This section is organised as follows. In Section 6.1 we describe *environments*, a general model for the behaviour of agents and their interaction. This model may be used in a variety of ways to ascribe a state of knowledge to the agents after a particular sequence of events has occurred. We focus here on just one of the possibilities, in which it is assumed that agents have *perfect recall* of their observations. In Section 6.2, we define broadcast environments, a special case of this general model that constrains all communication between agents to be by synchronous broadcast. Section 6.3 considers the special case of homogeneous broadcast environments, establishes the connection between the systems generated by these environments and hypercubes.

6.1 Environments

In the model of Halpern and Moses [Halpern and Moses 1990], a distributed system corresponds to a set of runs, where each run constitutes a history that identifies at each point of time a state of the environment and a local state for each agent. This model is perhaps overly general, since in practice one is interested in the particular sets of runs that are generated by executing a given program, or *protocol*, within a given communication architecture. A formal framework to capture this idea, *contexts*, was defined by Fagin et al. [Fagin et al. 1997; Fagin et al. 1995]. In this section we briefly recall a variant of this framework, *environments*, from [Meyden 1996a]. (We refer the reader to [Meyden 1996a; Fagin et al. 1995] for more extensive motivation and examples.) Compared to contexts, environments admit an additional degree of freedom by allowing knowledge to be interpreted in different ways in the same set of runs. We focus here on a particular interpretation, based on the assumption that agents have perfect recall. We describe how executing a protocol in an environment with respect to an interpretation of knowledge determines a Kripke structure that ascribes a state of knowledge to the agents after the occurrence of a particular sequence of events. In Section 6.2, we will present a special case of this model that defines a particular architecture in which agents communicate by synchronous broadcast.

For the definition of environment, we assume a set $A = \{0, 1, \ldots, n\}$ of agents. We also assume that for each agent $i \in A$, there is a non-empty set ACT_i, representing the set of *actions* that may be performed by agent i. A *joint action* is defined to be a tuple $\langle a_0, \ldots, a_n \rangle \in ACT_0 \times \cdots \times ACT_n$. We write ACT for the set of joint actions. As before, we assume a set *Atoms* of propositional variables of the language.

In the following definitions, agent 0 will play a role somewhat different from the other agents. Intuitively, it is intended that agent 0 be used to model aspects of the context, or communication architecture, within which the other agents operate. The actions of agent 0 correspond to nondeterministic behaviour of this context. In applications of the framework, the architecture is typically fixed, and one is interested in designing programs for the behaviour of agents $1 \ldots n$.

Definition 6.1 (Environment). *An* interpreted environment *is a tuple of the form $E = \langle S, I, P_0, \tau, O, V \rangle$ where the components are defined as follows:*

—S *is a set of* states of the environment. *Intuitively, states of the environment may encode such information as messages in transit, failure of components, etc.*

—I *is a subset of* S, *representing the possible* initial states *of the environment.*

—$P_0 : S \rightarrow \mathcal{P}(ACT_0)$ *is a function, called the* protocol of the environment, *mapping states to subsets of the set* ACT_0 *of actions performable by the environment. Intuitively,* $P_0(s)$ *represents the set of actions that may be performed by the environment when the system is in state* s.

—τ *is a function mapping joint actions* $\mathbf{j} \in ACT$ *to state transition functions* $\tau(\mathbf{j}) : S \rightarrow S$. *Intuitively, when the joint action* $\mathbf{j}$ *is performed in the state* s, *the resulting state of the environment is* $\tau(\mathbf{j})(s)$.

—O *is a function from* S *to* $\mathcal{O}^n$ *for some set* $\mathcal{O}$. *For each* $i = 1, \dots, n$, *the function* O_i *mapping* $s \in S$ *to the* i*th component of* $O(s)$, *is called the* observation function of agent i. *Intuitively,* $O_i(s)$ *represents the* observation *of agent* i *in the state* s.

—$V : S \times Atoms \rightarrow \{0, 1\}$ *is a valuation, assigning a truth value* $V(s, p)$ *in each state* s *to each atomic proposition* $p \in Atoms$.

A *trace* of an environment E is a *finite* sequence $s_0 \dots s_m$ of states such that $s_0 \in I$ and for all $k = 0 \dots m-1$ there exists a joint action $\mathbf{j} = \langle a_0, a_1, \dots, a_n\rangle$ such that $s_{k+1} = \tau(\mathbf{j})(s_k)$ and $a_0 \in P_0(s_k)$. We write *fin*(r) for the final state of a trace r.

Intuitively, the traces of an environment correspond to the finite histories that may be obtained from some behaviour of the agents in that environment.[4] Note that the nondeterministic choices of action made by the environment itself are constrained by the protocol of the environment, and that these choices are determined at each step from the state of the environment. On the other hand, the notion of trace assumes that the choices of action of agents $1 \dots n$ are unconstrained. In practice, we wish these agents to behave according to some program (perhaps nondeterministic), that determines their choice of next possible action as some function of the observations that they have made. The following definition captures this intuition.

DEFINITION 6.2 (PERFECT RECALL). *The* perfect recall local state *of agent* $i = 1 \dots n$ *in a trace* $r = s_0 \dots s_m$, *denoted* $\{r\}_i$, *is defined to be the sequence* $O_i(s_0) \dots O_i(s_m)$ *of observations made by the agent in the trace.*

A perfect recall protocol *for agent* $i = 1 \dots n$ *is a function* P_i *mapping each sequence of observations in* $\mathcal{O}^*$ *to a non-empty subset of* ACT_i. *A* joint perfect recall protocol *is a tuple* $\mathbf{P} = \langle P_1, \dots, P_n\rangle$ *consisting of a perfect recall protocol* P_i *for each agent* $i = 1 \dots n$. *We write* $\mathbf{P}_i$ *for* P_i *when* $\mathbf{P}$ *is given.*

Protocols specify the possible choices of next action of the agents, given a certain history of events, as follows. For each agent $i = 1 \dots n$, we say that an action $a_i \in ACT_i$ is *enabled* with respect to a protocol $\mathbf{P}$ at a trace r of E if $a_i \in \mathbf{P}_i(\{r\}_i)$.

[4] One could also define *runs* of the environment, which are infinite sequences of states satisfying the same constraint on state transitions. This would correspond more closely to the framework of [Fagin et al. 1995]. Runs are essential when one is interested in languages containing temporal operators, but there is a precise sense in which it suffices to work with traces when only modal operators for knowledge are of interest, as in the present paper. See the appendix of [Meyden 1998] for a discussion of this issue.

An action a_0 of the environment is enabled at r if $a_0 \in P_0(\mathit{fin}(r))$. A joint action is enabled at r with respect to a protocol $\mathbf{P}$ if each of its components is enabled at r.

We obtain the traces that result when agents execute a joint protocol in an environment as follows. Define a trace $s_0 \dots s_m$ of E to be *consistent* with a joint protocol $\mathbf{P}$ if for each $k < m$, there exists a joint action $\mathbf{j}$ enabled at $s_0 \dots s_k$ with respect to $\mathbf{P}$, such that $\tau(\mathbf{a})(s_k) = s_{k+1}$. We are now in a position to describe the frame that captures the agents' states of knowledge when they execute a protocol in an environment.

DEFINITION 6.3 (PERFECT RECALL FRAME DERIVED FROM A PROTOCOL AND ENVIRONMENT). *Given an environment E and a joint protocol $\mathbf{P}$, the* perfect recall frame *derived from E and $\mathbf{P}$ is the structure $F_{E,\mathbf{P}} = (W, \sim_1, \dots, \sim_n)$, where:*

—*W is the set of all traces of E consistent with $\mathbf{P}$,*

—*$\sim_i$ is the binary relation on W defined by $r \sim_i r'$ if $\{r\}_i = \{r'\}_i$, for each agent $i = 1 \dots n$.*

Intuitively, because W contains only traces of E consistent with $\mathbf{P}$, this frame encodes the assumption that it is common knowledge amongst the agents that the environment in which they are operating is E and that the protocol they are running is $\mathbf{P}$. Moreover, the accessibility relations $\sim_i$ expressing agents' knowledge are defined in a way that corresponds to assuming that agents have perfect recall of their observations. The relations $\sim_i$ could have been defined in many different ways: for example, it is meaningful to consider instead the relations $\approx_i$ defined by $r \approx_i r'$ if $O_i(\mathit{fin}(r)) = O_i(\mathit{fin}(r'))$. This would correspond to the assumption that agents are only aware of their most recent observation. The assumption of perfect recall we work with in this paper is frequently made in the literature because it amounts to assuming that agents make optimal use of the information to which they are exposed. This assumption is essential for the derivation of lower bounds and impossibility results and the synthesis of optimal protocols [Halpern and Moses 1990; Moses and Tuttle 1988; Halpern et al. 1990].

6.2 Broadcast environments

In this subsection we define *broadcast environments* (BE), a special case of the formalism described in Section 6.1. Broadcast environments model situations in which all communication is by synchronous broadcast. Examples of this are systems in which agents communicate by means of a shared bus, by writing tokens onto a shared blackboard [Nii 1986] and in face to face conversation. Other examples are classical puzzles such as the wise men, or muddy children puzzle [Moses et al. 1986], and a variety of games of incomplete information, including battleships, Stratego and Bridge. Broadcast environments have been considered previously in [Meyden 1996b].

To define broadcast environments, we need to impose a number of constraints on the components making up the definition of environments given in the previous section. We do so here in a way that slightly simplifies the model in [Meyden 1996b], eliminating some features that will be irrelevant in the context of homogeneous broadcast environments. The intuition we wish to capture is that each agent holds some private information, which is unobservable to all other agents. The actions taken by the agents will have two types of effects: they will update this private information, and simultaneously broadcast some information to all the other agents.

The actions performed by agents in broadcast environments have two components: an *internal* component and an *external* component. The internal component of an agent's action will affect only the agent's private state, and will be unobservable to the other agents. On the other hand, the external component will be observable to all agents, but it will affect only the state of the environment.

ASSUMPTION 6.1 (BE ACTIONS). *For each $i = 0 \dots n$ there exists a set A_i of* external actions *and a set B_i of* internal actions. *All the sets A_i contain the special "null" action ϵ. For $i = 0 \dots n$, the set ACT_i of actions of agent i consists of the pairs $a \cdot b$ where $a \in A_i$ and $b \in B_i$.*

The role of the null action is to allow for a uniform representation of initial states (see Assumption 6.3 below). It follows from Assumption 6.1 and the definitions of the previous section that the set of the joint actions ACT in a broadcast environment consists of the tuples of the form $\mathbf{j} = \langle a_0 \cdot b_0, a_1 \cdot b_1, \dots, a_n \cdot b_n \rangle$ where $a_i \cdot b_i \in ACT_i$ for each $i = 0 \dots n$. We define $\mathbf{a}(\mathbf{j})$ to be the component $\langle a_0, a_1, a_2, \dots, a_n \rangle$, and call this the *joint external component* of $\mathbf{j}$. We write $\mathbf{A}$ for the set $A_0 \times \dots \times A_n$ of joint external actions.

To represent the private information held by agents, we assume that for each agent $i = 0 \dots n$ there exists a set S_i of *instantaneous private states*. Intuitively, for $i = 1 \dots n$ the states S_i represent the information observable by agent i only. In the case $i = 0$, the states S_0 represent that part of the environment's state which is observable to no agent.

ASSUMPTION 6.2 (BE STATES). *The set of states S of a broadcast environment is required to consist of tuples of the form $\langle a_0, \dots, a_n; p_0, \dots, p_n \rangle$, where for each $i = 0 \dots n$, the component $a_i \in A_i$ is an external action of agent i and the component $p_i \in S_i$ is a private state of agent i.*

Intuitively, a tuple $\langle a_0, \dots, a_n; p_0, \dots, p_n \rangle$ models a situation in which each agent i is in the instantaneous private state p_i, and in which a_i is the most recent external action performed by agent i. If $s = \langle a_0, \dots, a_n; p_0, \dots, p_n \rangle$ is a state then we define the *joint external action at s*, denoted $\mathbf{a}(s)$, to be the tuple $\langle a_0, \dots, a_n \rangle$. We define the *joint private state at s* to be the tuple $\mathbf{p}(s) = \langle p_0, \dots, p_n \rangle$, and *agent i's private state at s*, denoted $\mathbf{p}_i(s)$, to be the private state p_i.

Clearly, in initial states it does not make sense to talk of a most recent external action. This motivates the following.

ASSUMPTION 6.3 (BE INITIAL STATES). *The set of initial states I of a broadcast environment contains only states $\langle a_0, \dots, a_n; p_0, \dots, p_n \rangle$ with $a_i = \epsilon$ for all $i = 0 \dots n$.*

The definition of broadcast environment allows the set I of initial states to be any nonempty set of states of this form. As we will see later, homogeneous broadcast environments restrict the possible sets of initial states.

In a broadcast environment, agents are aware of their own private state, and also of the external actions performed by all agents. All communication between agents will be by means of the external actions. This constraint is modelled by the definition of the agents' observations.

ASSUMPTION 6.4 (BE OBSERVATIONS). *For $i = 1 \dots n$, we require agent i's observation function O_i to be given by $O_i(\langle a_0, \dots, a_n; p_0, \dots, p_n \rangle) = \langle a_0, \dots, a_n; p_i \rangle$.*

That is, in a given state, an agent's observation consists of the external component of the joint action producing that state, and the agent's private state.[5]

It will be convenient in what follows to introduce an observation function for agent 0, similarly defined by $O_0(\langle a_0, \dots, a_n; p_0, \dots, p_n \rangle) = \langle a_0, \dots, a_n; p_0 \rangle$. Moreover, we obtain using this observation function an equivalence relation $\sim_0$ on traces, defined exactly as the relations $\sim_i$.

One of the effects of performing a joint action in a broadcast environment is that each agent updates its private state in a way that depends on its internal action and the joint external action simultaneously being performed. In addition to this, the joint external action will be recorded in the resulting state.

ASSUMPTION 6.5 (BE TRANSITIONS). *For each agent $i = 0 \dots n$ there exists a private action interpretation function $\tau_i : \mathbf{A} \times B_i \to (S_i \to S_i)$. The joint action interpretation function $\tau : ACT \to (S \to S)$ of a broadcast environment is obtained from the private action interpretation functions as follows. For each joint action $\mathbf{j} = \langle a_0 \cdot b_0, \dots, a_n \cdot b_n \rangle$, the transition function $\tau(\mathbf{j})$ maps a state $s = \langle a'_0, \dots, a'_n; p_0, \dots, p_n \rangle$ to*

$$\tau(\mathbf{j})(s) = \langle a_0, \dots, a_n; \tau_0(\mathbf{a}(\mathbf{j}), b_0)(p_0), \dots, \tau_n(\mathbf{a}(\mathbf{j}), b_n)(p_n) \rangle.$$

That is, for each joint external action $\mathbf{a} \in \mathbf{A}$ and internal action $b_i \in B_i$, the function $\tau_i(\mathbf{a}, b_i) : S_i \to S_i$ is a private state transition function, intuitively representing the effect on agent i's private states of performing the internal action b_i when the joint external action $\mathbf{a}$ is being simultaneously performed. The state of the environment resulting from a joint action records the external component of the joint action, and updates each agent's private state using its private action interpretation function.

Finally, we require that the protocol P_0 of the environment depend only upon its private state and the most recent external action.

ASSUMPTION 6.6 (BE PROTOCOL). *If s and t are states with $O_0(s) = O_0(t)$ then $P_0(s) = P_0(t)$.*

The propositional constants *Atoms* of a broadcast environment are allowed to describe any property of the global states S, so we do not make any assumption on the valuation V.

We may now state the main definition of this section.

DEFINITION 6.4 (BROADCAST ENVIRONMENT). *A* broadcast environment *is an environment satisfying assumptions 6.1—6.6.*

In broadcast environments, agents' mutual knowledge can be shown to have a particularly simple structure [Meyden 1996b]. Intuitively, this is because broadcast communication maintains a high degree of common knowledge. The following section provides an illustration of this point in a special case of broadcast environments.

[5]This is a slight simplification of the definition in [Meyden 1996b], eliminating an extra component that is incompatible with homogeneity.

EXAMPLE 6.1. *We illustrate the definitions presented so far by modelling a game theoretic example: a repeated game in which each of two players follows some predetermined strategy, which is not known to the other. We suppose that each player chooses in each round an action from a set A. A strategy defines a choice of an action based on the outcome of the previous rounds, i.e., a strategy is a mapping $\sigma : (A^2)^* \to A$. We suppose that T is a set of strategies, representing the possible types of players in the game. Applications of this model include reasoning about auctions that proceed in a number of rounds, where each player makes a bid in each round that is then revealed to the other. A strategy in this application would correspond to a policy for deciding the next bid based on the previous bids in the auction, and the player's personal valuation of the item for sale. As the game proceeds, each player acquires some information about the other's valuation and bidding policy.*

To model this situation as a broadcast environment, we may take the set of private states S_0 of the environment to be any singleton set $\{s_0\}$. We use the agents' instantaneous private states to capture their strategies, by taking $S_i = T$ for $i = 1, 2$. Since the players may choose to follow any strategy from T, we take the set of initial states of the environment to be given by $I = \{\langle \epsilon, \epsilon, \epsilon; s_0, \sigma_1, \sigma_2 \rangle \mid \sigma_1, \sigma \in T\}$.

Actions of the environment are irrelevant in this example, so we also take $A_0 = B_0 = \{\epsilon\}$. Thus, $ACT_0 = \{\epsilon \cdot \epsilon\}$. Clearly, this means that the environment's protocol must be the mapping defined by $P_0(s_e) = \{\epsilon \cdot \epsilon\}$. When $i = 1$ or $i = 2$, we model the internal actions of agent i trivially by the singleton set $B_i = \{\epsilon\}$. The external actions A_i we take to be the set $A \cup \{\epsilon\}$, where A is the set of moves that the agents can take in the game.

Thus, a state of the environment is given by a tuple $\langle \epsilon, a_1, a_2; s_0, \sigma_1, \sigma_2 \rangle$ where $a_1, a_2 \in A \cup \{\epsilon\}$ and $\sigma_1, \sigma_2 \in T$. Since a player's strategy is predetermined, and does not change during the game, its private action interpretation function is defined by taking each function $\tau_i(\mathbf{a}(\mathbf{j}), b_i)$ to be the identity function on T, when $i = 1, 2$. In the case $i = 0$, each $\tau_i(\mathbf{a}(\mathbf{j}), b_0)$ is the identity function on $\{s_0\}$. It follows that the effect of a joint action $\mathbf{j} = \langle \epsilon \cdot \epsilon, a_1 \cdot \epsilon, a_2 \cdot \epsilon \rangle$ is to map the state $\langle \epsilon, a_1', a_2'; s_0, \sigma_1, \sigma_2 \rangle$ to the state $\langle \epsilon, a_1, a_2; s_0, \sigma_1, \sigma_2 \rangle$.

At each step of the game, a player is aware of its own strategy and observes the result of the previous round. We note that this is in accordance with the observation functions $O_i(\langle \epsilon, a_1, a_2; s_0, \sigma_1, \sigma_2 \rangle) = \langle \epsilon, a_1, a_2; \sigma_i \rangle$ required by the definition of a broadcast environment.

Moreover, note that the information available to an agent i's perfect recall protocol after k rounds is a sequence $\Sigma = \langle \epsilon, \epsilon, \epsilon; \sigma_i \rangle \langle \epsilon, a_1^1, a_2^1; \sigma_i \rangle \ldots \langle \epsilon, a_1^k, a_2^k; \sigma_i \rangle$. This contains both the strategy σ_i and the sequence $(a_1^1, a_2^1) \ldots (a_1^k, a_2^k)$ required by the agent's strategy to compute its next move. Thus, we may define the agent's perfect recall protocol by

$$P_i(\Sigma) = \sigma_i((a_1^1, a_2^1) \ldots (a_1^k, a_2^k)).$$

This states that the player follows the strategy encoded in its private state.

In the sequel, we will make use of the following observation.

LEMMA 6.1. *Suppose E is a broadcast environment, and $\mathbf{P}$ is a joint perfect recall protocol. Let r and r' be traces of E consistent with $\mathbf{P}$ such that $r \sim_i r'$, where $i \in \{0 \ldots n\}$. Then every action of agent i that is enabled at r is also enabled at r'.*

The proof is immediate from the definitions. (In the case of agent 0, note that $r \sim_0 r'$ implies $O_0(r) = O_0(r')$ and use Assumption 6.6.)

6.3 Homogeneous broadcast environments

Homogeneous broadcast environments are a special case of broadcast environments. These environments satisfy the additional, and quite natural, constraint that agents start in a condition of ignorance about each others states, and the state of the environment. Thus, their initial state of knowledge is characterised by a hypercube system. We will show that, under the assumption that agents have perfect recall, their knowledge can also be characterised as a hypercube system at all subsequent times.

DEFINITION 6.5 (HOMOGENEOUS BROADCAST ENVIRONMENT). *A broadcast environment E is* homogeneous *if there exists for each agent $i = 0 \ldots n$ a set $I_i \subseteq P_i$ of* initial private states*, such that the set of initial states I of the environment E is the set of all states $\langle \epsilon, \ldots, \epsilon; p_0, \ldots p_n \rangle$, where $p_i \in I_i$ for $i = 0 \ldots n$.*

In other words, the set of initial states is isomorphic to the Cartesian product $I_0 \times \ldots \times I_n$. This system is a hypercube if we view agent 0 on a par with the other agents in the system; if we view the states of agent 0 as providing the environment component L_e, then this Cartesian product is a full system. In either event, agents are initially ignorant of each others' states and the state of the environment. The environment in Example 6.1 is a homogeneous broadcast system. Of the other examples mentioned in the previous section, battleships and Stratego satisfy this constraint, but the wise men puzzle, the muddy children puzzle and Bridge do not. (For example, the initial configurations of Bridge, i.e. after cards have been dealt but before bidding, do not form a hypercube because it is not possible for two players to simultaneously hold the same card.)

We may now introduce the main object of study in this section.

DEFINITION 6.6 (PERFECT RECALL HOMOGENEOUS BROADCAST FRAME). *A* perfect recall homogeneous broadcast frame *is any frame $F_{E,\mathbf{P}}$ obtained from a joint perfect recall protocol $\mathbf{P}$ in a homogeneous broadcast environment E.*

We are now in a position to state the main result of this section.

THEOREM 6.1. *Every perfect recall homogeneous broadcast frame is isomorphic to a frame obtained from a disjoint union of systems of the form $X_0 \times X_1 \times \ldots \times X_n$. In particular, every such frame is weakly-directed.*

This result establishes a close connection between perfect recall homogeneous broadcast frames and hypercube systems. In particular, it follows that the logic S5WD$_n$ is sound for this class of frames.

For the proof, it is convenient to introduce the following notions. If $r = s_0 s_1 \ldots s_m$ is a trace of a broadcast environment, we will write $\mathbf{a}(r)$ for the sequence $\mathbf{a}(s_0) \ldots \mathbf{a}(s_m)$ of joint external actions performed in r. If $s = \langle a_0, \ldots, a_n; p_0, \ldots, p_n \rangle$ and $t = \langle a_0, \ldots, a_n; q_0, \ldots, q_n \rangle$ are global states with the same joint external action component, and $i \in \{0 \ldots n\}$, define $s \bowtie_i t$ to be the state $\langle a_0, \ldots, a_n; p_0, \ldots, p_{i-1}, q_i, p_{i+1}, \ldots, p_n \rangle$, that is like s except that agent i has the private state it has in t.

Note that for $j \neq i$, we have $O_j(s \bowtie_i t) = \langle a_0, \ldots, a_n; p_j \rangle = O_j(s)$. Additionally, $O_i(s \bowtie_i t) = \langle a_0, \ldots, a_n; q_i \rangle = O_i(t)$. More generally, if $r_1 = s_0 s_1 \ldots s_m$ and

$r_2 = t_0t_1 \dots t_m$ are sequences of states of the same length with $\mathbf{a}(r_1) = \mathbf{a}(r_2)$ then we define $r_1 \bowtie_i r_2$ to be the sequence $(s_0 \bowtie_i t_0)(s_1 \bowtie_i t_1)\dots(s_m \bowtie_i t_m)$. The following result states a closure condition of the set of traces of a homogeneous broadcast environment.

LEMMA 6.2. *Let E be a homogeneous broadcast environment, and $\mathbf{P}$ a joint perfect recall protocol. If r_1 and r_2 are traces in $F_{E,\mathbf{P}}$ with $\mathbf{a}(r_1) = \mathbf{a}(r_2)$ then for any i we have that $r_1 \bowtie_i r_2$ is a trace in $F_{E,\mathbf{P}}$ with $(r_1 \bowtie_i r_2) \sim_i r_2$ and for $j \neq i$ we have that $(r_1 \bowtie_i r_2) \sim_j r_1$.*

PROOF. Note that $\mathbf{a}(r) = \mathbf{a}(r')$ implies that r and r' have the same length. It is immediate from the comments above that $(r_1 \bowtie_i r_2) \sim_j r_1$ for $j \neq i$ and $(r_1 \bowtie_i r_2) \sim_i r_2$. It therefore suffices to show that $r_1 \bowtie_i r_2$ is a trace. We do this by induction on the length of the trace r_1.

The base case is straightforward. If r_1 is a trace of length one, then it consists of an initial state s_1. Similarly, r_2 consists of an initial state s_2. It is immediate from the assumption that $I = \{\langle\epsilon, \dots, \epsilon\rangle\} \times I_0 \times \dots \times I_n$ that $r_1 \bowtie_i r_2 = s_1 \bowtie_i s_2$ is a trace.

Assume that the result has been established for traces of length m, and consider traces r_1 and r_2 of length $m+1$ with $\mathbf{a}(r_1) = \mathbf{a}(r_2)$. Write $r_1 = r_1' s_1 t_1$ where t_1 is the final state of r_1 and s_1 is the next-to-final state of r_1, and similarly write $r_2 = r_2' s_2 t_2$. By the induction hypothesis, $r = r_1' s_1 \bowtie_i r_2' s_2$ is a trace indistinguishable to agent i from $r_2' s_2$, and indistinguishable to all other agents from $r_1' s_1$.

Let $\mathbf{j}_1$ be a joint action enabled at $r_1' s_1$ such that $t_1 = \tau(\mathbf{j}_1)(s_1)$, and similarly, let $\mathbf{j}_2$ be a joint action enabled at $r_2' s_2$ such that $t_2 = \tau(\mathbf{j}_2)(s_2)$. Note that because state transitions record the joint external action component of a joint action in the resulting state, and because $\mathbf{a}(r_1) = \mathbf{a}(r_2)$, we have $\mathbf{a}(\mathbf{j}_1) = \mathbf{a}(t_1) = \mathbf{a}(t_2) = \mathbf{a}(\mathbf{j}_2)$. Write $\langle a_0, \dots, a_n\rangle$ for the common joint external action of these states and joint actions. Then we may also write $\mathbf{j}_1 = \langle a_0 \cdot b_0, \dots, a_n \cdot b_n\rangle$ and $\mathbf{j}_2 = \langle a_0 \cdot c_0, \dots, a_n \cdot c_n\rangle$. To show that $r_1 \bowtie_i r_2$ is a trace we show that the joint action

$$\mathbf{j} = \langle a_0 \cdot b_0, \dots, a_{i-1} \cdot b_{i-1}, a_i \cdot c_i, a_{i+1} \cdot b_{i+1}, \dots, a_n \cdot b_n\rangle$$

is enabled at r and satisfies $\tau(\mathbf{j})(s_1 \bowtie_i s_2) = t_1 \bowtie_i t_2$.

To show that $\mathbf{j}$ is enabled at r we show that each of its components is enabled at r. In the case of agents $j \neq i$, we need to show that the action $a_j \cdot b_j$ of agent j is enabled at r. This follows, using Lemma 6.1, from the fact that $a_j \cdot b_j$ is enabled for agent j at $r_1' s_1$, and from the fact that $r_1' s_1 \sim_j r$. For agent i, we need to show that the action $a_i \cdot c_i$ is enabled at r. This follows, again using Lemma 6.1, from the fact that $a_i \cdot c_i$ is enabled for agent i at $r_2' s_2$, and from the fact that $r_2' s_2 \sim_i r$.

It therefore remains to show that $\tau(\mathbf{j})(s_1 \bowtie_i s_2) = t_1 \bowtie_i t_2$. Note first that $\mathbf{a}(\tau(\mathbf{j})(s_1 \bowtie_i s_2)) = \mathbf{a}(\mathbf{j}) = \mathbf{a}(\mathbf{j}_1) = \mathbf{a}(t_1 \bowtie_i t_2)$. Thus, the states $\tau(\mathbf{j})(s_1 \bowtie_i s_2)$ and $t_1 \bowtie_i t_2$ record the same joint external action. We show that they also have the same private state for each agent. In case of agents $j \neq i$, we have

$$\begin{aligned}
\mathbf{p}_j(\tau(\mathbf{j})(s_1 \bowtie_i s_2)) &= \tau_j(\mathbf{a}(\mathbf{j}), b_j)(\mathbf{p}_j(s_1 \bowtie_i s_2)) \\
&= \tau_j(\mathbf{a}(\mathbf{j}), b_j)(\mathbf{p}_j(s_1)) \\
&= \mathbf{p}_j(t_1) \\
&= \mathbf{p}_j(t_1 \bowtie_i t_2)
\end{aligned}$$

In case of agent i, we have

$$\begin{aligned}\mathbf{p}_i(\tau(\mathbf{j})(s_1 \bowtie_i s_2)) &= \tau_i(\mathbf{a}(\mathbf{j}), c_i)(\mathbf{p}_i(s_1 \bowtie_i s_2)) \\ &= \tau_i(\mathbf{a}(\mathbf{j}), c_i)(\mathbf{p}_i(s_2)) \\ &= \mathbf{p}_i(t_2) \\ &= \mathbf{p}_i(t_1 \bowtie_i t_2)\end{aligned}$$

This completes the proof. □

Note that because agents observe the most recent joint external action, if r and r' are traces in $F_{\mathbf{P},E}$ with $r \sim_i r'$ then these traces were generated by the same sequence of joint external actions, i.e., $\mathbf{a}(r) = \mathbf{a}(r')$. It follows from this that if r and r' are in the same connected component of $F_{\mathbf{P},E}$ then we also have $\mathbf{a}(r) = \mathbf{a}(r')$. In fact, we have the following stronger result:

LEMMA 6.3. *For every trace r we have that:*

—*the connected component F of $F_{\mathbf{P},E}$ containing r consists of all traces r' in $F_{\mathbf{P},E}$ with $\mathbf{a}(r') = \mathbf{a}(r)$ and*

—*this connected component is isomorphic to the frame obtained from the Cartesian product $\Pi_{i=0\dots n}\{\{r'\}_i \mid r' \in F\}$.*

PROOF. 1) We prove that r is connected to r' if and only if $\mathbf{a}(r) = \mathbf{a}(r')$. Left to right is immediate from Definition 6.3. Right to left follows from Lemma 6.2.

2) For each agent $i = 0 \dots n$, let r_i be any trace of $F_{\mathbf{P},E}$ with $\mathbf{a}(r_i) = \mathbf{a}(r)$. To show that the connected component is a Cartesian product, we prove that there is an r' such that $r' \sim_i r_i$. In fact, define $r' = (\dots(r_1 \bowtie_2 r_2)\dots \bowtie_{n-1} r_{n-1}) \bowtie_n r_n$. By Lemma 6.2, r' is a trace of E, with $r' \sim_i r_i$ for all $i = 0 \dots n$ and $\mathbf{a}(r') = \mathbf{a}(r)$. It is immediate that all traces r' of $F_{\mathbf{P},E}$ with $\mathbf{a}(r') = \mathbf{a}(r)$ are connected, and that the component containing r is isomorphic to the Cartesian product $\Pi_{i=0\dots n}\{\{r'\}_i \mid r' \in F\}$. □

This lemma characterises the sense in which agents' states of knowledge at times other than time 0 in a homogeneous broadcast system are characterised by a hypercube system. Theorem 6.1 follows immediately from Lemma 6.3.

We now obtain a result that provides one final characterisation of the logic S5WD$_n$.

THEOREM 6.2. *The logic S5WD$_n$ is sound and complete for the class of all homogeneous broadcast frames.*

PROOF. Soundness is direct from Theorem 6.1. For completeness, suppose that ϕ is not a theorem of S5WD$_n$. Since S5WD$_n$ is complete for the class of all hypercubes, there exists a hypercube $H = L_e \times L_1 \times \dots L_n$, where L_e is a singleton, an interpretation π_H on H, and a world $w \in S$ such that $(F(H), \pi_H) \models_w \neg\phi$. We show that it is possible to construct a homogeneous broadcast environment E whose decomposition into a union of Cartesian products contains H as one of its components. Indeed H will be the component consisting of all the traces of length one, i.e., the component characterising the initial state of knowledge of the agents.

We define the environment $E = \langle S, I, P_0, \tau, O, V\rangle$ as follows. For each agent $i = 0 \dots n$, we take the both the set of external actions A_i and the set of internal actions B_i to be the set $\{\epsilon\}$. Thus, the set of actions of each agent is also a singleton, viz $\{\epsilon \cdot \epsilon\}$. The components of the environment are as follows:

—The set of states $S = \langle \epsilon, \ldots, \epsilon; p_0, \ldots, p_n \rangle$ where $(p_0, \ldots, p_n) \in H$. Thus, the set S_i of instantaneous private states of agent i is exactly the set of local states L_i of agent i in H.

—All states are initial, i.e. $I = S$.

—Since the set actions ACT_0 of agent 0, the environment, is a singleton, the protocol of the environment is the unique function $P_0 : S \to ACT_0$.

—The transition function τ is defined by $\tau(\mathbf{j})(s) = s$ for (the unique) joint action $\mathbf{j}$ and state s. (Thus, similarly, the local transition functions τ_i satisfy $\tau_i(\mathbf{a}, b_i)(p_i) = p_i$ for (the unique) joint external action $\mathbf{a}$, (the unique) internal action b_i and private state $p_i \in L_i$.)

—The definition of the observation function O is determined by the fact that E is a broadcast environment, i.e. $O_i(\langle \epsilon, \ldots, \epsilon; p_0, \ldots, p_n \rangle) = \langle \epsilon, \ldots, \epsilon; p_i \rangle$ for each $i = 1, \ldots, n$.

—The valuation V is defined by $V(\langle \epsilon, \ldots, \epsilon; p_0, \ldots, p_n \rangle, q) = \pi_H((p_0, \ldots, p_n), q)$.

This is a homogeneous broadcast environment by construction. It is now straightforward to establish that for *every* joint perfect recall protocol $\mathbf{P}$, the connected component of $F_{E,\mathbf{P}}$ consisting of all traces of length one is isomorphic to $(F(H), \pi_H)$. (We remark that our choice of action sets and transition function above are not actually relevant to this conclusion.) □

One way to understand Theorem 6.2 is that it states completeness of S5WD$_n$ with respect to a class of models, namely those models obtained by adding an interpretation to a homogeneous broadcast frame. In these models the interpretation could assign to a proposition a meaning at a trace that depends not just on the final state of the trace, but also on prior states and actions. The proof of Theorem 6.2 in fact establishes that S5WD$_n$ is complete for a smaller class of models with underlying homogeneous broadcast frames, in which the interpretation π is derived from the environment. Given an environment E with valuation V, define the interpretation π_E by $\pi_E(r, p) = V(\mathit{fin}(r), p)$ for traces r of E and propositions $p \in \mathit{Atoms}$.

THEOREM 6.3. *The logic S5WD$_n$ is sound and complete for the class of all models of the form $(F_{E,\mathbf{P}}, \pi_E)$, where E is a homogeneous broadcast environment and $\mathbf{P}$ is a joint protocol.*

PROOF. Similar to the proof of Theorem 6.2. Note that the construction of this proof uses only the initial component of the frame. The valuation of the environment may be chosen to operate as required on this initial component. □

These results are in some respects similar to results of Fagin et al. [Fagin et al. 1992; Fagin and Vardi 1986], who show that that there exist natural classes of systems with respect to which the logic of knowledge is not characterised by S5$_n$, but by a stronger logic ML_n^-, that consists of S5$_n$ plus the following axiom:

$$\beta \ \& \ K_i(\beta \Rightarrow \neg\alpha) \Rightarrow K_1\neg\alpha \vee \ldots \vee K_1\neg\alpha$$

where α is a *primitive state formula* (intuitively, describing the assignment associated with the current state, but not of agents knowledge), and β is a *pure knowledge formula* (intuitively, describing properties of the agents knowledge but not dealing with the assignment associated with the current state). We refer the reader to

[Fagin et al. 1992] for a precise explanation of these terms. In particular, one class of systems to which this result applies is a class of systems in which the assignment is static, agents communicate by unreliable synchronous message passing and have perfect recall [Fagin and Vardi 1986].

Theorem 6.2 provides another interesting and natural class of systems that requires additional axioms. In our result, agents also have perfect recall, but the class is otherwise quite different from those considered in [Fagin et al. 1992; Fagin and Vardi 1986] since our agents communicate by reliable broadcast, and we allow the assignment to vary from moment to moment. The axiom (WD) we need to capture such systems is also quite different from that used by Fagin et al.

We remark that it is possible to prove a variant of the results of this section that deal with full systems rather than hypercubes. For this variant, we modify the definition of homogeneity to state that the initial states of the environment form a full system (treating agent 0 as part of the environment). Moreover, instead of Assumption 6.6, we assume that for all states s and t with the same joint external action, i.e., $\mathbf{a}(s) = \mathbf{a}(t)$, and for all external actions a_0 of agent 0, there exists an internal action b_0 of agent 0 such that $a_0 \cdot b_0 \in P_0(s)$ iff there exists an internal action b_0' of agent 0 such that $a_0 \cdot b_0' \in P_0(s)$. (Informally, this means that an external action of agent 0 is enabled in s iff it is enabled in t.) The environment and protocol in Example 6.1 satisfy both these assumptions.

Under these assumptions, Lemma 6.2 holds provided we restrict i and j to range over agents 1 to n only (i.e., we exclude agent 0.) The proof is a trivial adaptation. Consequently, we also obtain an analogue of Lemma 6.3 stating that the connected components of $F_{E,\mathbf{P}}$ are full systems.

7. CONCLUSIONS

In this paper we have formally investigated several classes of interpreted systems that arise by considering the full Cartesian product of the local state spaces. We have argued that these interpreted systems provide an appropriate model for the initial configurations of many systems of interest. Moreover, we have shown that a similar constraint arises at all later configurations in the special case of homogeneous broadcast systems. By relating these classes of systems to several classes of Kripke frames, we have established that a single modal logic, $\mathrm{S5WD}_n$, provides a sound and complete axiomatisation in all these cases. On the conceptual level, this logic provides a well motivated example of interaction among agents' knowledge.

In conducting this work, we have identified the interesting class of WD equivalence frames that generates a complete and decidable logic whose satisfaction problem is NEXPTIME complete. 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 $\mathrm{S5}_n$. This raises the question of whether it is possible to prove similar results for weaker logics, such as $\mathrm{S4}_n$, which model agents that do not have negative introspection capabilities.

Our results leave open many other questions. It should be noted that the fact that the same logic $\mathrm{S5WD}_n$ axiomatises all the different classes of semantic structures we have studied is due in part to the limited expressive power of the language we have considered. It would be interesting to investigate more expressive languages containing operators such as distributed knowledge and common knowledge [Fagin et al. 1995]. In the former case we have already identified the axiom $\phi \Leftrightarrow D_A\phi$ as of interest with respect to equivalence I frames (Lemma 3.4).

It would also be of interest to determine which of the language extensions contemplated above maintain decidability of the logic. Finally, for the dynamic model we have considered, extensions of the language to include temporal operators are of interest. Indeed, consideration of the logic of knowledge and time in homogeneous broadcast systems is just one example of a range of unresolved issues concerning the knowledge of agents operating within specific communications models: a great deal of work remains to be done in the axiomatisation of logics of knowledge with respect to such models.

REFERENCES

AUMANN, R. J. 1976. Agreeing to disagree. *Annals of Statistics 4*, 6, 1236–1239.

BERGMANN, G. 1949. The finite representations of S5. *Methodos 1*, 217–219.

CATACH, L. 1988. Normal multimodal logics. In T. M. SMITH AND G. R. MITCHELL Ed., *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI'88)* (St. Paul, MN, Aug. 1988), pp. 491–495. Morgan Kaufmann.

CHELLAS, B. 1980. *Modal Logic*. Cambridge University Press, Cambridge.

DWORK, C. AND MOSES, Y. 1990. Knowledge and common knowledge in a byzantine environment. *Information and Computation 88*, 2, 156–186.

FAGIN, R., HALPERN, J. Y., MOSES, Y., AND VARDI, M. Y. 1995. *Reasoning about Knowledge*. MIT Press, Cambridge.

FAGIN, R., HALPERN, J. Y., MOSES, Y., AND VARDI, M. Y. 1997. Knowledge-based programs. *Distributed Computing 10*, 4, 199–225.

FAGIN, R., HALPERN, J. Y., AND VARDI, M. Y. 1992. What can machines know? On the properties of knowledge in distributed systems. *Journal of the ACM 39*, 2 (April), 328–376.

FAGIN, R. AND VARDI, M. Y. 1986. Knowledge and implicit knowledge in a distributed environment: Preliminary report. In J. Y. HALPERN Ed., *TARK: Theoretical Aspects of Reasoning about Knowledge* (San Francisco (CA), 1986), pp. 187–206. Morgan Kaufmann.

FINE, K. AND SCHURZ, G. 1996. Transfer theorems for stratified multimodal logics. In J. COPELAND Ed., *Logic and Reality, Essays in Pure and Applied Logic. In memory of Arthur Prior*, pp. 169–213. Oxford University Press.

FINGER, M. AND GABBAY, D. M. 1992. Adding a temporal dimension to a logic. *Journal of Logic, Language and Information 1*, 203–233.

GABBAY, D. 1998. *Fibring Logics*. Oxford University Press.

GABBAY, D. AND SHEHTMAN, V. 1998. Products of modal logics, part 1. *Logic Journal of the IGPL 6*, 1, 73–146.

GOLDBLATT, R. 1992. *Logics of Time and Computation, Second Edition, Revised and Expanded*, Volume 7 of *CSLI Lecture Notes*. CSLI, Stanford. Distributed by University of Chicago Press.

HALPERN, J., MEYDEN, R., AND VARDI, M. Y. 1997. Complete axiomatisations for reasoning about knowledge and time. Submitted.

HALPERN, J. AND MOSES, Y. 1990. Knowledge and common knowledge in a distributed environment. *Journal of the ACM 37*, 3, 549–587. A preliminary version appeared in *Proc. 3rd ACM Symposium on Principles of Distributed Computing*, 1984.

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., MOSES, Y., AND WAARTS, O. 1990. A characterization of eventual Byzantine agreement. In *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing* (Quebec City, Quebec, Canada, 22–24 Aug. 1990), pp. 333–346.

HUGHES, G. E. AND CRESSWELL, M. J. 1984. *A Companion to Modal Logic*. Methuen, London.

HUGHES, G. E. AND CRESSWELL, M. J. 1996. *A New Introduction to Modal Logic*. Routledge, New York.

JOHNSON, J. S. 1969. Nonfinitizability of classes of representable polyadic algebras. *Journal of Symbolic Logic 34*, 344–352.

KRACHT, M. AND WOLTER, F. 1991. Properties of independently axiomatizable bimodal logics. *Journal of Symbolic Logic 56*, 4, 1469–1485.

KURUCZ, A. 2000. On axiomatising products of kripke frames. *The Journal of Symbolic Logic*. To appear.

LADNER, R. E. AND REIF, J. H. 1977. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing 6*, 3, 467–480.

LANG, S. 1984. *Algebra* (2 ed.). Addison-Wesley, Reading.

LEMMON, E. J. 1977. *The "Lemmon Notes": An Introduction to Modal Logic*. Basil Blackwell, Oxford. Written in collaboration with D. Scott; edited by K. Segerberg. American Philosophical Quarterly Monograph Series. Monograph No. 11.

LOMUSCIO, A. 1999. *Knowledge Sharing among Ideal Agents*. Ph. D. thesis, School of Computer Science, University of Birmingham, Birmingham, UK.

LOMUSCIO, A. AND RYAN, M. 2000. A spectrum of modes of knowledge sharing between agents. In M. J. WOOLDRIDGE AND N. R. JENNINGS Eds., *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Volume 1757 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag, Heidelberg.

LYNCH, N. 1996. *Distributed Algorithms*. Morgan Kaufmann, San Fancisco.

MARX, M. 1999. Complexity of products of modal logics. *Journal of Logic and Computation 9*, 2, 221–238.

MCCARTHY, J. 1989. Artificial intelligence, logic and formalizing common sense. In R. THOMASON Ed., *Philosophical Logic and Artificial Intelligence*. Kluwer Academic.

MCKINSEY, J. 1941. A solution of the decision problem for the Lewis systems S2 and S4 with an application to topology. *The Journal of Symbolic Logic 6*, 117–134.

MEYDEN, R. 1996a. Knowledge based programs: On the complexity of perfect recall in finite environments. In S. Y Ed., *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge* (San Francisco, March 1996), pp. 31–50. Morgan Kaufmann Publishers.

MEYDEN, R. 1998. Common knowledge and update in finite environments. *Information and Computation 140*, 2 (1 Feb.), 115–157.

MEYDEN, R. V. D. 1996b. Finite state implementations of knowledge-based programs. In *Proceedings of the Conference on Foundations of Software Technology and Theoretical Computer Science*, Volume 1180 of *Lecture Notes in Computer Science* (Hyderabad, India, Dec. 1996), pp. 262–273.

MEYER, J.-J. C. AND HOEK, W. 1995. *Epistemic Logic for AI and Computer Science*, Volume 41 of *Cambridge Tracts in Theoretical Computer Science*. Cambdridge University Press.

MOSES, Y., DOLEV, D., AND HALPERN, J. 1986. Cheating husbands and other stories: A case study of knowledge, action, and communication. *Distributed Computing 1*, 3, 167–176.

MOSES, Y. AND TUTTLE, M. R. 1988. Programming simultaneous actions using common knowledge. *Algorithmica 3*, 121–169.

NII, H. P. 1986. Blackboard systems (part one): the blackboard model of problem solving and the evolution of blackboard architectures. *The AI Magazine Summer*, 38–53.

POPKORN, S. 1994. *First Steps in Modal Logic*. Cambridge University Press: Cambridge, England.

RUBENSTEIN, A. AND WOLINSKY, A. 1990. On the logic of 'agreeing to disagree' type results. *Journal of Economic Theory 51*, 184–193.

VENEMA, Y. 1995. Cylindric modal logic. *Journal of Symbolic Logic 60*, 2, 591–623.