Privacy Regulations in Dynamic Epistemic Deontic Logic*

Guillaume Aucher
University of Luxembourg
guillaume.aucher@uni.lu

Guido Boella
Università di Torino - Italy.
guido@di.unito.it

Leendert van der Torre
University of Luxembourg
leon.vandertorre@uni.lu

Abstract

Privacy policies are often defined in terms of permitted messages. Instead, in this paper we derive dynamically the permitted messages from static privacy policies defined in terms of permitted and obligatory knowledge. With this new approach, we do not have to specify the permissions and prohibitions of all message combinations explicitly. To specify and reason about such privacy policies, we extend a multi-modal logic introduced by Cuppens and Demolombe with update operators modeling the dynamics of both knowledge and privacy policies. We show also how to determine the obligatory messages, how to express epistemic norms, and how to check whether a situation is compliant with respect to a privacy policy.

1. Introduction

Privacy policies are often static and defined in terms of permitted messages, for example in traditional access control languages (2; 4; 9; 14). If policies were instead defined in terms of the permitted and forbidden knowledge of the resulting epistemic state of the recipient of information, then the permitted messages could be derived by combining and reasoning on this knowledge. This raises the following research problem studied in this paper:

> How to formally specify and reason about privacy policies in terms of permitted and forbidden knowledge?

The challenge in this research problem is that the exchange of messages changes the knowledge, and we therefore need a dynamic language which allows us to reason about these changes. Moreover, we impose the following requirements on languages for specifying and reasoning about privacy policies.

We should be able to distinguish between a permission to know and the permission to send a message. For example, you may be permitted to know your medical file, while it may not be permitted that someone not being a doctor sends your medical file. How do such distinctions allow for a more fined-grained account of classical problems of security such as the Chinese wall problem?

We must be able to specify and reason about the *order* in which messages can be sent. For example, it may be permitted to send some sensitive information only if a message has been sent before detailing how to deal with sensitive messages. In many cases it is more efficient or natural to specify that a given piece of information may not be distributed, than explicitly forbidding the different ways of communicating it.

We must be able to specify *obligations* in privacy policies. It might happen that some additional instructions must be sent to the user about the nature of the previous information he received. As (3) notices, privacy laws actually specify which counter measures should apply in case a situation is not compliant with a privacy policy. E.g., if personal information is disclosed inappropriately, the subject of information should be informed.

We must be able to express *meta-security policies*. These are regulations about how to access the regulation itself. For instance, in some applications there is a need for constraints of the form: "agents who play the role r_1 are forbidden to know that agents who play the role r_2 are permitted to know p"; these constraints may be fulfilled using "cover stories" to hide some data (11).

We use modal logic, since both knowledge and obligations and permissions are traditionally and naturally modeled in branches of modal logic called epistemic and deontic logic respectively. This is no new observation in the area of security: Frédéric Cuppens already introduced in 1993 a modal logic for a logical formalization of secrecy (10), and together with Robert Demolombe he developed a logic for reasoning about confidentiality (12) and a modal logical framework for security policies (13). This epistemic deontic logic is the basis of the formalism we introduce in this paper.

The Cuppens-Demolombe logic already got many things right. However, despite the strengths of the Cuppens-Demolombe logic, it is not able to specify or reason about the dynamics of knowledge and privacy policies, and it does not satisfy the four requirements we have posed above. They were ahead of their times, since in 1993 dynamics in modal logic was mainly restricted to propositional dynamic logic for reasoning about programs. In fact the dynamics of knowledge was studied mainly in the AGM paradigm of theory revision (1). In the meantime, much has changed. Dynamic epistemic logic has become a standard branch of modal logic, on which text books have been written (17),

*We thank the anonymous reviewers for helpful comments.

and which is taught at many universities. Our modal logic extends the Cuppens-Demolombe logic with dynamic update operators, to model both the dynamics of knowledge and of privacy policies. As Cuppens and Demolombe, we define privacy policies in terms of knowledge that the recipient of information is permitted/prohibited to have. The way we defined the dynamics of knowledge then allows us to derive the policies on messages. With this new dynamic feature, we can not only determine in a generic way the permitted sequence of messages in a given situation but also determine which change is needed in order to enforce a (possibly new) privacy policy.

2. Our Scenario of Privacy Regulations

In this paper, we consider a single agent (*sender*) communicating information from a knowledge base to another agent (*recipient*), with the effect that the *recipient* knows the information. The *sender* is subject to privacy regulations which restrict what he can communicate to *recipient*. We illustrate the distinction between norms of transmission of information and epistemic norms with an example:

Example 2..1 Consider a *sender* s, e.g., a web server, which is subject to a privacy regulation: he should not communicate the address a of a person to the *recipient* u: we could write this as a norm of transmission of information, regulating the sending of a message: $\neg P(s\ sends\ a)$, which denotes the permission that the *sender* sends message a.

Instead, in an epistemic norm perspective, this prohibition can be derived from the prohibition for the *sender* that the *recipient* comes to know the address: $K_u a$. This is expressed by a deontic operator addressed to the system and having as content the knowledge of the *recipient*: $\neg P_s K_u a$. $\lhd$

This distinction is bridged by modelling sending actions performed by the *sender* which update the knowledge of the *recipient*.

Example 2..2 The message sending action $[s\ sends\ a]$ expresses that the *sender* s is communicating to the *recipient* u the address a. The result of the action is that the *recipient* knows a: $K_u a$. Since $K_u a$ is not permitted by the epistemic norm $\neg P_s K_u a$, the *sender* during his decision process derives that also the action $[s\ sends\ a]$ is not permitted: $\neg P(s\ sends\ a)$. Analogously, all other possible actions leading to the forbidden epistemic state $K_u a$, if any, are prohibited too. E.g., the address is composed by street e, number n and town t: $e \wedge n \wedge t \leftrightarrow a$, thus the sequence of messages $[s\ sends\ e][s\ sends\ n][s\ sends\ t]$ leads to the forbidden epistemic state $K_u a$.

While we need to explicitly model the knowledge of the *recipient* resulting from the message, it is not necessary to have an explicit modality for the *sender*, since we have only one *sender* and we adopt his point of view. So a alone means that the *sender* knows the address.

This explains also why we talk about "knowledge" of the *recipient*: the *sender* never lies, so the result of his actions on the epistemic state of the *recipient* is knowledge rather than belief: $K_u a$ implies a, i.e., that the *sender* holds a as true. If instead we allowed the *sender* to lie to protect some secrets (as, e.g., (6) do), then the result of the action of sending messages would be a mere belief of the *recipient*: the result of $[s\ sends\ a]$ would be that the *recipient* believes a, but a - from the point of view of the *sender* - would not follow from this. $\lhd$

A logical approach to privacy provides a solution to the so-called inference problem: how further permissions propagate from permitted information:

Example 2..3 Assume it is prohibited to know the street where some person lives. Thus, it must be prohibited to know the address of this person. If $e \wedge n \wedge t \leftrightarrow a$, then $\neg P_s K_u e$ implies $\neg P_s K_u a$. Viceversa, if it is permitted to know the address, then it must be permitted to know the street. The same kind of reasoning is transferred at the level of norms of transmission of information: e.g., $\neg P(s\ sends\ e)$ implies $\neg P(s\ sends\ a)$, if it is prohibited to send the name of the street, it is prohibited to send the entire address. $\lhd$

Note that to attribute knowledge to the *recipient*, it is neither necessary to have user profiles nor to have any uncertainty. This stems from the assumption that the *sender* is the only source of information for the *recipient* from the knowledge base. The only knowledge that should be considered is the one derived from the past interaction between the two agents, i.e., the information already disclosed by the *sender*. Assuming for simplicity that the *sender* is rational and sends only information consistent with his previous communicative acts, there is no need of belief revision.

When the forbidden state is achieved by a sequence of messages, there is the possibility that each message of the sequence is permitted while the resulting state is prohibited: this is a new kind of the Chinese wall problem.

Example 2..4 (Website example) Consider the information about websites contacted by a user (the *recipient*), which are available on a server (the *sender*) logfile. The list of websites for each user is clearly a sensitive information which he would not like to disclose. However, knowing which websites have been visited is a valuable information, for example, for the configuration of a firewall, or to make statistics. Thus it has become anonym by replacing the names of the users with numbers by means of a hashcode (h). So even if one knows the list of users one cannot understand who contacted which website. However, from the association between users and numbers and between numbers and websites the original information can be reconstructed. Therefore the mappings from the users to the numbers (c) and from the numbers to the websites (e) can be distributed individually but not altogether since their association would allow to reconstruct the mapping from the users to the websites they visited (v): $\theta = c \wedge e \rightarrow v$

A solution to enforce this privacy policy could be to forbid the distribution of a mapping if the other one has been already distributed, using a language like the one proposed by Barth *et al.* (3), which is able to express policies about the flow of information referring to actions already performed. This solution, however, requires two rules corresponding to the possible permutations of communicative acts.

Moreover, this solution is not general, because there can be further ways of making the forbidden information available. E.g., by distributing the hash function h used. Expressing a flexible policy on all the alternative combinations of actions becomes soon unfeasible. Moreover, new ways of computing the forbidden information could be devised later, which would not be taken into account by the policy.

In this situation we have that it is permitted to know the individual pieces of information, but not what is implied by the conjunction of them:

$$P_s K_u c, P_s K_u e, \neg P_s K_u v$$

It states that it is permitted to 'know' the mapping between users and numbers ($PK_u c$), it is permitted to 'know' the mapping between numbers and websites visited ($PK_u e$) but it is not permitted to 'know' the mapping between users and their websites visited ($\neg PK_u v$).

We have the same situation from the point of view of permissions concerning actions: it is permitted to send the messages c and e individually, but not their combination: $P(s\ sends\ c) \wedge P(s\ sends\ e)$ but $\neg P(s\ sends\ (e \wedge c))$ otherwise the epistemic norm $\neg PK_u v$ would be violated. This means that after sending one of the two messages, the other one becomes prohibited: $[s\ sends\ e]\neg P(s\ sends\ c)$ and $[s\ sends\ c]\neg P(s\ sends\ e)$. $\lhd$

The possibility of nesting formulas with epistemic and deontic modalities allows us to express meta-security, i.e., policies concerning the disclosure of policies, as proposed, e.g., by (6):

Example 2..5 Sometimes, giving the *recipient* the information about the prohibition of sending some information leads him to infer something he should not know. For example, if the *recipient* asks whether a person is a secret agent (p), replying "I cannot tell this to you" to the question makes the *recipient* infer that the person is actually a secret agent, otherwise the answer would have been "no". To avoid this case, it should be prohibited to let the *recipient* know the policy that knowing p is prohibited:

$$\neg P_s K_u \neg P_s K_u p$$

In contrast, if a policy is permitted to be known, it can even be communicated to the *recipient*: if $P_s K_u P_s K_u p$ then it is permitted to send the message $P_s K_u p$: $P(s\ sends\ P_s K_u p)$. This illustrates also that policies can be the content of messages. $\lhd$

3. Privacy policies, compliance and messages

The logic for privacy regulation should reason about epistemic norms, obligations, permissions, knowledge, and information exchange. To deal with these notions altogether, we first extend in Section 3.1 the logic of Cuppens and Demolombe (13) to a more expressive and flexible logic. This logic is actually based on the well-known deontic logic of Castañeda. In Section 3.2, we then add dynamics to the picture. This allows us to have a more fine-grained account of privacy regulations and to solve the research problems that we mentioned in the introduction.

3.1 Static privacy policies

Epistemic Deontic Logic (EDL) Starting from a linguistic analysis, the insight of Castañeda's well known approach to deontic logic is to acknowledge the grammatical duality of expressions depending whether they are within or without the scope of deontic operators (7). We follow this approach and therefore split our language into two kinds of formulas: circumstances and epistemic practitions. The former cannot be (alone) in the scope of an obligation operator O whereas the latter are always within the scope of a deontic operator O. This yields the following language $\mathcal{L}_{EDL} = \mathcal{L}^{\varphi}_{EDL} \cup \mathcal{L}^{\alpha}_{EDL}$ whose formulas are denoted φ^*.

$$\mathcal{L}^{\varphi}_{EDL} : \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_u\varphi \mid O_s\alpha$$
$$\mathcal{L}^{\alpha}_{EDL} : \alpha ::= K'_u\varphi \mid \neg\alpha \mid \alpha \wedge \alpha$$

where p ranges over Φ^{φ}. Formulas of $\mathcal{L}^{\varphi}_{EDL}$ are called circumstances and formulas of $\mathcal{L}^{\alpha}_{EDL}$ are called epistemic practitions. $O_s\alpha$ reads 'it is obligatory for the *sender* that α'. $K_u\varphi$ reads 'the *recipient* knows that φ'. $P_s\alpha$ is an abbreviation for $\neg O\neg\alpha$ and reads 'α is permitted'. Pure circumstances are circumstances without obligation operators $O_s\alpha$.

As it turns out, this language is strictly more expressive than the language of Cuppens and Demolombe (13), even if the semantics is slightly different.

Definition 3..1 An *EDL-model* M is a tuple $M = (W, D, R_u, R'_u, V)$, where W is a non-empty set of possible worlds, $R_u : W \to 2^W$, $R'_u : W \to 2^W$ and $D : W \to 2^W$ are accessibility relations on W, D being serial and R_u, R'_u being reflexive[1]. V is a valuation. The truth conditions are given by:

$M, w \models p^*$	iff	$w \in V(p^*)$
$M, w \models \varphi^* \wedge \psi^*$	iff	$M, w \models \varphi^*$ and $M, w \models \psi^*$
$M, w \models \neg\varphi^*$	iff	not $M, w \models \varphi^*$
$M, w \models O_s\alpha$	iff	for all $v \in D(w)$, $M, v \models \alpha$.
$M, w \models K_u\varphi$	iff	for all $v \in R_u(w)$, $M, v \models \varphi$
$M, w \models K'_u\varphi$	iff	for all $v \in R'_u(w)$, $M, v \models \varphi$

$M \models \varphi$ iff for all $w \in W$, $M, w \models \varphi$. (M, w) is called a pointed EDL-model. If $\mathcal{P}$ is a set of formulas, we write $M, w \models c(\mathcal{P})$ iff $M, w \models \varphi$ for all $\varphi \in \mathcal{P}$. $\lhd$

Obviously, one can map epistemic practitions to circumstances. This mapping $t : \mathcal{L}^{\alpha}_{EDL} \to \mathcal{L}^{\varphi}_{EDL}$ is needed in order to check whether obligations are fulfilled: for example $O_s\alpha \wedge \neg t(\alpha)$ means that we are in a violation state. The mapping function $t : \mathcal{L}^{\alpha}_{EDL} \to \mathcal{L}^{\varphi}_{EDL}$ is defined inductively as follows:

$$\begin{array}{lcl} t(\neg\alpha) & = & \neg t(\alpha) \\ t(\alpha \wedge \alpha') & = & t(\alpha) \wedge t(\alpha') \\ t(K'_u\varphi) & = & K_u\varphi \end{array}$$

[1] An accessibility relation R is reflexive if and only if for all worlds w, $w \in R(w)$. An accessibility relation R is serial if $R(w) \neq \emptyset$ for all worlds w.

Theorem 3..2 *The semantics of $\mathcal{L}_{EDL}$ is* sound and complete *with respect to the logic L_{EDL} axiomatized as follows:*

A_1 *All propositional tautologies based on* Φ^{φ}
A_2 $\vdash O_s\alpha \to P_s\alpha$
A_3 $\vdash K_u\varphi \to \varphi$
A_4 $\vdash O_s(\alpha \to \alpha') \to (O_s\alpha \to O_s\alpha')$
A_5 $\vdash K(\varphi \to \psi) \to (K\varphi \to K\psi)$
R_1 *If* $\vdash \alpha$ *then* $\vdash O_s\alpha$
R_2 *If* $\vdash \varphi$ *then* $\vdash K\varphi$
R_3 *If* $\vdash \varphi^* \to \psi^*$ *and* $\vdash \varphi^*$ *then* $\vdash \psi^*$

where K stands for K_u or K'_u. L_{EDL} is also decidable

Proof. QED

It follows straightforwardly from the Sahlqvist correspondence theorem (5) because Axioms A_2 and A_3 are Sahlqvist formulas. To prove decidability, one can show that L_{EDL} has the finite model property by adapting the selection method (5).

Privacy policies and compliance in EDL As discussed by Barth *et al.* (3) in the theory of contextual integrity, privacy norms are relevant only in some context, usually defined by roles played by *sender* and *recipient*. This leads us to define the following notions.

Definition 3..3 An *epistemic norm* is a formula of the form $\varphi \to O_s\alpha$ or $\varphi' \to P\alpha'$ where φ, φ' are pure circumstances and α, α' are epistemic practitions. A *privacy policy* $\mathcal{P}$ is a consistent set of epistemic norms. We abusively write $\varphi \in \mathcal{P}$ if there is $\varphi \to O_s\alpha \in \mathcal{P}$, and in that case the corresponding α is written α_φ. ◁

Note that permissions concern the knowledge of the *recipient*. This fact should not let the reader think that a privacy policy concerns the behavior of the *recipient*. Indeed, the beliefs of the *recipient* are only modified by actions of the *sender*, so these policies regulate the behavior of the *sender* who might disclose information or not to the *recipient* depending on wether or not this disclosure is in conflict with the privacy policy.

Privacy policies are imposed to the decision maker (*sender*) from a hierarchical superior or set up by himself. They should be enforced in any case. However, this set of epistemic norms is not necessarily complete. As a result, the *sender* can perfectly add other epistemic norms as long as they are consistent with the privacy policy, depending on the particular situation at stake. This leads us to define the following notions of open and closed privacy policies.

Intuitively, an open privacy policy is a policy where only the permissions of the security policies hold, everything else being forbidden. A closed privacy policy is a policy where only the prohibitions of the security policy hold, everything else being permitted. These definitions are similar with the definitions of permissive and restrictive approach of Cuppens and Demolombe (13).

Definition 3..4 Let $\mathcal{P}$ be a privacy policy.

- The privacy policy $\mathcal{P}$ is *open* if for all EDL-model (M, w), if $\mathcal{E}(M, w) \cup \mathcal{P} \nvdash P_s\alpha$, then $M, w \models \neg P_s\alpha$.
- The privacy policy $\mathcal{P}$ is *closed* if for all EDL-model (M, w), if $\mathcal{E}(M, w) \cup \mathcal{P} \nvdash \neg P_s\alpha$, then $M, w \models P_s\alpha$.

where $\mathcal{E}(M, w) = \{\varphi \in \mathcal{L}^{\varphi}_{EL} \mid M, w \models \varphi\}$ represents the epistemic state of the *recipient*. ◁

Note that specifying whether a privacy policy $\mathcal{P}$ is closed or open specifies completely what is permitted and forbidden to know for the *recipient* in the pointed EDL-model (M, w). However, in the general case, the privacy policy $\mathcal{P}$ does not specify all the obligations that should hold in a situation (M, w). This leads us to define two notions of compliance. The first notion of compliance, simply called compliance, just checks wether the obligations $O_s\alpha_\varphi$ strictly following from the privacy policy $\mathcal{P}$ given the epistemic state $\mathcal{E}(M, w)$ are fulfilled. The second notion of compliance, called strong compliance, checks whether *all* the obligations are fulfilled.

Definition 3..5 Let (M, w) be a pointed EDL-model and $\mathcal{P}$ a privacy policy.

- The situation (M, w) is *compliant* with respect to $\mathcal{P}$ if $M, w \models c(\mathcal{P})$ and $M, w \models \varphi \to t(\alpha_\varphi)$ for all $\varphi \in \mathcal{P}$.
- The situation (M, w) is *strongly compliant* with respect to $\mathcal{P}$ if $M, w \models c(\mathcal{P})$ and $M, w \models O_s\alpha \to t(\alpha)$ for all $\alpha \in \mathcal{L}^{\alpha}_{EDL}$.

◁

The following proposition shows that the distinction between compliance and strong compliance is not relevant for closed privacy policies. It also gives a semantic counterpart to the syntactic notion of strong compliance: an epistemic state (represented by $R_u(w)$) is strongly compliant if there exists a corresponding ideal epistemic state (represented by $R'_u(v)$ for some $v \in D(w)$) containing the same information (i.e. R_uD-bisimilar).

Proposition 3..6 *Let (M, w) be a pointed EDL-model and $\mathcal{P}$ a privacy policy.*

- *If $\mathcal{P}$ is closed then (M, w) is compliant w.r.t. $\mathcal{P}$ if and only if (M, w) is strongly compliant w.r.t. $\mathcal{P}$.*
- *The situation (M, w) is strongly compliant w.r.t. $\mathcal{P}$ if and only if there exists $v \in D(w)$ such that $R_u(w)$ and $R'_u(v)$ are R_uD-bisimilar*[2].

Example 3..7 (Website example continued) Consider Example 2.4, where we have the mappings from the users to the numbers (c) and from the numbers to the websites (e), the related mapping from the users to the websites they visited (v) such that $\theta = c \wedge e \to v$.

The epistemic norm solution is to express the *privacy policy* $\mathcal{P}_1$ as:

$$\mathcal{P}_1 = \{P_sK'_uc, P_sK'_ue, \neg P_sK'_uv\}$$

[2]Two pointed models (M, v) and (M', v') are R_uD-bisimilar if there is a relation on $W \times W'$ satisfying the base condition for Φ^{φ} and the back and forth conditions for R_u and D (see Blackburn *et al.* (5) for details). If S is a set of worlds of M and S' a set of worlds of M', S and S' are R_uD-bisimilar if and only if for all $v \in S$ there is $v' \in S'$ such that (M, v) is bisimilar to (M', v'), and vice versa.

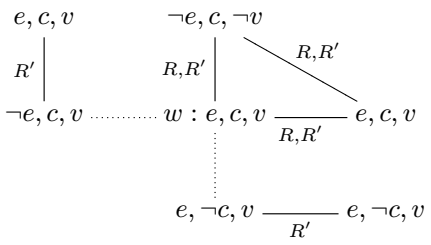

Figure 1: Website example

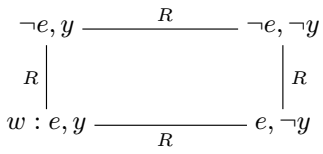

Figure 2: Spyware example

It states that it is permitted to 'know' the mapping between users and numbers ($PK'_u c$), it is permitted to 'know' the mapping between numbers and websites visited ($PK'_u e$) but it is not permitted to 'know' the mapping between users and their websites visited ($\neg PK'_u v$).

The pointed EDL-model (M, w) of Figure 1 represents semantically a situation which is *compliant* with respect to this privacy policy. The accessibility relations R and R' are indexed by R and R' respectively and the accessibility relation D is represented by dashed arrows. Reflexive arrows are omitted, which means that for all worlds $v \in M$ we also have that $v \in R_u(v)$, $v \in R'_u(v)$ and $v \in D(v)$. So we get: $M \models \theta$ ◁

Example 3..8 (Spyware example continued) Consider a situation where the list of websites mentioned is e and the fact that websites might contain risky softwares is y. The privacy policy is expressed by a unique epistemic norm:

$$\mathcal{P}_2 = \{y \wedge K_u e \rightarrow O_s K'_u y\}$$

Note that the condition of this epistemic norm contains an epistemic formula. In Figure 2 is depicted a situation compliant with this privacy policy. In this pointed EDL-model (M, w), the accessibility relation R is indexed by R and reflexive arrows are omitted, which means that for all $v \in M$, we have $v \in R_u(v)$ and $\{v\} = R'_u(v)$, $\{v\} = D(v)$. We do have that the situation is compliant with respect to the privacy policy $\mathcal{P}_2$. ◁

In fact, we can generalize this kind of policies to stronger policies where the *sender* has to inform the *recipient whether* some information has some property or not.

3.2 The dynamic turn

Dynamic Epistemic Deontic Logic (DEDL) We now want to add dynamics to the picture by means of messages sent to the *recipient*. The content of these messages can affect the situation in two ways: either it affects the epistemic realm (represented in a EDL-model by the relation R_u) or it affects the normative realm (represented in a EDL-model by the relations R'_u and D). This leads us to enrich the language $\mathcal{L}_{EDL}$ with two dynamic operators $[s\ sends\ \varphi]$ and $[s\ prom\ \alpha]$, yielding the language $\mathcal{L}_{DEDL}$, whose formulas are denoted φ^*:

$$\mathcal{L}^{\varphi}_{DEDL} : \varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid K_u\varphi \mid O_s\alpha \mid [s\ sends\ \varphi]\varphi \mid [s\ prom\ \varphi]\varphi$$

$$\mathcal{L}^{\alpha}_{DEDL} : \alpha ::= K'_u\varphi \mid \neg\alpha \mid \alpha \wedge \alpha \mid [s\ sends\ \varphi]\alpha \mid [s\ prom\ \alpha]\alpha$$

where p ranges over Φ^{φ}.

$[s\ sends\ \psi]\varphi$ reads 'after the *recipient* learns ψ, φ holds', and $[s\ prom\ \alpha]\varphi$ reads 'after the *sender* promulgates α, φ holds'. The semantics of these dynamic operators is inspired by Kooi (15) and defined as follows.

Intuitively, after learning ψ, the *recipient* restricts his attention to the worlds accessible from the current world which satisfy ψ, unless ψ is not true in this current world. In that case, the message is just ignored. But this second case actually never occurs here because we assume that *sender* only sends truthful messages. Likewise, after the promulgation of α, the ideal worlds are restricted to the worlds which satisfy α, unless the imperative α is not permitted.

Definition 3..9 Let $M = (W, D, R_u, R'_u, V)$ be an EDL-model, $\psi \in \mathcal{L}^{\varphi}_{EDL}$ and $\alpha \in \mathcal{L}^{\alpha}_{EDL}$. We define the EDL-models $M * \psi$ and $M * \alpha$ as follows.

- $M * \psi = (W, D, R^s_u, R'_u, V)$ where for all $w \in W$, $R^s_u(w) = \begin{cases} R_u(w) \cap ||\psi|| & \text{if } M, w \models \psi \\ R_u(w) & \text{otherwise.} \end{cases}$
- $M * \alpha = (W, D^s, R_u, R'_u, V)$ where for all $w \in W$, $D^s(w) = \begin{cases} D(w) \cap ||\alpha|| & \text{if } M, w \models P\alpha \\ D(w) & \text{otherwise.} \end{cases}$

where $||\varphi^*|| = \{v \in M \mid M, v \models \varphi^*\}$. The truth conditions:

$$M, w \models [s\ sends\ \psi]\varphi^* \quad \text{iff} \quad M * \psi, w \models \varphi^*$$
$$M, w \models [s\ prom\ \alpha]\varphi^* \quad \text{iff} \quad M * \alpha, w \models \varphi^*.$$

◁

Permitted and obligatory messages Obviously, given a privacy policy and a situation, some messages might not be permitted by the privacy policy because they might lead to a non-compliant situation.

Definition 3..10 Let $\varphi \in \mathcal{L}^{\varphi}_{DEDL}$, $\mathcal{P}$ be a privacy policy and (M, w) an EDL-model representing a given situation.

- The message φ is *permitted* with respect to $\mathcal{P}$ in (M, w), written $M, w \models P(s\ sends\ \varphi)$, if $(M * \varphi, w)$ is compliant with respect to $\mathcal{P}$.
- A message φ is *obligatory* with respect to $\mathcal{P}$ in (M, w), written $M, w \models O(s\ sends\ \varphi)$, if $M, w \models OK'_u\varphi \wedge \neg K_u\varphi \wedge P(s\ sends\ \varphi)$.

◁

Note also that if a message is obligatory in a situation then this situation is not *strongly* compliant.

Example 3..11 (Website example continued) In Example 3.7, we have:

$$M, w \models P(s \; sends \; c) \wedge P(s \; sends \; e).$$

So it is permitted to send the mappings from the users to the numbers (c) and it is permitted to send the mapping from the numbers to the web-sites (e). However, we also have

$$M, w \models [s \; sends \; e] \neg P(s \; sends \; c) \text{ and}$$

$$M, w \models [s \; sends \; c] \neg P(s \; sends \; e)$$

which means that after sending the mapping from the numbers to the web-sites (e) it is *not* permitted to send the mapping from the users to the numbers (c), and vice versa for the second conjunct. This is because in both cases we would violate the epistemic norm $\neg PK'_u v$:

$$M, w \models [s \; sends \; e][s \; sends \; c](K_u v \wedge \neg P_s K'_u v) \text{ and}$$

$$M, w \models [s \; sends \; c][s \; sends \; e](K_u v \wedge \neg P_s K'_u v).$$

So we have

$$M, w \models \neg P(s \; sends \; (e \wedge c)) \wedge \neg P(s \; sends \; (c \wedge e)).$$

◁

Our approach is very flexible because it is applicable in infinitely many other contexts than the one of the above example, once the privacy policy is fixed. E.g., assume that the hash function computing the mapping from users to numbers is now available (h) and that the *recipient* is able to apply it to get the mapping from numbers to users (c):

$$M \models h \rightarrow c.$$

Applying the same reasoning, we would get:

$$M, w \models [s \; sends \; e] \neg P(s \; sends \; h)$$

$$M, w \models \neg P(s \; sends \; (e \wedge h)).$$

and so without having to introduce explicitly new prohibitions or permissions on h.

Privacy policies do not only concern which information can be disclosed but also which information *should* be disclosed. We can express such policies due to the fact that our epistemic deontic logic can express obligations about knowledge:

Example 3..12 (Spyware Example continued) After sending the message e in the previous situation represented by the EDL-model (M, w) of Figure 2 we obtain the pointed EDL-model $(M * e, w)$ depicted in Figure 3. The corresponding situation $(M * e, w)$ is not compliant with respect to $\mathcal{P}'$. Therefore, it was forbidden to disclose e:

$$M, w \models \neg P(s \; sends \; e)$$

But it is now obligatory (with respect to $\mathcal{P}'$) to disclose s:

$$M * e, w \models O(s \; sends \; y)$$

So we have that

$$M, w \models [s \; sends \; e] O(s \; sends \; y)$$

$$M, w \models \neg P(s \; sends \; e) \wedge P(s \; sends \; (e \wedge y))$$

As it turns out, after sending the message y we reach a compliant situation. ◁

$$w : e, y \underset{R}{\longrightarrow} e, \neg y$$

Figure 3: Spyware example updated

The above example suggests that even if it is prohibited to send message e, it might still be permitted to send message e as long as it is followed by another message y. We leave the investigation of the permissibility of iterative messages for future work.

In privacy regulations, the permission to disclose the names of users also allows to disclose their family names (which are part of their name). This problem, discussed in Example 2.3, is known as the inference problem, and is in general difficult to model (see for instance Barth *et al.* (3)). In our logical framework it follows easily from the fact that the *recipient* has reasoning capabilities. Indeed, if we assume that the conditions of the epistemic norms of the privacy policy $\mathcal{P}$ are propositional then for all $\varphi, \varphi' \in \mathcal{L}^{\varphi}_{DEDL}$,

$$\varphi \rightarrow \varphi' \models^g P(s \; sends \; \varphi) \rightarrow P(s \; sends \; \varphi')$$

where $\models^g$ is the global consequence relation.

Example 3..13 (Website example continued) Assume we have a situation modeled by an EDL-model M such that $M \models v \rightarrow v'$: the association between the users' name and the web-sites they visited (v) induces the association between the users' *family* name and the web-sites they visited (v'). So if $M, w \models P(s \; sends \; v)$ then $M, w \models P(s \; sends \; v)'$: if it is permitted to disclose the name of the users in association with the websites they visited, it is also permitted to disclose their family name in association with the web-sites they visited. Dually, if $M, w \models v \rightarrow v'$, then $M \models \neg P(s \; sends \; v')$ implies $M, w \models \neg P(s \; sends \; v)$: if it is prohibited to disclose their family names in association with the web-sites they visited then it is also prohibited to disclose their names in association with the web-sites they visited. ◁

We have another interesting property connecting the notions of permitted and obligatory communicative acts. Let $\varphi, \varphi' \in \mathcal{L}^{\varphi}_{DEDL}$:

If $\vdash \varphi' \rightarrow \varphi$ then $\vdash O(s \; sends \; \varphi') \rightarrow \neg P(s \; sends \; \neg\varphi)$

This proposition states that if it is obligatory to disclose a fact then it is prohibited to disclose the opposite of any of its logical consequences. However, note that $O(s \; sends \; \varphi)$ and $P(s \; sends \; \varphi)$ are not dual operators:

$$\nvdash O(s \; sends \; \varphi) \leftrightarrow \neg P(s \; sends \; \neg\varphi).$$

This is intuitively correct: in Example 3.12 it is prohibited to disclose e but it does not entail that it is obligatory to disclose $\neg e$. Moreover, we know by the 'Web-site Example' that we have the following property:

$$\nvdash P(s \; sends \; \varphi) \wedge P(s \; sends \; \psi) \rightarrow P(s \; sends \; (\varphi \wedge \psi)).$$

Indeed, in Example 3.11 we have that $M, w \models P(s \; sends \; e) \wedge P(s \; sends \; c) \wedge \neg P(s \; sends \; (e \wedge c))$.

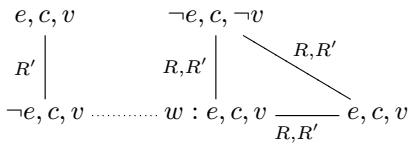

Figure 4: Website example

Enforcing privacy policies: $[s\ prom\ \varphi]$ The hierarchical superior of the *sender* or the *sender* himself might decide to change the policy privacy from $\mathcal{P}$ to $\mathcal{P}'$. As a result, the sender needs to enforce this new privacy policy $\mathcal{P}'$. This enforcement is captured in our formalism by $[s\ prom\ \psi]$.

Example 3..14 (Web-site Example) In case of attack by some hacker, the privacy policies can be made more strict. For example, the *sender* can decide to strengthen the privacy policy $\mathcal{P}_1$ of Example 3.7 to

$$\mathcal{P}_4 = \{P_s K'_u c, \neg P_s K'_u e, \neg P_s K'_u v\}$$

where $P_s K'_u e$ has been replaced by $\neg P_s K'_u e$: it is now prohibited to disclose the mapping from numbers to visited web-sites. This new privacy policy $\mathcal{P}_4$ can be enforced by the *sender* through the update $[s\ prom\ \neg K'_u e]$. We get the EDL-model $(M * \neg K'_u e, w)$ depicted in Figure 4 which is compliant with respect to $\mathcal{P}_4$. $\lhd$

4. Concluding remarks

Cuppens and Demolombe (12) extend the original framework (10) by using an epistemic deontic logic to model security in databases. They do not introduce the dynamics of their system, neither for beliefs nor for obligations, even if they recognize the importance of this dimension. We share many properties of their epistemic-deontic modalities, but we also extend them to permissions and obligations concerning actions and not only propositions, getting a more fine grained analysis, for example of the Chinese wall problem. Moreover, they do not introduce separately the epistemic and deontic operators but only combined ones, like (6) do, limiting the expressivity of the logic. Given the ability to nest epistemic and deontic operators we are able to model more complex formulas like those for meta-security or obligations to know whether something holds. Given that our approach is based on their approach, their solutions to several problems can naturally be transferred in our setting. They show for example that multi-level security policies which assign a degree of clearance l to formulae φ and which might be incomplete can be expressed in their framework by indexing the modality $PK_u\varphi$ with the degree of clearance l: $PK_{ul}\varphi$ reads 'an agent u cleared at level l is explicitly permitted to know that the database believes φ'.

We introduced a multi-modal logic to formally specify and reason about privacy policies in terms of permitted and forbidden knowledge. The logic satisfies the four requirements we gave in the introduction. A topic for further research is to deal with multi-agent scenarios involving more agents than just a *sender* and a *recipient*.

References

[1] C. Alchourrón and P. Gärdenfors and D. Makinson. On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic logic*, 50(2):510–530, 1985.

[2] A. Anderson *et al.* Extensible access control markup language (XACML) version 2.0. 2004.

[3] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum. Privacy and contextual integrity: Framework and applications. In *19th IEEE Symposium on Security and Privacy*, pages 184–198. IEEE Computer Society, 2006.

[4] M. Bishop. *Computer Security: Art and Science*. Addison Wesley Professional, 2003.

[5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Computer Science*. Cambridge University Press, 2001.

[6] P. Bonatti, S. Kraus and V. Subrahmanian. Foundations of Secure Deductive Databases. *IEEE Transactions on Knowledge Data and Engineering*, 7(3):406–422, 1995.

[7] H.-N. Castañeda. *The paradoxes of Deontic Logic: the simplest solution to all of them in one fell swoop*, pages 37–86. Synthese library. 1981.

[8] H.-N. Castañeda. Knowledge and epistemic obligation. *Philosophical perspectives*, 2:211–233, 1988.

[9] L. Cranor. *Web Privacy with P3P*. O'Reilly and Associates Inc., 2002.

[10] F. Cuppens. A Logical Formalization of Secrecy. In *6th IEEE Computer Security Foundations Workshop - CSFW'93*. IEEE Computer Society, 1993.

[11] F. Cuppens and R. Demolombe. Normative Conflicts in a Confidentiality Policy. In ECAI Workshop on Artificial Normative Reasoning. 1994.

[12] F. Cuppens and R. Demolombe. A Deontic Logic for Reasoning about Confidentiality. In Deontic Logic, Agency and Normative Systems, DEON '96: Third International Workshop on Deontic Logic in Computer Science. Springer, 1996.

[13] F. Cuppens and R. Demolombe. A Modal Logical Framework for Security Policies. In *Foundations of Intelligent Systems, 10th International Symposium, ISMIS '97*. Springer, pages 579-589, 1997.

[14] G. Karjoth and M. Schunter. A privacy policy model for enterprises. In *15th IEEE Computer Security Foundations Workshop*. IEEE Computer Society, 2002.

[15] B. Kooi. Probabilistic dynamic epistemic logic. *Journal of Logic, Language and Information*, 12(4):381–408, 2003.

[16] E. Pacuit and R. Parikh. The logic of knowledge based obligation. *Synthese*, 149(2), 2006.

[17] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*, volume 337 of *Synthese library*. Springer, 2007.