Mesh Models of Knowledge in Distributed Systems

Marek A. Suchenek
Henrietta Okeke
California State University Domingues Hills
Carson, CA 90747 U.S.A.
e-mail addr: Suchenek@filly.calstate.edu

Abstract

In this paper theoretical aspects of knowledge in distributed systems are investigated. A nonmonotonic multi-modal variant of logic S5, deductively complete with respect to a Kripke-style minimal-knowledge semantics, is introduced. It allows for formal verification and analysis of information and its flow within a distributed system, either by proof-theoretic methods or by computing models of system's knowledge. The introduced system is applied to Mr. Sum and Mr. Product Puzzle *which was not properly handled by existing methods of modal logic.*

Key words: *distributed knowledge, epistemic logic, minimal-knowledge semantics, multi-modal logic S5, nonmonotonic logic.*

AMS classification: **68T30**, 03C40, 03B45.

1 Introduction

Applications of modal logic in reasoning about knowledge represented in digital systems belong to main stream research in Artificial Intelligence. A number of formal systems have been proposed and studied in professional literature, e.g. in [MD80, Moo85, MaT91, Kam91]. They involve a modal operator K with the intentional interpretation "one knows that ...". The distinguishing property of these systems is the so called *nonmonotonicity*, a feature which admits refutation of old theorems when new axioms are added to the system. Indeed, if somebody's whole knowledge is fully represented by a set of axioms Σ then that person must be ignorant of every sentence φ which is not entailed by Σ. Under such interpretation, Σ should entail $\neg K\varphi$ for every φ not entailed by Σ, while $\Sigma \cup \{\varphi\}$ (a larger set) should not. Clearly, $\Sigma \cup \{\varphi\}$ should entail $K\varphi$, instead.

Because the single operator K does not allow for the adequate treatment of a distributed knowledge, a multimodal version of modal logic has been suggested as a formal tool for proper treatment of the distributed case. Each autonomous element A_n of a distributed system, the so called *agent*, has its own operator K_n with the intentional interpretation "agent A_n knows that...". So far, all published systems of this kind (cf. [HM90, FHV91, FHV92] for the most recent development in this subject) have been monotonic. This may be considered a serious disadvantage taking into account the nonmonotonic behavior of sentences of the form $\neg K_n\varphi$. In particular, none of these systems properly formalises general patterns of reasoning about one's ignorance.

In this paper, we introduce a nonmonotonic *multi-epistemic logic ME* which unifies two formal systems of reasoning about knowledge: a well known multimodal monotonic variant [Hin62] of logic S5, and a unimodal nonmonotonic variant [Suc92a] of S5. The proposed *ME* logic allows for reasoning about both knowledge *and* ignorance. It turns out that *ME* logic is sound and complete with respect to a certain restriction of Kripke-style semantics, which we call a *minimal-knowledge mesh semantics*. This completeness property allows for using interchangeably the proof-theoretic and the model-computing techniques, depending on which one is computationally feasible in a given circumstance.

Unlike most logical approaches which are mainly focused on highly abstract Artificial Intelligence applications, *ME* logic, because of the concept of the mesh model used, has a strong appeal for architecture and hardware oriented interpretations. We illustrate our approach with application of *ME* logic to the *Mr. Sum and Mr. Product Puzzle*, a non-trivial example which has been a hard nut to crack for monotonic multimodal logics.

2 System topology

We assume that the distributed system consists of N autonomous agents $A_1, ..., A_N$ related to a certain global world $\vec{S} = \langle S_1, ..., S_N \rangle$, with each agent A_n having a complete knowledge of the n-th component S_n of $\vec{S}$ (e.g., $\vec{S}$ may be interpreted as a global status register with local components S_n). All these agents are capable of reasoning about the global world on the basis of the information they possess. They may also broadcast what they know to all other agents.

Moreover, the agents admit the existence of certain possible worlds $\vec{S}'$ which together with $\vec{S}$ form a model $\mathcal{M}$ for their distributed knowledge. This model $\mathcal{M}$ represents all publicly known information in the system, so that every agent knows what the possible worlds are, and every agent knows that every agent knows what the possible worlds are, etc.... The agents also know that each of them knows the respective component of the global world $\vec{S}$. Except for this knowledge, and its logical consequences, there is nothing else that every agent knows, and that every agent knows that every agent knows, etc....

Broadcasting may increase the agents' knowledge, so that the model $\mathcal{M}$ for public knowledge may vary (as we will see, it may shrink).

3 Multimodal language L_{MI}

To express the agent's knowledge, and to reason about it, we use the multimodal language L_{MI} defined as follows. (cf. any textbook on modal logic, e.g. [Che80], for details about the unimodal case). L_{MI} has a finite collection I of constants $c_1, ..., c_{|I|}$, N status variables $\pi_1, ..., \pi_N$, the equality symbol $=$, usual logical connectives (including the falsehood symbol $\bot$ and the truth symbol $\top$), the common knowledge modality C, and the individual agents' modalities $K_1, ..., K_N$. Terms of L_{MI} are defined as constants and status variables. Formulas of L_{MI} are $\bot$, $\top$, $\tau_1 = \tau_2$, where $\tau_{1,2}$ are terms, $\neg\varphi$, $\varphi \vee \psi$, $\varphi \wedge \psi$, $C\varphi$, $K_1\varphi, ...,$ and $K_N\varphi$, where φ and ψ are formulas. Other connectives are treated as abbreviations, i.e. $\varphi \supset \psi$ is an abbreviation for $\neg\varphi \vee \psi$, $\vee_{c \in I}\varphi(c)$ is an abbreviation for $\varphi(c_1) \vee ... \vee \varphi(c_{|I|})$, and $\vec{\pi} = \vec{\tau}$ is an abbreviation for $\wedge_{1 \leq k \leq N}(\pi_k = \tau_k)$. We denote by $mPos$ the set of all modally positive formulas of L_{MI}, that is, formulas without occurrence of any of the modalities C, $K_1, .., K_N$ in a scope of negation. For example, $K_3\neg(\tau_2 = \tau_4)$ is an element of $mPos$ while $\neg K_3(\tau_2 = \tau_4)$ is not.

The formulas of L_{MI} are partitioned on groups with respect to the depth of modal negation, the so called *ranks*, according to the following definition.

If $\varphi \in mPos$ then $Rank(\varphi) = 0$;
$Rank(K_n\varphi) = Rank(\varphi)$;
$Rank(\varphi \vee \psi) = max\{Rank(\varphi), Rank(\psi)\}$;
$Rank(\varphi \wedge \psi) = max\{Rank(\varphi), Rank(\psi)\}$;
$Rank(\neg K_n\varphi) = 1 + Rank(\varphi)$;
if $\varphi \equiv \psi$ is a propositional tautology then
$Rank(\varphi) = Rank(\psi)$.

For instance, $Rank(\neg K_1 \neg K_2 \neg(\pi_1 = c)) = 2$. Routine induction shows that every formula of L_{MI} has a uniquely defined rank. Therefore, $Rank$ is a total function on L_{MI}, and defines a position of L_{MI} on non-empty sets $L_{MI}^{(0)}$, $L_{MI}^{(1)}$, ..., given by:

$$L_{MI}^{(k)} = \{\varphi \in L_{MI} \mid Rank(\varphi) = k\}.$$

Formulas of rank $= 0$ (modally positive formulas) are particularly important because of their regular semantic properties. They will be given special attention in this paper. For instance, modal-free formulas represent the objective knowledge. Formulas of rank > 0, e.g., $K_1 \neg K_2 \neg K_3(\pi_1 = c)$, are somewhat troublesome. In particular, the statement "and nothing else is objectively known" may not have a unique interpretation when applied to such formulas.

A formula φ is closed for modal operator K_n iff φ is a Boolean combination of formulas of the form $K_n\psi$.

4 Mesh models for L_{MI}

As we already mentioned, a model $\mathcal{M}$ for the agents' knowledge is a non-empty set of possible worlds, including the actual (but possibly unknown) one. We assume that each world $\mathcal{M}$ is an N-tuple $\vec{i} = [i_1, ..., i_N]$ of constants from I, so that $\mathcal{M} \subseteq I^N$. Each model $\mathcal{M}$ provides the semantics for terms and formulas of L_{MI}. The meaning of a term or a formula in model $\mathcal{M}$ is a function defined for each element $\vec{i}$ of $\mathcal{M}$ as follows. For constants c: $c[\vec{i}] = c$; for status variables: $\pi_n[\vec{i}] = i_n$. For atomic formulas: $\bot[\vec{i}] = \bot$; $\top[\vec{i}] = \top$; $(\tau_1 = \tau_2)[\vec{i}] = \top$ iff $\tau_1[\vec{i}] = \tau_2[\vec{i}]$, $\bot$ otherwise. The meaning of logical connectives are defined by routine induction, with $K_n\varphi[\vec{i}] = \top$ iff for every $\vec{j} \in \mathcal{M}$ with $i_n = j_n$, $\varphi[\vec{j}] = \top$, $\bot$ otherwise. ($C\varphi[\vec{i}]$ is defined later in this section.) In the case $\varphi[\vec{i}] = \top$, we apply the usual notation $\mathcal{M} \models \varphi[\vec{i}]$. Moreover, we use $\mathcal{M} \models \varphi$ iff $\mathcal{M} \models \varphi[\vec{i}]$ holds for every $\vec{i} \in \mathcal{M}$, and for the set Σ of formulas, $\mathcal{M} \models \Sigma$ iff $\mathcal{M} \models \varphi$ for all $\varphi \in \Sigma$.

For instance, if $\mathcal{M} = I^N$ and $N > 1$ then for every $\vec{i} \in \mathcal{M}$, $\mathcal{M} \models \vee_{c \in I} K_n(\pi_n = c)[\vec{i}]$ and $\mathcal{M} \models \neg \vee_{c \in I} K_n(\pi_m = c)[\vec{i}]$, where $n, m \leq N$, $n \neq m$.

Because every model $\mathcal{M}$ is a subset of the Cartesian product I^N, we may represent it as a mesh, e.g., (for $N = 2$ and $I = \{0, 1, 2, 3, 4, 5\}$):

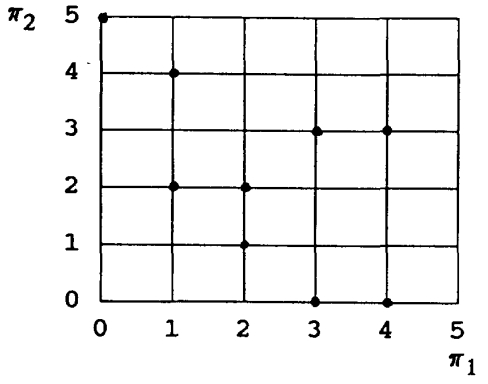

Fig. 1.

where joints • represent possible worlds. The above mesh visualizes model $\mathcal{M} = \{[0,5], [1,2], [1,4], [2,1], [2,2], [3,0], [3,3], [4,0], [4,3]\}$. If the actual world is [2,1] then A_1 knows that $(\pi_1 = 2) \wedge ((\pi_2 = 1) \vee (\pi_2 = 2))$ and A_2 knows that $(\pi_1 = 2) \wedge (\pi_2 = 1)$, i.e. $\mathcal{M} \models K_1((\pi_1 = 2) \wedge ((\pi_2 = 1) \vee (\pi_2 = 2)))[2,1]$ and $\mathcal{M} \models K_2((\pi_1 = 2) \wedge (\pi_2 = 1))[2,1]$. Consequently, if the actual world is [2,2], A_2 does not know $(\pi_1 = 2)$, but also A_1 does not know that A_1 does not know that $(\pi_1 = 2)$, i.e., $\mathcal{M} \models \neg K_2(\pi_1 = 2)[2,2]$ and $\mathcal{M} \models \neg K_1 \neg K_2(\pi_1 = 2)[2,2]$.

By connecting those joints of this mesh which lie on the same coordinate line ($\pi_1 = const$ or $\pi_2 = const$) one obtains a graph whose connected components provide the semantics for the common knowledge modality C. Namely, $\mathcal{M} \models C\varphi[\vec{i}]$ iff for each $\vec{j} \in \mathcal{M}$ which belongs to the same connected component as $\vec{i}$ does, $\mathcal{M} \models \varphi[\vec{j}]$. For $\mathcal{M}$ of Fig. 1, $\mathcal{M} \models C(\pi_1 = 1 \vee \pi_1 = 2)[1,4]$, but $\mathcal{M} \not\models C(\pi_1 = 1)[1,4]$.

If mesh $\mathcal{M}$ is implemented as a mesh network of autonomous processors, each of them representing a possible world, then the computation of logical values of formulas of L_M in $\mathcal{M}$ may be executed in parallel. For instance, to evaluate $K_2(\pi_1 = 2)[2,2]$, the formula $(\pi_1 = 2)$ would have to be sent along bus $\pi_2 = 2$ to all nodes of the form [i,2] (in the above case, to [1,2], and to [2,2]) for evaluation. Then $K_2(\pi_1 = 2)$ would be assumed true iff $(\pi_1 = 2)$ was found true in all these nodes, and false otherwise (in the above case, $(\pi_1 = 2)[1,2]$ would return false, hence $K_2(\pi_1 = 2)[2,2]$ would evaluate to false), that is, $\mathcal{M} \models \neg K_2(\pi_1 = 2)[2,2]$. Similarly, to evaluate $K_1 \neg K_2(\pi_1 = 2)[2,2]$, the formula $\neg K_2(\pi_1 = 2)$ would have to be sent along bus $\pi_1 = 2$ (in the above case, to nodes [2,1] and [2,2]), and each recipient node would subsequently send formula $(\pi_1 = 2)$ along bus $\pi_2 = 1$ or $\pi_2 = 2$, respectively (in our case, $\pi_1 = 2$ would have been sent to nodes [1,2], [2,1], [2,2], and evaluated to false in [2,1], thus making $K_1 \neg K_2(\pi_1 = 2)[2,2]$ false, that is $\mathcal{M} \models \neg K_1 \neg K_2(\pi_1 = 2)[2,2]$). Evaluating formulas $C\varphi[\vec{i}]$ would involve building and traversing a spanning tree of a connected component of $\mathcal{M}$ containing $\vec{i}$. For instance, $\mathcal{M} \models C\neg(\pi_2 = 3)[2,1]$.

$\mathcal{M}$ is called a *minimal-knowledge model* of set Σ of formulas iff $\mathcal{M} \models \Sigma$ and for every $\mathcal{N} \models \Sigma$, if $\mathcal{M} \subseteq \mathcal{N}$ then $\mathcal{M} = \mathcal{N}$, (i.e. $\mathcal{M}$ is maximal w.r.t. set-theoretical inclusion). Minimal-knowledge model of Σ is represented by a mesh with a maximal arrangement of joints. As we will see in Section 6, it may be calculated from an initial mesh by the successive elimination of those joints which violate the constraint imposed by Σ.

It may be easily verified by induction that for every formula φ of rank 0 (i.e., modally positive), and every two models $\mathcal{M}$, $\mathcal{N}$, with $\mathcal{M} \subseteq \mathcal{N}$, if $\mathcal{N} \models \varphi$ then $\mathcal{M} \models \varphi$, but not necessarily vice versa. Consequently, the larger the model the less objective formulas it satisfies, that is, the less objective knowledge it represents. Thus minimal-knowledge models explicate an implicit assumption "and nothing else is objectively known to the agents", which is customarily used in formal specifications of distributed knowledge, particularly when mathematical induction is used as the definition vehicle. Because I^N is finite, every set of formulas which has a model has a minimal-knowledge model as well. We call a set Σ of formulas *definite* iff it has at most one minimal-knowledge model. For instance, every set of modal-free formulas is definite.

Restriction of the semantics of Σ to minimal-knowledge models has a profound consequence. The logical entailment $\vdash$ defined by:

$$\Sigma \vdash \varphi \text{ iff every model of } \Sigma \text{ is a model of } \varphi$$

is no longer adequate, because a sentence φ which is true in all minimal-knowledge models of Σ, and therefore a consequence of Σ in the sense of this restricted semantics, may happen to be false in a certain non-minimal model of Σ and, therefore, not entailed by

Σ. This undesirable situation gives rise to the following relation of minimal-knowledge entailment $\vdash_{min}$ defined for every set of formulas Σ and every formula φ by:

> $\Sigma \vdash_{min} \varphi$ iff for every minimal-knowledge model $\mathcal{M}$ of Σ, $\Sigma \models \varphi$.

Minimal-knowledge entailment is the central concept in this paper. It allows for adequate reasoning about knowledge in a distributed system. As any entailment based on minimisation, it is nonmonotonic.

For example, if $\Sigma \vdash_{min} \varphi$ but $\Sigma \not\vdash \varphi$ then $\Sigma \cup \{\neg\varphi\}$ has a minimal-knowledge model, therefore $\Sigma \cup \{\neg\varphi\} \not\vdash_{min} \varphi$. In other words, there are Σ, Σ', with $\Sigma \subseteq \Sigma'$, such that for some formula φ, $\Sigma \vdash_{min} \varphi$ but $\Sigma' \not\vdash_{min} \varphi$. Logics based on nonmonotonic entailments, so called *nonmonotonic logics* are essentially different from those based on classic, monotonic entailments. Obviously, their expressive power is greater because of the narrower respective semantics. For instance, the empty set of sentences 0 with I^N as its only minimal-knowledge model expresses the fact that all elements for I^N are possible worlds; a statement which is not expressible by any set of objective formulas under the standard (i.e., non minimal one) semantics. Most remarkably, non-monotonic entailments are considerably more difficult for syntactic characterizations than monotonic ones. As one can see, axioms and familiar rules of inference can only define a monotonic entailment.

In the next section we will provide a complete syntactic characterization of $\vdash_{min}$ for modal-free Σ's in terms of provability within multimodal logic S5.

5 Multi-epistemic logic *ME*

First we define the monotonic part of *ME* logic, which is based on the multimodal variant of logic S5 (see [Hin62, Sat76] for details).

Axioms of S5

Pl: Propositional tautologies (e.g. $\varphi \vee \neg\varphi$)

T: $K_n\varphi \supset \varphi$

5: $\varphi \supset K_n\varphi$ (only for φ closed for K_n).

K: $K_n(\varphi_1 \supset \varphi_2) \supset (K_n\varphi_1 \supset K_n\varphi_2)$

Rules of inference

$$\text{MP}: \frac{\vdash \varphi_1, \vdash \varphi_1 \supset \varphi_2}{\vdash \varphi_2}$$

$$\text{RN}: \frac{\vdash \varphi}{\vdash K_n\varphi}$$

where $n = 1, ..., N$, and φ, φ_1, and φ_2 are arbitrary (except for axiom 5) formulas of L_M. In addition, we include the following axioms:

Equality axioms

$\tau = \tau$;

$\tau_1 = \tau_2 \supset \tau_2 = \tau_1$;

$\tau_1 = \tau_2 \wedge \tau_2 = \tau_3 \supset \tau_1 = \tau_3$;

where τ, τ_1, τ_2, τ_3 are terms;

Unique names axioms

$\neg(c = d)$

where c and d are different elements of I;

Domain closure axioms

$\vee_{c \in I}(\pi_n = c)$

where $n = 1, ..., N$;

Status knowledge axioms

$(\pi_n = c) \supset K_n(\pi_n = c)$

where $c \in I$ and $n = 1, ..., N$ (they assure that agent A_n knows the value of π_n).

We call the above system $MS5^+$ (*M*ultimodal *S*5 + the above axioms). We use Cn_{MS5+} as the consequence relation defined by $MS5^+$.

The nonmonotonic part of logic *ME* consists of the following rule of inference

$$\text{NRN}: \frac{\not\vdash (\pi_n = c) \supset \varphi}{\vdash (\pi_n = c) \supset \neg K_n\varphi},$$

where $n = 1, ..., N$, $c \in I$, and φ is an arbitrary formula of L_M.

The nonmonotonic consequence operation Cn_{ME} of logic *ME* is defined inductively for every $\Sigma \subseteq L_M$:

$$\Pi_0(\Sigma) = Cn_{MS5+}(\Sigma),$$

$$\Pi_{k+1}(\Sigma) = Cn_{MS5+}(\text{NRN}(\Pi_k(\Sigma))),$$

$$Cn_{ME}(\Sigma) = \bigcup_{k \in \omega} \Pi_k(\Sigma),$$

where $\text{NRN}(\Pi_k(\Sigma)) = \{(\pi_n = c) \supset \neg K_n\varphi \mid (\pi_n = c) \supset \varphi \in L_M^{(k)} \setminus \Pi_k(\Sigma), n = 1, ..., N, \text{ and } c \in I\}$. The above definition of Cn_{ME} restricts, in fact, the use of NRN rule to a *stratified* one, in the sense that the NRN rule which infers a formula of rank $k+1$ may be used only after stratum $\Pi_k(\Sigma)$ has been completed.

Here is the main result of the paper (in what follows, the common knowledge modality C is not allowed to occur in Σ or in φ).

The Completeness Theorem 5.1 For every set Σ of modal-free formulas of L_M, and every formula φ of L_M,

$$\Sigma \vdash_{min} \varphi \text{ iff } \varphi \in Cn_{ME}(\Sigma).$$

Proof. We prove by induction on $Rank(\varphi)$, that

$$\Sigma \vdash_{min} \varphi \text{ iff } \varphi \in \Pi_{Rank(\varphi)}(\Sigma). \quad (1)$$

(i) $Rank(\varphi) = 0$ (i.e., $\varphi \in mPos$). First, we prove that

$$\Sigma \vdash_{min} \varphi \text{ iff } \Sigma \cup Axioms \vdash_{S5} \varphi, \quad (2)$$

where $X \vdash_{S5} \varphi$ means that φ is true in all Kripke models of X, and *Axioms* is the set of all instances of the equality axioms, unique names axioms, domain closure axioms, and status knowledge axioms (all of them are elements of $mPos$). Because every mesh model satisfies *Axioms*, it suffices to demonstrate implication to the right in (2). Let $\Sigma \vdash_{min} \varphi$ and let a Kripke structure $\mathcal{K} \models \Sigma \cup Axioms$. Let mesh $\mathcal{M}$ be defined by:

$\bigwedge_{\vec{i} \in I^N} \vec{i} \in \mathcal{M}$ iff $\mathcal{K} \not\models \neg(\vec{\pi} = \vec{i})$.

Because $\mathcal{K} \models Axioms$, for every $w \in \mathcal{K}$ there is exactly one $\vec{i} \in I^N$ with $\mathcal{K} \models (\vec{\pi} = \vec{i})[w]$. Hence, for every modal-free ψ and every $v, w \in \mathcal{K}$, and every $\vec{i} \in I^N$, if $\mathcal{K} \models (\vec{\pi} = \vec{i})[w]$, $\mathcal{K} \models \psi[w]$, and $\mathcal{K} \models (\vec{\pi} = \vec{i})[v]$, then $\mathcal{K} \models \psi[v]$. In particular, $\mathcal{M}$ and $\mathcal{K}$ satisfy the same modal-free formulas. Thus $\mathcal{M} \models \Sigma$, and $\mathcal{M} \models \varphi$. Moreover, because $\mathcal{K}$ satisfies the status knowledge axioms, if $w \equiv_n v$ in $\mathcal{K}$ then there is $c \in I$ with $\mathcal{K} \models (\pi_n = c)[w]$ and $\mathcal{K} \models (\pi_n = c)[v]$. From that, straightforward induction shows that for every $\psi \in mPos$, if $\mathcal{M} \models \psi$ then $\mathcal{K} \models \psi$. Hence, $\mathcal{K} \models \varphi$, which proves (2). By the completeness theorem of multimodal $S5$ (see [Sat76]), (2) implies $\Sigma \vdash_{min} \varphi$ iff $\varphi \in Cn_{S5}(\Sigma \cup Axioms)$, that is, $\Sigma \vdash_{min} \varphi$ iff $\varphi \in Cn_{MS5+}(\Sigma)$, which yields (1).

(ii) If $\varphi = \neg K_n\psi$ and $Rank(\psi) = k$ then we have $\Sigma \vdash_{min} \neg K_n\psi$ **iff** [by the definiteness of Σ] $\mathcal{M} \models \neg K_n\psi$ **iff** $\bigwedge_{c \in I} \mathcal{M} \models (\pi_n = c) \supset \neg K_n\psi$ **iff** $\bigwedge_{c \in I} \bigwedge_{\vec{i} \in \mathcal{M}} \mathcal{M} \models (\pi_n = c)[\vec{i}]$ or $\mathcal{M} \models \neg K_n\psi[\vec{i}]$ **iff** $\bigwedge_{c \in I} \bigwedge_{\vec{i} \in \mathcal{M}: i_n = c} \mathcal{M} \models \neg K_n\psi[\vec{i}]$ **iff** $\bigwedge_{c \in I} \bigwedge_{\vec{i} \in \mathcal{M}: i_n = c} \bigvee_{\vec{j} \in \mathcal{M}: j_n = i_n} \mathcal{M} \models \neg\psi[\vec{j}]$ **iff** $\bigwedge_{c \in I} \bigvee_{\vec{j} \in \mathcal{M}: j_n = c} \mathcal{M} \not\models \psi[\vec{j}]$ **iff** $\bigwedge_{c \in I} \not\models (\pi_n = c) \supset \psi$ **iff** [by induction hypothesis, because $Rank((\pi_n = c) \supset \psi) = Rank(\psi) = k$] $\bigwedge_{c \in I} (\pi_n = c) \supset \psi \notin \Pi_k(\Sigma)$ **iff** $\bigwedge_{c \in I} (\pi_n = c) \supset \neg K_n\psi \in \mathrm{NRN}(\Pi_k(\Sigma))$ **iff** [by the domain closure axiom] $\neg K_n\psi \in Cn_{MS5+}(\mathrm{NRN}(\Pi_k(\Sigma)))$ **iff** [by definition of $\Pi_{k+1}(\Sigma)$] $\neg K_n\psi \in \Pi_{k+1}(\Sigma)$.

(iii) If $\varphi = \psi_1 \wedge \psi_2$ then the inductive step is obvious.

(iv) If $\varphi = \psi_1 \vee \psi_2$ then $\Sigma \vdash_{min} \psi_1 \vee \psi_2$ **iff** $\mathcal{M} \models \psi_1 \vee \psi_2$ **iff** $\bigwedge_{\vec{i} \in \mathcal{M}} \mathcal{M} \models \psi_1[\vec{i}]$ or $\mathcal{M} \models \psi_2[\vec{i}]$ **iff** there is a partition of I^N onto two sets X_1 and X_2 with $\bigwedge_{\vec{i} \in \mathcal{M} \cap X_1} \mathcal{M} \models \psi_1[\vec{i}]$ and $\bigwedge_{\vec{i} \in \mathcal{M} \cap X_2} \mathcal{M} \models \psi_2[\vec{i}]$ **iff** there is such a partition with $\Sigma \cup \{\neg(\vec{\pi} = \vec{i}) \mid \vec{i} \in I^N \setminus X_1\} \vdash_{min} \psi_1$ and $\Sigma \cup \{\neg(\vec{\pi} = \vec{i}) \mid \vec{i} \in I^N \setminus X_2\} \vdash_{min} \psi_2$ **iff** [by inductive hypothesis] $\psi_1 \in \Pi_k(\Sigma \cup \{\neg(\vec{\pi} = \vec{i}) \mid \vec{i} \in I^N \setminus X_1)$ and $\psi_2 \in \Pi_k(\Sigma \cup \{\neg(\vec{\pi} = \vec{i}) \mid \vec{i} \in I^N \setminus X_2\})$ **iff** [by the domain closure axiom, because Π_k is closed under Cn_{MS5+}] $\psi_1 \vee \psi_2 \in \Pi_k(\Sigma)$. Finally, we conclude that $\Sigma \vdash_{min} \varphi$ **iff** $\varphi \in \bigcup_{k \in w} \Pi_k(\Sigma)$ **iff** [by definition of Cn_{ME}] $\varphi \in Cn_{ME}(\Sigma)$.

(v) If $\varphi = K_n\psi$ then $\Sigma \vdash_{min} K_n\psi$ **iff** [by definiteness of Σ] $\mathcal{M} \models K_n\psi$ **iff** $\mathcal{M} \models \psi$ **iff** $\Sigma \vdash_{min} \psi$ **iff** [by inductive hypothesis] $\psi \in \Pi_{Rank(\psi)}(\Sigma)$ **iff** $K_n\psi \in \Pi_{Rank(\varphi)}(\Sigma)$. □

Theorem 5.1 gives this corollary.

Corollary 5.2 For every modal-free Σ and every $\varphi, \psi \in L_M$,

$$\Sigma \vdash_{min} \varphi \supset C\psi \text{ iff } \bigwedge_{A \in \mathbf{K}^{|I^N|-1}} \varphi \supset A\psi \in Cn_{AE}(\Sigma),$$

where $\mathbf{K}^k$ denotes the set of all sequences of elements of $\{K_1, ..., K_N\}$ of length k.

Proof follows from the fact that any mesh $\mathcal{M}$ cannot have a simple path longer than $|I^N| - 1$. □

6 Updates of public knowledge

Information broadcasted by an agent to the other agents may change the public knowledge. Specifically, if the knowledge before broadcast φ by agent A_n was represented by a mesh $\mathcal{M}$ then after this broadcast, all worlds $\vec{i}$ of $\mathcal{M}$ in which A_n doesn't know φ, that is, $\mathcal{M} \not\models K_n\varphi[\vec{i}]$, must be eliminated from $\mathcal{M}$, so that the $\mathcal{M} \restriction K_n\varphi$ represents the knowledge after the broadcast. Formally, $\mathcal{M} \restriction \psi$ is defined by:

$$\mathcal{M} \restriction \psi = \{\vec{i} \in \mathcal{M} \mid \mathcal{M} \models \psi[\vec{i}]\}.$$

In the case of a sequence of broadcasts by various agents (a broadcast dialogue), a mesh $\mathcal{M}$ which represents public knowledge shrinks as the dialogue proceeds and the amount of public knowledge in the system grows. The following result comes in handy when evaluating the final public knowledge in the system by proof-theoretic means, without calculating a sequence of mesh models.

Let φ, ψ be formulas of L_M. *Relativization* of ψ to φ (notation: ψ^φ) is defined by induction.

If ψ is modal-free then $\psi^\varphi = \psi$;
$(\neg\psi)^\varphi = \neg(\psi^\varphi)$;
$(\psi_1 \wedge \psi_2)^\varphi = \psi_1^\varphi \wedge \psi_2^\varphi$; $(\psi_1 \vee \psi_2)^\varphi = \psi_1^\varphi \vee \psi_2^\varphi$;
$(K_n\psi)^\varphi = K_n(\varphi \supset \psi^\varphi)$.

Theorem 6.1 For every formula φ, ψ of L_{MI} and every model $\mathcal{M}$, and every $\vec{i} \in \mathcal{M}$,

$$\mathcal{M} \restriction \varphi \models \psi[\vec{i}] \text{ iff } \mathcal{M} \models (\varphi \wedge \psi^{\varphi})[\vec{i}].$$

Proof by induction on the length of ψ.

(i) Assume ψ is modal-free. Then for every $\vec{i} \in \mathcal{M}$, $\mathcal{M} \restriction \varphi \models \psi[\vec{i}]$ iff $\vec{i} \in \mathcal{M} \restriction \varphi$ and $\mathcal{M} \models \psi[\vec{i}]$ iff [by definition of $\mathcal{M} \restriction \varphi$] $\mathcal{M} \models \varphi[\vec{i}]$ and $\mathcal{M} \models \psi[\vec{i}]$ iff $\mathcal{M} \models (\varphi \wedge \psi)[\vec{i}]$ iff (by definition of ψ^{φ}) $\mathcal{M} \models (\varphi \wedge \psi^{\varphi})[\vec{i}]$.

(ii) Assume $\psi = K_n\vartheta$. Then for every $\vec{i} \in \mathcal{M}$, $\mathcal{M} \restriction \varphi \models \psi[\vec{i}]$ iff $\vec{i} \in \mathcal{M} \restriction \varphi$ and for every $\vec{j}$ in $\mathcal{M} \restriction \varphi$ with $j_n = i_n$, $\mathcal{M} \restriction \varphi \models \vartheta[\vec{j}]$ iff $\mathcal{M} \models \varphi[\vec{i}]$ and for every $\vec{j}$ with $\mathcal{M} \models \varphi[\vec{j}]$ and with $j_n = i_n$, [by the induction hypothesis] $\mathcal{M} \models (\varphi \wedge \vartheta^{\varphi})[\vec{j}]$ iff $\mathcal{M} \models \varphi[\vec{i}]$ and for every $\vec{j}$ with $\mathcal{M} \models \varphi[\vec{j}]$ and with $j_n = i_n$, $\mathcal{M} \models \vartheta^{\varphi}[\vec{j}]$ iff $\mathcal{M} \models \varphi[\vec{i}]$ and for every $\vec{j}$ with $j_n = i_n$, $\mathcal{M} \models (\varphi \supset \vartheta^{\varphi})[\vec{j}]$ iff $\mathcal{M} \models \varphi[\vec{i}]$ and $\mathcal{M} \models K_n(\varphi \supset \vartheta^{\varphi})[\vec{i}]$ iff $\mathcal{M} \models (\varphi \wedge \psi^{\varphi})[\vec{i}]$.

(iii) Cases of $\psi = \neg\vartheta$, $\psi = \vartheta_1 \vee \vartheta_2$, and $\psi = \vartheta_1 \wedge \vartheta_2$ are straightforward. □

Theorem 6.1 yields the following

Corollary 6.2 For every formula φ, ψ of L_{MI}, and every model $\mathcal{M}$,

$$\mathcal{M} \restriction \varphi \models \psi \text{ iff } \mathcal{M} \models \varphi \supset \psi^{\varphi}.$$

Proof. $\mathcal{M} \restriction \varphi \models \psi$ iff $\bigwedge_{\vec{i} \in \mathcal{M}|\varphi} \mathcal{M} \restriction \varphi \models \psi[\vec{i}]$ iff [by theorem 6.1] $\bigwedge_{\vec{i} \in \mathcal{M}|\varphi} \mathcal{M} \models (\varphi \wedge \psi^{\varphi})[\vec{i}]$ iff [by definition of $\mathcal{M} \restriction \varphi$] $\bigwedge_{\vec{i} \in \mathcal{M}}$ if $\mathcal{M} \models \varphi[\vec{i}]$ then $\mathcal{M} \models (\varphi \wedge \psi^{\varphi})[\vec{i}]$ iff $\bigwedge_{\vec{i} \in \mathcal{M}} \mathcal{M} \models (\varphi \supset (\varphi \wedge \psi^{\varphi}))[\vec{i}]$ $\mathcal{M} \models (\varphi \supset \psi^{\varphi})$. □

The above result allows for purely syntactic (that is, without explicit reference to any model) evaluation of public knowledge after the dialogue $\varphi_1, ..., \varphi_n$ took place, as states the following

Theorem 6.3 For every set Σ of modal-free formulas of L_{MI}, the following conditions are equivalent:
for every minimal-knowledge model $\mathcal{M}$ of Σ,

$$((\mathcal{M} \restriction \varphi_1)...) \restriction \varphi_n) \models \psi \tag{3}$$

and

$$\Sigma \vdash_{min} \varphi_1 \supset (\varphi_2 \supset ... \supset \varphi_{n-1}(\varphi_n \supset \psi^{\varphi_n})^{\varphi_{n-1}}...)^{\varphi_1}. \tag{4}$$

Proof. By Corollary 6.2, (3) is equivalent to

$$\mathcal{M} \models \varphi_1 \supset (\varphi_2 \supset ... \supset \varphi_{n-1}(\varphi_n \supset \psi^{\varphi_n})^{\varphi_{n-1}}...)^{\varphi_1}. \tag{5}$$

Application of Theorem 5.1 yields (4). □

7 Mr. Sum and Mr. Product Puzzle

We use the following classic example to illustrate how these two approaches work.

There are two logically omniscient and fully introspective agents S and P, for whom a possible world is a pair $(a+b, a \times b)$ where a and b are natural numbers with $2 \leq a, b \leq 100$. Agent S knows the value of sum $a+b$, and agent P knows the value of product $a \times b$. This is all that is known in this system. (In particular, agent S knows that agent P knows the product and no more than the product, etc.)

The following dialogue takes place.

S: "I don't know the value of $a \times b$, but P doesn't know the value of $a+b$, either".

P: "But now, I do know the value of $a+b$".

S: "Now, I know the value of $a \times b$, too".

What are these values?

Unlike most examples used in Artificial Intelligence articles, this puzzle is far from being trivial, i.e., not every math major will be able to solve it. (We suggest that the reader tries it before proceeding). The formal specification of the initial situation is

$$\Sigma = \{\neg((\pi_S = c) \wedge \pi_P = d) \mid x + y = c \text{ and } x \times y = d \text{ has no solutions}\}.$$

It is clear that Σ is definite, that is, it has a unique minimal-knowledge model, say, $\mathcal{M}$. The dialogue between S and P is formalized as follows.
S's first statement:
$\varphi = K_S \neg \vee_{4 \leq c \leq 200} K_P(\pi_S = c)$.
P's first statement:
$\psi = \vee_{4 \leq c \leq 200} K_P(\pi_S = c)$.
S's second statement:
$\vartheta = \vee_{4 \leq c \leq 10,000} K_S(\pi_P = c)$.
The problem is to find x, y which satisfy:
$\mathcal{M} \models \varphi[x, y]$,
$\mathcal{M} \restriction \varphi \models \psi[x, y]$,
$(\mathcal{M} \restriction \varphi) \restriction \psi \models \vartheta[x, y]$.
The last condition is equivalent to
$[x, y] \in ((\mathcal{M} \restriction \varphi) \restriction \psi) \restriction \vartheta$.

One can calculate (using a suitable program) models $\mathcal{M}$, $\mathcal{M} \restriction \varphi$, $(\mathcal{M} \restriction \varphi) \restriction \psi$ and $((\mathcal{M} \restriction \varphi) \restriction \psi) \restriction \vartheta$. They are visualized on Figures 2, 3, 4, and 5. Fig. 2 visualizes a lower left corner of the initial mesh $\mathcal{M}$. Joints $\vec{i}$ with $\mathcal{M} \models \psi[\vec{i}]$ (those for which P knows up front the value of $a+b$) are encapsulated in squares. Rows of $\mathcal{M}$ containing a squared asterisk (and consequently, those

for which S cannot predict that P does not know $a+b$) have been crossed out with solid horizontal lines.

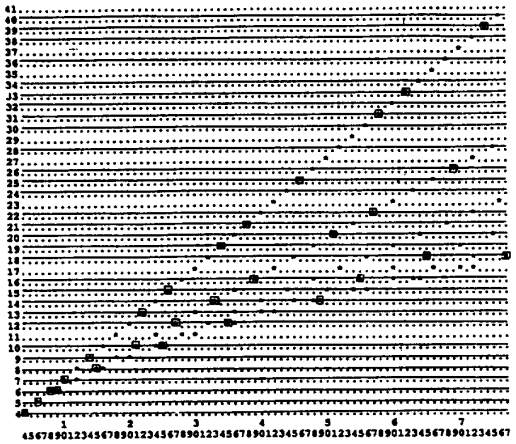

Fig. 2.

Fig. 3 visualizes a lower left corner of the mesh $\mathcal{M} \restriction \varphi$ (after S's first statement). It is the result of the removal from $\mathcal{M}$ all the joints which lie on the solid horizontal lines of Fig. 2. Columns of $\mathcal{M} \restriction \varphi$ containing more that one asterisk (those for which P does not know $a + b$) have been crossed out with solid vertical lines.

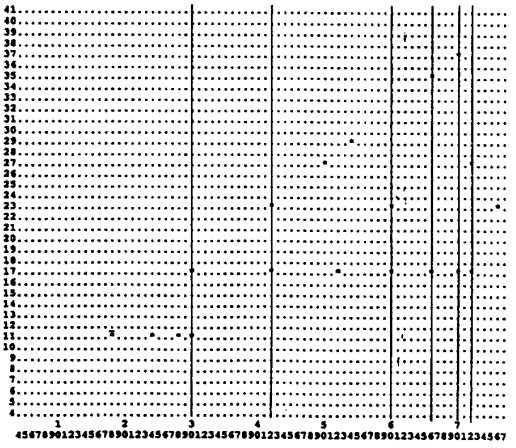

Fig. 3.

Fig. 4 visualizes a lower left corner of the mesh $(\mathcal{M} \restriction \varphi) \restriction \psi$. It is the result of removal from $\mathcal{M}$ all the joints which lie on the solid vertical lines of Fig. 3. Rows of $(\mathcal{M} \restriction \varphi) \restriction \psi$ containing more than one asterisk (those for which S does not know $a \times b$) have been crossed out with solid horizontal lines. After removing all the worlds which lie on these lines, one obtains mesh $((\mathcal{M} \restriction \varphi) \restriction \psi) \restriction \vartheta$, visualised in Fig. 5.

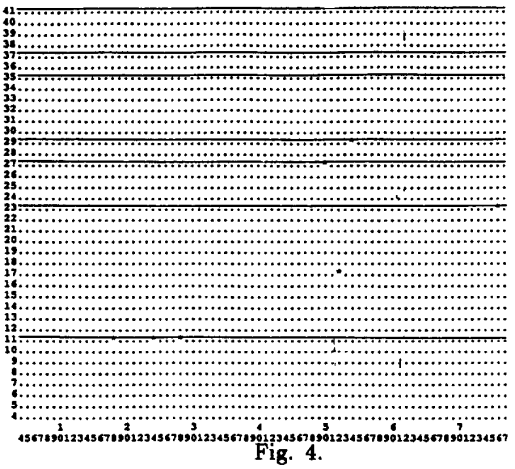

Fig. 4.

Because the last mesh contains only one joint, pair [17,52] (that is, $a = 4$, $b = 13$), we conclude that it is the solution of the puzzle.

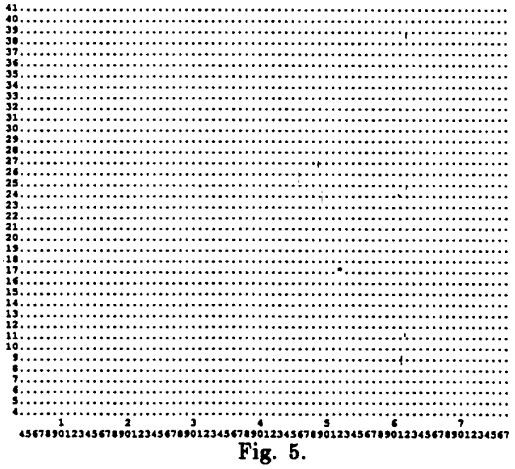

Fig. 5.

One can easily evaluate the common knowledge between S and P using the above meshes.

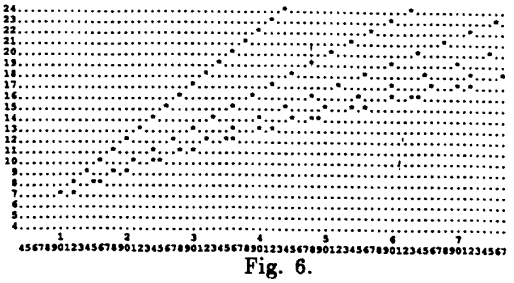

Fig. 6.

As we noted in Section 4, this problem reduces to traversing the connected component $C(\vec{i})$ of $\mathcal{M}$ containing a given joint $\vec{i}$. Then $\mathcal{M} \models C\varphi[\vec{i}]$ holds iff $C(\vec{i}) \models \varphi$. For instance, one can easily see from Fig. 6, which visualises a lower left corner of the connected component of $\mathcal{M}$ containing $[17,52]$, that $\mathcal{M} \models C(\pi_S \geq 7) \wedge \neg C(\pi_S \geq 8)[17,52]$ (a known fact from [Pan92]). Similarly, $\mathcal{M} \restriction \varphi \models C(\pi_S \geq 11) \wedge \neg C(\pi_S \geq 12)[17,52]$, and obviously, $(\mathcal{M} \restriction \varphi) \restriction \psi \models C(\pi_S = 17)[17,52]$.

In the case where mesh is too large to be effectively computed (for instance, if upper limit on a and b is lifted), a method of direct evaluation in the minimal-knowledge model of Σ leads to a faster verification of a solution. It follows from Theorem 6.1 that $[x,y]$ is a solution of the puzzle iff

$\mathcal{M} \models \varphi \wedge \psi^{\varphi} \wedge (\vartheta^{\psi})^{\varphi}[x,y]$,

that is, after simplification

$\mathcal{M} \models \varphi[x,y]$,

$\mathcal{M} \models K_P(\varphi \supset \pi_S = x)[x,y]$, and

$\mathcal{M} \models K_S(K_P(\varphi \supset \pi_S = x) \supset \pi_P = y)[x,y]$.

It is a matter of straightforward although tedious calculations (a program can do that) to check that [17,52] satisfies the above conditions.

Because for every $\vec{i} \in \mathcal{M}$ and $\rho \in L_M$, $\mathcal{M} \models \rho[\vec{i}]$ is equivalent to $\mathcal{M} \models (\vec{\pi} = \vec{i}) \supset \rho$, it follows from Theorem 5.1 that the same can be done by means of a proof within *ME* logic. Examples of such proofs were presented in [Suc92b] and [SO92].

The nonmonotonicity of *ME* logic was necessary to assure the proper treatment of the above puzzle. For instance, let

$$\Sigma' = \{\rho \mid \rho \text{ is modal-free and } \mathcal{M} \restriction \varphi \models \rho\}.$$

Obviously, $\Sigma \subseteq \Sigma'$ and $\Sigma \vdash_{min} (\vec{\pi} = [17,52]) \supset \neg\psi$, but $\Sigma' \not\vdash_{min} (\vec{\pi} = [17,52]) \supset \neg\psi$. This means that the increase in public knowledge invalidated some of its earlier consequences.

References

[Che80] Brian F. Chellas. *Modal Logic: An Introduction.* Cambridge University Press, 1980.

[FHV91] Ronald Fagin, Joseph Y. Halpern, and Moshe Y. Vardi. A model-theoretic analysis of knowledge. *JACM*, 38(2):382–428, April 1991.

[FHV92] Ronald Fagin, Joseph Y. Halpern, and Moshe Y. Vardi. What can machines know? On the properties of knowledge in distributed systems. *JACM*, 39:328–376, April 1992.

[Hin62] Jaakko Hintikka. *Knowledge and Belief.* Cornell University Press, Ithaca, N. Y., 1962.

[HM90] Joseph Y. Halpern and Yoram Moses. Knowledge and common knowledge in a distributed environment. *JACM*, vol. 37(3):549–587, July 1990.

[Kam91] Michael Kaminski. Embedding a default system into nonmonotonic logics. *Fundamenta Informaticae*, 14:345–353, 1991.

[MaT91] Wiktor Marek and Mirosław Truszczyński. Autoepistemic logic. *JACM*, 38(3):588–619, 1991.

[MD80] Drew McDermott and Jon Doyle. Nonmonotonic logic I. *Artificial Intelligence*, 13(1–2):41–72, 1980.

[Moo85] Robert C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25(1):75–94, 1985.

[Pan92] Giovanni Panti. Solution of a number theoretic problem involving knowledge. *International Journal of Foundations of Computer Science*, 1992. To appear.

[Sat76] M. Sato. A study of Kripke-type models for some modal logics. Research Institute for Mathematical Science, Kyoto University, Kyoto, Japan, 1976.

[SO92] Marek. A. Suchenek and Henrietta Okeke. Calculating Kripke models for distributed systems. In *3rd Annual CSU Artificial Intelligence Symposium*, pages 69–78. California State University, June 18 - 19 1992.

[Suc92a] Marek A. Suchenek. Notes on nonmonotonic autoepistemic logic. Submitted for publication, 1992.

[Suc92b] Marek A. Suchenek. Sieves of Eratosthenes, Kripke models, and distributed knowledge bases. In *3rd Annual Ulam Mathematics Conference*, West Palm Beach, FL, March 19-20 1992. Invited lecture.