Update Semantics of Security Protocols

Arjen Hommersom (arjenh@cs.kun.nl)
*Nijmegen Institute for Computing and Information Sciences,
University of Nijmegen*

John-Jules Meyer (jj@cs.uu.nl)
*Institute of Information and Computing Sciences,
University of Utrecht*

Erik de Vink (evink@win.tue.nl)
*Department of Mathematics and Computer Science,
Technische Universiteit Eindhoven
Leiden Institute of Advanced Computer Science,
Leiden University*

Abstract. We present a model-theoretic approach for reasoning about security protocols, applying recent insights from dynamic epistemic logics. This enables us to describe exactly the subsequent epistemic states of the agents participating in the protocol, using Kripke models and transitions between these based on updates of the agents' beliefs associated with steps in the protocol. As a case study we will consider the SRA Three Pass protocol and discuss the Wide-Mouthed Frog protocol.

1. Introduction

In today's world of e-commerce and the Internet, the role of security protocols is getting increasingly important. The design of these security protocols is difficult and error-prone (Lowe, 1996; Schneier, 2000; Anderson, 2001), which makes (automatic) verification of protocols of crucial importance. Since the late eighties, one line of research, amongst others, for reasoning about security protocols is based on the use of the so-called BAN logic, proposed by Burrows, Abadi and Needham in (Burrows et al., 1990). This is an epistemic logic augmented by constructs that are relevant for reasoning about security, such as the property of having the disposal of a cryptographic key to be able to decode a message and therefore to know its contents. Although many useful results have been reported (e.g., (Kessler and Neumann, 1998; Agray et al., 2001; Stubblebine and Wright, 2002)), due to their complexity and their semantic underpinning the use of BAN logics to prove the correctness of security protocols has so far been of limited success (cf. (Abadi and Tuttle, 1991; Bleeker and Meertens, 1997; Wedel and Kessler, 1996)).

In this paper we will apply insights from dynamic epistemic logics as recently developed by Gerbrandy (Gerbrandy, 1997; Gerbrandy, 1999), Baltag and Moss (Baltag et al., 1998; Baltag, 2002; Baltag and Moss, 2004), Van Ditmarsch (van Ditmarsch, 2000; van Ditmarsch, 2001), and Kooi (Kooi, 2003). Moreover, contrary to the traditional BAN logic approach, our approach is semantic or model-theoretic. We use Kripke models to represent the epistemic state of the agents involved in a protocol, similarly to the S5 preserving approach of Van Ditmarsch to analyze certain kinds of games involving knowledge. From the action models of Baltag and Moss we import the idea to describe belief updates of the agents by semantic operators transforming the Kripke models at hand by copying and deleting parts of these models, although we use traditional Kripke models rather than action models. To this end, we need also operations for unfolding models, which is in its turn inspired by Gerbrandy's work on possibilities. The difference being that in our approach only *partial* unfolding is called for. We furthermore propose a language to express belief updates in the context of security protocols as well as properties of these updates, and give a semantics of this language in terms of the models mentioned and the operators on them. Since our approach is model-theoretic, we believe that it may serve as a starting point for the automatic verification of (properties of) security protocols.

As a case study illustrating our approach we will consider the so-called SRA Three Pass protocol. It is not our intention to prove that the protocol is completely secure (as it is not in full generality), but we will prove that if the agents participating in the protocol are honest, then an intruder watching the communication does not learn anything about the plain-text messages in a single run. Furthermore, we show what the intruder is able to learn about the agents participating. We also discuss the Wide-Mouthed Frog protocol to illustrate the operation developed in the sequel for updating the beliefs of agents.

Acknowledgements

We are grateful to the anonymous reviewers for their constructive comments and to Simona Orzan for various discussions on the subject. We are also indebted to Peter Lucas for his support to this research.

2. Preliminaries

In this section we briefly discuss some preliminaries and background regarding the semantic updates we will handle and the epistemic models we will use. First, we define the notion of an objective formula and introduce so-called o-seriality. The set of propositional variables in a model is denoted as $\mathcal{P}$.

Definition 2.1. The class of objective formulas is the smallest class such that

- all propositional variables and atoms $p \in \mathcal{P}$ are objective;
- if ϕ is objective, then $\neg\phi$ is objective;
- if ϕ_1 and ϕ_2 are objective, then $\phi_1 \wedge \phi_2$ is objective.

So, objective formulas do not involve beliefs. For our purposes it is important that every agent, at every world, distinguishes a world with the same 'objective' information. This leads to the notion of an o-serial model. The operations on Kripke structures discussed in the sequel degenerate for models that are not o-serial.

Definition 2.2. A model $M = \langle S, \pi, R_1, ..., R_m \rangle$ is *o-serial* iff for all agents i and $w \in S$, there exists $v \in S$ such that $R_i(w, v)$ and for all objective formulas ϕ it holds that $(M, w) \models \phi \Leftrightarrow (M, v) \models \phi$.

We use a, b, c, etc. and i, j as typical agents, taken from a class $\mathcal{A}$. Furthermore, B is used as a doxastic modal operator. For example, $B_a\phi$ should be read as 'a believes ϕ'. We interpret formulas on standard Kripke models $(M, s) = (\langle S, \pi, R_1, ..., R_m \rangle, s)$, where $(M, s) \models B_i\phi$ iff $\forall t \in S\colon R_i(s, t) \rightarrow (M, t) \models \phi$.

We require the relations R_i to be o-serial, transitive and euclidean. This yields a class of models that we will call Kt45, a proper subset of the class of models of the well-known doxastic logic *KD45*. The lower case t refers to the axiom

$$B_i\phi \Rightarrow \phi, \tag{t}$$

where the formula ϕ ranges over objective formulas. The system Kt45 is sound with respect to the class of o-serial, transitive and euclidean models (Hommersom, 2003). (We conjecture that Kt45 is complete as well for this class.) We will show that the operations we introduce preserve Kt45. The point is that in worlds of Kt45 models, we cannot both have $B_i\phi$ and $\neg\phi$, for an objective formula ϕ. This is reasonable from the assumption that agents are conscious about the protocol. Therefore, they will not infer objective falsehoods. This objectivity is captured locally for each state. As a consequence, the operations that we introduce can restrict the set of states without destroying objective information.

For the analysis of security protocols below, we assume that we are omniscient about the values of the variables in different runs of a protocol. For example, the program variable p in a protocol run has the value $[\![p]\!]$. In the real world it is, obviously, always true that $p = [\![p]\!]$. However, it is cumbersome to keep track of what is the real world in the operations on Kripke structures that we employ below. Therefore, we assume that an interpretation $[\![\cdot]\!]$ is given, that provides the 'real' values (not necessarily boolean) of the program variables when needed. It might very well be the case that $p \neq [\![p]\!]$ in a certain state. Often, we will abbreviate $p = [\![p]\!]$ to p on (thus transforming a program expression into a propositional variable). Similarly, $\neg p$ is an abbreviation of $p \neq [\![p]\!]$. For example, agent a that learns $B_b p \vee B_b \neg p$, learns that agent b has assigned a value to the program variable p.

The types of updates we consider are

- public announcement of a variable,
- private learning of a variable, and
- private learning about the knowledge of other agents.

The first type of update typically runs as follows: In an open network, agent a sends a message to agent b. From a security perspective, it is customary in the so-called Dolev-Yao framework (Dolev and Yao, 1983), to assume that all agents in the network can read this message too. In contrast, the second type of update, describes private learning. For example, agent b receives a message $\{x\}_k$ from agent a. (Here, we use the notation $\{x\}_k$ to denote a message x encrypted with the cryptographic key k.) If b possesses the key k, then b privately learns the message content x. The final type of update is probably the most interesting. It is realistic to assume that the steps in a protocol run are known to all agents. Therefore, observing that an agent receives a message will increase the knowledge of the other agents. For example, if agent a sends a message $\{x\}_k$ to agent b, then agent c learns that b has learned the information contained in the message $\{x\}_k$, but typically, c does not learn x if c does not possess the key k.

Stronger types of updates we do not consider here. For example, we will not update the beliefs of an honest agent such that it learns that an intruder has learned about others. In the present paper, we restrict ourselves to beliefs about objective formulas and the updating of such beliefs.

3. Update constructions

In this section we describe various types of updates in detail. We will start by defining an update for propositions in Subsection 3.1. In Subsection 3.2 we will define a belief update for agents that learn something about the belief of others. We do this in two slightly different ways by varying in the operations that describe a side-effect for an agent.

3.1. Objective updates

The belief update of objective formulas we will use is based on the work reported in (Roorda et al., 2002; Baltag et al., 1998). The construction works as follows: We will make copies of the states of the model such that the *old* worlds correspond to the information in the original model and the *new* worlds correspond to the new information.

Definition 3.1. Let a model $(M, w) = (\langle S, \pi, R_1, ..., R_m \rangle, w)$, a group of agents $\mathcal{B}$, and an objective formula ϕ be given. Then the model $\textsc{update}_{(\phi,\mathcal{B})}(M, w)$, the update of (M, w) for agents in $\mathcal{B}$ and formula ϕ, is given by

$$\textsc{update}_{(\phi,\mathcal{B})}(M, w) = (\langle S', \pi', R'_1, ..., R'_m \rangle, w')$$

where

- $S' = \{\ old(s),\ new(s) \mid s \in S\ \}$
- $w' = new(w)$
- for all $p \in \mathcal{P}$: $\pi'(old(u))(p) = \pi'(new(u))(p) = \pi(u)(p)$
- for $a \in \mathcal{A}$, the binary relation R'_a on S' is minimal such that

$$\begin{array}{lcll} R'_a(old(u), old(v)) & \Leftrightarrow & R_a(u, v) & \\ R'_a(new(u), new(v)) & \Leftrightarrow & R_a(u, v) \wedge (M, v) \models \phi & \text{if } a \in \mathcal{B} \\ R'_a(new(u), old(v)) & \Leftrightarrow & R_a(u, v) & \text{if } a \notin \mathcal{B} \end{array}$$

In order to distinguish the two copies of the states, the tagging function *old* and *new* are used. In the new part of the model, agents in $\mathcal{B}$ will only consider possible worlds that verify ϕ. Therefore, states can become unreachable from the actual world $new(w)$, and can be dropped.

The following example shows how this works on a concrete model.

Example 3.1. (updating) Consider the model (M, s) in Figure 1.a, where we have $\pi(s)(p) = \texttt{true}$ and $\pi(t)(p) = \texttt{false}$. The operation we execute is $\textsc{update}_{(p,\{b\})}$, i.e. b learns p. This results in the model (M, u) in Figure 1.b, where $\pi(u)(p) = \pi(v)(p) = \texttt{true}$ and $\pi(w)(p) = \texttt{false}$ and $\mathit{new}(s) = u, \mathit{old}(s) = v$ and $\mathit{old}(t) = w$. The world $\mathit{new}(t)$ is unreachable and is omitted.

We can see that the belief of agent a has not changed: it still considers its old worlds possible. The belief of agent b, however, has changed. It now only considers the state u possible where p holds. Note that agent b is aware that agent a does not know about p.

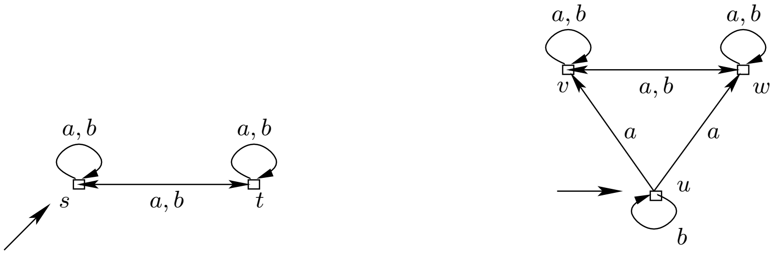

a. (M, s) b. (M, u)

Figure 1.

The update operation $\textsc{update}_{(\phi,\mathcal{B})}$ is based on a formula ϕ and a set of agents $\mathcal{B}$. Roorda *et al.* (Roorda et al., 2002) propose a characterization of the formulas that are altered by such an operation with a single learning agent. Here, we extend their definition for multi-agent purposes.

Definition 3.2. An update function $(\cdot)[\phi, \mathcal{B}]$ is called proper if

$$\begin{array}{lll}
(M, w)[\phi, \mathcal{B}] \models p & \Leftrightarrow & (M, w) \models p \\
(M, w)[\phi, \mathcal{B}] \models \alpha \wedge \beta & \Leftrightarrow & (M, w)[\phi, \mathcal{B}] \models \alpha \text{ and } (M, w)[\phi, \mathcal{B}] \models \beta \\
(M, w)[\phi, \mathcal{B}] \models \neg\alpha & \Leftrightarrow & (M, w)[\phi, \mathcal{B}] \not\models \alpha \\
(M, w)[\phi, \mathcal{B}] \models B_a\alpha & \Leftrightarrow & (M, w) \models B_a\alpha \text{ if } a \notin \mathcal{B} \\
(M, w)[\phi, \mathcal{B}] \models B_a\alpha & \Leftrightarrow & \forall v\colon ((R_a(w, v) \text{ and } (M, v) \models \phi) \Rightarrow \\
 & & \qquad (M, v)[\phi, \mathcal{B}] \models \alpha) \text{ if } a \in \mathcal{B}
\end{array}$$

for every model M and state w.

Following (Roorda et al., 2002) we have that $\textsc{update}_{(\phi,\mathcal{B})}$ is proper (cf. (Roorda et al., 2002, Proposition 3.2)). Moreover, $\textsc{update}_{(\phi,\mathcal{B})}$ is uniquely characterized by Definition 3.2 upto elementary equivalence, i.e., if $(\cdot)[\phi, \mathcal{B}]$ is a proper update function, then $(M, w)[\phi, \mathcal{B}]$ and $\textsc{update}_{(\phi,\mathcal{B})}(M, w)$ are elementary equivalent. We collect the following properties of $\textsc{update}_{(\phi,\mathcal{B})}$.

Theorem 3.1.

(a) For any objective formula ϕ and set of agents $\mathcal{B}$, it holds that

$$\text{UPDATE}_{(\phi,\mathcal{B})}(M,w) \models B_b\phi$$

for all $b \in \mathcal{B}$.

(b) If (M,w) satisfies the Kt45 properties, the formula ϕ is objective and $(M,w) \models \phi$, then $\text{UPDATE}_{(\phi,\mathcal{B})}(M,w)$ satisfies the Kt45 properties as well.

(c) Updating is commutative, i.e. the models

$$\begin{array}{l}(M^1,w^1) = \text{UPDATE}_{(\psi,\mathcal{C})}(\text{UPDATE}_{(\phi,\mathcal{B})}(M,w)), \textit{ and}\\ (M^2,w^2) = \text{UPDATE}_{(\phi,\mathcal{B})}(\text{UPDATE}_{(\psi,\mathcal{C})}(M,w))\end{array}$$

for objective formulas ϕ,ψ and sets of agents $\mathcal{B},\mathcal{C}$, are bisimilar.

(d) Update is idempotent, i.e., the two models

$$\begin{array}{l}(M^1,w^1) = \text{UPDATE}_{(\phi,\mathcal{B})}(\text{UPDATE}_{(\phi,\mathcal{B})}(M,w)), \textit{ and}\\ (M^2,w^2) = \text{UPDATE}_{(\phi,\mathcal{B})}(M,w)\end{array}$$

for an objective formula ϕ and a set of agents $\mathcal{B}$, are bisimilar.

Proof. We prove parts (a) to (c). The proof of part (d) is similar to that of part (c). For part (a) we need to prove, that for any objective ϕ, set of agents $\mathcal{B}$ and $b \in \mathcal{B}$, it holds that $(M',w') = \text{UPDATE}_{(\phi,\mathcal{B})}(M,w) \models B_b\phi$. Take $b \in \mathcal{B}$. Since w' is $\mathit{new}(w)$, we have, by definition, for all v, if $R_b(w,v)$ then $(M',v) \models \phi$. Hence, $(M',w') \models B_b\phi$.

We prove part (b) by checking each of the properties for a Kt45 model. Assume that $(M,w) = \langle S, R_1, ..., R_m, \pi\rangle$ is transitive, and that

$$\text{UPDATE}_{(\phi,\mathcal{B})}(M,w) = (\langle S, R'_1, ..., R'_m, \pi'\rangle, w')$$

is not. Then there is, for some agent i, $(s,t) \in R'_i$, $(t,u) \in R'_i$ and $(s,u) \notin R'_i$. Because, by definition, there are no arrows from *old* to *new* states and no i has both a relation from *new* to *new* states and from *new* to *old* states, there are only three cases. Firstly, suppose $\mathit{old}(s)$, $\mathit{old}(t)$ and $\mathit{old}(u)$. Then, by definition, we have $(s,t) \in R_i$, $(t,u) \in R_i$ and by transitivity $(s,u) \in R_i$. Then $(\mathit{old}(s), \mathit{old}(u)) \in R'_i$. Contradiction. Now suppose $\mathit{new}(s)$, $\mathit{old}(t)$ and $\mathit{old}(u)$. Then $(s,t) \in R_i$, $(t,u) \in R_i$ and $i \notin \mathcal{B}$. Again, this implies $(\mathit{new}(s), \mathit{old}(u)) \in R'_i$, which is a contradiction. Finally, suppose $\mathit{new}(s), \mathit{new}(t), \mathit{new}(u)$, then $(s,t) \in R_i$, $(t,u) \in R_i$ and $(M,u) \models \phi$. Thus, $(\mathit{new}(s), \mathit{new}(u)) \in R'_i$, which is again a contradiction and completes the proof for transitivity. For proving that R'_i is euclidean (for all agents i), we need to consider exactly the same cases and for all these cases a similar argument can be made. For o-seriality, the proof obligation is to show that for all $s \in S'$, there is a $t \in S'$ such that $(s,t) \in R_i \wedge (M,s) \models \psi \leftrightarrow (M,t) \models \psi$, for all objective ψ. Suppose $i \notin \mathcal{B}$, then this follows directly from the definition, since we will have some t such that $(s, \mathit{old}(t)) \in R_i$. Suppose $i \in \mathcal{B}$, then by assumption of M being o-serial, there is a t such that $(s,t) \in R_i$ which agree on the objective formulas. In particular, they agree on ϕ. Suppose $(M,s) \models \phi$, then $(M,t) \models \phi$ and therefore $(\mathit{new}(s), \mathit{new}(t)) \in R_i$. Suppose $(M,s) \not\models \phi$, then both s and t will be unreachable from w' and will thus have no corresponding state in M'. Thus, in all cases o-seriality is preserved.

For part (c), we observe that 4 copies of the original states are made. In both cases we have an original copy (call them the *old* ones), a new copy that is made in the first step is called *middle*, and the last copy that is made (so a copy from both the *old* and *middle*) is called *new*. Furthermore, the predicates have a superscript of 1 or 2 depending on the model they belong to. In addition, the elements of the model will have superscripts according to their model. We construct a bisimulation $\mathcal{R} \subseteq S^1 \times S^2$ such that it is minimal with respect to

$$\begin{array}{lcl} \mathcal{R}(\mathit{old}^1(s), \mathit{old}^2(s')) & \Leftrightarrow & s = s' \\ \mathcal{R}(\mathit{middle}^1(s), \mathit{new}^2(s')) & \Leftrightarrow & s = s' \\ \mathcal{R}(\mathit{new}^1(s), \mathit{middle}^2(s')) & \Leftrightarrow & s = s' \end{array}$$

(referring to equality in the original model).

1. $\mathcal{R}$ satisfies forward-choice: Suppose $\mathcal{R}(s,s') \wedge R^1_i(s,t), s,t \in S^1, s' \in S^2$, then there is a $t' \in S^2$ such that $\mathcal{R}(t,t'), (s',t') \in R^1$. If $\mathit{old}^1(s)$, it is trivially satisfied (take a $\mathit{old}^2(t')$, and it is satisfied, since $\mathit{old}(t)$). Say $\mathit{middle}^1(s), \mathit{new}^2(s')$. If $i \in \mathcal{C}$, then t is apparently a copy from the old^1 states to the middle^1 states. But then, it is also copied to the new^2 states and it is reachable for the same reason that is was reachable in the middle^1 state. If $i \in \mathcal{B}$, the same argument applies as for $\mathcal{C}$. If $i \notin \mathcal{C}$ and $i \notin \mathcal{B}$, then one can go back to the *old* world again as these worlds are still considered possible for these agents i.

2. $\mathcal{R}$ satisfies backward-choice: The reasoning is similar to the case of forward choice.

3. For all $s \in S$, $s' \in S'$, $(\mathcal{R}(s,s') \Rightarrow \pi(s) = \pi'(s'))$: from the definition of $\mathcal{R}$ and the definition of $\textsc{update}_{(\phi,\mathcal{B})}(\cdot,\cdot)$ this is trivial.

□

The update operation of Definition 3.1 is restricted to objective formulas. In principle, one can do the same constriction for non-objective ones. However, for a non-objective formula $B_i\phi$, it can happen that, unintendedly, an agent increases the objective knowledge encapsulated by the formula ϕ. This is illustrated by the next example.

Example 3.2. (updating of non-objective formula) Suppose we are interested in agent a learning the formula $B_bp \vee B_b\neg p$, but not the property p itself. So, agent a learns that agent b knows about p without getting information about p itself. Consider the Kripke model (M, s) in Figure 2.a, where $\pi(s)(p) = \pi(u)(p) =$ `true` and $\pi(t)(p) =$ `false`. This models the state where b knows that p is true. Agent a does not know p or $\neg p$, and it does not know if b knows p.

If we apply the definition of the update operation, it results in the model (M', v) from Figure 2.b, where $\pi'(v)(p) = \pi'(s)(p) = \pi'(u)(p) =$ `true` and $\pi'(t)(p) =$ `false`. In (M', v) it holds that B_ap, as $\pi'(v)(p) =$ `true`, since v was copied from the state s in M. Figure 2.b illustrates that a has learned $B_bp \vee B_b\neg p$, but also that a has learned p itself, which we wanted to avoid. The reason that it turns out like this, is because the only state in M where $B_bp \vee B_b\neg p$ holds, is the state s. Thus, all the other states have no corresponding *new* states.

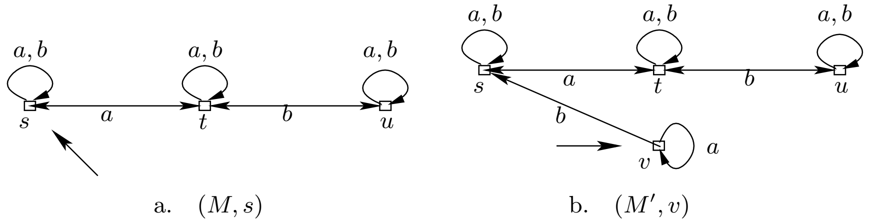

a. (M, s) b. (M', v)

Figure 2.

In the next subsection we will define a side-effect function such that a will learn about others, but does not learn any objective formulas itself.

3.2. Side-effects

The main reason that an update of $B_bp \vee B_b\neg p$ for agent a has undesired consequences, is that it actually does not include the right arrows between the copies of the original states. The construction, in the case of the non-objective formula $B_bp \vee B_b\neg p$, deletes arrows of a to gain the states that satisfy the updating formula. However, for the rest, we want a to keep all the states it considers possible. Moreover, we do not want to change the knowledge of the other agents. In this subsection we define the functions that accomplish these requirements.

A technical obstacle is that states can be shared among agents. It is obvious that if we change a state with the intention to change the belief of one agent, then the belief of the other agents that consider this state possible, is changed as well. Therefore, the first thing to do, is to separate the states of learning agents from the states of agents that do not learn. This procedure will be called *unfolding*. The tag $\mathit{new}_{\mathcal{B}}$ is a generalization of *new* and *old* from the previous section; the tag *orig* is only used for the point of the model, i.e. the actual world.

Definition 3.3. Given a model (M, w) with $M = \langle S, \pi, R_1, ..., R_m \rangle$, and a partitioning $\mathcal{X}$ of $\mathcal{A}$, we define the operation $\textsc{unfold}_{\mathcal{X}}(M, w)$, the unfolding of (M, w) with respect to $\mathcal{X}$, by $\textsc{unfold}_{\mathcal{X}}(M, w) = (\langle S', \pi', R'_1, ..., R'_m \rangle, w')$, where

- $S' = \{ \mathit{new}_{\mathcal{B}}(s) \mid s \in S,\ \mathcal{B} \in \mathcal{X} \} \cup \{\mathit{orig}(w)\}$
- $w' = \mathit{orig}(w)$
- $\pi'(\mathit{new}_{\mathcal{B}}(s))(p) = \pi(s)(p)$ and $\pi'(\mathit{orig}(w))(p) = \pi(w)(p)$ for all $s \in S$, $p \in \mathcal{P}, \mathcal{B} \in \mathcal{X}$
- for $a \in \mathcal{A}$, the binary relation R'_a on S' is minimal such that

$$\begin{array}{lcl} R'_a(\mathit{new}_{\mathcal{B}}(s), \mathit{new}_{\mathcal{B}}(t)) & \Leftrightarrow & R_a(s, t) \\ R'_a(\mathit{orig}(w), \mathit{new}_{\mathcal{B}}(s)) & \Leftrightarrow & R_a(w, s) \text{ and } a \in \mathcal{B} \end{array}$$

  where $\mathcal{B}$ ranges over $\mathcal{X}$.

So, for every group of agents $\mathcal{B}$ there is copy of the original states (viz. $\mathit{new}_{\mathcal{B}}(s)$ for every $s \in S$). The unfold operation does indeed preserve our Kt45 properties and it models the same knowledge, which is captured by the following theorem.

Theorem 3.2.

(a) If (M, w) is a Kt45 model and $\mathcal{X}$ a partitioning of $\mathcal{A}$, then it holds that $\textsc{unfold}_{\mathcal{X}}(M, w)$ is a Kt45 model too.

(b) For every model (M, w) and partitioning $\mathcal{X}$, it holds that (M, w) and $\textsc{unfold}_{\mathcal{X}}(M, w)$ are bisimilar.

Proof. Part (a) First, we prove R'_a is euclidean under the assumption that R_a is euclidean, for any $a \in \mathcal{A}$. Assume that $R'_a(s', t') \wedge R'_a(s', u'), s', t', u' \in S'$. The proof obligation is that $R'_a(t', u')$ where

s', t' and u' are either in one of the partitions or in *orig*. Suppose $s' = \mathit{new}_{\mathcal{W}}(s)$, $t' = \mathit{new}_{\mathcal{Y}}(t)$ and $u' = \mathit{new}_{\mathcal{Z}}(u)$ for $\mathcal{W}, \mathcal{Y}, \mathcal{Z} \in \mathcal{X}$. From the definition of $R'_a(s', t')$, it follows that $R_i(s, t) \wedge \mathcal{W} = \mathcal{Y}$, from $R'_a(s', u')$ follows that $R_a(s, u) \wedge \mathcal{W} = \mathcal{Z}$. Since R_a is euclidean, we have $R_a(t, u)$. From $\mathcal{W} = \mathcal{Y}$ and $\mathcal{W} = \mathcal{Z}$ we have $\mathcal{Y} = \mathcal{Z}$. Thus we conclude $R'_a(\mathit{new}_{\mathcal{Y}}(t), \mathit{new}_{\mathcal{Z}}(u)) = R'_a(t', u')$. The only other case is that $s' = \mathit{orig}(w)$, $t' = \mathit{new}_{\mathcal{W}}(t)$ and $u' = \mathit{new}_{\mathcal{W}}(u)$ for some $\mathcal{W} \in \mathcal{X}$. Then $R_a(w, u)$ and $R_a(w, u)$. Thus $R'_a(\mathit{new}_{\mathcal{W}}(t), \mathit{new}_{\mathcal{W}}(u))$. The proof that R'_a is transitive is similar to the euclidean proof. O-seriality can be proven directly, by observing that each $\mathit{new}_{\mathcal{W}}$ is a copy of the original model, so the property is preserved inside a partition. Since $\mathcal{A} \neq \emptyset$, it also holds in *orig*, because the world which is a copy of the *orig* in each partition is accessible (which, clearly, have the same valuation).

Part (b) Construct a bisimulation $\mathcal{R} \subseteq S \times S'$ such that, for $u, v \in S$ and $\mathcal{W} \in \mathcal{X}$,

$$\mathcal{R}(u, \mathit{new}_{\mathcal{W}}(v)) \Leftrightarrow u = v \text{ and } \mathcal{R}(u, \mathit{orig}(w)) \Leftrightarrow u = w$$

We check the various properties.

$\mathcal{R}$ satisfies forward-choice: Suppose $\mathcal{R}(s, s')$ and $R_i(s, t)$, $s, t \in S$, $s' \in S'$, $a \in \mathcal{A}$. If $s' = \mathit{new}_{\mathcal{W}}(s)$, then by the definition of $\textsc{unfold}_{\mathcal{X}}(\cdot, \cdot)$ and $R_a(s, t)$ we have $R'_a(\mathit{new}_{\mathcal{W}}(s), \mathit{new}_{\mathcal{W}}(t))$. If $s' = \mathit{orig}(s)$, we have $R'_a(\mathit{orig}(w), \mathit{new}_{\mathcal{W}}(t))$ for some $\mathcal{W} \in \mathcal{X}$ with $a \in \mathcal{W}$. Furthermore, for all cases of t' we get $\mathcal{R}(t, t')$.

$\mathcal{R}$ satisfies backward-choice: Suppose $\mathcal{R}(s, s')$ and $R_a(s', t')$, $s \in S$, s' and $t' \in S'$. For all cases of $R'_a(s', t')$, we immediately get $R_a(s, t)$. Also, for all cases of t' we have $\mathcal{R}(t, t')$.

$\mathcal{R}(s, s') \Rightarrow \pi(s) = \pi'(s')$ for all $s \in S$, $s' \in S'$: This is immediate, from the definition of $\mathcal{R}$ and the definition of $\textsc{unfold}_{\mathcal{X}}(\cdot, \cdot)$. □

Example 3.3. (Unfolding) Consider the Kripke model (M, s) in Figure 3.a with $\pi(s)(p) = \pi(u)(p) = \mathtt{true}$, $\pi(t)(p) = \mathtt{false}$. So, b knows that p is true, while a does not. Furthermore, a does not know if b knows p. Now the operation we perform is $\textsc{unfold}_{\{\{a\},\{b\}\}}(M, s)$, thus $\{a, b\}$ is split into $\{a\}$ and $\{b\}$, which results in the model (M', s) in Figure 3.b.

So, we have separated the knowledge of a and b. In Figure 3.b, the state s is the original state, the primed states model a's knowledge and the double primed states model b's knowledge. Thus, the upper half of the model represents the knowledge of a, and the lower half represents the knowledge of b. Note that no states are shared, in particular because the point of the model is not reflexive.

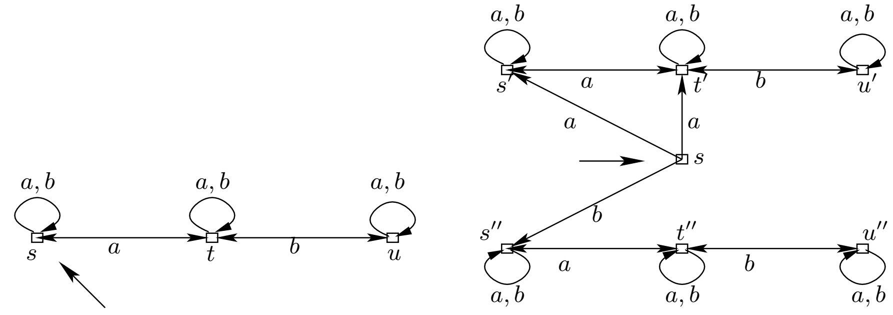

a. (M, s) b. (M', s)

Figure 3.

Now, we give some preparatory definitions leading to the formulation of a side-effect in Definition 3.9. First, we define the notion of a partial submodel.

Definition 3.4. A model $M = \langle S, \pi, R_1, ..., R_m \rangle$ is a partial submodel of $M' = \langle S', \pi', R'_1, ..., R'_m \rangle$, notation $M \sqsubseteq M'$, iff $S \subseteq S'$, $\pi(s)(p) = \pi'(s)(p)$ for all $s \in S$, $p \in \mathcal{P}$ and $R_i \subseteq R'_i$.

Note that a partial submodel is not pointed. Our notion of a partial submodel is slightly more liberal compared to the standard notion of a submodel, as here we allow to drop arrows. It is for technical reasons, viz. the handling of the atom split operation and the operation for side-effects below, that we have occasion to consider partial submodels here.

Next, we construct a partial submodel that represents the knowledge of a group of agents $\mathcal{B}$.

Definition 3.5. Given a model (M, w) such that

$$(M, w) = (\langle S, \pi, R_1, ..., R_m \rangle, w) = \text{UNFOLD}_{\mathcal{B}, \mathcal{A} \backslash \mathcal{B}}(M', w')$$

for some (M', w'), define $\text{SUB}_{\mathcal{B}}(M)$, the submodel of M for $\mathcal{B}$, by $\text{SUB}_{\mathcal{B}}(M) = \langle S', \pi', R'_1, ..., R'_m \rangle$ where

- $S' = \{\, new_{\mathcal{B}}(s) \mid s \in S \,\} \cup \{orig(w)\}$
- $\pi'(s)(p) \Leftrightarrow \pi(s)(p)$ for all $s \in S'$, $p \in \mathcal{P}$,
- for all $a \in \mathcal{A}$, $R'_a(s, t) \Leftrightarrow R_a(s, t) \wedge s, t \in S'$.

Clearly a $\mathcal{B}$-submodel is a partial submodel in the sense of Definition 3.4. The restmodel is the complementary part of the model that is the complement with respect to the accessibility relation of a given partial submodel.

Definition 3.6. Given a model $M = \langle S, \pi, R_1, ..., R_m \rangle$ and a partial submodel $N = \langle S'', \pi'', R_1'', ..., R_m'' \rangle$ of M, define $\textsc{rest}_N$, the restmodel of N in M, by $\textsc{rest}_N(M) = \langle S', \pi', R_1', ..., R_m' \rangle$ where

- $s \in S' \Leftrightarrow s \in S \wedge \exists a \in \mathcal{A} \exists (u,v) \in R_a' \ (u = s \vee v = s)$
- $\pi'(s)(p) \Leftrightarrow \pi(s)(p)$ for all $s \in S'$, $p \in \mathcal{P}$
- for all $a \in \mathcal{A}$, $R_a'(s,t) \Leftrightarrow R_a(s,t) \wedge \neg R_a''(s,t)$

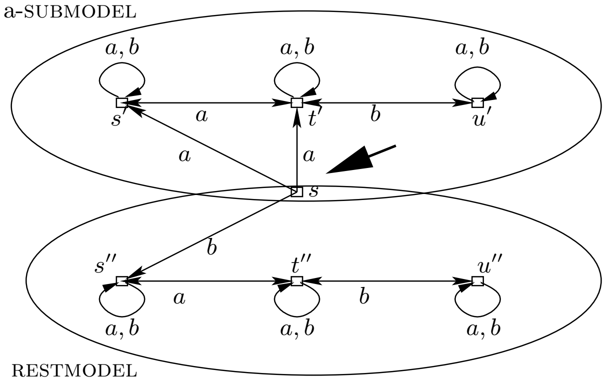

Figure 4. partial submodel and restmodel

We can see the partial submodel and restmodel definitions in action by taking the model of example 3.3 and applying the above definitions (see Figure 4). This exactly corresponds to the idea of two submodels that represent the belief of different agents.

Now, we would like to update the belief of some agents. To this end, we want to replace the submodel that represents their belief by a new model. We will apply the following definition.

Definition 3.7. Given a model $N = \langle S, \pi, R_1, ..., R_m \rangle$, a model M and a model N' such that $N \sqsubseteq N' \sqsubseteq M$ and

$$\textsc{rest}_{N'}(M) = \langle S', \pi', R_1', ..., R_m' \rangle$$

we define the operation $\textsc{replace}_{N'}(N, M)$, the replacement of N' by N in M, by $\textsc{replace}_{N'}(N, M) = \langle S'', \pi'', R_1'', ..., R_m'' \rangle$ where

- $s \in S'' \Leftrightarrow s \in S \vee s \in S'$,
- $\pi''(s)(p) \Leftrightarrow \pi(s)(p)$ for all $s \in S''$, $p \in \mathcal{P}$
- for all $a \in \mathcal{A}$, $R''_a(s,t) \Leftrightarrow (s,t) \in R_a \vee (s,t) \in R'_a$.

The idea is, that once the belief is completely separated, we can not only safely change the belief of certain agents, but also preserve the Kt45 properties. The operation $\textsc{atomsplit}_{(\phi,\mathcal{B})}$ removes the arrows for agents in the group $\mathcal{B}$ between states that have a different valuation for the objective formula ϕ.

Definition 3.8. Given a model $M = \langle S, \pi, R_1, ..., R_m \rangle$, an objective formula ϕ and a set of agents $\mathcal{B}$, we define the operation $\textsc{atomsplit}$ by

$$\textsc{atomsplit}_{(\phi,\mathcal{B})}(M) = \langle S', \pi', R'_1, ..., R'_m \rangle$$

where

- $S = S'$,
- $\pi'(s)(p) \Leftrightarrow \pi(s)(p)$ for all $s \in S'$, $p \in \mathcal{P}$,
- for $a \in \mathcal{B}$, $R'_a(s,t) \Leftrightarrow R_a(s,t) \wedge M, s \models \phi \Leftrightarrow M, t \models \phi$,
- for $a \notin \mathcal{B}$, $R'_a(s,t) \Leftrightarrow R_a(s,t)$.

Finally we are in a position to define the actual side-effect function that ties these things together.

Definition 3.9. For a model (M', w'), a set of agents $\mathcal{B}$ and an objective formula ϕ such that $(M', w') = \textsc{unfold}_{\{\mathcal{B}, \mathcal{A} \setminus \mathcal{B}\}}(M, w)$ and $N = \textsc{sub}_{\mathcal{B}}(M')$ we define the operation $\textsc{side-effect}_{(\phi,\mathcal{B},\mathcal{C})}(M, w)$, the side-effect for agents in $\mathcal{B}$ with respect to the agents in $\mathcal{C}$ and the formula ϕ, by

$$\textsc{side-effect}_{(\phi,\mathcal{B},\mathcal{C})}(M, w) = (\textsc{replace}_N(\textsc{atomsplit}_{(\phi,\mathcal{C})}(N), M'), w').$$

Note, that the formula ϕ in Definition 3.9 is required to be objective (cf. Example 3.2).

Example 3.4. We continue Example 3.3. Consider the a-submodel of M. We now apply $\textsc{atomsplit}_{(p,b)}$ on this model which results in the model (M'', s) in Figure 5. The arrow (t', u') has disappeared, since $\pi(t)(p) \neq \pi(u)(p)$. Therefore, u is not reachable anymore, and can be dropped. Notice, that a believes $B_b p \vee B_b \neg p$, while a has learned nothing about p itself, as was the case for Example 3.2.

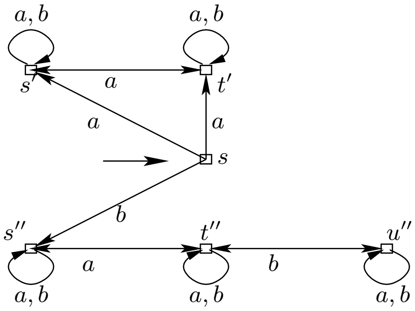

Figure 5. (M'', s)

A typical application of the side-effect function is one of the form $\text{SIDE-EFFECT}_{(p,\mathcal{A},b)}(M, w)$ where all the agents collectively learn that agent b knows about the atom p.

We introduce the notion of interconnection of relations, that comes in handy for a proof of the preservation of the Kt45 properties by the side-effect operation. Two binary relations A and B are called interconnected iff there is a $(w, v) \in A$ and $(s, t) \in B$ such that $w = s$, $w = t$, $v = s$ or $v = t$. If two binary relations are not interconnected, we call them separated. Separateness is useful because of its following properties.

Lemma 3.3.

(a) If binary relations A and B are separated and are both Kt45, then $A \cup B$ is also Kt45.

(b) If A and B are separated and $A \cup B$ has the Kt45 properties then both A and B have the Kt45 properties.

Proof. Part (a). We restrict ourselves only to the proof that union preserves the euclidean property. Assume $C = A \cup B$ not euclidean. Then there is $(s, t) \in C$ and $(s, u) \in C$ and $(t, u) \notin C$. Observe that (1) (s, t) and (s, u) cannot both come from A or both come from B, since those relations were both euclidean, and that would mean $(t, u) \in C$ and (2) (s, t) and (s, u) cannot come from the distinct subsets since that would contradict the interconnection property. So, in conclusion (s, t) and (s, u) cannot be elements of C. This directly contradicts the assumption. Hence C is is euclidean.

Part (b). Because of symmetry we only have to prove this for A. Suppose $(s, t) \in A$, $(s, u) \in A$. Since $C = A \cup B$ is euclidean, C must contain (t, u). But because A and B are not interconnected, (t, u) must be part of A. Therefore, A is euclidean. Proofs for transitivity and o-seriality are similar. □

We have seen that sets that are separated can be split and joined together without changing the Kt45 properties. In the next lemma we apply this for the replace operation.

Lemma 3.4. *Let $f\colon \mathcal{M} \to \mathcal{M}$ be an operation on the class of models such that (i) $f(M) \sqsubseteq M$ for any model $M \in \mathcal{M}$ and (ii) preserves $Kt45$ properties. Suppose model $(M, w) = \textsc{unfold}_{\{\mathcal{B},\mathcal{A}\backslash\mathcal{B}\}}(M', w')$ and $N = \textsc{sub}_{\mathcal{B}}(M)$. Then the operation defined by $\textsc{replace}_N(f(N), M)$ preserves the Kt45 properties too.*

Proof. From the definition of UNFOLD, it is quite easy to see, that N is separated (for all R_a, $a \in \mathcal{A}$) with the rest, since for a there's only a relation between *orig* and the partition where a is in, and for the other agents the relation from *orig* to the partition of a does not exist. So from Lemma 3.3 and the fact that M has the Kt45 properties, we must conclude that both the $\mathcal{B}$-submodel (N) and the restmodel have the Kt45 properties. Since f preserves the Kt45 properties, $f(N)$ also has the Kt45 properties too. By observing that doing a $\textsc{replace}_N$ is the same as doing a union of the accessibility-relations of the $\mathcal{B}$-restmodel with the new replacement, we can now use lemma 3.3, and conclude that $\textsc{replace}_N(f(N), M)$ has the Kt45 properties as well. □

In order to apply the above lemma we check that splitting preserves Kt45.

Lemma 3.5. *Given a Kt45-model $(M, w) = (\langle S, \pi, R_1, ..., R_m\rangle, w)$, then, for an objective formula ϕ and subset of agents $\mathcal{B}$, the model $\textsc{atomsplit}_{(\phi,\mathcal{B})}(M) = \langle S', \pi', R'_1, ..., R'_m\rangle$ has all the Kt45 properties too.*

Proof. For proving the new model is euclidean, suppose R_i is euclidean and $(s,t) \in R'_i \wedge (s,u) \in R'_i$. From $(s,t) \in R'_i$ follows that $\pi(s)(p) = \pi(t)(p)$ and $(s,t) \in R_i$. From $(s,u) \in R'_i$ follows that $\pi(s)(p) = \pi(u)(p)$ and $(s,u) \in R_i$. Thus, $\pi(t)(p) = \pi(u)(p)$ and therefore if $(t,u) \in R_i$, then $(t,u) \in R'_i$. Hence, $(t,u) \in R'_i$. The proof for transitivity is similar. For preservation of o-seriality we suppose $s \in S'$. By definition $s \in S$ and there is some $t \in S$ such that $(s,t) \in R_i$ and $(M,s) \models \varphi \leftrightarrow (M,t) \models \varphi$, where φ is objective. In particular then it holds that $\pi(s)(p) = \pi(t)(p)$. By definition $(s,t) \in R'_i$. □

We are now in a position to prove a number of properties of the side-effect operation.

Theorem 3.6.

(a) If (M, w) is a Kt45-model, then $\textsc{side-effect}_{(\phi,\mathcal{B},\mathcal{A})}(M, w)$*, for any objective formula ϕ and sets of agents $\mathcal{B}, \mathcal{C}$, is a Kt45-model as well.*

(b) (commutativity of side-effect) Given a model (M, w), sets of agents $\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}$ and two formulas ϕ, ψ, it holds that

$$\textsc{side-effect}_{(\phi,\mathcal{D},\mathcal{E})}(\textsc{side-effect}_{(\psi,\mathcal{B},\mathcal{C})}(M, w))$$

and $\textsc{side-effect}_{(\psi,\mathcal{B},\mathcal{C})}(\textsc{side-effect}_{(\phi,\mathcal{D},\mathcal{E})}(M, w))$ *are bisimilar.*

(c) (swapping update and side-effect) Given a model (M, w), sets of agents $\mathcal{B}, \mathcal{C}, \mathcal{D}$, a formula ϕ and an objective formula ψ, it holds that the models

$$\textsc{side-effect}_{(\phi,\mathcal{C},\mathcal{D})}(\textsc{update}_{(\psi,\mathcal{B})}(M, w))$$

and

$$\textsc{update}_{(\psi,\mathcal{B})}(\textsc{side-effect}_{(\phi,\mathcal{C},\mathcal{D})}(M, w))$$

are bisimilar.

(d) (idempotency of side-effect) Given a model (M, w), sets of agents $\mathcal{B}, \mathcal{C}, \mathcal{D}$ and a formula ϕ, it holds that

$$\textsc{side-effect}_{(\phi,\mathcal{B},\mathcal{C})}(\textsc{side-effect}_{(\phi,\mathcal{B},\mathcal{C})}(M, w))$$

and $\textsc{side-effect}_{(\phi,\mathcal{B},\mathcal{C})}(M, w)$ *are bisimilar.*

Proof. (a) Clearly, $\textsc{atomsplit}_{(p,\mathcal{C})}(M) \sqsubseteq M$ holds. Therefore, the statement follows from Lemma 3.4 and Lemma 3.5.

(b) There are several tagged states. We have *orig*, $\mathit{new}_{\mathcal{B}}$, $\mathit{new}_{\mathcal{A}\setminus\mathcal{B}}$, $\mathit{new}_{\mathcal{D}}$, $\mathit{new}_{\mathcal{A}\setminus\mathcal{D}}$ in both models.

Construct a bisimulation $\mathcal{R} \subseteq S \times S'$ such that:

$$\begin{array}{lcl}
\mathcal{R}(\mathit{orig}(w), \mathit{orig}(v)) & \Leftrightarrow & w = v \\
\mathcal{R}(\mathit{new}_{\mathcal{B}}(u), \mathit{new}_{\mathcal{B}}(v)) & \Leftrightarrow & u = v \\
\mathcal{R}(\mathit{new}_{\mathcal{A}\setminus\mathcal{B}}(u), \mathit{new}_{\mathcal{A}\setminus\mathcal{B}}(v)) & \Leftrightarrow & u = v \\
\mathcal{R}(\mathit{new}_{\mathcal{D}}(u), \mathit{new}_{\mathcal{D}}(v)) & \Leftrightarrow & u = v \\
\mathcal{R}(\mathit{new}_{\mathcal{A}\setminus\mathcal{D}}(u), \mathit{new}_{\mathcal{A}\setminus\mathcal{D}}(v)) & \Leftrightarrow & u = v
\end{array}$$

We can now follow the reasoning from Theorem 3.1 to see that this bisimulation has all the desired properties for bisimulation. For example, $\mathcal{R}$ satisfies forward-choice: suppose $\mathcal{R}(s, s') \wedge R_a(s, t)$, $s, t \in S, s' \in S'$. Suppose *orig*(s), then t could be $\mathit{new}_{\mathcal{B}}$ or $\mathit{new}_{\mathcal{A}\setminus\mathcal{B}}$. But this copy of

t is present in S' as well, since that was created after the first execution of SIDE-EFFECT on the original model. And indeed, there's a $orig(w)$ such that $(w, t) \in S'$. The checks for all the other options are similar.

(c) Construct a bisimulation $\mathcal{R} \subseteq S \times S'$ such that:

$$\begin{array}{lcl}
\mathcal{R}(orig(w), orig(v)) & \Leftrightarrow & w = v \\
\mathcal{R}(old(u), old(v)) & \Leftrightarrow & u = v \\
\mathcal{R}(new(u), new(v)) & \Leftrightarrow & u = v \\
\mathcal{R}(new_{\mathcal{B}}(u), new_{\mathcal{B}}(v)) & \Leftrightarrow & u = v \\
\mathcal{R}(new_{\mathcal{A}\backslash\mathcal{B}}(u), ew_{\mathcal{A}\backslash\mathcal{B}}(v)) & \Leftrightarrow & u = v
\end{array}$$

Checking all the properties is similar to (b).

(d) The result of applying the same side-effect operation twice is the same model as applying it just once modulo a number of unreachable states. Again, construct a bisimulation $\mathcal{R}$ similar to all the previous proofs such that the relation exists if they are a copy of each other and reachable from the point of the model. The unprimed variables belong to $\text{SIDE-EFFECT}_{(\phi,\mathcal{B},\mathcal{C})}(M, w)$ and the primed ones to $\text{SIDE-EFFECT}_{(\phi,\mathcal{B},\mathcal{C})}(\text{SIDE-EFFECT}_{(\phi,\mathcal{B},\mathcal{C})}(M, w)$. Now, $\mathcal{R}$ satisfies forward-choice: suppose $\mathcal{R}(s, s') \wedge R_a(s, t)$, $s, t \in S, s' \in S'$. Suppose $new_{\mathcal{B}}(s)$, then we indeed have such a copy $new_{\mathcal{B}}(t')$ with $\mathcal{R}(t, t')$, since no more arrows inside $new'_{\mathcal{B}}$ were deleted, because only the states from $new_{\mathcal{B}}$ are reachable in $new'_{\mathcal{B}}$. For the other cases it it trivial. It can also be shown that it satisfies the other properties that are required. □

Next, we consider how the formulas are altered by the side-effect operation. We will partially answer this by presenting a few interesting formulas that hold in the resulting model. We distinguish

1. the group of agents $\mathcal{B}$ that learn about other agents, ranged over by b;

2. the group of agents $\mathcal{C}$ that is learned about, ranged over by c;

3. other agents in the group $\mathcal{D}$, ranged over by d.

The fact that the agents in $\mathcal{B}$ are the only agents that learn at all, is clear. The other agents consider exactly (copies of) their old worlds possible; their belief has not changed. With this in mind, we present a few properties of the side-effect operation. Below $C_{\mathcal{BCD}}$ expresses common knowledge of agents in the sets $\mathcal{B}$, $\mathcal{C}$ and $\mathcal{D}$.

Lemma 3.7. *Given a model (M, w), disjoint sets of agents $\mathcal{B}, \mathcal{C}, \mathcal{D}$ and a formula ϕ, put $(M', w') = \text{SIDE-EFFECT}_{(\phi,\mathcal{B},\mathcal{C})}(M, w)$. Then it holds that* *(a)* $(M', w') \models B_b(B_c\phi \vee B_c\neg\phi)$*;*

(b) $(M', w') \models B_b\psi$ *iff* $(M, w) \models B_b\psi$ *for any objective formula* ψ*;*

(c) $(M', w') \models B_bC_{\mathcal{BCD}}(B_c\phi \vee B_c\neg\phi)$*;*

(d) $(M', w') \models B_a\psi$ *iff* $(M, w) \models B_a\psi$ *for any formula* ψ *and any agent* $a \notin \mathcal{B}$*.*

Proof. We prove the typical cases of part (b) and (c).

(b) The unfolded model of (M, w) is bisimilar with (M', w'), and no arrows of b were deleted afterwards. Hence, the knowledge of b about objective formulas has not changed.

(c) We have already seen that in every state of $\text{SUB}_{\mathcal{B}}(M')$ it holds that $B_c\phi \vee B_c\neg\phi$. Now, what we need to prove is that the path $w' \xrightarrow{b} s_1 \xrightarrow{i} s_2 \xrightarrow{j} ..$ with $i, j, .. \in \mathcal{A}$ is a path to a state where it holds that $B_bp \vee B_b\neg p$. But since $s_1 \in new_{\mathcal{B}}$, and since there are no arrows from $new_{\mathcal{A}}$ to other partitions, all s_k are elements of $new_{\mathcal{B}}$. Thus any s_k is part of $\text{SUB}_b(M')$. By the construction of $\text{SIDE-EFFECT}_{(\phi,\mathcal{B},\mathcal{C})}(M, w)$, in any state in $new_{\mathcal{B}}$, we either have all arrows of c to a world where ϕ holds (at least one, by o-seriality) or to a world where $\neg\phi$ holds. Furthermore, $B_c\phi$ or $B_c\neg\phi$ holds. Hence, $B_c\phi \vee B_c\neg\phi$ holds. □

In part (a) of the above lemma, an agent b obtains derived knowledge of an agent c. Part (b) states that no objective knowledge is learned. Part (c) phrases that an agent b considers the rest of the agents as smart itself. Finally, part (d) captures that other agents do not learn.

Property (c) may or may not be a reasonable assumption of b about the other agents. If one agent believes that another agent knows the value of ϕ, then it is reasonable to assume that another agent will believe the same. On the other hand common knowledge might be too strong to assume.

Next, we address the issue that an agent b shares its belief about an agent c with only some other agents. We represent this by linking b's beliefs of those other agents back to the original (unmodified) states. We distinguish four different type of groups of agents.

1. the groups $\mathcal{B}$ and $\mathcal{C}$ are as before;
2. the group $\mathcal{D}$ of agents of which agents in $\mathcal{B}$ believe they have learned in common about agents in group $\mathcal{C}$, ranged over by d;
3. the group $\mathcal{E}$ of agents of which agents in $\mathcal{B}$ believe they have learned nothing about, ranged over by e.

We define the new side-effect operation 0-UNFOLD that handles this refinement. Here, 0 refers to zero-knowledge for the group of agents $\mathcal{E}$. As before, we define an unfolding operation first.

Definition 3.10. Given a model (M, w), with $M = \langle S, \pi, R_1, ..., R_m \rangle$, and a partitioning $\mathcal{X} = \{\mathcal{B}, \mathcal{C}, \mathcal{D}, \mathcal{E}\}$ of the set of agents $\mathcal{A}$, we define the operation 0-UNFOLD$_{\mathcal{X}}(M, w)$, the zero-knowledge unfolding of (M, w) with respect to $\mathcal{X}$, by 0-UNFOLD$_{\mathcal{X}}(M, w) = (\langle S', \pi', R'_1, ..., R'_m \rangle, w')$, where

- $S' = \mathit{new}_{\mathcal{B}}(S) \cup \mathit{new}_{\mathcal{CDE}}(S) \cup \{\mathit{orig}(w)\}$
- $w' = \mathit{orig}(w)$
- $\pi'(\mathit{new}_{\mathcal{Y}}(v))(p) = \pi(v)(p)$ and $\pi'(\mathit{orig}(w))(p) = \pi(w)(p)$ for all $p \in \mathcal{P}$, $\mathcal{Y} \in \{\{\mathcal{B}\}, \{\mathcal{CDE}\}\}$
- for $a \in \mathcal{A}$, R'_a on S' is the minimal binary relation such that

$$\begin{array}{lcl}
R'_a(\mathit{orig}(w), \mathit{new}_{\mathcal{B}}(v)) & \Leftrightarrow & R_a(w, v) \ \wedge \ a \in \mathcal{B} \\
R'_a(\mathit{orig}(w), \mathit{new}_{\mathcal{CDE}}(v)) & \Leftrightarrow & R_a(w, v) \ \wedge \ a \notin \mathcal{B} \\
R'_a(\mathit{new}_{\mathcal{B}}(u), \mathit{new}_{\mathcal{CDE}}(v)) & \Leftrightarrow & R_a(u, v) \ \wedge \ a \in \mathcal{E} \\
R'_a(\mathit{new}_{\mathcal{B}}(u), \mathit{new}_{\mathcal{B}}(v)) & \Leftrightarrow & R_a(u, v) \ \wedge \ a \notin \mathcal{D} \\
R'_a(\mathit{new}_{\mathcal{CDE}}(u), \mathit{new}_{\mathcal{CDE}}(v)) & \Leftrightarrow & R_a(u, v)
\end{array}$$

So, instead of completely separating the knowledge of agents in $\mathcal{B}$ with the other agents, we share this knowledge with the other agents. Since the other agents do not learn anything, agents in $\mathcal{B}$ does not gain knowledge about $\mathcal{E}$. We present a theorem similar to Theorem 3.2.

Theorem 3.8.

(a) If (M, w) is a Kt45 model, then so is 0-UNFOLD(M, w).

(b) For every model (M, w), it holds that the models (M, w) and 0-UNFOLD(M, w) *are bisimilar.*

Proof. (a) The accessibility relations of every agent other than those in $\mathcal{E}$ are constructed exactly the same way as in Definition 3.3. Since that operation preserves the Kt45 properties, we do not have to prove it for those agents. A proof for the agents in $\mathcal{E}$ now follows. For transitivity assume $(s, t) \in R'_a, (t, u) \in R'_a$. By the definition, we immediately see that u is of the form $\mathit{new}_{\mathcal{CDE}}(\cdot)$. Also we see that if there is a $R_a(s, u)$ then there's a $R'_a(s, u)$. Thus, we can conclude $R'_a(s, u)$. The proof that R'_a euclidean is similar. For o-seriality observe that every arrow of an agent in $\mathcal{E}$ ends in a state that used to be its old state. Thus, o-seriality is preserved.

(b) Construct a bisimulation $\mathcal{R} \subseteq S \times S'$ such that

$$\mathcal{R}(u, \mathit{new}_{\mathcal{Y}}(v)) \Leftrightarrow u = v \text{ and } \mathcal{R}(u, \mathit{orig}(v)) \Leftrightarrow u = v.$$

We prove that

$\mathcal{R}$ satisfies forward-choice: Suppose $\mathcal{R}(s, s')$ and $R_a(s, t)$, $s, t \in S, s' \in S'$. Suppose $\mathit{orig}(s)$, then we have $\mathit{new}_{\mathcal{B}}(t)$ if $a \in \mathcal{B}$, else we have $\mathit{new}_{\mathcal{BCD}}(t)$ if $a \notin \mathcal{B}$. Suppose $\mathit{new}_{\mathcal{CDE}}(s)$, then we immediately have $\mathit{new}_{\mathcal{CDE}}(t)$ for all agents a. Suppose $\mathit{new}_{\mathcal{B}}(s)$, then we have $\mathit{new}_{\mathcal{B}}(t)$ if $a \notin \mathcal{E}$, else $\mathit{new}_{\mathcal{CDE}}(t)$.

$\mathcal{R}$ satisfies backward-choice: Suppose $\mathcal{R}(s, s')$ and $R_a(s', t')$, $s \in S, s', t' \in S'$. We have that $R_a(s', t')$ immediately implies $R_a(s, t)$ for some the s, t such that s', t' are copies from s, t. This implies $\mathcal{R}(t, t')$.

$\mathcal{R}(s, s') \Rightarrow \pi(s) = \pi'(s')$ for $s \in S$, $s' \in S'$: Direct from the definition of $\mathcal{R}$. □

Due to the case distinction for arrows leaving the point $\mathit{orig}(w)$, in Definition 3.10 above, it holds that the knowledge of $\mathcal{B}$-agents about $\mathcal{C}$-agents is separated, with the knowledge of other agents about agents in $\mathcal{C}$ or of $\mathcal{C}$-agents themselves. So we can 'cut out' the submodel containing the c arrows from the belief of b. In this definition we can omit the point of the model in our partial submodel, which makes it slightly easier.

Definition 3.11. Let $(M', w') = \langle S', \pi', R'_1, ..., R'_m \rangle$, ϕ be a model, ϕ an objective formula and $\mathcal{X}$ a partitioning of $\mathcal{A}$ (as in Definition 3.10) such that $(M', w') = 0\text{-}\textsc{unfold}_{\mathcal{X}}(M, w)$, for some model (M, w). We define $\textsc{sub}_{\mathcal{B}}(M') = \langle S'', \pi'', R''_1, ..., R''_m \rangle$ where

- $S'' = \{ \mathit{new}_{\mathcal{B}}(s) \mid s \in S \}$
- $\pi''(s)(p) \Leftrightarrow \pi(s)(p)$ for all $s \in S'$, $p \in \mathcal{P}$
- $R''_c(s, t) \Leftrightarrow R'_c(s, t)$ for $s, t \in S''$
- $R''_a = \emptyset$ $(a \notin \mathcal{C})$.

The operation $0\text{-}\textsc{side-effect}_{(\phi,\mathcal{X})}$, the zero-knowledge side-effect of ϕ with respect to the partitioning $\mathcal{X}$, is then given by

$$0\text{-}\textsc{side-effect}_{(\phi,\mathcal{X})}(M, w) = (\textsc{replace}_N(\textsc{atomsplit}_{(\phi,\mathcal{C})}(N), M'), w')$$

where $N = \textsc{sub}_{\mathcal{B}}(M')$.

The operation $0\text{-}\textsc{side-effect}$ has algebraic properties comparable to those of the previous side-effect operation (cf. Theorem 3.6). We also have to following result, corresponding to Lemma 3.7.

Lemma 3.9. *Let (M, w) be a model, $\mathcal{X}$ a partitioning of $\mathcal{A}$ and ϕ an objective formula. Put $(M', w') = 0\text{-}\textsc{side-effect}_{(\phi,\mathcal{X})}(M, w)$. Then it holds that*

(a) $(M', w') \models B_b(B_c\phi \vee B_c\neg\phi)$

(b) $(M', w') \models B_b\psi$ *iff* $(M, w) \models B_b\psi$ *for any objective formula* ψ

(c) $(M', w') \models B_bC_{\mathcal{BCD}}(B_c\phi \vee B_c\neg\phi)$

(d) $(M', w') \models B_a\psi$ *iff* $(M, w) \models B_a\psi$ *for* $a \notin \mathcal{B}$ *for any formula* ψ

(e) $(M', w') \models B_bB_e\psi$ *iff* $(M, w) \models B_bB_e\psi$ *for any formula* ψ*.*

Proof. We provide the proof for (e). ($\Rightarrow$) Assume $(M', w') \models B_bB_e\psi$, then we have for all paths $w \overset{b}{\to} s \overset{e}{\to} t$, $t \models \psi$. But, since all arrows of e in $new_{\mathcal{B}}$ point back to (a copy of) the original model, we have that $(M, t) \models \psi$. ($\Leftarrow$) Assume $(M, w) \models B_bB_e\psi$. Now, we have a path in the original model $w \overset{b}{\to} s \overset{e}{\to} t$. In the construction of (M, w) no arrows of b and e are ever added nor deleted. □

Parts (a) to (d) are correspond to the properties given in Lemma 3.7. Part (e) states that the knowledge of agent b about agent e has not changed, which is exactly as desired.

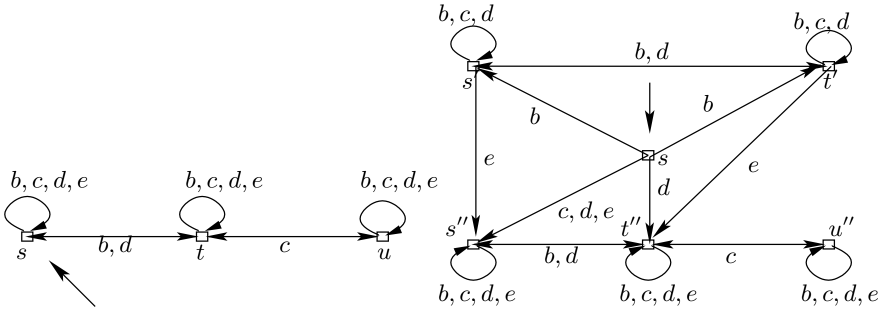

a. (M, s) b. (M', s)

Figure 6.

Example 3.5. (alternative side-effect function) Recall from Example 3.3 the model (M, s) (Figure 3.a). We now present this model with four agents $\{b, c, d, e\}$ in Figure 6.a, with a partitioning into singletons, such that $\pi(s)(p) = \pi(u)(p) = \texttt{true}$ and $\pi(t)(p) = \texttt{false}$. Now, apply $0\text{-}\textsc{side-effect}_{(p,b,c)}(M, w)$ and we gain the model (M', s) from Figure 6.b such that $\pi(s)(p) = \pi(s')(p) = \pi(s'')(p) = \pi(u')(p) = \texttt{true}$, $\pi(t)(p) = \pi(t')(p) = \texttt{false}$. Note that in this model, b still knows exactly the same about e as it did before.

3.3. Comparison with the action model approach

Baltag and Moss (Baltag and Moss, 2004; Baltag et al., 1998) propose a framework for describing epistemic actions using action models. Similar to Kripke models that describe the uncertainty of agents about which world they are in, they use Kripke models to describe the uncertainty of agents about the action that is being performed.

Formally, an epistemic action model is a triple $\Sigma = \langle \Sigma, \rightarrow \mathcal{A}, \mathrm{pre} \rangle$, where Σ is a set of *simple actions*, $\rightarrow \mathcal{A}$ is an accessibility relation of agents on actions and the precondition pre is a mapping $\mathrm{pre}\colon \Sigma \rightarrow \Phi$ with Φ being the collection of all epistemic propositions. The central operation of updating an epistemic model $M = \langle S, R_1, \ldots, R_m, \pi \rangle$ as we have used so far with such an action model Σ is defined as $M \otimes \Sigma = \langle S \otimes \Sigma, R'_1, \ldots, R'_m, \pi' \rangle$, where

- $S \otimes \Sigma = \{(s, \sigma) \in S \times \Sigma \mid (M, s) \models \mathrm{pre}(\sigma)\}$
- $R'_i((s, \sigma), (s', \sigma'))$ iff $R_i(s, s')$ and $\sigma \rightarrow \mathcal{A}\sigma'$
- $\pi'((s, \sigma))(p)$ iff $\pi(s)(p)$

In (Baltag and Moss, 2004), Baltag and Moss provide, based on this notion of updating, illuminating examples and study several interesting applications of this idea such as the public and private learning of, what we would call, objective formulas. This poses the question if we can describe the more complicated updates as well.

In Figure 7 the ovals depict a precondition with φ some objective formula. Intuitively, this action model corresponds to the operation of 0-side-effect.

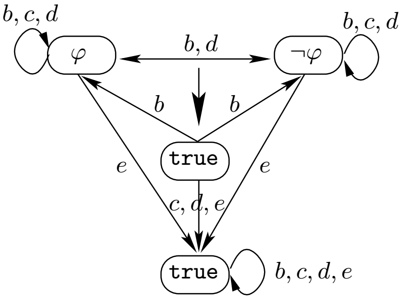

Figure 7. 0-side-effect action model

Indeed, the update product of the Kripke structure of Example 3.5 in Figure 6.a and the action model of Figure 7 results in the model of Figure 6.b. Also, omitting the e-arrows from the action model in Figure 7 yields an action model that for the concrete examples discussed above corresponds to the side-effect operation of Definition 3.9.

However, currently we have no proof that such a correspondence holds in general.

4. A logical language for security protocols

In this section we exploit the ideas of the previous section for a logical language to reason about security protocols. The UPDATE and SIDE-EFFECT operations are used for its semantics. We introduce so-called transition rules for the modeling of security protocols, that we discuss in the next section.

Definition 4.1. Fix a set of proposition $\mathcal{P}$, ranged over by p, and a set of agents $\mathcal{A}$ of m elements, ranged over by i, j. The language $\mathcal{L}_C$ is given by

$$\phi ::= p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid B_i\phi \mid [\sigma]\phi$$
$$\sigma ::= \mathit{Priv}(i \rightarrow j, p) \mid \mathit{Pub}(i, p) \mid \sigma; \sigma'$$

where $\mathcal{C}$ is a collection of so-called transition rules.

The σ symbol denotes a (possibly composed) communication action. The action $\mathit{Priv}(i \rightarrow j, p)$ is a private or peer-to-peer message p from i to j; the action $\mathit{Pub}(i, p)$ means a public announcement or a broadcast by i about p. In the latter, every agent on the network learns p, whereas in the former, only j learns p. The bracket operator $[\sigma]\phi$ has the interpretation that after executing the communication action σ, ϕ holds.

The subscript in $\mathcal{L}_C$ refers to a set of so-called transition rules $\mathcal{C}$. The transition rules capture the updates, i.e. the expansions and side-effects, necessary for the interpretation of indirect effect of the constructs $\mathit{Priv}(i \rightarrow j, p)$ and $\mathit{Pub}(i, p)$. The transitions rules enforce consistency among the propositions that hold. For example, if an agent believes that the value of a message m is $[\![m]\!]$ and possesses a key k, then it must believe that the value of the encryption $\{m\}_k$ of m has a value that corresponds with $[\![m]\!]$.

A transition rule has the form $B_i p \Rightarrow \beta$. The condition $B_i p$ expresses that p must be believed by agent i. The body β of a transition rule is a sequence of actions $\alpha_1; \ldots; \alpha_n$. Actions come in three flavours, viz. $L_\mathcal{B} p$, $S_{\mathcal{B},\mathcal{C}} p$ and $S^0_{\mathcal{B},\mathcal{C},\mathcal{D}} p$. Here, $L_\mathcal{B} p$ expresses that p is learned among the agents in the set $\mathcal{B}$ and corresponds to belief expansion, whereas $S_{\mathcal{B},\mathcal{C}} p$ expresses the side-effect that the agents in the set $\mathcal{B}$ have learned that the agents in the set $\mathcal{C}$ now know about p. Similarly, $S^0_{\mathcal{B},\mathcal{C},\mathcal{D}} p$ is used for the side-effect where agents in $\mathcal{B}$ assume that the agents in $\mathcal{D}$ have learned as well.

As an example, we will have the transition rule $B_b\{x\}_k \Rightarrow L_{ab}x$, when agents a and b share the key k, and a sends b the message $\{x\}_k$. In the situation described above, agent a sends the message x to agent b and agent b returns the message $\{x\}_k$. Since it is shared, a already can compute $\{x\}_k$ itself, so the delivery of $\{x\}_k$ does not teach a anything about this value. However, the transition rule expresses that a and b commonly learn, and, in particular, a learns that b knows the message x.

The semantics for the language $\mathcal{L}_C$, provided in the next definition, follows the set-up of, e.g., (Baltag et al., 1998; van Ditmarsch, 2000). Definition 4.2 is organized in three layers. First, there is the layer of the actions of the language. The defining clauses make use of an auxiliary operation $\triangleright_p$. This operation helps in the processing of the relevant transition rules. The set $\mathit{Mod}(M, w, p)$ collects all the transition rules that will change the model. The next layer of the definition concerns the body of a transition rule. The last part of Definition 4.2 concerns the validity of the formulas of $\mathcal{L}_C$.

Definition 4.2. Let $\mathcal{C}$ be a finite set of transition rules. For $\sigma \in \mathcal{L}_C$ the relation $[\![\sigma]\!]$ on models for $\mathcal{A}$ over $\mathcal{P}$ is given by

$$\begin{aligned}
&(M,w)[\![\mathit{Priv}(i \to j,p)]\!](M',w') \\
&\qquad \Leftrightarrow (M,w) \models B_i p \Rightarrow (\textsc{update}_{p,j}(M,w) \triangleright_p (M',w')) \\
&(M,w)[\![\mathit{Pub}(i,p)]\!](M',w') \\
&\qquad \Leftrightarrow (M,w) \models B_i p \Rightarrow (\textsc{update}_{p,\mathcal{A}}(M,w) \triangleright_p (M',w')) \\
&(M,w)[\![\sigma;\sigma']\!](M',w') \\
&\qquad \Leftrightarrow (M,w)[\![\sigma]\!](M'',w'')[\![\sigma']\!](M',w') \text{ for some model } (M'',w'')
\end{aligned}$$

$$\begin{aligned}
&(M,w) \triangleright_p (M',w') \\
&\qquad \Leftrightarrow \texttt{if}\ \ (x \Rightarrow \beta) \in \mathit{Mod}(M,w,p) \\
&\qquad\quad\ \texttt{then}\ \ (M,w)\ \langle\beta\rangle\ (M'',w'') \triangleright_p (M',w') \text{ for some } (M'',w'') \\
&\qquad\quad\ \texttt{else}\ \ (M,w) = (M',w')\ \texttt{end}
\end{aligned}$$

$$\begin{aligned}
&(M,w)\ \langle\rangle\ (M',w') \\
&\qquad \Leftrightarrow (M,w) = (M',w') \\
&(M,w)\ \langle L_{\mathcal{B}}p;\beta\rangle\ (M',w') \\
&\qquad \Leftrightarrow \textsc{update}_{(p,\mathcal{B})}(M,w)\langle\beta\rangle(M',w') \\
&(M,w)\ \langle S_{\mathcal{B},\mathcal{C}}p;\beta\rangle\ (M',w') \\
&\qquad \Leftrightarrow \textsc{side-effect}_{(p,\mathcal{B},\mathcal{C})}(M,w)\ \langle\beta\rangle\ (M',w') \\
&(M,w)\ \langle S^0_{\mathcal{B},\mathcal{C},\mathcal{D}}p;\beta\rangle\ (M',w') \\
&\qquad \Leftrightarrow 0\text{-}\textsc{side-effect}_{(p,\mathcal{X})}(M,w)\langle\beta\rangle \\
&\qquad\qquad \text{where } \mathcal{X} = \{\mathcal{B},\mathcal{C},\mathcal{D},\mathcal{A} \setminus (\mathcal{B} \cup \mathcal{C} \cup \mathcal{D})\}
\end{aligned}$$

$$
\begin{aligned}
&(M,w) \models p\\
&\qquad \Leftrightarrow \pi(w)(p) = \texttt{true}\\
&(M,w) \models \neg\phi\\
&\quad \Leftrightarrow (M,w) \not\models \phi\\
&(M,w) \models \phi \wedge \psi\\
&\qquad \Leftrightarrow (M,w) \models \phi \text{ and } (M,w) \models \psi\\
&(M,w) \models B_i\phi\\
&\qquad \Leftrightarrow (M,v) \models \phi \text{ for all } v \text{ such that } R_i(w,v)\\
&(M,w) \models [\sigma]\phi\\
&\qquad \Leftrightarrow (M',w') \models \phi \text{ if } (M,w)\llbracket\sigma\rrbracket(M',w')
\end{aligned}
$$

where

$$
\begin{aligned}
Mod(M,w,p) = \{\, B_ip \Rightarrow \beta \,\in \mathcal{C} \mid (M,w) \models B_ip, (M,w)\, \langle\beta\rangle\, (M',w'),\\
\exists\phi\colon (M',w') \models \phi \not\leftrightarrow (M,w) \models \phi \,\}
\end{aligned}
$$

The private and public communication of p can only be executed under the condition B_ip. The communication has the effect that the agent j, respectively all agents get informed about p. Next, the transition rules are invoked as a consequence of some parties learning p, as expressed by the operator $\rhd_p$. The 'modifiers', the transition rules in the set $Mod(M,w,p)$ are those rules that match the learning of p and, moreover, will transform (M,w) into a different model, i.e. some formula ϕ will have changed its truth value. As a consequence of the algebraic properties of the update and side-effect operators of Section 3, the order of in which the transition rules are processed does not matter and every 'modifying' rule gets applied at most once (by idempotency). So, no rules are applied over and over again. Apart from this, the above definition also works for general formulas instead of objective ones (cf. (Baltag et al., 1998; van Ditmarsch, 2000)).

5. Examples

In this section we discuss how the machinery developed above works out for a concrete example. Preparatory for this, in order to keep the models within reasonable size, we employ two helpful tricks. The first one is the disregarding of propositions not known to any agent. Thus, if a proposition is not part of the model, then the interpretation is that no agent has any knowledge about it. What we then have to specify is how to add a fresh proposition to the model. We accomplish this by making two copies of the original states. One of them we assign 'positive' and the other 'negative'. In the positive states, the proposition will be `true`, and in the negative states, the proposition will be `false`.

Definition 5.1. Given a model $(M, w) = \langle S, \pi, R_1, ..., R_m \rangle$ and a fresh proposition p, we define the operation $\textsc{addatom}_p$ such that $(M', w') = \textsc{addatom}_p(M, w) = \langle S', \pi', R'_1, ..., R'_m \rangle$ where

- $S' = pos(S) \cup neg(S)$
- $\pi'(pos(s))(q) =$ if $p = q$ then `true` else $\pi(s)(q)$
- $\pi'(neg(s))(q) =$ if $p = q$ then `false` else $\pi(s)(q)$
- $R'_i(s, t) \Leftrightarrow R_i(s, t)$ for any agent i
- $w' = pos(w)$

We suppress straightforward technicalities regarding the restriction of the domain of the valuation π or expansion of the set of proposition $\mathcal{P}$.

We have the following property.

Lemma 5.1. *Given a model (M, w) and a fresh proposition p such that $(M', w') = \textsc{addatom}_p(M, w)$ it holds that*

(a) $(M', w') \models p$;

(b) $(M, w) \models \phi \Leftrightarrow (M', w') \models \phi$ *for* $p \notin \phi^*$ *with* ϕ^* *the closure under subformulas of* ϕ;

(c) $(M', w') \not\models B_i p$ *for all agents* i.

Proof. (a) Trivial, since we make the new point the positive copy of the old point.

(b) We prove the stronger property $(M, s) \models \phi, p \notin \phi^* \Leftrightarrow (M', s')$ where s' is a copy of s. Proof by induction on complexity of ϕ. Suppose ϕ is objective then it's trivial (since $p \notin \phi^*$). Now suppose $\phi = B_i\psi$, then each state that is reachable in the resulting model is reachable if and only if it was reachable from a copy in the original model. By induction, we have the property for ψ, so it follows that we have it for ϕ.

(c) It holds that $w' = pos(w)$ and we have some $w'' = neg(w)$, such that $R_a(w', w'')$. □

The second trick helps to short-cut the application of rules which helps keeping the model in a reasonable size.

Lemma 5.2. *Given a model (M, w), an agent $i \in \mathcal{A}$, a set of agents $\mathcal{B} \subseteq \mathcal{A}$ it holds that the model (M', w') such that*

$$(M, w)\langle L_i p; S_{\mathcal{B},i} p\rangle x \langle L_{\mathcal{A}} p\rangle (M', w')$$

for some model x, and the model (M'', w'') such that

$$(M, w)\langle L_{\mathcal{A}}p\rangle(M'', w'')$$

are bisimilar.

Proof. The proof is a corollary of the following two properties:

1. $L_i p; L_{\mathcal{A}}p = L_{\mathcal{A}}p$
2. $S_{\mathcal{B},i}p; L_{\mathcal{A}p} = L_{\mathcal{A}}p$

The proof of these properties is similar to the ones we have seen before. Here we proof (1). Because in both operations the last operation is a $L_{\mathcal{A}}$ we can easily that only *new* states are reachable. Bisimulation is constructed by associating links between states that are copies of each other in the original model. Checking bisimulation now is trivial. For (2) it is similar. Intuitively for the first operation, again in the last step, only the *new* are reachable. The removal of a in the knowledge of $\mathcal{B}$ in the first step is redundant because of the removal that happens in any case in the second step. □

That is to say, if an agent i learns p and then all other agents learn about i that it has learned p, followed by the action where everyone learns p (commonly), then it is equivalent to say that they have just learned p commonly.

5.1. The SRA Three Pass protocol

Shamir, Rivest and Adleman have suggested the three-pass protocol (Clark and Jacob, 1997) for the transmission of a message under the assumption of a commutative cipher. It is known to be insecure and various attacks have been suggested. However, it serves an illustrative purpose here. The protocol has the following steps:

1. $a \rightarrow b : \{x\}_{k_a}$
2. $b \rightarrow a : \{\{x\}_{k_a}\}_{k_b}$
3. $a \rightarrow b : \{x\}_{k_b}$

Here, both agent a and b have their own symmetric and unshared encryption key, k_a and k_b, respectively. Agent a wants to send message x to agent b through an insecure channel and therefore wants to send x encrypted to b. It does this by sending x protected with its own key. Next, b will encrypt this message with b's key and sends this back. Since the encryption is assumed to be commutative, a can now decrypt this message and sends the result to b. Finally, b can decrypt the message it has just received and learn the value of x.

In our modeling, we consider three agents $\{a, b, c\}$. It is assumed that all agents can see the activity of the network. In particular, they see messages been sent out and received. We are interested in what agent c can learn during a run of this protocol between agents a and b. We use the notation $m_{\mathcal{K}}$, for a possibly empty set of agents $\mathcal{K}$, to denote the message m encrypted with the keys of all agents in the set $\mathcal{K}$. We have, e.g., $m_{\{a,b\}} = \{\{m\}_{k_a}\}_{k_b}$. Since the cipher is commutative, this is well defined. Also, we write $S_{*,a}\phi$ instead of $S_{\mathcal{A},\{a\}}\phi$, to express that all agents learn that agent a knows about formula ϕ.

Next, we define the transition rules. The first transition rule models the fact that agents can encrypt with their own key:

$$B_j m_{\mathcal{K}} \Rightarrow L_j m_{\mathcal{K}+j}; S_{*,j} m_{\mathcal{K}+j}.$$

For simplicity, it is assumed that $j \notin \mathcal{K}$. Thus, if an agent j happens to learn the value of a message encrypted with the keys of the agents in the set $\mathcal{K}$, then agent j can encrypt the message received with its own key added (provided it was not used already). Moreover, as the other agents have seen that agent j has received the message, the agents collectively learn that agent j knows about the result after adding its key, as expressed by $S_{*,j} m_{\mathcal{K}+j}$.

Complementary, we have the transition rule

$$B_j m_{\mathcal{K}} \Rightarrow L_j m_{\mathcal{K}-j}; S_{*,j} m_{\mathcal{K}-j}$$

for every agent $j \in \mathcal{A}$. Now, it is assumed that the agent j is among the agents in $\mathcal{K}$. By commutativity of the cipher, agent j can then (partially) decrypt the message $m_{\mathcal{K}}$ and learn $m_{\mathcal{K}-j}$ whereas all others know that agent j can do this.

In the modeling, we limit ourselves by defining the list of useful propositions. The propositions we want to consider here are $\mathcal{P} = \{m, m_a, m_b, m_{ab}\}$ where m_a abbreviates $\{x\}_{k_a} = [\![\{x\}_{k_a}]\!]$ and m_{ab} abbreviates $\{\{x\}_{k_a}\}_{k_b} = [\![\{\{x\}_{k_a}\}_{k_b}]\!]$. Recall, that $[\![y]\!]$ denotes the *real* value of y, i.e. the value of the expression y in the point of the model.

Next, we must represent the initial knowledge of the agents, i.e. their knowledge before the run of the protocol. We will assume that a is the only agent that knows m and m_a. Furthermore, we will assume that the other agents know this about a. The corresponding Kripke structure is in Figure 8.a. The SRA protocol can be captured by three public announcement $Pub(a, m_a)$, $Pub(b, m_{ab})$ and $Pub(a, m_b)$. We are curious whether at the end of the protocol

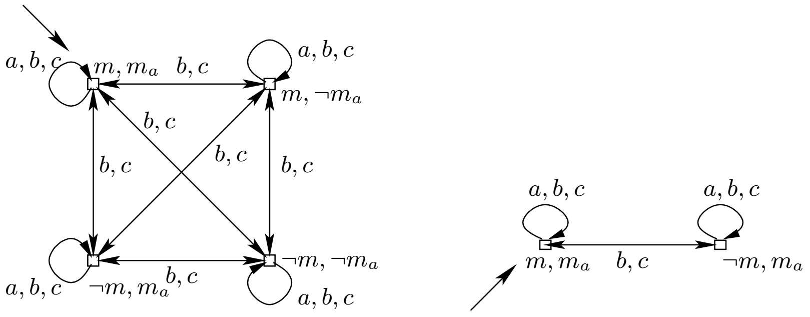

a. starting point   b. after $Pub(a, m_a)$

Figure 8.

(i) agent b will know m;

(ii) agent a will know that agent b knows m;

(iii) agent c only knows that agents a and b know about m.

The first step is executed. That is, m_a is propagated on the network, so all agents will learn its value. So, we execute the action $Pub(a, m_a)$. If we discard the states that become unreachable, this results in the model of Figure 8.b. Note that in this model $B_b m_a$ holds. This the condition of one of the transition rules, that is, it triggers $B_b m_a \Rightarrow L_b m_{ab}; S_{*,b} m_{ab}$ since its antecedent holds in the point now.

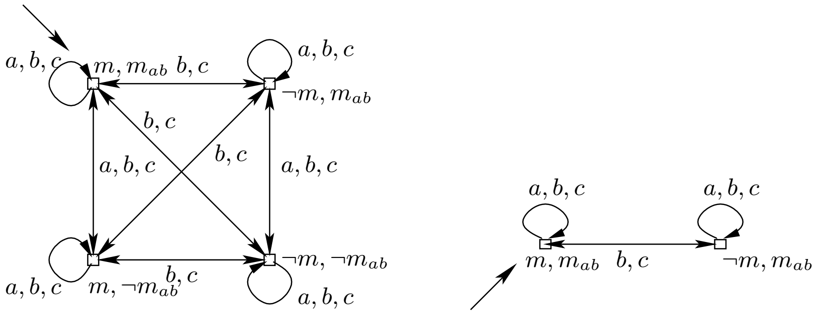

a. added m_{ab} b. after $Pub(b, m_{ab})$)

Figure 9.

For processing the operation $L_b m_{ab}$, we notice that m_{ab} is not modeled yet, so this is the first thing to do. We will not repeat m_a in the figure since this holds in any state of the model. The operation ADDATOM$_{m_{ab}}$ results in the model of Figure 9.a. Instead of applying the body of the transition rule $L_b m_{ab}; S_{*,b} m_{ab}$ to this model, we observe that in the next step of the protocol $L_{\mathcal{A}} m_{ab}$ is executed, as the result of the action $Pub(b, m_{ab})$, since the message is being transmitted to all agents on the network. So, with an appeal to Lemma 5.2, it is justified to skip the operation that are required by the transition rules and perform $L_{\mathcal{A}}m_{ab}$ only. We arrive at the model in Figure 9.b.

In turn, this results in the triggering of the transition rule: $B_a m_{ab} \Rightarrow L_a m_b; S_{*,a} m_b$. This is in fact a completely similar case to the previous step of the protocol. Since the next action of the protocol is $Pub(a, m_b)$ the value of m_b will be learned by all agents, anyway. Again, we dismiss the m_{ab} proposition since every agent has learned this.

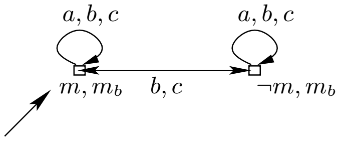

Figure 10. after $Pub(a, m_b)$

We introduce the proposition m_b and execute $Pub(a, m_b)$. So, we end up with the model in Figure 10. The last transition rule that is triggered is $B_b m_b \Rightarrow L_b m; S_{*,b} m$. Again, we discard the proposition that holds in every state, which is m_b, and focus on the most interesting proposition m. First b learns m, as dictated by the operation $L_b m$, which results in the model in Figure 11.a.

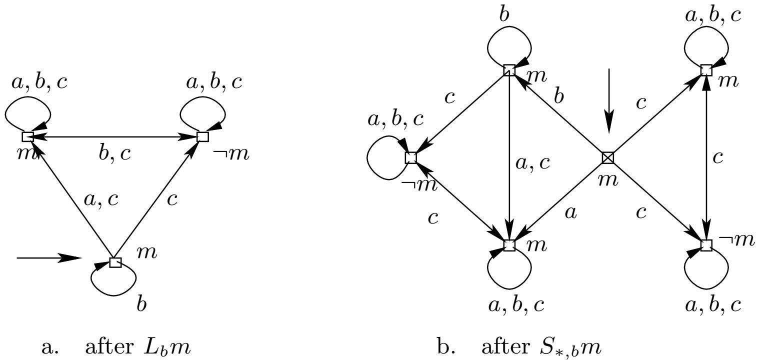

a. after $L_b m$ b. after $S_{*,b} m$

Figure 11.

Next, the second action for the transition rule is that all learn $B_b m \vee B_b \neg m$. If we execute this, we get the model (M', w') which is depicted in Figure 11.b. Recall that this model (M', w') is obtained from the initial model (M, w) by application of the actions $Pub(a, m_a); Pub(b, m_{ab}); Pub(a, m_b)$ and associated transition rules. Moreover, in the resulting model (M', w') it holds that (i) $B_b m$, (ii) $C_{ab} m$, and $\neg(B_c m \vee B_c \neg m)$.

5.2. The Wide-Mouthed Frog protocol

The next example we address to illustrate the update machinery developed above, is the well-known Wide-Mouthed Frog protocol (see, e.g., (Burrows et al., 1990; Abadi and Gordon, 1999)). The protocol exchanges a session key k from the agent a to another agent b via a server s. Then, agent a sends agent b a message protected with the session key k. It is assumed, that the agents a and b share each a symmetric key, k_{as} and k_{bs} say, with the server. The protocol can be described by

1. $a \to s : \{k\}_{k_{as}}$
2. $s \to b : \{k\}_{k_{bs}}$
3. $a \to b : \{m\}_k$

The keys k_{as} and k_{bs} are shared among a and s, and among b and s, respectively. The key k is fresh and initially only known to agent a, as is message m.

In the analysis we want to focus on the session key k and the message m it protects. Therefore the protocol is represented by the sequence of actions

$$\mathit{Priv}(a \to s, k); \mathit{Priv}(s \to b, k); \mathit{Pub}(a, \{m\}_k).$$

Thus, the security of the channel, based on the server keys k_{as} and k_{bs} is expressed by private rather than public communication. We assume that the 'ports' of the channel from a to b can be observed, but the ones for the communication with the server are not visible to other parties.

The initial knowledge is depicted in Figure 12.a. As transition rule we adopt

$$B_b\{m\}_k \Rightarrow L_b m; S_{*,b} m; S_{*,b} k$$

i.e., after b has received the encrypted message $\{m\}_k$ it can learn its content m and everybody learns that b knows about it. Moreover, if agent b is known to know about the message m, then it must know about the session key k as well. Note that, there are no transition rules dealing with the communication with the server.
Execution of the first action $\mathit{Priv}(a \to s, k)$ leads to an update of the knowledge of s. This is represented by a model with six (reachable) states in Figure 12.b. Similarly, but more complicated, the execution of the second action $\mathit{Priv}(s \to b, k)$ induces the model in Figure 13.a. The model gets more involved because, by assumption, the learning of messages exchanged with the server is private. For example, agent a

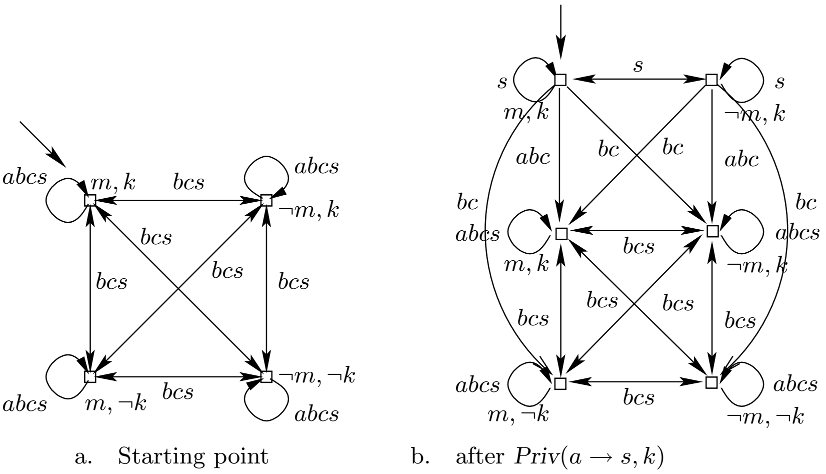

Figure 12.

is not aware of agent b learning about the key k. So far, no transition rules have been triggered.

Next, we execute the last step of the protocol, viz. the public communication $Pub(a, \{m\}_k)$. For this we need to add the atom m_k abbreviating $\{m\}_k = [\![\{m\}_k]\!]$. This doubles the number of states of the models. However, since m_k will be known to all agents, its negative part can be discarded. Now, the transition rules gets activated. So, agent b learns the content m and the other agents learn that b knows about m and k, resulting in the final model in Figure 13.b. Since agent b is learning twice in a row, the difference between Figure 13.a and 13.b are the absence of b-arrows to $\neg m$-states and between states with different values for about m and k.

Typical properties of this model include $B_b(m \wedge k)$, agent b knows the values of the message m and session key k, $\neg C_{ab}m$, m is not commonly known by agents a and b, and, $\neg B_c B_s k$ agent c does not know that the server knows the session key k. That agents a and b do not share the knowledge about the session key, is debatable. One way out is to modify the transition rules and have the operation $L_{ab}m$ instead of L_b.

6. Conclusion

Inspired by recent work on dynamic epistemic logics, we have proposed a logical language for describing (properties of) runs of security protocols. The language contains constructs for the three basic types of epistemic actions that happen during such runs. The semantics of the language is based on traditional Kripke models representing the

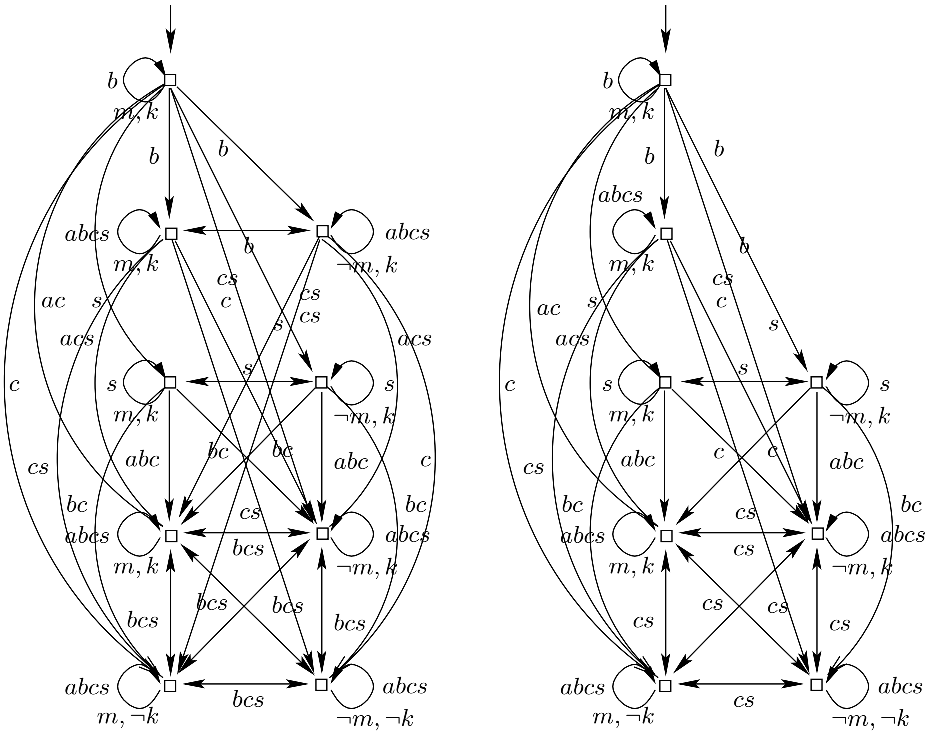

a. after *Priv*$(s \to b, k)$ b. after *Pub*$(a, \{m\}_k)$

Figure 13.

epistemic state of the agents. Changes in the epistemic state of the agent system as a result of the execution of a protocol are described by means of transition rules that precisely indicate what belief updates happen under certain preconditions. These belief updates give rise to modifications of the models representing the agents' epistemic state in a way that is precisely given by semantic operations on these models. We have illustrated our approach for two well-known security protocols, viz. the SRA Three Pass protocol and the Wide-Mouthed Frog protocol.

The semantic updates we used, operate on traditional Kripke models as opposed to updates in the approaches of Gerbrandy and Baltag et al. We believe that this will make it less troublesome to integrate these updates into existing model checkers, which hopefully will lead to better and new tools for verifying properties of security protocols. However, for the development of the theory, it is important to establish the precise connection of the explicit approach followed here and the approach based on action models as advocated in (Baltag et al., 1998; van Ditmarsch, 2000; Baltag, 2002; Baltag and Moss, 2004). A first step into this direction has been presented here, but many others will have to follow. Nevertheless, it points to a promising opportunity to establish a firm relationship between logical theory and security protocol analysis, to the benefit of the latter.

Although future research will have to justify this, we are confident that our method, preferably with some form of computer assistance, can be employed for a broad class of verification problems concerning security protocols because of the flexibility of our approach using transition rules for epistemic updates.

References

Abadi, M. and A. Gordon: 1999, 'A Calculus for Cryptographic Protocols: The Spi Calculus'. *Information and Computation* **148**, 1–70.

Abadi, M. and M. Tuttle: 1991, 'A Semantics for a Logic of Authentication'. In: *Proc. PODC'91*. pp. 201–216, ACM.

Agray, N., W. van der Hoek, and E. de Vink: 2001, 'On BAN logics for industrial security protocols'. In: B. Dunin-Keplicz and E. Nawarecki (eds.): *From Theory to Practice in Multi-Agent Systems*. pp. 29–38, LNAI 2296.

Anderson, R.: 2001, *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley.

Baltag, A.: 2002, 'A Logic for Suspicous Players: Epistemic Actions and Belief-Updates in Games'. *Bulletin of Economic Research* **54**, 1–46.

Baltag, A. and L. Moss: 2004, 'Logics for Epistemic Programs'. *Synthese: Knowledge, Rationality and Action* **139**, 165–224.

Baltag, A., L. Moss, and S. Solecki: 1998, 'The logic of public announcements, common knowledge and private suspicions'. In: I. Gilboa (ed.): *Proc. TARK'98*. pp. 43–56.

Bleeker, A. and L. Meertens: 1997, 'A Semantics for BAN Logic'. In: *Proceedings DIMACS Workshop on Design and Formal Verification of Protocols*. DIMACS, Rutgers University, `http://dimacs.rutgers.edu/Workshops/Security`.

Burrows, M., M. Abadi, and R. Needham: 1990, 'A logic of authentication'. *ACM Transactions on Computer Systems* **8**, 16–36.

Clark, J. and J. Jacob: 1997, 'A Survey of Authentication Protocols 1.0'. Technical report, University of York.

van Ditmarsch, H.: 2000, 'Knowledge games'. Ph.D. thesis, ILLC Dissertation Series 2000-06. University of Groningen.

van Ditmarsch, H.: 2001, 'The semantics of concurrent knowledge actions'. In: M. Pauly and G. Sandu (eds.): *Proc. ESSLLI Workshop on Logic and Games*. Helsinki.

Dolev, D. and A. Yao: 1983, 'On the security of public-key protocols'. *IEEE Transaction on Information Theory* **29**, 198–208.

Gerbrandy, J.: 1997, 'Dynamic Epistemic Logic'. Technical Report LP–97–04, ILLC.

Gerbrandy, J.: 1999, 'Bisimulations on Planet Kripke'. Ph.D. thesis, ILLC Dissertation Series 1999-01. University of Amsterdam.

Hommersom, A.: 2003, 'Reasoning about security'. Master's thesis, Universiteit Utrecht.

Kessler, V. and H. Neumann: 1998, 'A sound logic for analyzing electronic commerce protocols'. In: J.-J. Quisquater, Y. Deswarte, C. Meadows, and D. Gollman (eds.): *Proc. ESORICS'98*. pp. 345–360, LNCS 1485.

Kooi, B.: 2003, 'Knowledge, Chance, and Change'. Ph.D. thesis, ILLC Dissertation Series 2003-01. University of Groningen.

Lowe, G.: 1996, 'Breaking and Fixing the Needham-Schroeder Public-Key Protocol Using FDR'. *Software – Concepts and Tools* **17**, 93–102.

Roorda, J.-W., W. van der Hoek, and J.-J. Meyer: 2002, 'Iterated Belief Change in Multi-Agent Systems'. *Proceedings of the first International Joint Conference on Autonomous Agents and Multi-Agent Systems: Part 2*.

Schneier, B.: 2000, *Secrets and Lies: Digital Security in a Networked World*. Wiley.

Stubblebine, S. and R. Wright: 2002, 'An authentication logic with formal semantics supporting synchronization, revocation and recency'. *IEEE Transactions on Software Engineering* **28**, 256–285.

Wedel, G. and V. Kessler: 1996, 'Formal Semantics for Authentication Logics'. In: E. Bertino, H. Kurth, G. Martello, and E. Montolivo (eds.): *Proc. ESORICS'96*. pp. 219–241, LNCS 1146.