

Composing Models

Jan van Eijck, Yanjing Wang, and Floor Sietsma

Centrum Wiskunde en Informatica,
P.O. Box 94079, NL-1090 GB Amsterdam, The Netherlands

Abstract. We study a new composition operation on (epistemic) multiagent models and update actions that takes vocabulary extensions into account. This operation allows to represent partial observational information about a large model in a small model, where the small models can be viewed as representations of the observational power of agents, and about their powers for changing the facts of the world. Our investigation provides ways to check relevant epistemic properties on small components of large models, and our approach generalizes the use of ‘locally generated models’.

Keywords: dynamic epistemic logic, multi-agent models, model composition, agent observational power, action model update, epistemic model checking, interpreted systems, reduction techniques.

1 Introduction

The initial model for a muddy children scenario with n children has 2^n nodes. In this paper we show how such a model can be viewed as a composition of n two-node models, each talking only about the muddiness of a single child. For this, we introduce restricted multi-agent models, i.e., models with a limited set of ‘relevant’ proposition letters, we define a composition operation on restricted models and show that it is nicely behaved. Next, we adjust the definition of action model update to restricted models, by incorporating vocabulary expansion in the update process.

The intuition behind building models from components is that each component contains partial observational information about the whole model.

Dynamic epistemic logic [Ger99,BMS99,BvEK06,DvdHK07] has a somewhat monolithic architecture, and the decomposition perspective that we propose in this paper brings the framework closer to that of interpreted systems [FHMV95] (see [BH⁺09] and [KP10] for discussion). The intuition guiding interpreted systems is that each agent is following its own computational procedure, and that the internal states of the agents at different moments in time are given by local states. We claim that model components with restricted vocabularies can play a similar role in dynamic epistemic logic.

Our investigation is encouraging for epistemic model checking with dynamic epistemic logic, for it suggests ways to check relevant epistemic properties on small components of large models. An extended version of the present paper can be found in Chapter 5 of the forthcoming PhD thesis [Wanng], available upon request from the author.

2 Composing Static Models

Let P be a finite set of proposition letters. A static multi-agent model for P is a quadruple (W, I, R, V) with W a set of worlds, I a finite set of agents, R a function that assigns to each $i \in I$ a binary relation R_i on W , and V a function that assigns to each $w \in W$ a subset of P . (In computer science static multi-agent models are often called labelled transition systems.)

A *vocabulary* is a subset Q of P . A model over a vocabulary Q is a multi-agent model (W, I, R, V) where V is a valuation satisfying $V(w) \subseteq Q$ for each $w \in W$.

A *restricted* multi-agent model is a quintuple (W, I, R, V, Q) such that (W, I, R, V) is a model over vocabulary Q .

The unit model \mathcal{E} for I is the restricted model $(\{e\}, I, \{(e, e) \mid i \in I\}, e \mapsto \emptyset, \emptyset)$. In a picture:



The parallel composition $\mathcal{M} \oplus \mathcal{N}$ of two restricted multi-agent models with the same agent set I is given by $(W, I, R, V, Q_M \cup Q_N)$, where

$$W = \{(w, v) \mid w \in W_M, v \in W_N, V_M(w) \cap Q_N = V_N(v) \cap Q_M\},$$

and


$$(w, v)R_i(w', v') \text{ iff } wR_{i_M}w' \text{ and } vR_{i_N}v',$$

and


$$V(w, v) = V_M(w) \cup V_N(v).$$

The new accessibility relations are defined as the product of the relations on the components, in the usual way, restricted to the pairs of worlds where the old valuations agree on the respective restricted vocabularies. Note that $V(w, v)$ agrees with $V_M(w)$ on Q_M and with $V_N(v)$ on Q_N .

As a first example, here is a ‘compositional version’ of the muddy children scenario. Consider the following models, where each pair $m_i \xleftrightarrow{i} \overline{m}_i$ represents a child that does not know whether it is muddy. We assume that each model $m_i \xleftrightarrow{i} \overline{m}_i$ is restricted to $\{m_i\}$, and we leave out reflexive arrows (present for all agents).




Composing with a third model gives:



And so on, for composition of multidimensional hypercubes with more and more children present.

Note that it does not generally hold that $\mathcal{M} \oplus \mathcal{M} \simeq \mathcal{M}$ (using \simeq to indicate the existence of a total bisimulation). In other words, the \oplus operation is not idempotent. To see this, consider the following model:



It is clear that (t, u) is in the composed model $\mathcal{M} \oplus \mathcal{M}$, but according to the definition of \oplus , (t, u) cannot reach a non- p world in the composed model. Therefore, (t, u) is not bisimilar to any world in \mathcal{M} . Still, Kripke models with restricted vocabularies form a commutative monoid:

Theorem 1. *Restricted multi-agent models form a commutative monoid under the \oplus operation, with bisimilarity as the appropriate equality notion. In particular, we have (using \simeq to indicate the existence of a total bisimulation between two models):*

$$\begin{aligned} \mathcal{E} \oplus \mathcal{M} &\simeq \mathcal{M} \\ \mathcal{M} \oplus \mathcal{E} &\simeq \mathcal{M} \\ \mathcal{M} \oplus (\mathcal{N} \oplus \mathcal{K}) &\simeq (\mathcal{M} \oplus \mathcal{N}) \oplus \mathcal{K} \\ \mathcal{M} \oplus \mathcal{N} &\simeq \mathcal{N} \oplus \mathcal{M} \end{aligned}$$

This yields the well-known algebraic preordering \leq on the set of restricted multi-agent models:

$$\mathcal{M} \leq \mathcal{N} \text{ iff there is a } \mathcal{K} \text{ with } \mathcal{M} \oplus \mathcal{K} = \mathcal{N}.$$

We proceed to give a structural characterization of this relation. For this, let a left-simulation between two restricted static models \mathcal{M} and \mathcal{N} be a bisimulation with the invariance condition restricted to proposition letters in the vocabulary of \mathcal{M} , and without the zig condition (see the section on simulation and safety in [BRV01]). Formally, a left-simulation between \mathcal{M} and \mathcal{N} is a relation $C \subseteq W_M \times W_N$ such that wCv implies that the following hold:

Restricted invariance $p \in V_M(w)$ iff $p \in V_N(v)$ for all $p \in Q_M$,

Zag If for some $i \in I$ there is a $v' \in W_N$ with $v \xrightarrow{i} v'$ then there is a $w' \in W_M$ with $w \xrightarrow{i} w'$ and $w' C v'$.

We will use $\mathcal{M}, w \Leftarrow \mathcal{N}, v$ to indicate that there is a left-simulation that connects w and v , and $\mathcal{M} \Leftarrow \mathcal{N}$ to indicate that there is a total left-simulation between \mathcal{M} and \mathcal{N} (and we also write $\mathcal{N} \Rightarrow \mathcal{M}$ for $\mathcal{M} \Leftarrow \mathcal{N}$).

Theorem 2. *If $\mathcal{M} \leq \mathcal{N}$ then $\mathcal{M} \Leftarrow \mathcal{N}$.*

Proof. Assume $\mathcal{M} \leq \mathcal{N}$. Then there is a restricted model \mathcal{K} with $\mathcal{M} \oplus \mathcal{K} \simeq \mathcal{N}$. Let Z be a total bisimulation between $W_{\mathcal{M} \oplus \mathcal{K}}$ and W_N . Define C as $w C v$ iff there is some world $x \in W_{\mathcal{K}}$ with $(w, x) Z v$. C is easily seen to be a total left-simulation between \mathcal{M} and \mathcal{N} . The restricted invariance property follows from the definition of the valuation on $\mathcal{M} \oplus \mathcal{K}$. The zag property follows from the definition of the accessibility relations on $\mathcal{M} \oplus \mathcal{K}$. Thus, $\mathcal{M} \Leftarrow \mathcal{N}$. \square

Note that the converse of Theorem 2 does not hold without restrictions. For example, let \mathcal{M} and \mathcal{N} be the following two S5 models:



It is clear that $\mathcal{M} \Leftarrow \mathcal{N}$. Now suppose towards a contradiction that there exists a \mathcal{K} with $\mathcal{M} \oplus \mathcal{K} \simeq \mathcal{N}$. Since there is a pq world in \mathcal{N} there must be a world s in \mathcal{K} with q true in s and s compatible with the top-right world in \mathcal{M} (call this world t). Then it is easy to see that (t, s) must be in $\mathcal{M} \oplus \mathcal{K}$, with $V_{\mathcal{M} \oplus \mathcal{K}}(t, s) = \{p, q\}$. However, according to the definition of \oplus , (t, s) cannot reach an \bar{p} world in one step, and therefore (t, s) cannot be bisimilar to any world in \mathcal{N} .

We will prove the converse of Theorem 2 for propositionally differentiated models. Call a model \mathcal{M} propositionally differentiated if it holds for all worlds w, w' of \mathcal{M} that if w and w' have the same valuation then $w \simeq w'$.

Theorem 3. *Let \mathcal{M} be a propositionally differentiated model. Then $\mathcal{M} \Leftarrow \mathcal{N}$ implies $\mathcal{M} \leq \mathcal{N}$.*

Proof. We assume without loss of generality that \mathcal{M} is bisimulation minimal. This, together with the fact that \mathcal{M} is propositionally differentiated, gives that different worlds of \mathcal{M} have different valuations. Let C be a left-simulation between \mathcal{M} and \mathcal{N} . Note that since C is total and \mathcal{M} is propositionally differentiated and bisimulation minimal, for each $v \in W_N$ there is exactly one $w \in W_M$ such that $V_M(w) = V_N(v) \cap Q_M$ and $w C v$.

We will show that $\mathcal{M} \oplus \mathcal{N} \simeq \mathcal{N}$. Let the relation Z between $\mathcal{M} \oplus \mathcal{N}$ and \mathcal{N} be defined as

$$(w, v) Z v' \text{ iff } v = v'$$

We claim Z is a total bisimulation. Note that it follows from the fact that C is total and the definition of Z that Z is total.

Suppose $(w, v) Z v$. By construction of $\mathcal{M} \oplus \mathcal{N}$, $V_{\mathcal{M} \oplus \mathcal{N}}((w, v)) = V_M(w) \cup V_N(v) = V_N(v)$. This proves the invariance property.

Suppose $(w, v)Zv$ and $(w, v)R_i(w', v')$. By construction of $\mathcal{M} \oplus \mathcal{N}$ this means wR_iw' and vR_iv' . By construction of Z , $(w', v')Zv'$. This proves the zig property.

Suppose $(w, v)Zv$ and vR_iv' . Since $(w, v) \in \mathcal{M} \oplus \mathcal{N}$, $V_M(w) = V_N(v) \cap Q_M$. So then w is the unique element of W_M that has that property for v , and wCv . Then because C is a left-simulation, there must be some w' such that wR_iw' and $w'Cv'$. Since $w'Cv'$, $V_M(w') = V_N(v') \cap Q_M$ so $(w', v') \in \mathcal{M} \oplus \mathcal{N}$. Since wR_iw' and vR_iv' , $(w, v)R_i(w', v')$ and by definition of Z , $(w', v')Zv'$. This proves the zag property. \square

Van Ditmarsch and French [DF09] prove that for static models \mathcal{M}, \mathcal{N} : $\mathcal{M} \Leftarrow \mathcal{N}$ iff there is an action model A with $\mathcal{M} \otimes A \simeq \mathcal{N}$. From this it follows that for all \mathcal{M} and A : $\mathcal{M} \Leftarrow \mathcal{M} \otimes A$. Combining this and Theorem 3 we get:

Theorem 4. *For any propositionally differentiated static model \mathcal{M} and action model A : $\mathcal{M} \leq \mathcal{M} \otimes A$.*

As a side remark we mention that ‘erasing’ valuation information can sometimes be used for decomposing models, as illustrated in the following example:

$$\begin{array}{ccc}
 \begin{array}{c} m_1 m_2 \xrightarrow{1} \overline{m_1 m_2} \\ 2 \left| \quad \quad \right| 2 \\ \overline{m_1 \overline{m_2}} \xrightarrow{1} \overline{m_1 \overline{m_2}} \\ 1 \end{array} & \Rightarrow \text{erase } m_1 \text{ info} \Rightarrow & \begin{array}{c} m_2 \xrightarrow{1} m_2 \quad m_2 \\ 2 \left| \quad \quad \right| 2 \Leftrightarrow \left| 2 \right. \\ \overline{m_2} \xrightarrow{1} \overline{m_2} \quad \overline{m_2} \\ 1 \end{array}
 \end{array}$$

Let Q^I be the universal ignorance model for Q , i.e. $Q^I = (W, I, R, V, Q)$ with $W = \mathcal{P}(Q)$, $R_i = W^2$, $V = \text{id}$. If $\mathcal{M} = (W, I, R, V, Q)$ is a restricted static model and Q_1 is a set of proposition letters, then we define the expanded model for the larger vocabulary $Q \cup Q_1$ as follows: $\mathcal{M} \triangleleft Q_1^I = \mathcal{M} \oplus Q_1^I$.

Here is an example of expanding with a single new proposition letter m_2 . Note: **Here and henceforth, worlds are i -linked if there is an i -path in the picture.**

$$\begin{array}{ccc}
 \begin{array}{c} m_1 \xrightarrow{1} \overline{m_1} \\ \oplus \quad \left| 1, 2 \right. \\ \overline{m_2} \end{array} & \Leftrightarrow & \begin{array}{c} m_1 m_2 \xrightarrow{1} \overline{m_1 m_2} \\ 1, 2 \left| \quad \quad \right| 1, 2 \\ \overline{m_1 \overline{m_2}} \xrightarrow{1} \overline{m_1 \overline{m_2}} \\ 1 \end{array}
 \end{array}$$

Model expansion to a larger vocabulary will be used in the definition of action model update for restricted models, in Section 3.

If we let p range over P and i over I , then the PDL language over P, I , notation L_{PI} , is given by:

$$\begin{aligned}
 \phi &::= \top \mid p \mid \neg\phi \mid \phi \vee \phi \mid \langle \alpha \rangle \phi \\
 \alpha &::= i \mid ?\phi \mid \alpha; \alpha \mid \alpha \cup \alpha \mid \alpha^*
 \end{aligned}$$

We employ the usual abbreviations. In particular \perp abbreviates $\neg\top$. The semantics for L_{PI} is defined as usual. (We leave the development of a Kleene style three valued logic [Kle50] for interpretation of L_{PI} in restricted models for a future occasion.)

The diamond fragment of L_{PI} is given by the formulas of the syntactic form of ϕ in the following definition:

$$\begin{aligned}\psi &::= \top \mid p \mid \neg\psi \mid \psi \vee \psi \\ \alpha &::= i \mid ?\phi \mid \alpha; \alpha \mid \alpha \cup \alpha \mid \alpha^* \\ \phi &::= \psi \mid \langle \alpha \rangle \phi \mid \phi \vee \phi \mid \phi \wedge \phi.\end{aligned}$$

It is well-known that diamond formulas are preserved from the larger to the smaller model under simulation. The following theorem generalizes this to cases where the vocabularies of the two models may be different.

Theorem 5. *If $\mathcal{M}, w \Leftarrow \mathcal{N}, v$ then all formulas ϕ over the vocabulary Q_M in the diamond fragment of L_{PI} are preserved from right to left under left simulation: if $\mathcal{N} \models_v \phi$ then $\mathcal{M} \models_w \phi$.*

Proof. Let C be a left simulation with wCv . We prove the property by induction on the construction of ϕ . If ϕ has the form ψ and is a proposition letter p , then p is in the vocabulary of \mathcal{M} , and the result holds by the restricted invariance property of C . Purely boolean combinations of ϕ are obvious. So the property holds for all boolean formulas ψ . As an example of the reasoning for $\langle \alpha \rangle$ we give the case of $\langle i \rangle \phi$. Suppose the property holds for ϕ , and assume $\mathcal{N} \models_v \langle i \rangle \phi$. Then there is a v' with $v \xrightarrow{i} v'$ and $\mathcal{N} \models_{v'} \phi$. By the zag property of C , there is a w with $w \xrightarrow{i} w'$ and $w'Cv'$. By the induction hypothesis, $\mathcal{M} \models_{w'} \phi$, and therefore $\mathcal{M} \models_w \langle i \rangle \phi$. \square

We can get rid of the vocabulary constraint. Since $\mathcal{M}, w \Leftarrow \mathcal{N}, v$ implies

$$\mathcal{M} \triangleleft_{Q_N, (w, V(v) \cap (Q_N - Q_M))} \Leftarrow \mathcal{N}, v,$$

we have:

Corollary 1. *If $\mathcal{M}, w \Leftarrow \mathcal{N}, v$ then all formulas ϕ in the diamond fragment of L_{PI} are preserved from right to left under left simulation: if $\mathcal{N} \models_v \phi$ then $\mathcal{M} \triangleleft_{Q_N, (w, V(v) \cap (Q_N - Q_M))} \models \phi$.*

The box fragment of L_{PI} is defined similarly to the diamond fragment. Box formulas are preserved in the other direction:

Theorem 6. *If $\mathcal{M}, w \Leftarrow \mathcal{N}, v$ then all formulas ϕ over the vocabulary Q_M in the box fragment of L_{PI} are preserved from left to right under left simulation: if $\mathcal{M} \models_w \phi$ then $\mathcal{N} \models_v \phi$.*

3 Composing Action Models

An action model is like a static model, but with valuations replaced by precondition formulas taken from an appropriate language.

An action model over P for agent set I is a quadruple (U, I, S, T) where U is a set of events, S is a function that assigns every $i \in I$ a binary relation S_i on U , and T is a function that assigns to every $u \in U$ an L_{PI} formula, the so-called precondition of u .


We will now give a version of action model update \otimes (see [BMS99]) for restricted models. A restricted action model for P, I is a quintuple (U, I, S, T, Q) where (U, I, S, T) is an action model, Q is a subset of P , and the formulas in $\{T(u) \mid u \in U\}$ use only proposition letters in Q .

Model expansion to a larger vocabulary X by means of $M \triangleleft X$ is used in the definition of product update to ensure that we get a result model for a vocabulary consisting of the union of the vocabulary of the static input model and the vocabulary of the action model.


Let $M = (W, I, R, V, Q)$ be a restricted static model and $A = (U, I, S, T, Q_1)$ a restricted action model for the same agent set I . Let X be the new vocabulary, i.e., $X = Q_1 - Q$. Then $M \otimes A$ is the static model (W', I, R', V', Q') given by $(M \triangleleft X) \otimes (U, I, S, T)$, where \otimes is the usual update product. This definition boils down to the following:

1. $W' = \{(w, X', u) \mid w \in W, u \in U, X' \subseteq X, M \triangleleft X \models_{(w, X')} T(u)\}$,
2. $(w, X', u)R'_i(w', X'', u')$ iff wR_iw' and uS_iu' ,
3. $V'(w, X', u) = V(w) \cup X'$.
4. $Q' = Q \cup X$.


Here is an example of update with a public announcement “At least one of the two children is muddy”, or $!(m_1 \vee m_2)$. This is represented by an event model with a single world and precondition $m_1 \vee m_2$. Note that the update involves model expansion:



Here is an update of the other component of the two-muddy-children model:



And here is the outcome of composing the two update results:



This is the same as the result of public announcement of $m_1 \vee m_2$ (“at least one of you is muddy”) on the composition on $m_1 \xleftrightarrow{1} \overline{m}_1$ and $m_2 \xleftrightarrow{2} \overline{m}_2$. The following theorem expresses that this outcome is not accidental: update of a composed model yields the same result (modulo bisimulation) as composing the updates of its components, provided the action model satisfies certain conditions. Call an action model propositionally differentiated if all its preconditions are boolean, and mutually exclusive.

Theorem 7. *If A is propositionally differentiated then:*

$$(\mathcal{M} \oplus \mathcal{N}) \otimes A \simeq (\mathcal{M} \otimes A) \oplus (\mathcal{N} \otimes A).$$

Proof. Now define a relation $Z \subseteq W_{(M \oplus N) \otimes A} \times W_{(M \otimes A) \oplus (N \otimes A)}$, as follows:

$$((w, v, X), s) Z ((w', X_1, s'), (v', X_2, s'')) \text{ iff } w = w', v = v', s = s' = s'', X = X_1 \cap X_2.$$

It is not hard to see that Z is total: Given $((w, v, X), s) \in W_{(M \oplus N) \otimes A}$, let $X_1 = X \cup (V_N(v) - V_M(w))$ and $X_2 = X \cup (V_M(w) - V_N(v))$, then $((w, X_1, s), (v, X_2, s)) \in W_{(M \otimes A) \oplus (N \otimes A)}$ and $X_1 \cap X_2 = X$. On the other hand, given $((w, X_1, s), (v, X_2, s)) \in W_{(M \otimes A) \oplus (N \otimes A)}$ we have $((w, v, X_1 \cap X_2), s) \in W_{(M \oplus N) \otimes A}$.

To check that Z is a bisimulation, assume that $((w, v, X), s) Z ((w, X_1, s), (v, X_2, s))$ and note that:

$$\begin{aligned} V_{(M \oplus N) \otimes A}(w, v, X), s) &= V_{M \oplus N}(w, v, X) \\ &= V_M(v) \cup V_N(w) \cup X \\ &= (V_M(v) \cup X_1) \uplus (V_N(w) \cup X_2) \\ &= V_{(M \otimes A) \oplus (N \otimes A)}((w, X_1, s), (v, X_2, s)). \end{aligned}$$

This proves invariance.

For zig, assume $((w, v, X), s) \rightarrow ((w', X', v'), s')$. Then $w \rightarrow w'$ and $v \rightarrow v'$ and $s \rightarrow s'$. and (w', X'_1) and (v', X'_2) satisfy the precondition of s' , where $X'_1 = X' \cup (V_N(v') - V_M(w'))$ and $X'_2 = X' \cup (V_M(w') - V_N(v'))$. It follows that $((w, X_1, s), (v, X_2, s)) \rightarrow ((w, X_1, s), (v, X_2, s))$, and we get $((w', X', v'), s') Z ((w, X_1, s), (v, X_2, s))$ by definition of Z .

For zag, assume $((w, X_1, s), (v, X_2, s)) \rightarrow ((w', X'_1, s'), (v', X'_2, s''))$. Then from the fact that the preconditions of A are mutually exclusive, we know that $s' = s''$, and we get from the definition of \oplus that $(w, X_1, s) \rightarrow (w', X'_1, s')$ and $(v, X_2, s) \rightarrow (v', X'_2, s')$. Since both (w', X'_1) and (v', X'_2) satisfy the precondition of s' , we can join the vocabulary extensions X'_1 and X'_2 to a single extension X' of the joint vocabulary of w' and v' , and conclude $((w, v, X), s) \rightarrow ((w', X', v'), s')$. Again we get that Z relation between $((w', X', v'), s')$ and $((w, X_1, s), (v, X_2, s))$ by definition of Z . \square

Action models are very similar to static models, and it turns out that parallel composition on action models can be defined in a natural way.

The parallel composition $A \oplus B$ of two restricted action models A and B is given by (U, I, S, T, Q) , where $U = \{(w, v) \mid w \in U_A, v \in U_B\}$, S is given by

$$(w, v) S_i(w', v') \text{ iff } w S_{i_A} w' \text{ and } v S_{i_B} v',$$

T by $T(w, v) = T_A(w) \wedge T_B(v)$, and Q by $Q = Q_A \cup Q_B$.

Updating with a composite action model should yield the same outcome as updating with its components and then composing the results. The following theorem says that it does.

Theorem 8. $M \otimes (A \oplus B) \Leftrightarrow (M \otimes A) \oplus (M \otimes B)$.

Proof. Let the relation $Z \subseteq W_{M \otimes (A \oplus B)} \times W_{(M \otimes A) \oplus (M \otimes B)}$ be given by

$$(w, X, (s, t))Z((w', X_1, s'), (w'', X_2, t')) \text{ iff } w = w' = w'', s = s', t = t' \text{ and } X = X_1 \cup X_2$$

To see that $(w, X, (s, t))$ is in $W_{M \otimes (A \oplus B)}$ implies $((w, X_1, s), (w, X_2, t))$ (with $X_1 = X \cap Q_A$ and $X_2 = X \cap Q_B$) is in $W_{(M \otimes A) \oplus (M \otimes B)}$, notice that

$$(w, X, (s, t)) \in W_{M \otimes (A \oplus B)}$$

\Rightarrow

$$M \triangleleft (Q_A \cup Q_B) \models_{w, X} T_A(s) \wedge T_B(t)$$

\Rightarrow

$$M \triangleleft Q_A \models_{w, X \cap Q_A} T_A(s) \text{ and } M \triangleleft Q_B \models_{w, X \cap Q_B} T_B(t)$$

\Rightarrow

$$(w, X_1, s) \in W_{M \otimes A} \text{ and } (w, X_2, t) \in W_{M \otimes B}$$

\Rightarrow

$$((w, X_1, s), (w, X_2, t)) \in W_{(M \otimes A) \oplus (M \otimes B)}.$$

On the other hand, if $((w, X_1, s), (w, X_2, t))$ is in $W_{(M \otimes A) \oplus (M \otimes B)}$ then it is not hard to see that $(w, X_1 \cup X_2, (s, t))$ is in $W_{M \otimes (A \oplus B)}$. This proves the totality of Z .

The invariance, zig and zag properties of Z are immediate based on the totality of Z . \square

4 Composing Change Models

A *PI* substitution is an expression of the form $\{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$ with the $p_k \in P$, all different, and each $\phi_k \in L_{PI}$. If $\sigma = \{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$ is a substitution, we call $\{p_1, \dots, p_n\}$ its domain, notation $\text{dom}(\sigma)$. *PI* substitutions can be used to express factual change (see [BvEK06]). An action model with change (henceforth: a change model) is a quintuple (U, I, S, T, C) where (U, I, S, T) is an action model, and C is a function that assigns to every $u \in U$ a *PI* substitution. The substitutions express instructions for changing the facts of the world.

A substitution $\sigma = \{p_1 \mapsto \phi_1, \dots, p_n \mapsto \phi_n\}$ is over a vocabulary Q if each $p_k \in Q$ and moreover the set of proposition letters in each ϕ_k is contained in Q . A change model (U, I, S, T, C) is over a vocabulary Q if (U, S, T) is over vocabulary Q , and each $C(u)$ is a substitution over Q . A restricted change model is a quintuple (U, I, S, T, C, D) such that (U, I, S, T, C) is a change model over vocabulary D .

The definition of \otimes for update of restricted models with change models is similar to that of \otimes for action models.

As an example, consider an epistemic model for the 100 prisoners and a lightbulb riddle [DvEW10]. This has 100 proposition letters and 2^{100} possible states. The simplest protocol for solving the riddle has an agent acting as counter who only is aware of a single boolean variable (for signaling whether the lightbulb is on or off) and of the value of a counting register (for keeping track of how many times she has switched the light off). Here is how the counter component of the model gets updated:

$$\begin{array}{c} \text{light}, c = n \\ \hline \end{array} \otimes \boxed{\text{light}, \{\text{light} := \perp, c := c + 1\}} \Leftrightarrow \overline{\text{light}}, c = n + 1$$

This models an update with the event of finding the light on and carrying out part of the protocol: switch the light off and increment the counting register.

Next, consider what happens if one of the non-counting prisoners gets interrogated, seen from the point of view of the counter. If the prisoner has not touched the switch before and finds the light off he will switch it on. This can be viewed as an update with

$$\boxed{\top, \{\text{light} := \text{done}_i \rightarrow \text{light}, \text{done}_i := \text{light} \rightarrow \text{done}_i\}}$$

where done_i indicates that prisoner i has done his signalling deed. From the counter perspective we get a simplified view of this. Suppose we start out from a situation where the counter knows that the light is off. Next, prisoner i gets interrogated, and light is set to $\text{done}_i \rightarrow \text{light}$. Since the counter cannot make observations about done_i , the result is that afterwards she does not know what has happened to the light.

The definition of $A \oplus B$ for change models is as before, but with the new change component C of the composition result given by:

$$\begin{array}{ll} C(w, v) \text{ binds } p \text{ to } C_A(w)(p) \wedge C_B(v)(p) & \text{if } p \in \text{dom}(C_A(w)) \cap \text{dom}(C_B(v)) \\ C(w, v) \text{ binds } p \text{ to } C_A(w)(p) & \text{if } p \in \text{dom}(C_A(w)), p \notin \text{dom}(C_B(v)) \\ C(w, v) \text{ binds } p \text{ to } C_B(v)(p) & \text{if } p \notin \text{dom}(C_A(w)), p \in \text{dom}(C_B(v)) \end{array}$$

Theorems 7 and 8 extend to change models.

5 Characterizing Epistemic Models in Terms of Composition

Our framework suggests new questions about composition: is it possible to build every multimodal S5 model from components with at most two worlds? The answer turns out to be ‘no’, and the follow-up question is: what do the multimodal S5 models that can be built from components with at most two worlds look like? Similar questions can be asked about KD45 models. If a model is built from S5 components using \oplus , then the result is again S5. How about building from KD45 components: will the result again be KD45?

There are also more general questions: for which (finite) \mathcal{M} is it possible to find a decomposition of \mathcal{M} (modulo bisimulation) such that each component has strictly smaller vocabulary than \mathcal{M} ? Or more general still: Is there a *normal form* of \mathcal{M} by composition and relativization (public announcement)? There is a connection between this question and the prime factor product decomposition theorem in graph theory (see, e.g., [IK00], theorem 5.21).

As an example result in this area, we give a characterization of models that can be decomposed into two-world building blocks.

Theorem 9. *An S5-model can be generated from two-world S5 components iff it is bisimilar to a model satisfying the following conditions:*

1. all worlds have different valuations
2. for any proposition p , either all worlds agree on the valuation of p or p is in some subset of propositions P such that
 - (a) half of the worlds have one valuation V of the propositions in P and the other half of the worlds have the opposite valuation \bar{V} of the propositions in P
 - (b) there is a bisimulation Z between the model restricted to worlds with valuation V of P and the model restricted to worlds with valuation \bar{V} of P , if we limit the invariance conditions to propositions not in P
 - (c) there is a fixed set of agents J such that if w and w' differ on the valuation of P , wZu and $u'Zw'$ then the set of agents that relate w and w' is the intersection of J with the set of agents that relate w and u' or, equivalently, w' and u .

Proof. \Leftarrow : For any model satisfying above conditions, take a set of propositions P satisfying conditions 2a, 2b and 2c. We can decompose the model into two smaller models: a two-world model with vocabulary P of which one world has valuation V and the other valuation \bar{V} ; and one of the two models that are related by the bisimulation Z , restricted to propositions not in P . This latter model also satisfies the above conditions, so we can also decompose that model resulting in another model satisfying the conditions, and so on. Since the conditions imply that the number of worlds in the model is a power of two, we will end up with a sequence of two-world models.

\Rightarrow : Zero-world, single-world and two-world models trivially satisfy conditions 1 and 2.

Suppose we have a model \mathcal{M} satisfying conditions 1 and 2 and we compose it with the two-world model \mathcal{K} to get a new model \mathcal{N} . Suppose \mathcal{K} satisfies condition 1. Then clearly this property is preserved in \mathcal{N} . Suppose the two worlds of \mathcal{K} have identical valuations, then \mathcal{N} is a model consisting of two ‘copies’ of each world of \mathcal{M} that matches the valuation of the worlds of \mathcal{K} . We can remove one of each two copies to get a model bisimilar to \mathcal{N} satisfying condition 1. This operation does not effect the satisfaction of condition 2. So we can assume \mathcal{N} satisfies condition 1. Now we will show that \mathcal{N} satisfies condition 2. Let p be a proposition. Clearly, if either all worlds of \mathcal{M} or both worlds of \mathcal{K} agree on the valuation of p then \mathcal{N} satisfies condition 2. Suppose otherwise. Let P , V , \bar{V} and J be the corresponding subset of propositions, valuations and set of agents justifying satisfaction of condition 2 for p in \mathcal{M} . There are three possibilities:

- Neither V nor \bar{V} match the valuation of any worlds of \mathcal{K} . In this case \mathcal{N} is the zero-world model.
- V matches the valuation of one world of \mathcal{K} and \bar{V} matches the valuation of no worlds of \mathcal{K} , or the other way around. In this case all worlds in \mathcal{N} will agree on the valuation of p .
- Both V and \bar{V} match the valuation of one world of \mathcal{K} (clearly, these worlds must be different). Let W be the set of worlds with (partial) valuation V in \mathcal{M} that match a world of \mathcal{K} , and let W' be the set of worlds with (partial) valuation \bar{V} in \mathcal{M} that are bisimilar to W . Clearly this is exactly the set of worlds with (partial) valuation \bar{V} that match a world of \mathcal{K} . Now let U and U' be the sets of worlds in \mathcal{N} that result from the worlds in W and W' and let P' be the result of removing from P all propositions for which both worlds of \mathcal{K} have the same valuation. Then P' is a subset of propositions containing p satisfying conditions 2a, 2b and 2c in \mathcal{N} . U is one half of the worlds of \mathcal{N} satisfying $V \cap P'$ and U' is the other half of worlds satisfying $\bar{V} \cap P'$, and the appropriate set of agents is the intersection of J with the set of agents relating the worlds of \mathcal{K} . \square

As an immediate application of this theorem, we see that any initial muddy children model can be built from two-world components, but the model that results from the father’s announcement ‘At least one of you is muddy’ cannot.

Our composition approach holds promise for epistemic model checking with dynamic epistemic logic. The following theorem gives examples of epistemic properties that can be checked on small components of large models.

Theorem 10 (Preservation). *If a pointed model (\mathcal{M}, s) is decomposable into reflexive pointed models $(\mathcal{M}_0, s_0), \dots, (\mathcal{M}_n, s_n)$ with disjoint vocabularies Q_0, Q_1, \dots, Q_n , then for any epistemic formula ϕ based on $Q_i : \mathcal{M}_i, s_i \models \phi \iff \mathcal{M}, s \models \phi$.*

Proof. Let Z_i be the relation between the worlds of \mathcal{M} and the worlds of \mathcal{M}_i given by $\bar{i}Z_it$ iff $\bar{i}[i] = t$ (where $\bar{i}[i]$ is the i th component of the vector \bar{i}). We show that Z_i is a Q_i -restricted bisimulation (a bisimulation with the invariance condition restricted to Q_i). Assume $\bar{i}Z_it$. Then $V(\bar{i}) \cap Q_i = V_i(t)$, by the definition of parallel composition and the fact that Q_i is disjoint from the other vocabularies. Thus, Q_i -restricted invariance holds. Next suppose $\bar{i}R_k\bar{u}$. Then by the definition of the accessibility relations on \mathcal{M} , $\bar{i}[i]R_k\bar{u}[i]$, whence, by definition of Z_i , there is a u in the domain of \mathcal{M}_i with $\bar{u}Z_it$. It follows that the zig condition holds. Finally, assume there is a u in the domain of \mathcal{M}_i with tR_ku . Consider the state \bar{u} given by $\bar{u}[i] = u$ and $\bar{u}[j] = \bar{i}[j]$ for $j \neq i$. Then by reflexivity of the component models and the fact that tR_ku , $\bar{i}R_k\bar{u}$. From the definitions of \bar{u} and Z_i we get that $\bar{u}Z_it$, i.e., the zag condition holds. \square

For an example application, consider a muddy children model of 2^n components. This can be viewed as built from n reflexive two-component models, each with its own vocabulary for talking about the muddiness of a single child. Any epistemic statement that talks about the muddiness of a single child in the big model can be checked by evaluation in a single two-world component.

The following definition is from the interpreted systems literature (cf. [EvdMM98]). A basic proposition $p \in Q$ is i -local for $i \in I$ in a model $\mathcal{M} = (W, I, \sim, V, Q)$, if $w \sim_i v$ implies that $p \in V(w) \iff p \in V(v)$. A model $\mathcal{M} = (W, I, \sim, V, Q)$ is said to be *locally generated* if for any $i \in I$, $w \sim_i v \iff (Q_i \cap V(w) = Q_i \cap V(v))$ where Q_i is the set of all the i -local propositions in \mathcal{M} . (We leave the logical characterization of this property for a future occasion.)

Intuitively, the i -local propositions are the *atomic observables* of agent i . A model is locally generated if those atomic observables determine the epistemic relations in the model. The muddy children model is a typical example of a locally generated model. But note that the property is not preserved under bisimulation: the fact that \mathcal{M}, s is locally generated does not imply that \mathcal{M}', s' is locally generated, even if $\mathcal{M}, s \simeq \mathcal{M}', s'$, for unconnected states harm.

Theorem 11 (Decomposition of locally generated models). *If $\mathcal{M} = (W, I, \sim, V, Q)$ is locally generated, then there are partial models $\mathcal{M}_1, \dots, \mathcal{M}_n$ and \mathcal{M}_0 such that $|I| = n$ and $\mathcal{M}_i = (W_i, I, \sim, V_i, Q_i)$ such that:*


- $Q_i = Q_i$ for $i \leq 0$ and $Q_0 = Q$;
- $\mathcal{M} \simeq (\mathcal{M}_0 \oplus \mathcal{M}_1 \oplus \dots \oplus \mathcal{M}_n)$.

Proof. For $i > 0$, let \mathcal{M}_i be (the bisimulation contraction of) (W, I, \sim, V_i, Q_i) where V_i is the restriction of V to Q_i . Let \mathcal{M}_0 be the universal ignorance model with the same state space: (W, I, \sim', V, Q) where \sim'_i is the universal relation on $S \times S$ for each $i \in I$. \square


Given a model $\mathcal{M} = (W, I, \sim, V, \mathcal{Q})$, let $\phi_{\mathcal{M}}$ be a DNF formula listing all the $V(w)$ in \mathcal{M} , i.e., $\phi_{\mathcal{M}} = \bigvee_{w \in W} (\bigwedge_{\{p \in V(w)\}} p \wedge \bigwedge_{\{p \notin V(w)\}} \neg p)$. We can then also rephrase the theorem as $\mathcal{M} \simeq (\mathcal{M}_0 \oplus \mathcal{M}_1 \oplus \dots \oplus \mathcal{M}_n) \upharpoonright_{\phi_{\mathcal{M}}}$ such that \mathcal{M}_0 is a universal ignorance model w.r.t to $\mathcal{Q}_0 = \mathcal{Q} - \bigcup_{i \in I} \mathcal{Q}_i$.

This theorem says that we can decompose a locally-generated model according to the observables of each agent and the *states of affairs* considered in the model. Muddy Children is an example of such decomposition where all the propositions are local for some agents.

On the other hand, there are models which are not locally generated but decomposable in a non-trivial way. The following model \mathcal{M} that can be decomposed into two three-world-models is an example:



It is clear that the p is local for the line agent 1 and q local for the dot agent 2, but \mathcal{M} is not locally generated. Nevertheless, \mathcal{M} can be decomposed as follows:



If we take the boldface states as the real worlds in these two models respectively, then the two models capture the situations where agent 2 is not sure about whether 1 knows that 2 knows q and agent 1 is not sure about whether 2 knows that 1 knows p .

6 Conclusions, Connections, and Future Work

We intend to extend the epistemic model checker DEMO [Eij07] with model composition operations, to investigate the practical usefulness of the approach. Our approach of composing epistemic models from small components differs in an interesting way from the decomposition by symmetry reduction technique of [CDLQ09].

In [DHKW08] ‘cooperative boolean games’ are studied: games where agents cooperatively can achieve goals stated as propositional formulas. In the present framework, the variables under the control of an agent can be taken to be the variables that are in the domain of a substitution in component models representing the perceptive and control abilities of agents. This points the way towards extending cooperative boolean games with an epistemic dimension, and for building a logical framework for the study of cooperative epistemic games.

Van Ditmarsch and French [DF09] study an operator $G\phi$ with semantics $\mathcal{M} \models_w G\phi$ iff for all \mathcal{M}', w' with $\mathcal{M}, w \preceq \mathcal{M}', w'$ it holds that $\mathcal{M}' \models_{w'} \phi$. Their definition of \preceq differs from ours in that they do not work with vocabulary restrictions on their models. In case $\mathcal{M}, w \preceq \mathcal{M}', w'$ they call \mathcal{M}', w' a refinement of \mathcal{M}, w , and they prove that product updates are refinements. This should be compared to our Theorem 2. There is also an obvious connection to the dynamics of awareness, as studied in [BVQ09].

Finally, the combination of communicative actions and vocabulary expansion deserves separate study. A first task here could be to axiomatize the strong Kleene logic of public announcement $!\phi$ and vocabulary expansion $\#p$, where $\#p$ is interpreted as the model changing operation $\mathcal{M} \mapsto \mathcal{M} \triangleleft \{p\}$.

Acknowledgements The second author was supported by the NWO project VEMPS, the third author by the NWO project KISEN. This support is gratefully acknowledged. We thank two anonymous LOFT referees for useful feedback.

References

- [BH⁺09] J. van Benthem, T. Hoshi, J. Gerbrandy, and E. Pacuit. Merging frameworks for interaction. *Journal of Philosophical Logic*, 38(5):491–526, 2009.
- [BMS99] A. Baltag, L.S. Moss, and S. Solecki. The logic of public announcements, common knowledge, and private suspicions. Technical Report SEN-R9922, CWI, Amsterdam, 1999.
- [BRV01] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2001.
- [BvEK06] J. van Benthem, J. van Eijck, and B. Kooi. Logics of communication and change. *Information and Computation*, 204(11):1620–1662, 2006.
- [BVQ09] J. van Benthem and F. Velázquez-Quesada. Inference, promotion, and the dynamics of awareness. Technical Report 43, ILLC, Amsterdam, 2009. Available from www.illc.uva.nl/Publications/ResearchReports/PP-2009-43.text.pdf.
- [CDLQ09] Mika Cohen, Mads Dam, Alessio Lomuscio, and Hongyang Qu. A symmetry reduction technique for model checking temporal-epistemic logic. In *IJCAI'09: Proceedings of the 21st international joint conference on Artificial intelligence*, pages 721–726, San Francisco, CA, USA, 2009. Morgan Kaufmann Publishers Inc.
- [DF09] H.P. van Ditmarsch and T. French. Simulation and information. In J. Broersen and J.-J. Meyer, editors, *Knowledge Representation for Agents and Multi-Agent Systems (Proceedings of KRAMAS 2008, Sydney, Australia)*, LNAI 5605, pages 51–65. Springer, 2009.
- [DHKW08] P. E. Dunne, W. van der Hoek, S. Kraus, and M. Wooldridge. Cooperative boolean games. In Padgham, Parkes, Müller, and Parsons, editors, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pages 1015–1022, May 2008.
- [DvdHK07] H.P. van Ditmarsch, W. van der Hoek, and B.P. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*. Springer, 2007.
- [DvEW10] Hans van Ditmarsch, Jan van Eijck, and William Wu. One hundred prisoners and a lightbulb – logic and computation. In *Proceedings of KR2010*. AAAI, 2010.

- [Eij07] Jan van Eijck. DEMO — a demo of epistemic modelling. In Johan van Benthem, Dov Gabbay, and Benedikt Löwe, editors, *Interactive Logic — Proceedings of the 7th Augustus de Morgan Workshop*, number 1 in Texts in Logic and Games, pages 305–363. Amsterdam University Press, 2007.
- [EvdMM98] Kai Engelhardt, Ron van der Meyden, and Yoram Moses. Knowledge and the logic of local propositions. In *TARK '98: Proceedings of the 7th conference on Theoretical aspects of rationality and knowledge*, pages 29–41, San Francisco, CA, USA, 1998. Morgan Kaufmann Publishers Inc.
- [FHMV95] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT Press, 1995.
- [Ger99] J. Gerbrandy. *Bisimulations on planet Kripke*. PhD thesis, ILLC, Amsterdam, 1999.
- [IK00] Wilfried Imrich and Sandi Klavár. *Product Graphs*. Wiley Interscience Series in Discrete Mathematics and Optimization. Wiley, New York etc., 2000.
- [Kle50] S.C. Kleene. *Introduction to Metamathematics*. Van Nostrand, Princeton, NJ, 1950.
- [KP10] B. Kooi and E. Pacuit. Logics of rational interaction. In *Dynamic Epistemology: Contemporary Perspectives*. 2010.
- [Wanng] Yanjing Wang. *Epistemic Modelling and Protocol Dynamics*. PhD thesis, Universiteit van Amsterdam, 2010 (forthcoming).