The Muddy Children Puzzle and Action Models

Alexandru Dragomir

Graduate student, University of Bucharest

January 10, 2012

Abstract

This paper's aim is twofold: first, to present Baltag, Moss and Solecki's *Action Model Logic* (see [2] and [8]) as a kind of Dynamic Epistemic Logic. Second, to present the intuitive, informal, solution to the muddy children puzzle and a formal solution using the logical apparatus of Action Model Logic.

1 The Logical Apparatus

In the following subsections I will present the innovative approach to dynamizing epistemic logic of Baltag, Moss and Solecki in [2]. Because this paper is intended to be (somewhat) self-contained, I will also introduce the reader to the very basics of epistemic logic.

1.1 Syntax

The well-formed formulas of a language $\mathfrak{L}$ of epistemic logic can be specified by the following Nackus-Baur form:

$\phi ::= p \mid \phi \mid \neg\phi \mid \phi \wedge \phi \mid K_a\phi$, for A a set of agents.

Formula $K_a\phi$ will be read *agent a knows that* ϕ. It may be that the truth value of ϕ is not interesting, but that a knows whether it is true or false! We'll formalize this by: $K_a\phi \vee K_a\neg\phi$.

As a tool for reasoning about knowledge, the $S5$ system is the most used:

(K) $K_a(\phi \rightarrow \psi) \rightarrow (K_a\phi \rightarrow K_a\psi)$ (*normality* or the *distribution* of K_a over $\rightarrow$)

(T) $K_a\phi \rightarrow \phi$ (*veridicity*)

(4) $K_a\phi \rightarrow K_aK_a\phi$ (*positive introspection*)

(5) $\neg K_a\phi \rightarrow K_a\neg K_a\phi$ (*negative introspection*)

1.2 Kripke Frames

A Kripke frame is a structure: $\mathfrak{F} = (W, \{R_i\}_{i\in I})$, where W is a finite, non-empty set of possible worlds (nodes, points, states), and $\{R_i\}_{i\in I}$ is a set of binary *accesibility relations* between elements of W and indexed by the agents in a set I: $R_i \subseteq W \times W$, for $i \in I$.

1.3 Kripke models

A Kripke model is a structure $\mathfrak{M} = (\mathfrak{F}, \pi) = (W, \{\longrightarrow_a |a \in A\}, \pi)$, where: $\mathfrak{F}$ is a frame and π is *the valuation function*, the function that assigns sets of worlds to propositional variables: $\pi : Var(\mathfrak{L}) \longrightarrow 2^W$.

The relation of satisfaction can be defined as: $\models$: $\models \subseteq (\mathfrak{M}, w) \times Prop(\mathfrak{L})$, the smallest relation that respects the following conditions (for a model $\mathfrak{M} = (W, \longrightarrow, \pi)$):

R1. $\mathfrak{M}, w \models p$ iff $w \in \pi(p)$, for $p \in W$ and $p \in Var(\mathfrak{L})$.

R2. $\mathfrak{M}, w \models \neg\phi$ iff $\mathfrak{M}, w \not\models \phi$, for $\phi \in Prop(\mathfrak{L})$

R3. $\mathfrak{M}, w \models \phi \wedge \psi$ iff $\mathfrak{M}, w \models \phi$ and $\mathfrak{M}, w \models \psi$, $\phi, \psi \in Prop(\mathfrak{L})$

R4. $\mathfrak{M}, w \models \phi \vee \psi$ iff $\mathfrak{M}, w \models \phi$ or $\mathfrak{M}, w \models \psi$, $\forall \phi, \psi \in Prop(\mathfrak{L})$

R5. $\mathfrak{M}, w \models \phi \rightarrow \psi$ iff $\mathfrak{M}, w \not\models \phi$ or $\mathfrak{M}, w \models \psi$, $\forall \phi, \psi \in Prop(\mathfrak{L})$

R6. $\mathfrak{M}, w \models \Diamond\phi$ iff $\exists u : Rwu \wedge \mathfrak{M}, u \models \phi$, $\forall \phi \in Prop(\mathfrak{L})$

R7. $\mathfrak{M}, w \models \Box\phi$ iff $\forall u : Rwu \Rightarrow \mathfrak{M}, u \models \phi$, $\forall \phi \in \mathfrak{L}$

1.4 Action models

This subsection will be based on: [8], [1], [2].

Definition. An *action model* is an **S5** model:

$\mathsf{M} = (\mathsf{S}, \longrightarrow_a, \mathsf{pre})$, where:

1. S is a domain of actions (action points or states), analogous to the possible worlds of an ordinary Kripke model,

2. $\longrightarrow_a$ is a accessibility relation for an agent $a \in A$ (where A is our agents set),

3. pre is a function that assigns to each action a precondition: $\mathsf{pre} : \mathsf{S} \to \mathfrak{L}$. As we can see, a precondition is an atom or a whole proposition of our language $\mathfrak{L}$.

Until now, we have defined two kinds of models: *epistemic models*, used to represent the knowledge of agents and *action models*, used to describe actions with respect to their preconditions.

Definition. Let $(\mathfrak{M}, s)$ be an epistemic state (epistemic model and actual world) of $\mathfrak{M} = (S, \longrightarrow, \pi)$, let M=$(\mathsf{S}, \longrightarrow, \mathsf{pre})$ be an action model. Then, $\mathfrak{M}' = (\mathfrak{M} \otimes \mathsf{M}) = (S', \longrightarrow', \pi')$ is defined as:

1. $S' := \{(s, \mathsf{s}) \mid s \in S, \mathsf{s} \in \mathsf{S}, \mathfrak{M}, s \models \mathsf{pre}(\mathsf{s})\}$

2. $(s, \mathsf{s}) \longrightarrow'_a (t, \mathsf{t})$ iff $s \longrightarrow_a t$ and $\mathsf{s} \longrightarrow_a \mathsf{t}$

3. $(s, \mathsf{s}) \in \pi'(p)$ iff $s \in \pi(p)$

Let's note that actions are epistemic, not ontic: they affect only the knowledge state of an agent, but not the world - or the state of affairs presupposed in our multi-agent setting. After executing an action at an epistemic state, what we obtain is another epistemic state, not an action state ! The underlying intuition is that if an agent has a certain set of knowledge or beliefs, after executing certain actions on that set, what is obtained is another set of knowledge or beliefs, modified accordingly to what that action meant. Turning a little bit more technical, any result-state of a restricted modal product (or the resulting state of an execution of an action at an epistemic state) is a tuple composed of a possible world and an action point: (s, s), $s \in W$, $\mathsf{s} \in \mathsf{S}$.

Also, one should note that any resulting state: (s, s) is constructed so as to respect the following property: $\mathfrak{M}, s \models \mathsf{pre}(\mathsf{s})$. Informally, a resulting state is composed of an epistemic state and action that can be executed at that state, meaning that the precondition of action s is (a formula) satisfied at epistemic state s. Now we have a better understanding of why the resulting model is called an *restricted modal product*.

A semantic rule stating that after executing an action (S, s) in the epistemic model (S, s) with the effect that the resulting model is (S', s')[1] is the following:

$(S, s)\,[[\mathsf{S}, \mathsf{s}]]\,(S', s')$ ddacă $S, s \models \mathsf{pre}(\mathsf{s})$ şi $S', s' = (S \otimes \mathsf{S}, (s, \mathsf{s}))$

Also, the semantic rule for the evaluation of an execution of a program at any state is the following: (for $\mathfrak{M} = (S, \longrightarrow, \pi)$ şi $\mathsf{M} = (\mathsf{S}, \longrightarrow, \mathsf{pre})$):

$\mathfrak{M}, s \models [\mathsf{M}, \mathsf{s}]\phi$ ddacă $\mathfrak{M}, s \models \mathsf{pre}(\mathsf{s}) \Rightarrow (\mathfrak{M} \otimes \mathsf{M}, (s, \mathsf{s})) \models \phi$

This rule can be easily read as: at state s, in model $\mathfrak{M}$, it is true that ϕ after the execution of an action s of the action model M, iff: if at world s the precondition of action s is satisfied, then, in the newly obtained model $\mathfrak{M}' = \mathfrak{M} \otimes \mathsf{M}$, at world (s, s) the formula ϕ is satisfied.

The composition of two action models

Definition. Let's consider two action models: $\mathsf{M} = (\mathsf{S}, \longrightarrow, \mathsf{pre})$, $\mathsf{M}' = (\mathsf{S}', \longrightarrow'$ $, \mathsf{pre}')$. Their composition will be:

$(\mathsf{M}; \mathsf{M}') = (\mathsf{S}'', \longrightarrow'', \mathsf{pre}'')$ such that:

1. $\mathsf{S}'' = \mathsf{S} \times \mathsf{S}'$

[1] van Ditmarsch et al., *Dynamic Epistemic Logic*, p. 151

2. $(\mathsf{s},\mathsf{s}') \longrightarrow''_a (\mathsf{t},\mathsf{t}'')$ iff $\mathsf{s} \longrightarrow_a$ şi $\mathsf{s}' \longrightarrow;_a \mathsf{t}'$

3. $\mathsf{pre}''((\mathsf{s},\mathsf{s}')) = \langle \mathsf{M},\mathsf{s}\rangle \mathsf{pre}'(\mathsf{s}')$

Proposition. If the preconditions of two actions are not modalized (they do not contain epistemic operators), then the following is true:

$(\mathfrak{M} \otimes \mathsf{M}_1) \otimes \mathsf{M}_2 \simeq (\mathfrak{M} \otimes \mathsf{M}_2) \otimes \mathsf{M}_1$

Proof. See [4].

1.5 Public Announcements in an Action Model Logic setting

Proposition (see [8], p.150). Public announcements are actions executed at epistemic states. Equivalently, there is an action, pub (whose precondition is ϕ, the formula to be announced), such that after an execution of this action in an epistemic state $(\mathfrak{M}, w)$, will yield an epistemic state $(\mathfrak{M}', w')$ that represents the epistemic state after the announcement of ϕ.

Proof. Formally, the execution of an announcement can be modelled as follows:

$((\mathsf{S}, \longrightarrow, \mathsf{pre}), \mathsf{s} \in \mathsf{S}) := ((\{\mathsf{pub}\}, \longrightarrow, \mathsf{pre}), \mathsf{pub})$

$\mathsf{pre}(\mathsf{pub}) = \phi$

$\mathsf{pub} \longrightarrow_a \mathsf{pub}, \forall a \in A$

Let $\mathsf{Pub} = (\{\mathsf{pub}\}, \longrightarrow, \mathsf{pre})$ be the action model that represents the action of performing an announcement.

Also, let's assume that $\mathfrak{M}, s \models \phi$. Then, the execution of Pub in model $\mathfrak{M}$ is

$\mathfrak{M} \otimes \mathsf{Pub}$, defined as:

1. *The domain* $S' \in (\mathfrak{M} \otimes \mathsf{Pub}) = (\mathfrak{M} \otimes \{\mathsf{pub}\})$ is composed of all (s, pub) such that: $\mathfrak{M}, s \models \mathsf{pre}(\mathsf{pub})$ (equivalently, in this case: $\mathfrak{M}, s \models \phi$).

Note that the domain of the newly obtained epistemic model is composed of worlds at which the formula ϕ is satisfied (the model is restricted to the worlds at which it is true that ϕ).

2. *Accessibility relations*. Following the definition of the restricted modal product (or the *update product*), we have that:

$(s, \mathsf{pub}) \longrightarrow'_a (t, \mathsf{pub})$ iff $s \longrightarrow_a t$ and $\mathsf{pub} \longrightarrow_a \mathsf{pub}$

Because Pub is an **S5** model, we have that $\mathsf{pub} \longrightarrow \mathsf{pub}$. So :

$(s, \mathsf{pub}) \longrightarrow'_a (t, \mathsf{pub})$ iff $s \longrightarrow_a t$

It's easy to see that the above relation is restricted to states at which the announced formula is satisfied: just observe that (s, s) has the property: $\mathfrak{M}, s \models \phi$, and that $\longrightarrow'_a$ links only worlds that validate ϕ.

3. The *valuation* of model $\mathfrak{M} \otimes \mathsf{Pub}$ is (nothing surprising here):

$s \in \pi(p) \Leftrightarrow (s, \mathsf{pub}) \in \pi'(p)$

Which is equivalent to (using the notion of satisfaction in a model):

$\mathfrak{M}, s \models p \Leftrightarrow (\mathfrak{M} \otimes \mathsf{S})(s, \mathsf{s}) \models p$

Examples: See the pictures below and the explanation:

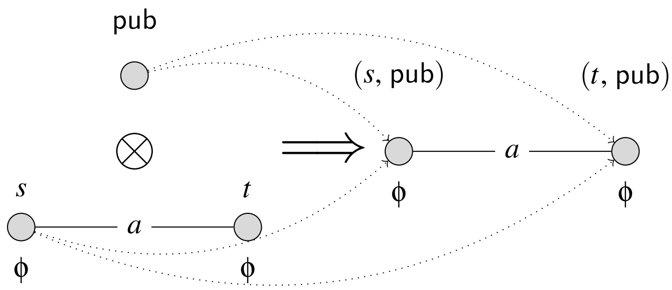

Figure 1 represents the execution of action $\mathsf{Pub} = (\{\mathsf{pub}\}, \longrightarrow, \mathsf{pre})$ in $\mathfrak{M} = (\{s,t\}, \longrightarrow, \pi)$. Because at both s and t it is true that ϕ, $\mathfrak{M} \otimes \mathsf{Pub}$ has two nodes: (s, pub) and (t, pub). Both satisfy ϕ.

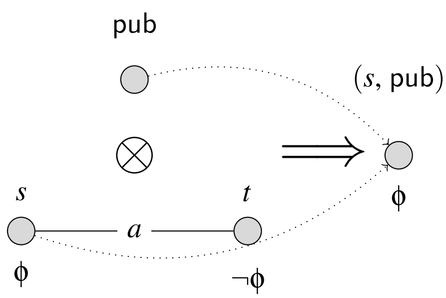

Figure 2 presents the transformation of a model in which only one world verifies the precondition of the executed action. If the precondition of pub is ϕ, and ϕ is true only at s, then after executing it we obtain an epistemic model with only a single world: (s, pub), a world that satisfies ϕ.

2 The Muddy Children Puzzle

For starters, I will offer a presentation (based on [8], [6] [2], [7]) of the muddy children puzzle. Suppose Anne, Bob and Charlie play outside and two of them

(Anne and Bob) get mud on their foreheads. Father tells them:

At least one of you has mud on his forehead.

Afterwards, their father asks them:

If you know whether you have mud on your forehead or not, raise your arm!

But none of them raises an arm. Father asks them again and... surprise ! Anne, Bill and Charlie raise their arms ! Why is that ?

2.1 An intuitive solution.

Please note that the children are perfect reasoners: they're able to make instantly (meaning that it requires them no time at all to logically pass from one sentence to another) all kinds of inferences allowed in an epistemic logic . Suppose Anne and Bob are muddy. Anne reasons as follows: after father's announcement that at least one of them is muddy, if she's not muddy, Bob will know he's muddy, because he sees a clean Anne and a clean Charlie. But after father asks the children who know whether they're muddy or not to raise an arm and Bob doesn't, she infers that she must be muddy. So, the second time father asks the children to raise an arm if they know their status, she raises her arm.

2.2 A formal solution (using the logical apparatus presented above)

To give the formal solution using action models I will simply follow the solution in [8], [6] [2] and [7], with the exception that the public announcements will be modeled as actions with preconditions. First, let's fix our formal language with a few useful formulas and their abbreviations:

Our set of agents will be $D = \{a, b, c\}$, each of them corresponding to our Ann, Bob and Charlie.

ma := a is muddy.

mb := b is muddy.

mc := c is muddy.

$\neg(K_a ma \vee K_a \neg ma)$:= a does not know whether she is muddy or not.

$father := \bigvee\{mi \mid i \in \{a,b,c\}\} = ma \vee mb \vee mc$

Meaning: At least one of them (a, b and c) is muddy .

$ann2 := \bigvee\{K_i mi \vee K_i \neg mi \mid i \in \{a,b,c\}\}$

Meaning: At least one of our agents knows its state.

$\neg ann2 := \neg(\bigvee\{K_i mi \vee K_i \neg mi \mid i \in \{a,b,c\}\}) = \neg((K_a ma \vee K_a \neg ma) \vee (K_b mb \vee K_b \neg mb) \vee (K_c mc \vee K_c \neg mc))$.

Meaning: Basically, it's the negation of $ann2$ and it says that none of them knows its state.

$ann3 := \bigwedge\{K_i mi \vee K_i \neg mi \mid i \in \{a,b\}\} = (K_a ma \vee K_a \neg ma) \wedge (K_b mb \vee K_b \neg mb)$

Meaning: a and b whether they're muddy or not .

$everyone := \bigwedge\{K_i mi \vee K_i \neg m_i \mid i \in \{a,b,c\}\} = (K_a ma \vee K_a \neg ma) \wedge (K_b mb \vee K_b \neg mb) \wedge (K_c mc \vee K_c \neg mc)$

Meaning: a, b and c know their states.

How can we model the uncertainty of our agents using epistemic models ? The model that represents the knowledge and uncertainty of our agents is the epistemic cube (named *Cube* hereafter) in the figure below. Of course, the model is not arbitrarily constructed: we have 3 agents, each of them being either muddy or clean (1 or 0, correspondingly), so we have 2^3 possibilities, so our model has to have 8 worlds. For convenience, the name of each node will represent the state of our agents at that world: for example, if the first agent is muddy but the rest are clean, the world will be named 100. Evidently, at that node (and only at that one) the formula $ma \wedge \neg mb \wedge mc$ will be satisfied. Generally, the name of each node will be: $\overline{xyz}$, such that $x, y, z \in \{0, 1\}$.

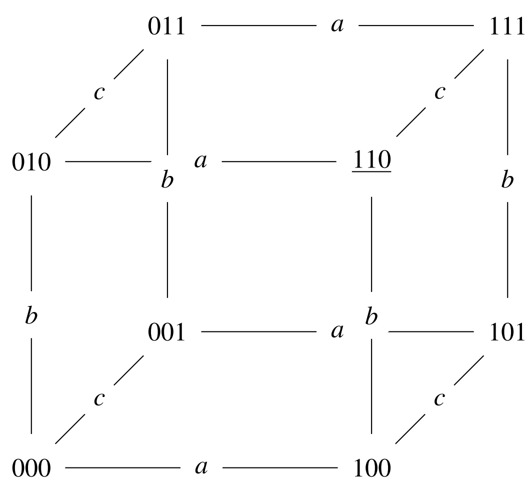

One should know that in the following we will employ a pragmatic view on what an announcement is to count. What *is* announced is the information that flows within our system, between our agents and not what their father literally announces. So how can we model the first announcement ? Using the logical apparatus presented above, the model after an announcement is the restricted modal product of an epistemic model and an action model. In this case, the action is the execution of an announcement. Therefore, the domain of this action model (let's name it Pub1) will be a singleton model, containing a single action point: pub1: $\mathsf{Pub1} = (\{\mathsf{pub1}\}, \longrightarrow, \mathsf{pre})$. The action pub1 has as precondition the formula $father$:

$$\mathsf{pre}(\mathsf{pub1}) = anunt1$$

The result of the product $Cube \otimes \mathsf{Pub1}$ will be a model model $Cube'$ whose all nodes satisfy $father$. Below, there is a graphic representation of the execution of Pub1 in $Cube$. Following step by step the construction of $Cube'$, it's none other than our Cub of which we eliminated the node at which it was false that $father$, 000, and the accessibility relations linking any other node to 000.

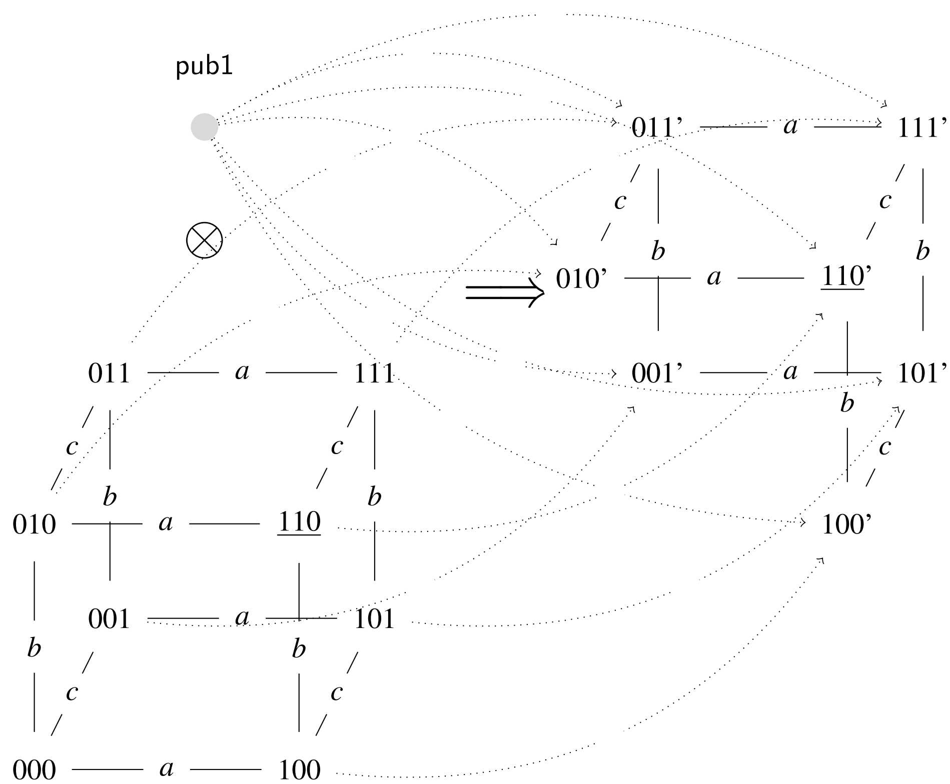

Figure represents the execution of action $\mathsf{pub1}$ in the epistemic model *Cube*. The result, $Cube'$, is composed of the following nodes: $011' = (011, \mathsf{pub1})$, $111' = (111, \mathsf{pub1})$, $010' = (010, \mathsf{pub1})$, $\underline{110'} = \underline{(110, \mathsf{pub1})}$, $001' = (001, \mathsf{pub1})$, $101' = (101, \mathsf{pub1})$, $100' = (100, \mathsf{pub1})$.

Let's write this in the language of Action Model Logic:

(**1**) $Cube \otimes \mathsf{Pub1}, (\underline{110, \mathsf{pub1}}) \models C_{\{a,b,c\}} father$

Meaning that in the resulting model it is common knowledge that at least one of the children is muddy. Equivalently:

(**1'**) $Cube, \underline{110} \models [\mathsf{Pub1}, \mathsf{pub1}] C_{\{a,b,c\}} father$

In the initial model, *Cube*, it is true that after the execution of (the public announcement action) $(\mathsf{Pub}, \mathsf{pub})$ it becomes common knowledge that $father$.

Now, let's consider father's first "raise your hands if you know"-command. No one knows her own state but everyone knows the state of anyone else. In the actual state of affairs, $\underline{110}$, every agent is uncertain: a, who is *de facto* muddy, can access a state at which he is clean: 010, b, who is muddy, considers possible a state at which she is clean: 100, and c, the clean one, considers that she may be muddy: 111. Since no child raises an arm, we can consider that what is announced in the group (the information that flows inside the group) is that no one knows whether she is muddy or not, which is exactly formula $\neg ann2$. To give a logical representation of this situation, let's assume a new action, $\mathsf{Pub2} = (\{\mathsf{pub2}\}, \longrightarrow, \mathsf{pre})$, such that $\mathsf{pre}(\mathsf{pub2}) = \neg ann2$. The execution of this public announcement transforms our $Cube'$ model to the restricted modal product, $Cube''$, in which there is no world at which everyone is uncertain, eliminating all the states at which only one of our agents is muddy:100, 010, 001. Observe below $Cueb'' = Cube' \otimes \mathsf{Pub2}$:

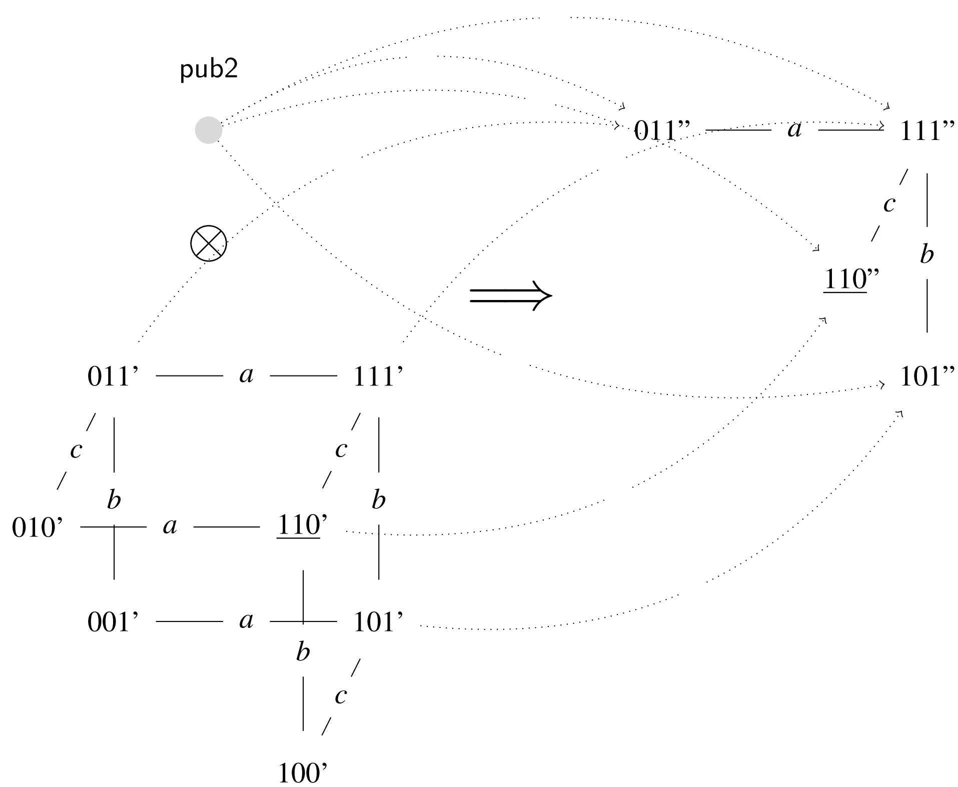

Figure presents the execution of action pub2 in *Cube*$'$. The result of composing the two models: *Cube*$'$ and *Pub*2 is *Cube*$''$, with nodes: $011'' = (011', \mathsf{pub2}) = ((011, \mathsf{pub1}), \mathsf{pub2})$, $111'' = (111', \mathsf{pub2}) = ((111, \mathsf{pub1}), \mathsf{pub2})$, $\underline{110''} = \underline{(110', \mathsf{pub2})} = \underline{((110, \mathsf{pub1}), \mathsf{pub2})}$, $011'' = (101', \mathsf{pub2}) = ((101, \mathsf{pub1}), \mathsf{pub2})$.

This time, let’s note that we have the following:

(**2**) $(Cube \otimes \mathsf{Pub1}) \otimes \mathsf{Pub2}, (\underline{110, \mathsf{pub1}, \mathsf{pub2}}) \models ann2$

(**2**) says that in the model resulting after the product of *Cube* with $(\mathsf{Pub1}, \mathsf{pub1})$ and, succesively, $(\mathsf{Pub2}, \mathsf{pub2})$ it is true that at least one child knows whether she is muddy or not (*ann*2). Equivalently:

(**2'**) $Cube, \underline{110} \models [\mathsf{Pub1}, \mathsf{pub1}][\mathsf{Pub2}, \mathsf{pub2}]ann2$

Formula (**2'**) expresses the fact that in the initial model, after the execution of the first two announcements, it becomes true that $ann2$.

The last announcement of their father coincides with the following: in $Cube''$, at $\underline{110}$, we have that a and b know their status: $Cube'', \underline{110} \models K_a ma \wedge K_b mb$. Because they know their status, they will raise their arms. This equivalates to performing a public announcement of formula $ann3$ by executing an action $\mathsf{Pub3} = (\{\mathsf{pub3}\}, \longrightarrow, \mathsf{pre})$, where $\mathsf{pre}(\mathsf{pub3}) = ann3$. Below, check the graphic representation of executing $\mathsf{Pub3}$ in $Cube''$ and obtaining $Cube''' = Cube'' \otimes \mathsf{Pub3}$.

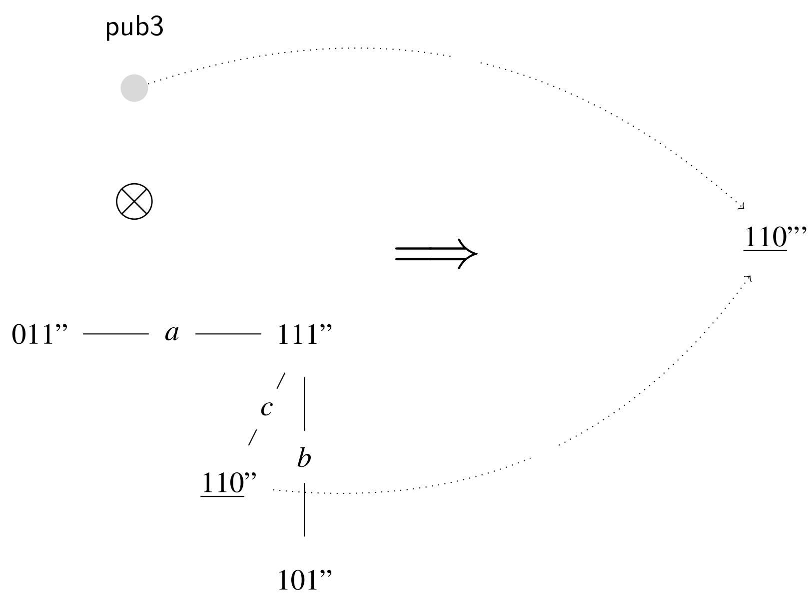

Figure prezintă rezultatul compunerii modelelor $Cube''$ şi $\mathsf{pub3}$. Modelul rezultat, $Cube'''$ conţine doar starea: $\underline{110'''} = \underline{(110'', \mathsf{pub3})} = \underline{((110', \mathsf{pub2}), \mathsf{pub3})} = \underline{(((110, \mathsf{pub1}), \mathsf{pub2}), \mathsf{pub3})}$.

Because the resulting model has a singleton as domain, all agents have certainty, so *c* also knows her status and raises her arm. Finally, they all know whether they're muddy or not:

(**3**) $(((Cube \otimes \mathsf{Pub1}) \otimes \mathsf{Pub2}) \otimes \mathsf{Pub3}), (\underline{110, \mathsf{pub1}, \mathsf{pub2}, \mathsf{pub3}}) \models everyone$

Formula (**3**) expresses that in the model resulting after three succesive products, it is true that *everyone*.

(**3'**) $Cube, \underline{110} \models [\mathsf{Pub1}, \mathsf{pub1}][\mathsf{Pub2}, \mathsf{pub2}][\mathsf{Pub3}, \mathsf{pub3}]everyone$

The same, formula (**3'**) says that in the model *Cube* it becomes true that all children know their status after the three public announcements.

References

[1] A. Baltag, H. P. Van Ditmarsch, and L. H. Moss. *Epistemic logic and information update*, pages 361–455. Elsevier, 2008.

[2] Solecki, S., Baltag, A., Moss, L.S.,. The Logic of Public Announcements, Common Knowledge and Private Suspicions. Technical report, Indiana University, 1999.

[3] de Rijke, M., Blackburn, P., Venema, Y.,. *Modal Logic*. Cambridge University Press, 2002.

[4] G.R. Renardel de Lavalette and H.P. van Ditmarsch. Epistemic actions and minimal models, 2002.

[5] M. Dumitru. *Modalitate si incompletitudine*. Paideia, Bucharest, 2001.

[6] Groeneveld, W. Gerbrandy, J.,. Reasoning about Information Change. *Journal of Logic, Language and Information, Volume 6, Number 2,*, 6(2):147–169, 1997.

[7] van Ditmarsch, H.P., Kooi, B. The Secret of My Success. *Synthese*, 151(2), 2006.

[8] van Ditmarsch, H.P., Kooi, K., van der Hoek, W. *Dynamic Epistemic Logic*. Springer, 2006.

[9] Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M. *Reasoning about Knowledge*. The MIT Press, Cambridge, Massachusetts, 1995.