Interpreting Product Update as Reasoning about Observations and Meta-Observations

Jan Broersen[1]

Abstract. In this brief note, I would like to suggest that it makes sense to reinterpret product update, as introduced by Baltag, Moss and Solecki, as a system to account for observations and meta-observations, where a meta-observation is an observation of an observation. Under this interpretation we also take products of action models with meta-action models. I deliberate on some possible consequences of this extension to the interpretation of product update.

1 Introduction

Product update, as defined by Baltag, Moss and Solecki [1, 2], is about updating multi-agent epistemic models by modeling the assimilation of new information as a (restricted) modal product with a multi-agent epistemic action model. This paper reports two observations[2] concerning this framework. The first observation is that the mechanism defined by taking products only fits with an interpretation of the actions in action models as *observations*. The second, related observation is that action models themselves might be viewed as resulting from products of meta-action models representing meta-observations. These (admittedly preliminary) ideas may give rise to new action languages for talking about epistemic action models.

2 Product updates model observations

That the possible worlds resulting from a product update are a Cartesian product of the original worlds and the actions, is intuitively explained by the observation that in principle any of the epistemic actions in the action model can be performed from any state in the static epistemic model. We get a *restricted* product by considering the preconditions of the actions that prevent some actions to be performed from certain states. For the uncertainty relation in the product models the intersection of the uncertainty relations of the epistemic model and the action model is taken.

Surprisingly, in the literature not much effort is spend on explaining why it is that we have to take the intersection of the uncertainty relations originating from the static epistemic model and the epistemic action model. Baltag and Moss [2], in their most recent account of product update, say the following:

> "We model the update of a state by an action as a partial update operation, given by a restricted product of the two structures: the uncertainties present in the given state and the given action are multiplied, while the impossible combinations of states and actions are eliminated (by testing the actions preconditions on the state). The underlying intuition is that, since the agents uncertainties concerning the state and the ones concerning the simple action are mutually independent, the two uncertainties must be multiplied, except that we insist on applying an action only to those states which satisfy its precondition."

The quote explains that the intersection reflects multiplication of *independent* uncertainties. But the quote does not explain why we can assume this independency, nor does it explain what kind of actions actually ensure independency under the constraints imposed by the preconditions.

I will approach the question concerning the reason for taking an intersection from a slightly different angle. Prima facie, one might consider taking an intersection surprising: if an agent performs the same action from states he cannot distinguish, it will also not be able to distinguish the result states. And if an agent does not distinguish two actions from a state it does distinguish, again two or more indistinguishable effect states will result. This would then rather suggest that we should take the *union* instead of the intersection. So why is it that the intersection is the right choice? The rough answer is: because the actions of action models are 'testing' or 'observation' actions. Such actions always aim at *reducing* uncertainty. Furthermore, what these actions observe, must be true in the state where their precondition holds. So there is just exactly only *one* way in which observation actions can result in uncertainty: from an uncertain state it must be uncertain whether the observation has taken place. That explains the intersection.

This also sheds light on the question above, concerning the reason for the independence of the uncertainties involved. The independence is explained by the reasonable assumption that observations themselves do not interact with the conditions observed[3].

The term 'observation' should not be taken too literally here. Actually, from the level of abstraction we are looking at information transfer, 'observation', 'testing', 'learning' and 'announcing' are all alike. The difference between these concepts can only become more clear if we can differentiate between sending and receiving agents, their motivations for transferring information, and their strategies for dealing with new information. The present setting, without notions like 'agency' or 'intention' is simply too abstract for that.

3 Product update and meta-observation

Many standard examples are explicitly about observations. A well-known one is the following [3, p.130]:

Example 1 *In a situation with three colored cards and three agents*

---

[1] University of Utrecht, The Netherlands, email: broersen@cs.uu.nl
[2] As far as I know, these observations have not been reported before.

[3] Note that this assumption conflicts with the uncertainty principle from quantum physics.

knowing their own card, agent 1 publicly observes a card of agent 3 (for instance, because 3 shows it to him).

The action model distinguishes three different observations: '1 observes 3 has red' ($1 : r@3$), '1 observes 3 has white' ($1 : w@3$) and '1 observes 3 has blue' ($1 : b@3$). Agent 1 and 3 distinguish these actions, agent 2 does not. Below are the pictures for the static initial epistemic model, the action model, and the product model.

```
      rwb — 2 — bwr
      /             \
    1                 3
   /                   \                    1:b@3
 rbw ———————— 2 ———— wbr       ⊗          /    \
   \         1  3       /                 2      2
    3       / \        1                 /        \
      \    /   \      /              1:r@3 - 2 - 1:w@3
      brw — 2 — wrb
```

```
              rwb — 2 — bwr
                          \
                           3
                            \
 :=    rbw ———————— 2 ——— wbr
          \              /
           3            3
            \            \
             brw — 2 — wrb
```

Figure 1. Agent 1 publicly observes agent 3's card

Now note that the description of example 1 also says that the observation itself is publicly observed. This is the first sign that something like 'meta'-observations are relevant for the example. In the following we will suggest that these meta-observations can be modeled as action models themselves. We will refer to these models as 'meta-action models'.

But now let us first extend the above example in order to make more clear what we mean.

Example 2 *Agent 3 has left the table, leaving his card on the table. After coming back he suspects that 1 has taken a look at his card, which, in fact is indeed what happened, and it happened publicly. Agent 3 publicly announces his suspicion.*

Figure 2 gives the right product model resulting from taking the product with the appropriate action model for this example. The model contains both the epistemic model of the initial situation and the epistemic model resulting from the previous example, and agent 3 hesitates between these two models. But what is the action model that produces this product model? Of course, it is not too difficult to find the right action model. However, below we show we can decompose this problem into two problems: finding the appropriate action model and finding the appropriate meta-action model.

It is rather clear that in this example there are at least two levels of observation. First there is the level where individual agents get to know the results of the card showing action. This level is exactly
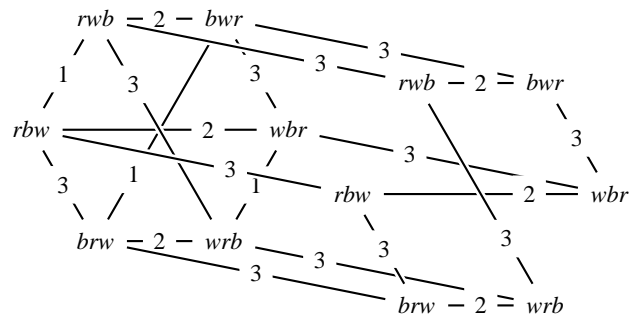
```
      rwb — 2 — bwr
      /    \          _____ 3 — 3
    1       3      3          rwb — 2 = bwr
   /        \     /                       \
 rbw ———————— 2 — wbr              3         3
   \        1   3    _____                 \
    3      /    1     rbw ———————— 2 = wbr
     \    /    /
      brw = 2 — wrb ___ 3 — 3
                  \____ 3          3
                       brw — 2 = wrb
```

Figure 2. Agent 3 suspects agent 1 has seen his card

the same as in the first example. Therefore, in figure 3, that gives the action models for the observation and the meta-observation levels, the observation level action model is the same as the one in figure 1 for the first example (with the exception of the non-connected 'skip' world, which we discuss later). Indeed, we might say that the only difference between the two examples is on the level of meta-observations: in the first the meta-observation is like a public announcement, and in the second the meta-observation action model contains at least two meta-observation actions '3 observes that 1 takes a look at 3's cards' and '3 observes that 1 observes nothing'. Agent 3 does not distinguish between these actions (because he does not know whether the looking action actually took place). But agents 1 and 2 can distinguish between the two. Note that this meta-level action model models the suspicion of agent 3 as the hesitation about whether or not some observation has taken place on the observation level. Maybe the hesitation and suspicion originates from agent 3 not being sure whether or not he saw that agent 1 was taking a look at his card. Also note that the meta-level contains a third action: the meta-action of agent 3 not meta-observing anything at all ('skip'). To make the view using meta-levels of observation work, for all agents in the system we have to add such 'non-observation/skip' actions at any level and meta-level. Note that in the meta-level action model of figure 3, I only give the non-action for agent 3. Actually, to prevent the picture from being flooded by states that are inessential for explaining the idea I do not give any of the meta-level observation actions of agents other than agent 3. This is why the figure says 'etc.' in the meta-level action model. In particular, as long as at the direct meta-level there is no uncertainty about inaction, the non-actions can be neglected. For instance, note that using the observation action model of figure 3 in stead of the one in figure 1 to solve the first example, does not make a difference. In particular, if we stick to the original set-up, with only one level of action models, adding a non-observation action to the action model does not make a difference as long as there are no uncertainty relations linking it to other actions. Finally, note that the product model resulting from the product of the observation model and the meta-observation model, when multiplied with the static epistemic model, yields the product model of figure 2, as required.

Now, what are the preconditions for the actions in this meta-observation action model? And how exactly do we define the product of the meta-observation model and the observation model? The preconditions for the meta-observations are straightforward. Just like for the observation action models they just say that we can only observe what is actually true. For the action 'observing nothing' this

```
    1:b@3                3:(1-obs)
   /   \                    |
  2     2                   3
 /       \                  |
1:r@3 - 2 - 1:w@3        3:(1-skip)


                  ⊗


1-skip      2-skip            3-skip



    3-skip                    etc.



              1:b@3
             / | \
            2  |  2
           /   |   \
:=     1:r@3 - 2 - 1:w@3
            \  ¦  /
           3  3  3
            \ | /
           1-skip
```

Figure 3.   The meta-observation as a meta-product

means we get ⊤ as a precondition, because it is an action that can always be done. And as expected, since this is all about observation, for the product of an action model and a meta-action model we also take the intersection of uncertainty relations. Finally, we get the 'restricted' product by checking the preconditions of the meta-action model on the action model. In the specific example we deal with here, we have three meta-observations. The meta-observations 3:(1−obs), 3:(1−skip) and 3−skip. The first meta-action has a precondition that is true on all actions of the observation model where agent 1 observes something. The second has a precondition that is true only on the action in the action model where agent 1 observes nothing. Finally, the precondition of 3−skip is true on all actions in the actions model. But this part of the meta-product only yields an unconnected copy of the action model.

In example 2 above, there is also a third meta-level of observation. Because of the involvement of this level, we say that agent 3 announces its suspicion. This ensures that the third meta-level is a public announcement. Actually, for any example we come up with, in principle infinitely many layers of observation play a role. Any observation is itself an action for which we can ask whether it is observed. So, for any observation potentially infinitely many actions are performed at the same time: every observation possibly has a meta-observation. In the examples above, this infinite hierarchy is broken down by the first meta-level of observation that is a public announcement. Actually, for any well-specified example, the highest relevant meta-level of observation is always a public announcement. If not, the example is under-specified, and leaves room for a non-deterministic interpretation of the update. Actually, in most examples in the literature, implicitly a closure assumption is applied: if nothing is said about the meta-levels of observation, it is assumed that they are public observations closed under meta-observations.

4   Future Research

The setting raises several questions. I briefly mention a few of them. The first is that products and meta-products are *not associative*. This is quite easy to see from the example above. The meta-products should always be taken first. In particular, if we first take the product of the static model and the observation level action model, it is not clear how to take the product of the resulting product model with the meta-level action model. Performing products in this order is not even well-defined. But also it is clear that we throw away information by first taking the product with the action model instead of the product with the meta-action model. We cannot recover this information. A possible problem is that associativity might be important for certain examples. For instance, what if only at a later moment an agent learns that his looking at the cards was observed or suspected by another agent. Since we do note have associativity, this can only be modeled by keeping track of all action and meta-action models over time.

Another interesting question is how we can add *agency* to the picture. Actually, viewing the actions as observations of specific agents, as we did in the above examples, is a first step in this direction. For every observation action it is important to describe whose observation it is, since on the meta-observation level this information is used. And for each observation action it is important to describe whose action is observed. Above we used an ad-hoc notation to describe actions. An obvious route for investigation is to try to turn this notation into an action language for observations and meta-observations.

One of the principles that suggest themselves is that agents always observe their own observations. Note that the setting we sketched actually does enable us to model situations where this is not the case. However, methodologically it would be strange to allow this. We do not want to get into the territory where agents make 'sub-conscious' observations.

The current set-up also enables us to speculate about a view where *all* knowledge is viewed as observation[4], even at the 'lowest' level. It is not too difficult to translate a standard static epistemic model into an action model containing the same information. This is accomplished by also seeing the knowledge of each agent at the lowest level as an observation. For instance, in the cards example, the static model is equivalent to the action model where each agent observes his own card.

I want to finish with a comment. It seems not right to claim that the setting we sketched adds nothing new only *because* the meta-products will simply return an action model of the form already available in the original setting of Baltag, Moss and Solecki. If that would be a good argument against the present proposal, the original proposal could be attacked with the same argument: product update adds nothing new, because it can be seen as a system that only specifies a complicated way to describe standard epistemic models.

REFERENCES

[1] L. S. Moss A. Baltag and S. Solecki, 'The logic of public announcements and common knowledge', in *Proceedings of TARK 1998*, (1998).
[2] A. Baltag and L. S. Moss, 'Logics for epistemic programs', *Synthese*, **139**, 165–224, (2004).
[3] Hans van Ditmarsch, Wiebe van der Hoek, and Barteld Kooi, *Dynamic Epistemic Logic*, volume 337 of *Synthese Library*, Springer, New York, NY, USA, 2007.

---

[4] Like in logical positivism.