Topology and Knowledge of Multiple Agents

Bernhard Heinemann

Fakultät für Mathematik und Informatik, FernUniversität in Hagen
58084 Hagen, Germany
Bernhard.Heinemann@FernUni-Hagen.de

Abstract. A multi-agent version of Moss and Parikh's logic of knowledge and effort is developed in this paper. This is done with the aid of particular modalities identifying the agents involved in the system in question. Throughout the paper, special emphasis is placed on foundational issues.

Keywords: knowledge of agents, epistemic logic, topological reasoning, spatial logics.

1 Introduction

The development of spatial calculi in AI (like, eg, RCC [1]) has directed the interest of researchers back to one of the earliest interpretations of modal logic, the *topological* one [2].[1] In modern terminology, a formula $\Box\alpha$ holds at some point x of a topological space iff there exists an open neighbourhood of x in which α is everywhere valid; cf [4]. Thus the modal box stands for a kind of nested quantification in this spatial context, viz an existential one concerning (open) sets followed by a universal one concerning points.

Moss and Parikh [5] discovered that breaking up this interlocking of quantifiers brings to light a certain relationship of topology and *knowledge*. In fact, regarding points as *system states* and sets as *knowledge states of some agent,* the above quantification over sets induces an S4-like modality of *shrinking* while the point quantifier, being S5-like, directly describes knowledge. But shrinking a knowledge state means gaining knowledge so that the first modality corresponds to some knowledge acquisition procedure and was called the *effort operator* thus. The semantic domains of the language underlying the Moss-Parikh system are obtained by releasing the class of all admissible structures. Now, every *subset space* can be taken for that, i.e., every triple $(X, \mathcal{O}, V)$ consisting of a non-empty set X of states, a set $\mathcal{O}$ of subsets of X representing the knowledge states of the agent,[2] and a valuation V determining the states where the atomic propositions are true. The *knowledge operator,* K, then quantifies over some knowledge state $U \in \mathcal{O}$, whereas the effort operator, denoted $\Box$ as well, quantifies 'downward' over $\mathcal{O}$.

[1] See [3] for bringing the precise connection with RCC about.
[2] Since these sets are of topological origin they are still called *the opens* sometimes.

H. Geffner et al. (Eds.): IBERAMIA 2008, LNAI 5290, pp. 1–10, 2008.

Following that initial work, a rich theory on the connection between knowledge and topology has emerged; see Ch. 6 of the recent handbook [6] for an overview of the state of the art. The original class of all topological spaces could be characterized within the broader framework, in particular; see [7]. In addition, several extensions including temporal or dynamic aspects of knowledge have been proposed. An essential shortcoming of the system remains up to now though: Regarding its applicability to real-life scenarios in computer science or AI, a corresponding multi-agent version is still missing. (This is at least true regarding the usual semantics as the paper [8] is based on a different one.)

This deficiency is rectified in the present paper. At first glance, we have two options in doing so. On the one hand, we could assume that there is one subset space $(X, \mathcal{O}, V)$ comprising the knowledge states of all the agents. If $\mathcal{O}$ is unstructured, then the agents are semantically indistinguishable. This means that they share all of their knowledge states. Thus the resulting theory reduces to that of a single agent, which is rather uninteresting for our problem. On the other hand, if a separate subset space is associated with every agent, then we get, among other things, into trouble when trying to capture the interplay of the knowledge of the agents. The latter leads us to a possible solution nevertheless: Adressing a particular agent will here be arranged by an appropriate modality distinguishing the agent. This entails that, unlike the usual logic of knowledge (see [9]) we do no longer have a knowledge operator K_A in the formal language, for every agent A. Instead, K_A is 'decomposed' into two modalities, K and the new one belonging to A. Thus the agents are 'super-imposed' on the original system, which, therefore, can essentially be preserved. This idea is carried out in the technical part of this paper.

On looking more carefully, representing an agent by a modal operator seems to be quite natural since it is thereby indicated that agents are acting entities rather than indices to knowledge operators. But most notably, the new approach enables us to present a multi-agent version of Moss and Parikh's system in a smooth and satisfactory way.

The body of the paper is organized as follows. In the next section, we precisely define the spatio-epistemic multi-agent language indicated above. Moreover, we reason about its expressiveness there. Section 3 contains a list of basic subset space validities and a discussion referring to this. Section 4 deals with the question of completeness. We also discuss effectivity issues, in Section 5. Concluding the paper, we sum up and point to variants, extensions, and future research assignments.

2 Multi-agent Subset Spaces

In this section, we define the extended language, $\mathcal{L}$. The syntax of $\mathcal{L}$ is based on a denumerable set $\mathrm{Prop} = \{p, q, \ldots\}$ of symbols called *proposition letters*. Let $m \in \mathbb{N}$. Then, the set Form_m of all *formulas* over Prop is given by the rule $\alpha ::= p \mid \neg\alpha \mid \alpha \wedge \alpha \mid K\alpha \mid \Box\alpha \mid [A_i]\,\alpha$, where $i \in \{1, \ldots, m\}$. The letters A_i should represent the agents involved in the scenario under discussion. While K

and $\Box$ denote the modalities of knowledge and effort as above, the new operators $[A_i]$ identify the respective agent in a way to be made precise in a minute. The duals of K, $\Box$ and $[A_i]$ (i.e., the operators $\neg K \neg$, $\neg \Box \neg$ and $\neg [A_i] \neg$) are denoted L, $\Diamond$ and $\langle A_i \rangle$, respectively. The missing boolean connectives are treated as abbreviations, as needed. – We now fix the relevant semantic domains. It turns out that every agent induces a substructure of the set of all opens in a natural way. Let $\mathcal{P}(X)$ designate the powerset of a given set X.

Definition 1 (Multi-agent structures)

1. *Let X be a non-empty set, $\mathcal{O} \subseteq \mathcal{P}(X)$ be a set of subsets of X such that $X \in \mathcal{O}$, and $\mathcal{A} : \{1, \ldots, m\} \to \mathcal{P}(\mathcal{P}(X))$ be a mapping such that the union $\bigcup \mathrm{Im}(\mathcal{A})$ of all the image sets $\mathcal{A}(i)$, where $i = 1, \ldots, m$, is contained in $\mathcal{O}$. Then $\mathcal{S} := (X, \mathcal{O}, \mathcal{A})$ is called a* multi-agent subset frame.
2. *Let $\mathcal{S} = (X, \mathcal{O}, \mathcal{A})$ be a multi-agent subset frame. Then the elements of the set $\mathcal{N}_\mathcal{S} := \{(x, U) \mid x \in U \text{ and } U \in \mathcal{O}\}$ are called the* neighbourhood situations of *the frame $\mathcal{S}$.*
3. *Let $\mathcal{S} = (X, \mathcal{O}, \mathcal{A})$ be a multi-agent subset frame and $V : \mathrm{Prop} \to \mathcal{P}(X)$ be a mapping. Then V is called an $\mathcal{S}$*-valuation.
4. *Let $\mathcal{S} = (X, \mathcal{O}, \mathcal{A})$ be a multi-agent subset frame and V be an $\mathcal{S}$-valuation. Then $\mathcal{M} := (X, \mathcal{O}, \mathcal{A}, V)$ is called a* multi-agent subset space, *or, in short, a* MASS *(*based on $\mathcal{S}$*).*

Some points are worth mentioning here. First, the function $\mathcal{A}$ is added to subset spaces in such a way that the set of all opens is not necessarily exhausted by the knowledge states of the agents. Having in mind certain spatial settings, this more general approach is quite reasonable. Second, neighbourhood situations are the atomic semantic objects of our language. They will be used for evaluating formulas. In a sense, the set component of a neighbourhood situation measures the uncertainty about the associated state component at any one time. And third, we are mainly interested in *interpreted systems,* which are here formalized by MASSs. The assignment of sets of states to proposition letters by means of valuations is in accordance with the usual logic of knowledge; cf [9] again.

Our next task is defining the relation of satisfaction. This is done with respect to a MASS $\mathcal{M}$. Thus satisfaction, which should hold between neighbourhood situations of the underlying frame and formulas from Form_m, is designated $\models_\mathcal{M}$. In the following, neighbourhood situations are written without brackets.

Definition 2 (Satisfaction and validity). *Let $\mathcal{M} = (X, \mathcal{O}, \mathcal{A}, V)$ be a MASS based on $\mathcal{S} = (X, \mathcal{O}, \mathcal{A})$, and let $x, U \in \mathcal{N}_\mathcal{S}$ be a neighbourhood situation. Then*

$$\begin{array}{lcl} x, U \models_\mathcal{M} p & :\Longleftrightarrow & x \in V(p) \\ x, U \models_\mathcal{M} \neg\alpha & :\Longleftrightarrow & x, U \not\models_\mathcal{M} \alpha \\ x, U \models_\mathcal{M} \alpha \wedge \beta & :\Longleftrightarrow & x, U \models_\mathcal{M} \alpha \textit{ and } x, U \models_\mathcal{M} \beta \\ x, U \models_\mathcal{M} \mathsf{K}\alpha & :\Longleftrightarrow & \textit{for all } y \in U : y, U \models_\mathcal{M} \alpha \\ x, U \models_\mathcal{M} \Box\alpha & :\Longleftrightarrow & \forall U' \in \mathcal{O} : \textit{if } x \in U' \subseteq U, \textit{ then } x, U' \models_\mathcal{M} \alpha \\ x, U \models_\mathcal{M} [A_i]\,\alpha & :\Longleftrightarrow & U \in \mathcal{A}(i) \textit{ implies } x, U \models_\mathcal{M} \alpha, \end{array}$$

for all $p \in \mathrm{Prop}$, $i \in \{1, \ldots, m\}$ *and* $\alpha, \beta \in \mathrm{Form}_m$. *In case* $x, U \models_{\mathcal{M}} \alpha$ *is true we say that* α is valid in $\mathcal{M}$ at *the neighbourhood situation* x, U. *Furthermore, a formula* α *is called* valid in $\mathcal{M}$ *iff it is valid in* $\mathcal{M}$ *at every neighbourhood situation. (Manner of writing:* $\mathcal{M} \models \alpha$.)

Note that the meaning of proposition letters is independent of the opens by definition, hence 'stable' with respect to $\Box$. This fact will find expression in the logical system considered later on. Additionally, note that the operator $[A_i]$ conditions the validity of formulas at the actual neighbourhood situation on agent A_i.

The rest of this section is concerned with some aspects of the expressiveness of $\mathcal{L}$. First of all, we show that the usual knowledge operator associated with an agent A is definable in the new language. In fact, the formula $\langle A_i \rangle K\alpha$ exactly says that A_i knows α at the actual neighbourhood situation (where $i \in \{1, \ldots, m\}$). This means that, with regard to subset spaces, $\mathcal{L}$ is at least as expressive as the usual language for knowledge of agents.

Actually, we have a little more expressive power. To see this, note first that the modalities $[A_i]$ remind one of the binding procedures which are known from *hybrid logic;* cf [10], Sec. 6, or [11], Ch. 14. Though there is a rather weak connection only, a kind of naming system is constituted by these modalities for one component of the semantics (the opens) nevertheless. For instance, the formula $\langle A_i \rangle \top$ means that the actual open represents a knowledge state of agent A_i.

It can easily be inferred from Definition 2 that the knowledge operator and the agent operators commutate in the following sense.

Proposition 1. *Let* $\mathcal{M}$ *be any MASS. Then, for all* $i \in \{1, \ldots, m\}$ *and* $\alpha \in \mathrm{Form}_m$, *we have* $\mathcal{M} \models [A_i] K\alpha \leftrightarrow K [A_i] \alpha$.

Thus it is known that the validity of α is conditioned on agent A_i if and only if it is conditioned on A_i that α is known. This is an example of a basic MASS-validity which will be part of the axiomatization of MASSs we propose in the next section. Concluding this section, we give a concrete example.

Example 1. Let be given a two-agent scenario with three states, x_1, x_2, x_3. Let $X = \{x_1, x_2, x_3\}$. Assume that a knowledge acquisition procedure P_1, which, for $k = 1, 2, 3$, step-by-step eliminates x_k from the set of all possible alternatives, is available to agent A_1. Thus $\mathcal{A}(1) = \{\{x_1, x_2, x_3\}, \{x_2, x_3\}, \{x_3\}\}$. On the other hand, assume that a corresponding procedure P_2, which successively eliminates x_{3-k}, is available to A_2 so that $\mathcal{A}(2) = \{\{x_1, x_2, x_3\}, \{x_1, x_2\}, \{x_1\}\}$. Then, each agent can obviously reach an 'exclusive state of complete knowledge'. In other words, for every $i, j \in \{1, 2\}$ such that $i \neq j$, there is a knowledge state $\in \mathcal{A}(i) \setminus \mathcal{A}(j)$ in which every valid formula α is known by A_i. A specification of this fact reads

$$\bigwedge_{i,j \in \{1,2\},\, i \neq j} L\Diamond \left(\langle A_i \rangle (\alpha \to K\alpha) \wedge [A_j] \bot \right),$$

which is valid at the neighbourhood situation x_k, X, for every $k \in \{1, 2, 3\}$.

3 Axiomatizing Multi-agent Subset Spaces

Our starting point to this section is the system of axioms for the usual logic of subset spaces from [12]. We later add several schemata involving the agent modalities. After that we define the arising multi-agent logic, MAL. – The axioms from [12] read as follows:

1. All instances of propositional tautologies.
2. $K(\alpha \to \beta) \to (K\alpha \to K\beta)$
3. $K\alpha \to (\alpha \wedge KK\alpha)$
4. $L\alpha \to KL\alpha$
5. $\Box(\alpha \to \beta) \to (\Box\alpha \to \Box\beta)$
6. $(p \to \Box p) \wedge (\Diamond p \to p)$
7. $\Box\alpha \to (\alpha \wedge \Box\Box\alpha)$
8. $K\Box\alpha \to \Box K\alpha$,

where $p \in \mathrm{Prop}$ and $\alpha, \beta \in \mathrm{Form}_m$. In this way, it is expressed that, for every Kripke model M validating these axioms,

– the accessibility relation $\stackrel{K}{\longrightarrow}$ of M belonging to the knowledge operator is an equivalence,
– the accessibility relation $\stackrel{\Box}{\longrightarrow}$ of M belonging to the effort operator is reflexive and transitive,
– the composite relation $\stackrel{\Box}{\longrightarrow} \circ \stackrel{K}{\longrightarrow}$ is contained in $\stackrel{K}{\longrightarrow} \circ \stackrel{\Box}{\longrightarrow}$ (this is usually called the *cross property*[3]), and
– the valuation of M is constant along every $\stackrel{\Box}{\longrightarrow}$-path, for all proposition letters (see the remark right after Definition 2).

We now present the axioms containing the operators $[A_i]$:

9. $[A_i]\,(\alpha \to \beta) \to ([A_i]\,\alpha \to [A_i]\,\beta)$
10. $\alpha \to [A_i]\,\alpha$
11. $K\,[A_i]\,\alpha \to [A_i]\,K\alpha$,

where $i \in \{1, \ldots, m\}$ and $\alpha, \beta \in \mathrm{Form}_m$. – Some comments are appropriate here. Item 9 contains the usual distribution schema being valid for every normal modality. The next schema, 10, says that the accessibility relation $\stackrel{A_i}{\longrightarrow}$ belonging to the operator $[A_i]$ has *height 1;* that is, $\forall\, s, t$: if $s \stackrel{A_i}{\longrightarrow} t$, then $s = t$ (where s and t are any points of the frame under discussion). This fact can be proved in the standard way. Finally, item 11 is one half of the commutation relation between the knowledge operator and the agent operators stated in Proposition 1. Note the formal similarity between the schemata 8 and 11. – It is somewhat surprising at first glance that no other interaction axioms between the modalities are needed.

[3] The cross property is caused by Axiom 8, which is, therefore, called the *Cross Axiom;* see [12].

By adding two of the standard rule schemata of modal logic, viz *modus ponens* and *necessitation with respect to each modality* (cf [13], Sec. 1.6), we obtain a logical system denoted MAL. The following proposition about MAL is quite obvious.

Proposition 2 (Soundness). *If $\alpha \in \mathrm{Form}_m$ is MAL-derivable, then α is valid in all MASSs.*

The proof of the opposite assertion, i.e., the completeness of MAL with respect to the class of all MASSs, is highly non-trivial. We turn to it in the next section.

4 Completeness

A completeness proof for MAL must clearly use the associated canonical model in some way. We fix several notations concerning that model first. Let $\mathcal{C}$ be the set of all maximal MAL-consistent sets of formulas. Furthermore, let $\stackrel{K}{\longrightarrow}$, $\stackrel{\Box}{\longrightarrow}$ and $\stackrel{A_i}{\longrightarrow}$ be the accessibility relations induced on $\mathcal{C}$ by the modalities K, $\Box$ and $[A_i]$, respectively. Let $\alpha \in \mathrm{Form}_m$ be non-MAL-derivable. We attain to a multi-agent subset space falsifying α by an infinite 'multi-dimensional' step-by-step construction. In every step, an approximation to the claimed model, which will be the 'limit' of the intermediate ones, is defined. In order to ensure that the final structure behaves as desired, several requirements on the approximations have to be kept under control during the process.

Suppose that $\Gamma_0 \in \mathcal{C}$ is to be realized (i.e., $\neg\alpha \in \Gamma_0$). We choose a denumerably infinite set of points, Y, fix an element $x_0 \in Y$, and construct inductively a sequence of quintuples $(X_n, P_n, i_n, a_n, t_n)$ such that, for every $n \in \mathbb{N}$,

1. $X_n \subseteq Y$ is a finite set containing x_0,
2. P_n is a finite set partially ordered by $\leq$ and containing $\bot$ as a least element,
3. $i_n : P_n \to \mathcal{P}(X_n)$ is a function such that $p \leq q \iff i_n(p) \supseteq i_n(q)$, for all $p, q \in P_n$,
4. $a_n : \{1, \ldots, m\} \to \mathcal{P}(\mathcal{P}(X_n))$ is a function satisfying $\bigcup\limits_{1 \leq i \leq m} a_n(i) \subseteq \mathrm{Im}(i_n)$,
5. $t_n : X_n \times P_n \to \mathcal{C}$ is a partial function such that, for all $x, y \in X_n$ and $p, q \in P_n$,
   (a) $t_n(x, p)$ is defined iff $x \in i_n(p)$; in this case it holds that
      i. if $y \in i_n(p)$, then $t_n(x, p) \stackrel{\mathsf{K}}{\longrightarrow} t_n(y, p)$,
      ii. if $q \geq p$, then $t_n(x, p) \stackrel{\Box}{\longrightarrow} t_n(x, q)$,
      iii. for all $i \in \{1, \ldots, m\}$: if $i_n(p) \in a_n(i)$, then $t_n(x, p) \stackrel{A_i}{\longrightarrow} t_n(x, p)$,
   (b) $t_n(x_0, \bot) = \Gamma_0$.

The next five conditions say to what extent the structures $(X_n, P_n, i_n, a_n, t_n)$ approximate the final model. Actually, it will be guaranteed that, for all $n \in \mathbb{N}$,

6. $X_n \subseteq X_{n+1}$,

7. P_{n+1} is an *end extension* of P_n (i.e., a superstructure of P_n such that no element of $P_{n+1} \setminus P_n$ is strictly smaller than any element of P_n),
8. $i_{n+1}(p) \cap X_n = i_n(p)$ for all $p \in P_n$,
9. for all $i \in \{1, \ldots, m\}$ and $U \in a_{n+1}(i) : U \notin a_n(i)$ or $U \cap X_n \in a_n(i)$,
10. $t_{n+1} \mid_{X_n \times P_n} = t_n$.

Finally, the construction complies with the following requirements on existential formulas: For all $n \in \mathbb{N}$,

11. if $L\beta \in t_n(x,p)$, then there are $n < k \in \mathbb{N}$ and $y \in i_k(p)$ such that $\beta \in t_k(y,p)$,
12. if $\Diamond\beta \in t_n(x,p)$, then there are $n < k \in \mathbb{N}$ and $p \leq q \in P_k$ such that $\beta \in t_k(x,q)$,
13. for all $i \in \{1, \ldots, m\}$: if $\langle A_i \rangle \beta \in t_n(x,p)$, then there is some $n < k \in \mathbb{N}$ such that $i_k(p) \in a_k(i)$ and $\beta \in t_k(x,p)$.

With that, the desired model can easily be defined. Furthermore, a relevant *Truth Lemma* (cf [13], 4.21) can be proved for it, from which the completeness of MAL follows immediately. Thus it remains to construct $(X_n, P_n, i_n, a_n, t_n)$, for all $n \in \mathbb{N}$, in a way meeting all the above requirements. This makes up the core of the proof. The case $n = 0$ is still easy. In the induction step, some existential formula contained in some maximal MAL-consistent set $t_n(x,p)$, where $x \in X_n$ and $p \in P_n$, must be realized. We confine ourselves to the case of the operator $\langle A_i \rangle$ here, where $i \in \{1, \ldots, m\}$.[4]

So let $\langle A_i \rangle \beta \in t_n(x,p)$. The first three components of the approximating structure are unaltered in this case, i.e., $X_{n+1} := X_n$, $P_{n+1} := P_n$, and $i_{n+1} = i_n$. From the *Existence Lemma* of modal logic (cf [13], 4.20) we know that there is some point s of $\mathcal{C}$ such that $t_n(x,p) \xrightarrow{A_i} s$ and $\beta \in s$. Axiom 10 now implies that $t_n(x,p) = s$. We therefore define $a_{n+1}(i) := a_n(i) \cup \{i_n(p)\}$. This already determines a_{n+1} since the agents with index $j \neq i$ are not affected in this case. Finally, we define $t_{n+1} := t_n$.

We must now check that the properties 1 – 10, and 13, remain valid (11 and 12 are irrelevant to the present case). Apart from 5 and 13, all items are more or less obvious from the construction. Concerning item 13, see the remark at the end of this proof outline.

Thus the verification of the property 5 is left. Since 5 (b) too is evident we concentrate on 5 (a). First, the condition on the domain of t_{n+1} is obviously satisfied. Second, (i) and (ii) are clear from the validity of this condition for n. Hence only for (iii) some arguments are needed. If $j \neq i$, where i is from above, then the property is valid because of the induction hypothesis, for $a_{n+1}(j)$ equals $a_n(j)$ then. So let $j = i$ and assume that $i_{n+1}(q) \in a_{n+1}(i)$, where $q \in P_{n+1}$. We distinguish two cases. First, let $i_{n+1}(q) \notin a_n(i)$. Then both $q = p$ and $i_{n+1}(p) = i_n(p)$ must hold due to the construction. It follows from Axiom 11 (and Axiom 10 as well) that $t_{n+1}(y,p) \xrightarrow{A_i} t_{n+1}(y,p)$, where $y \in X_{n+1}$. Second,

[4] Note that the axioms not mentioned below are used for the other cases.

let $i_{n+1}(q) \cap X_n \in a_n(i)$. Then the above construction step forces that $i_{n+1}(q) = i_n(q)$, and the assertion again follows from the induction hypothesis.

In order to ensure that *all* possible cases are eventually exhausted, processing has to be suitably scheduled with regard to each of the modalities involved. This can be done by means of appropriate enumerations. Regarding this and the construction in case of a modality of the usual subset space logic, the reader is referred to the paper [12] for further details. – Summarizing this section, we can state the first of the main results of this paper:

Theorem 1 (Completeness). *If $\alpha \in \mathrm{Form}_m$ is valid in all MASSs, then α is* MAL*-derivable.*

5 Decidability

In this final technical section of the paper, we prove that the set of all MAL-derivable formulas is decidable. Since the finite model property does not apply to the usual subset space logic with respect to the class of all subset spaces (see [12], Sec. 1.3), the same is true for MAL with respect to the class of all MASSs. Thus we have to make a little detour in order to obtain the desired decidability result. We shall single a certain subclass out of the class of all Kripke models and prove that MAL satisfies the finite model property with respect to *that* class of structures. This gives us the decidability of MAL in a standard fashion. – In the following definition, R and K, S and $\Box$, and T_i and $[A_i]$, respectively, correspond to each other ($i \in \{1, \ldots, m\}$).

Definition 3 (MACA-model). *Let $M := (W, R, S, T_1, \ldots, T_m, V)$ be a multi-modal model, where $R, S, T_1, \ldots, T_m \subseteq W \times W$ are binary relations and V is a valuation. Then M is called a* multi-agent cross axiom model *(or, in short, a* MACA-model*) iff the following conditions are satisfied:*[5]

1. *R is an equivalence relation, and S is reflexive and transitive,*
2. *for all $i \in \{1, \ldots, m\}$, the relations T_i have height 1,*
3. *$S \circ R \subseteq R \circ S$ and, for all $i \in \{1, \ldots, m\} : T_i \circ R \subseteq R \circ T_i$ (where $\circ$ denotes composition of relations),*
4. *for all $w, v \in W$ and $p \in \mathrm{Prop}$: if $w\,T\,w'$, then $w \in V(p) \iff v \in V(p)$.*

Note that, by taking neighbourhood situations as points, every subset space induces a semantically equivalent MACA-model.

It is easy to see that all the axioms from Section 3 are sound with respect to the class of all MACA-models. Moreover, the canonical model of MAL belongs to this class of structures. These facts imply the following theorem.

Theorem 2 (Kripke completeness). *The logical system* MAL *is sound and complete with respect to the class of all MACA-models.*

[5] The term 'cross axiom model' was introduced in [12], Sec. 2.3.

In the following, we use the method of *filtration* in order to prove the finite model property of MAL with respect to the class of all MACA-models; cf [14], Sec. 4. For a given MAL-consistent formula $\alpha \in \mathrm{Form}_m$, we define a filter set $\Sigma \subseteq \mathrm{Form}_m$ as follows. We first let $\Sigma_0 := \mathrm{sf}(\alpha) \cup \{\neg\beta \mid \beta \in \mathrm{sf}(\alpha)\}$, where $\mathrm{sf}(\alpha)$ designates the set of all subformulas of α. Second, we form the closure of Σ_0 under finite conjunctions of pairwise distinct elements of Σ_0. Third, we close under single applications of the operator L. And finally, we form the set of all subformulas of elements of the set obtained last. Let Σ then denote the resulting set of formulas.

Now, the *smallest* filtrations of the accessibility relations $\stackrel{K}{\longrightarrow}$, $\stackrel{\Box}{\longrightarrow}$ and $\stackrel{A_i}{\longrightarrow}$ of the canonical model are taken, where $i = 1, \ldots, m$; cf [14], Sec. 4. Let $M := (W, R, S, T_1, \ldots, T_m, V)$ be the corresponding filtration of a suitably generated submodel of the canonical model, for which the valuation V assigns the empty set to all proposition letters not occurring in Σ. Then we have the following lemma.

Lemma 1. *The structure M is a finite MACA-model of which the size computably depends on the length of α.*

Proof. Most of the assertion is clear from the definitions and the proof of [12], Theorem 2.11. Only item 2 and the second part of item 3 from Definition 3 have to be checked. – For item 2, let Γ, Θ be two points of the canonical model such that $[\Gamma]\ T_i\ [\Theta]$, where the brackets $[\ldots]$ indicate the respective classes ($i \in \{1, \ldots, m\}$). Since we are working with smallest filtrations, there are $\Gamma' \in [\Gamma]$ and $\Theta' \in [\Theta]$ such that $\Gamma' \stackrel{A_i}{\longrightarrow} \Theta'$. From this we conclude that $\Gamma' = \Theta'$ since $\stackrel{A_i}{\longrightarrow}$ has height 1. It follows that $[\Gamma] = [\Theta]$. Thus T_i too has height 1. – For item 4, we remember the formal similarity of the axiom schemata 8 and 11 (see Sec. 3). Thus the cross property for K and T_i can be established on the filtrated model in the same way as it was established for K and $\Box$ there; see [12], Lemma 2.10. Note that the special form of Σ is used for exactly that purpose.

Since the model M realizes α according to the *Filtration Theorem* (cf [13], 2.39), the decidability result we strived for follows readily from Theorem 2 and Lemma 1.

Theorem 3 (Decidability). *The set of all* MAL*-derivable formulas is decidable.*

6 Concluding Remarks

We developed a multi-agent version of Moss and Parikh's topological logic of knowledge. To this end, we introduced appropriate modalities addressing the knowledge states of the agents in question. The main issues of the paper are corresponding soundness, completeness and decidability results.

The generalization put forward here is quite natural in many respects. First, the original approach is preserved to a large extent. Second, the new system is open to variations and extensions depending on the applications one has in mind. For example, it is easy to characterize axiomatically those MASSs for which the union of all sets of knowledge states of the agents coincides with the set of all opens. Moreover, further-reaching concepts from the usual logic of knowledge can be incorporated.

We are particularly interested in topological spaces. A corresponding extension of the system presented above as well as the treatment of complexity problems are postponed to future research.

References

1. Randell, D.A., Cui, Z., Cohn, A.G.: A spatial logic based on regions and connection. In: Nebel, B., Rich, C., Swartout, W. (eds.) Principles of Knowledge Representation and Reasoning (KR 1992), San Mateo, CA, pp. 165–176. Morgan Kaufmann, San Francisco (1992)
2. McKinsey, J.C.C.: A solution to the decision problem for the Lewis systems S2 and S4, with an application to topology. Journal of Symbolic Logic 6, 117–141 (1941)
3. Nutt, W.: On the translation of qualitative spatial reasoning problems into modal logics. In: Burgard, W., Christaller, T., Cremers, A.B. (eds.) KI 1999. LNCS (LNAI), vol. 1701, pp. 113–125. Springer, Heidelberg (1999)
4. Aiello, M., van Benthem, J., Bezhanishvili, G.: Reasoning about space: The modal way. Journal of Logic and Computation 13, 889–920 (2003)
5. Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. In: Moses, Y. (ed.) Theoretical Aspects of Reasoning about Knowledge (TARK 1992), Los Altos, CA, pp. 95–105. Morgan Kaufmann, San Francisco (1992)
6. Aiello, M., Pratt-Hartmann, I.E., van Benthem, J.F.A.K.: Handbook of Spatial Logics. Springer, Heidelberg (2007)
7. Georgatos, K.: Knowledge theoretic properties of topological spaces. In: Masuch, M., Polos, L. (eds.) Logic at Work 1992. LNCS (LNAI), vol. 808, pp. 147–159. Springer, Heidelberg (1994)
8. Pacuit, E., Parikh, R.: The logic of communication graphs. In: Leite, J.A., Omicini, A., Torroni, P., Yolum, p. (eds.) DALT 2004. LNCS (LNAI), vol. 3476, pp. 256–269. Springer, Heidelberg (2005)
9. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Computer Science. Cambridge Tracts in Theoretical Computer Science, vol. 41. Cambridge University Press, Cambridge (1995)
10. Blackburn, P.: Representation, reasoning, and relational structures: a hybrid logic manifesto. Logic Journal of the IGPL 8, 339–365 (2000)
11. Blackburn, P., van Benthem, J., Wolter, F.: Handbook of Modal Logic. Studies in Logic and Practical Reasoning, vol. 3. Elsevier, Amsterdam (2007)
12. Dabrowski, A., Moss, L.S., Parikh, R.: Topological reasoning and the logic of knowledge. Annals of Pure and Applied Logic 78, 73–110 (1996)
13. Blackburn, P., de Rijke, M., Venema, Y.: Modal Logic. Cambridge Tracts in Theoretical Computer Science, vol. 53. Cambridge University Press, Cambridge (2001)
14. Goldblatt, R.: Logics of Time and Computation, 2nd edn. CSLI Lecture Notes, vol. 7. Center for the Study of Language and Information, Stanford (1992)