

A deontic logic for reasoning about confidentiality

Frédéric Cuppens and Robert Demolombe
ONERA-CERT, 2 Av. E. Belin
31055, Toulouse Cedex, France
Fax: + 33 62 25 25 93
email: {cuppens,demolomb}@tls-cs.cert.fr

Abstract

This paper presents a deontic logic Σ for reasoning about permission or prohibition to know some parts of the database content in the context of a multilevel confidentiality policy.

The most important logical features in the definition of a multilevel policy are that each confidentiality level is defined by a set of sentences and that, when the policy is designed, the permission to know is not necessarily the complement of the prohibition to know. These concepts are formalized in a modal logic where deontic modalities, doxastic modalities and confidentiality levels are interpreted by non-standard modal models. The corresponding axiomatics is also presented in the paper and its soundness and completeness have been proved. A limitation of the Σ logic is that sentences in the scope of modalities are sentences of Propositional Calculus.

Finally, it is shown how the logic can be used to express constraints to guarantee the consistency of a policy or to prevent the existence of inference channels. That is, the possibility to infer sentences that are not permitted to know from other sentences that are permitted to know. Both deductive and abductive channels are considered.

Keywords: knowledge representation, deontic logic, database security.

1 Introduction

To guarantee confidentiality with regard to accesses to a database or, more generally, to an information system, a standard approach is to define multilevel security policies. A given multilevel policy assigns a clearance level to users and a classification level to sentences of the language used to represent the database content. This assignment is not necessarily complete for sentences. There may be sentences whose level is not defined.

With a multilevel policy a user is permitted to access sentences stored in the database whose classification is lower than or equal to his clearance and he is forbidden to access sentences whose classification is greater than his clearance. One may immediately notice that, as the permission and prohibition of a given user is derived from two independently defined sets – namely, the sentences which are respectively lower than or equal to user’s clearance and the sentences greater than user’s clearance –, the permission is not necessarily the complement of prohibition. Moreover, there are sentences which are not assigned a classification. This means that the policy may be incomplete in the sense that there are sentences a given user may be neither permitted nor prohibited to know.

As a matter of simplification, we only consider here two security levels: the “confidential” level and the “secret” level. We consider that the assignment of a classification is defined independently of the database content and independently of the truth value value of sentences in a given situation. Therefore, it is, in one sense, syntactical. Also, the set of sentences to which a given classification is assigned should not be confused with the set of sentences actually stored in the database. For instance, the set of confidential sentences may be: $\{p, \neg p, q\}$, and the set of confidential sentences stored in the database may be: $\{p\}$ at a given time and $\{\neg p, q\}$ at another time.

In this context there are several issues related to confidentiality. The first is to find a formal definition for the “consistency” of a confidentiality policy. A confidentiality policy is consistent iff we cannot infer from this policy that there exists a sentence a given user is both permitted and prohibited to know. We can very briefly explain the reasons why the formalization of this definition is not a trivial problem. The two sets of sentences classified respectively at level confidential or secret, are intended to characterize the set of sentences a confidential user u is permitted to know or is forbidden to know. The user u is permitted to know all the sentences that are in the set of confidential sentences. He is also permitted to know the consequences of **each** consistent subset of this set but he does not necessarily have the permission to know, in a given situation, the conjunction of several sentences that are consequences of distinct subsets. For instance, if the set of confidential sentences is: $\{p, \neg p\}$, we can infer, for every sentence q , that user u has the permission to know $p \vee q$ and also to know $\neg p \vee q$. However, u does not have the permission to know $(p \vee q) \wedge (\neg p \vee q)$. Indeed, if this were the case, he would also be permitted to know the consequence q of $(p \vee q) \wedge (\neg p \vee q)$, whatever the sentence q may be. Reasoning about the consequences of the fact that a user is forbidden to know a sentence p is slightly less intuitive, as will be shown in the rest of the paper. Thus we can easily imagine that consistency checking of a confidentiality policy requires very clear definitions of the concepts of permission and prohibition to know and the associated inference rules have to be carefully defined.

The second issue is to characterize the “inference problem”. There is an inference problem, for a given confidentiality policy and a given database state, if it is possible that a confidential user can infer from a set of sentences he has retrieved from the database a sentence he is not permitted to know. This problem can be refined into two sub-problems. In the first, we consider inferences based on deductive reasoning. In this case we say that there exists a “deductive channel”. In the second, we consider abductive reasoning and say that there exists an “abductive channel”. By abductive reasoning, we mean the kind of

reasoning that allows the confidential user to deduce, by using an additional missing sentence that this user may abduce, a secret sentence. We want to be able to characterize, in our logical framework, the constraints that prevent the existence of deductive or abductive channels.

Most current research work (for instance [1, 2]) is restricted to the detection of deductive inference channels. There are few publications which are interested in abductive channels [3, 4]. Unfortunately, a formal definition of what is an abductive channel is not provided.

Moreover, in all these publications, the multilevel database is identified with a set of atomic facts and it is generally assumed that the classification procedure only applies to the database content, i.e., a given data may be assigned a classification only if it is a fact explicitly stored in the database. This leads to two important restrictions:

- first, it is not possible to specify, for a given sentence p_1 that both p_1 and $\neg p_1$ may be assigned a (possibly different) classification;
- second, it is not possible to specify that not only a given piece of information p_1 is classified as secret but also the disjunction $p_1 \vee p_2$ for some sentence p_2 . Therefore, it is not possible to represent a situation where not only the “exact” disclosure of p_1 is forbidden but also a “partial” disclosure of p_1 , represented by the formula $p_1 \vee p_2$, is to be prevented.

The next section presents the formal definition of the logic. The main reason we did not adopt a standard deontic logic and a standard doxastic logic is the necessity of dealing with complete representations of the regulation – in the sense that the regulation defines all the sentences and *only* the sentences that are permitted or prohibited to know – and complete sets of agent knowledge – i.e sets used to represent all the sentences and *only* the sentences that an agent know (see [5] for similar motivations). Another reason is that, the concept of permission is not defined as the complement of prohibition. This enables two different attitudes to be modeled [6]. In the first attitude, the multilevel database should only provide the user with data this user is explicitly permitted to know. We call **restrictive** this first attitude. With respect to the second attitude, the multilevel database should only provide the user with data he is not explicitly prohibited to know. We call **permissive** this second attitude. Finally, section 3 presents examples of practical application of the logic to regulation consistency and to inference channels.

2 Formal definition of the Σ logic

2.1 Basic Assumptions and Notations

We assume that the database content is represented by a consistent set db of sentences of a language L of Propositional Calculus. It is not assumed that sentences in db are true in the world. This means that db is considered a set of beliefs. By querying the database db , the users will know that the database believes some sentence. Therefore, the fact that a user u knows that db believes some sentence is represented by the modality KB_u , and the fact that u knows that db believes the sentence p is represented by $KB_u(p)$, $KB_u p$ for short. The

modality KB_u can be viewed as the combination of an epistemic modality and a doxastic modality. This means that the modality KB_u can be formalized by the KD logic [7].

Let us now turn to the representation of a multilevel confidentiality policy. In this case, each user is associated with a clearance level and some sentences of the propositional language L are associated with a classification level. If we consider a user u cleared at a given level i , it is possible to associate this user with two sets of sentences: the set of sentences whose classification level l is such that $l \leq i$ and the set of sentences whose classification level l is such that $l > i$. These two sets are used to define the set of sentences a user cleared at level i is permitted or prohibited to know.

For the sake of simplicity, we only consider two security levels. Level 1 is called “confidential” and level 2 is called “secret”. The set of sentences such that $l \leq 1$ is called C and the set of sentences such that $l > 1$ is called S . We also only consider the case of a user cleared at confidential (level 1). Generalization to more than two security levels is trivial; we have simply to replace level 1 by level i .

It is assumed that some sentences may not be classified, that is, we do not necessarily have: $C \cup S = L$. These two sets are formally represented in the language by two sentences of the form: $[C](p_1, \dots, p_m)$, and $[S](p'_1, \dots, p'_n)$, where $\{p_1, \dots, p_m\}$ (resp. $\{p'_1, \dots, p'_n\}$) is the set of **all** the sentences classified at the confidential (resp. secret) level.

It is worth noting that for sentences that are neither in C nor in S , it is not explicitly defined whether a user cleared at confidential is permitted or prohibited to know these sentences. In some sense, the regulation is not complete. However, when a confidential user is querying the multilevel database, we need to complete the regulation. This is because the multilevel database has to decide whether or not the answer to the query has to be delivered to the user. We can actually propose two different attitudes to complete the regulation.

The first attitude is to consider that a user cleared at confidential is only permitted to know sentences belonging to C . This attitude can be qualified as restrictive because it means that a confidential user, when he is querying a multilevel database should only be provided with an answer he is explicitly permitted to know.

The second attitude is to consider that this user is only forbidden to know sentences belonging to S . This second attitude can be qualified as permissive because it means that a confidential user, when he is querying a multilevel database, should only be provided with an answer he is not explicitly forbidden to know.

We consider that it is the responsibility of the security administrator to choose one of these two attitudes. For this purpose, our model enables these two attitudes to be formally represented. That is why, in the model, the concepts of permission and prohibition are **independent** in the sense that what is permitted is not necessarily the complement of what is forbidden. These two concepts are formalized by the two modalities PKB_c and FKB_c . The fact that some confidential users are permitted to know that the database believes p is represented by $PKB_c(p)$, PKB_cp for short, and the fact that it is forbidden for some confidential users to know that the database believes p is represented by

$\text{FKB}_c(p)$, FKB_{cp} for short.¹

2.1.1 Language Definition

Let L be a Propositional Calculus language defined on a finite set VAR of propositional variables.

Let L' be a language defined from L as follows.

- If $p \in L$ then $p \in L'$.
- If $p_1 \in L, \dots, p_n \in L$ then $[C](p_1, \dots, p_n) \in L'$ and $[S](p_1, \dots, p_n) \in L'$.
- If $p \in L$ then $\Box p \in L'$, $\text{KB}_u(p) \in L'$, $\text{PKB}_c(p) \in L'$ and $\text{FKB}_c(p) \in L'$.
- If $p \in L'$ then $\neg p \in L'$.
- If $p \in L'$ and $q \in L'$ then $p \vee q \in L'$.

There is no other sentence in L' than the sentences defined by the above rules. As usual conjunction and implication are defined from negation and disjunction.

2.1.2 Semantics

Let Σ be the logic we are defining.

A structure M of Σ is a tuple $M = \langle W, R_u, D_{cr}, D_{cp}, R_{xcr}, R_{xcp}, C, S, P \rangle$, where:

- W is a set of worlds.
- P is a function from VAR to 2^W .
- $R_u, D_{cr}, D_{cp}, R_{xcr}$ and R_{xcp} are relations on $W \times W$.
- C and S are two functions from W to 2^{2^W} .

M is a model of Σ iff it satisfies the following constraints:

C1: if p is a satisfiable² sentence of L then $P(p) \neq \emptyset$.

C2: relations R_u, R_{xcr} and R_{xcp} are serial.

We shall use the notation $\text{OB}_{xcr}(p)$, $\text{OB}_{xcp}(p)$, $\text{DB}(C(w))$ and $\text{DB}(S(w))$ which are formally defined by:

- $M, w' \models \text{OB}_{xcr}(p)$ iff $\forall w''(w'R_{xcr}w'' \text{ iff } M, w'' \models p)$
- $M, w' \models \text{OB}_{xcp}(p)$ iff $\forall w''(w'R_{xcp}w'' \text{ iff } M, w'' \models p)$
- $\|p\|$ will denote the set of worlds $\{w : M, w \models p\}$

¹In what follows, we consider only what a confidential user is permitted or prohibited to know. We shall ignore obligations with regard to a secret user.

²Here “satisfiable” means satisfiable in Classical Propositional Calculus (CPC).

- $DB(C(w)) = \{ X : \exists X_1 \dots \exists X_m$
 $(X_1 \in C(w) \text{ and } \dots \text{ and } X_m \in C(w) \text{ and}$
 $X = X_1 \cap \dots \cap X_m \text{ and } X \neq \emptyset)\}$

- $DB(S(w)) = \{ X : X \subseteq W \text{ and } \forall Y (Y \in S(w) \text{ implies } X - Y \neq \emptyset)\}$

C3: $\forall p (\exists w' (M, w' \models OB_{xcr}(p)) \text{ if } \exists w (\|p\| \in DB(C(w)))$

C4: $\forall w \forall w' (wD_{cp} w' \text{ iff } \exists p (\|p\| \in DB(C(w)) \text{ and } M, w' \models OB_{xcr}(p))$

C5: $\forall p (\exists w' (M, w' \models OB_{xcp}(p)) \text{ if } \exists w (\|p\| \in DB(S(w)))$

C6: $\forall w \forall w' (wD_{cp} w' \text{ iff } \exists p (\|p\| \in DB(S(w)) \text{ and } M, w' \models OB_{xcp}(p))$

Comments:

These constraints may seem to be unnecessarily complicated. The following comments are intended to explain their intuitive meaning and their justifications.

Functions C and S define the regulation in each world, in other words the set of sentences which are respectively assigned the confidential level or the secret level. Each sentence p is represented by its truth set $\|p\|$.

A consequence of constraint C1 is that there is at least one world in W for each possible truth value assignment to propositional variables, in other words for each possible state of the world.

Relation R_u represents, for each state of the world, the confidential user's knowledge about the database beliefs.

Each set of worlds X_i in $C(w)$ can be represented by a sentence p_i such that $\|p_i\| = X_i$, and each set of worlds X in $DB(C(w))$ can be represented by a sentence p such that $\|p\| = X$ and there exists a consistent subset of the p_i s that implies p .

The elements of $DB(C(w))$, and the corresponding sentences p , can be interpreted as a set of mental states of a user who is only permitted to know the sentences defined by $C(w)$ or their consequences. In this case permission is defined according to a restrictive attitude, since what is prohibited is the complement of what is explicitly permitted.

Each set of worlds X_j in $S(w)$ can be represented by a sentence p_j such that $\|p_j\| = X_j$, and each set of worlds X in $DB(S(w))$ can be represented by a sentence p such that $\|p\| = X$ and p implies none of the p_j s.

The elements of $DB(S(w))$, and the corresponding sentences p , can be interpreted as a set of mental states of a user who is **only** permitted to know any set of sentences that does not allow to infer any sentence defined by $S(w)$. In this case permission is defined according to a permissive attitude, since what is permitted is the complement of what is explicitly prohibited.

Our deontic is then based on the concept of role [8, 9]. Intuitively, each individual is associated with a set of roles which represents the behavior the individual is playing in a given situation. Each role defines the permissions, obligations and prohibitions laid upon the role-holder. In our model of a multilevel confidentiality policy, a confidential user is associated with two different roles: the role cr (resp. cp) which represents the behavior of a user cleared at confidential with respect to the restrictive (resp. permissive) attitude.

The deontic accessibility relation D_{cr} (resp. D_{cp}) characterizes the set of ideal [10, 11] mental states of a user xcr (resp. xcp) who plays the role of a user cleared at confidential when a restrictive (resp. permissive) attitude is adopted. Accessibility relation R_{xcr} (resp. R_{xcp}) characterizes xcr 's beliefs (resp. xcp 's beliefs).

Constraint C3 imposes that for each element in $DB(C(w))$ there exists a world where xcr believes and **only** believes a sentence p that corresponds to this element. Constraint C4 imposes that the worlds accessible via D_{cr} are the worlds where xcr 's mental states are exactly those ones represented by the elements in $DB(C(w))$ (see figure 1). That is the xcr 's permitted mental states and elements in $DB(C(w))$ are in a one to one correspondence.

Constraint C5 imposes that for each element in $DB(S(w))$ there exists a world where xcp believes and only believes a sentence p that corresponds to this element. Constraint C6 imposes that the worlds accessible via D_{cp} are the worlds where xcp 's mental states are exactly those ones represented by the elements in $DB(S(w))$ (see figure 2). That is the xcp 's permitted mental states and elements in $DB(C(w))$ are in a one to one correspondence.

Satisfiability Conditions

$M, w \models p$ means that the sentence p of L' is true in the world w of the model M . The satisfiability conditions are:

$$\begin{aligned}
& \text{If } p \in \text{VAR} : M, w \models p \text{ iff } w \in P(p). \\
& M, w \models \neg p \text{ iff } M, w \not\models p. \\
& M, w \models p \vee q \text{ iff } M, w \models p \text{ or } M, w \models q. \\
& M, w \models \Box p \text{ iff } \forall w' \in W \ M, w' \models p. \\
& M, w \models [C](p_1, \dots, p_n) \text{ iff } C(w) = \{\|p_1\|, \dots, \|p_n\|\}. \\
& M, w \models [S](p_1, \dots, p_n) \text{ iff } S(w) = \{\|p_1\|, \dots, \|p_n\|\}. \\
& M, w \models KB_{up} \text{ iff } \forall w' (wR_u w' \Rightarrow M, w' \models p). \\
& M, w \models PKB_{cp} \text{ iff } \exists w' (wD_{cp} w' \text{ and } M, w' \models KB_{xcr} p). \\
& M, w \models FKB_{cp} \text{ iff } \forall w' (wD_{cp} w' \Rightarrow M, w' \models \neg KB_{xcp} p).
\end{aligned}$$

where:

$M, w' \models KB_{xcr} p$ is an abbreviation for: $\forall w'' (w'R_{xcr} w'' \Rightarrow M, w'' \models p)$, and
 $M, w' \models \neg KB_{xcp} p$ is an abbreviation for: $\exists w'' (w'R_{xcp} w'' \text{ and } M, w'' \models \neg p)$.




Figure 1: Permission to know.




Figure 2: Prohibition to know.

Valid Sentences

$M \models p$ denotes the fact that p is true in the model M . We have:

$$M \models p \text{ iff } \forall w(M, w \models p).$$

$\models p$ denotes the fact that p is a valid sentence. We have:

$$\models p \text{ iff } \forall M(M \models p).$$

Properties of the Σ Models

Property 1: If $\|p\| \in C(w)$ then $M, w \models \text{PKB}_c p$.

Property 2: If $\|p\| \in S(w)$ then $M, w \models \text{FKB}_c p$.

Property 3: $\models p \rightarrow q \Rightarrow \models \text{PKB}_c p \rightarrow \text{PKB}_c q$.

Property 4: $\models p \rightarrow q \Rightarrow \models \text{FKB}_c p \rightarrow \text{FKB}_c q$.

Properties 1 and 2 show that the definitions of PKB_c and FKB_c work out as they should for sentences represented by $C(w)$ and by $S(w)$. Properties 3 and 4 show that the set of sentences a confidential user is permitted to know is extended by their logical consequences whereas the set of sentences that a confidential user is forbidden to know is extended by their logical implicants.

We have the following corollaries of properties 3 and 4:

Corollary 1: $\models \text{PKB}_c(p \wedge q) \rightarrow \text{PKB}_c p \wedge \text{PKB}_c q$.

Corollary 2: $\models \text{PKB}_c p \rightarrow \text{PKB}_c(p \vee q)$.

Corollary 3: $\models \text{FKB}_c p \rightarrow \text{FKB}_c(p \wedge q)$.

Corollary 4: $\models \text{FKB}_c(p \vee q) \rightarrow \text{FKB}_c p \wedge \text{FKB}_c q$.

Corollary 5: $\models p \leftrightarrow q \Rightarrow \models \text{PKB}_c p \leftrightarrow \text{PKB}_c q$.

Corollary 6: $\models p \leftrightarrow q \Rightarrow \models \text{FKB}_c p \leftrightarrow \text{FKB}_c q$.

Corollaries 5 and 6 show that in the definition of a regulation, formally represented by $[C](p_1, \dots, p_n)$ and $[S](p'_1, \dots, p'_n)$, each sentence p_i can be replaced by any logically equivalent sentence.

Finally, we have the following properties.

Property 5: $\not\models \text{PKB}_c p \wedge \text{PKB}_c q \rightarrow \text{PKB}_c(p \wedge q)$.

Property 6: $\not\models \text{FKB}_c(p \wedge q) \rightarrow (\text{FKB}_c p \vee \text{FKB}_c q)$.

The intuitive meaning of property 5 is that it may happen that p can be derived from a given permitted mental state of xcr and that q can be derived from another permitted mental state of xcr , whereas there is no permitted mental state of xcr that allows us to derive $p \wedge q$.

For instance, if, in a model M , we have: $C(w) = \{\|a\|, \|a \rightarrow p\|, \|\neg a\|, \|\neg a \rightarrow q\|\}$, then we have $M, w \models \text{PKB}_c p$, because there exists a world w' accessible from w via D_{cr} where we have $M, w' \models \text{OB}_{xcr}(a \wedge (a \rightarrow p))$ and we have $M, w \models \text{PKB}_c q$, because there exists a world w' accessible from w via D_{cr} where we have $M, w' \models \text{OB}_{xcr}(\neg a \wedge (\neg a \rightarrow p))$. However, we do not have $M, w \models \text{PKB}_c(p \wedge q)$ because there exists no world w' accessible from w via D_{cr} where we have $M, w' \models \text{OB}_{xcr}(a \wedge (a \rightarrow p) \wedge \neg a \wedge (\neg a \rightarrow p))$ because $a \wedge \neg a$ is inconsistent.

The intuitive meaning of property 6 is that it is possible that a user u , cleared at level c , is forbidden to know $p \wedge q$, whereas u is neither forbidden to know p nor forbidden to know q . In fact, it may happen that neither p nor q independently allows the derivation of a secret sentence, whereas their conjunction allows to derive it.

For instance, if, in a model M , we have: $S(w) = \{\|a \wedge b\|\}$, then we have $M, w \models \text{FKB}_c(a \wedge b)$, whereas we have neither $M, w \models \text{FKB}_c(a)$ nor $M, w \models \text{FKB}_c(b)$. This corresponds to a ‘‘Chinese wall’’ type of regulation [12], where the conjunction of two sentences may have a classification level strictly greater than each term in the conjunction.

2.1.3 Axiomatics

The axiomatics of the Σ logic is defined by the following axiom schemas and inference rules.

- Axiom schemas of CPC plus Modus Ponens.

- For the modality KB_u we have:

$$(K) \text{KB}_u p \wedge \text{KB}_u(p \rightarrow q) \rightarrow \text{KB}_u q$$

$$(D) \neg \text{KB}_u(\text{false})$$

$$\frac{\vdash p}{\vdash \text{KB}_u p}$$

- For the modality $\text{PKB}_c(p)$ we have:

$$(P) [C](p_1, \dots, p_n) \rightarrow$$

$$\left(\bigvee_{\langle i_1, \dots, i_j \rangle \text{ in } \langle 1, \dots, n \rangle} \diamond(p_{i_1} \wedge \dots \wedge p_{i_j}) \wedge \Box(p_{i_1} \wedge \dots \wedge p_{i_j} \rightarrow p) \right) \leftrightarrow \text{PKB}_c(p)$$

where $\langle i_1, \dots, i_j \rangle$ in $\langle 1, \dots, n \rangle$ means that i_1, \dots, i_j is a subsequence of $1, \dots, n$ ³.

- For the modality $\text{FKB}_c(p)$ we have:

$$(F) \quad [S](p_1, \dots, p_n) \rightarrow$$

$$\left(\bigvee_{i \in [1, n]} \Box(p \rightarrow p_i) \right) \leftrightarrow \text{FKB}_c(p)$$
- for the modality \Box we have:

$$\vdash \Box p \text{ iff } \vdash p$$

$$\vdash \neg \Box p \text{ iff } \not\vdash p$$

The intuitive meaning of axiom schema (P) (resp. (F)) is that we can infer $\text{PKB}_c(p)$ (resp. $\text{FKB}_c(p)$) iff p is the consequence of a consistent subset $\{p_{i_1}, \dots, p_{i_j}\}$ of $\{p_1, \dots, p_n\}$ (resp. p is the implicant of a sentence p_i in $\{p_1, \dots, p_n\}$).

Notice that the following sentences are theorems of the Σ logic:

- $\Box(p \rightarrow q) \rightarrow (\text{PKB}_c p \rightarrow \text{PKB}_c q)$, and
- $\Box(p \rightarrow q) \rightarrow (\text{FKB}_c q \rightarrow \text{FKB}_c p)$.

Theorem 1. The Σ logic is sound and complete.

The proof technique of Theorem 1 is based on a canonical model where the set of worlds is the set of maximal consistent sets of sentences of Σ (see [7]).

3 Security Constraints expressed in the Σ logic

This section defines, in the framework of the Σ logic, constraints which characterize acceptable states of a database in regard to a given confidentiality policy. These constraints depend on the choice of a security policy and we do not pretend that they have to be accepted in every context. Nevertheless, they show the expressive power and the flexibility of the Σ logic to formalize security policies.

A given state is defined by a tuple $\langle C, S, \text{cdb} \rangle$, where C represents the set of confidential sentences, S represents the set of secret sentences, and cdb represents the set of sentences that user u has retrieved from the database.

A state $\langle U, S, \text{cdb} \rangle$ is represented in the language L' by the three assumptions: $[C](p_1, \dots, p_n)$, $[S](p'_1, \dots, p'_m)$, $\text{KB}_u(\text{cdb})$.

We shall use the notations: $R = [C](p_1, \dots, p_n) \wedge [S](p'_1, \dots, p'_m)$ and $\text{CDB} = \text{KB}_u(\text{cdb})$. R represents a confidentiality policy called the “regulation” and CDB will be called the “confidential database”.

We first present a security constraint to guarantee the consistency of a confidentiality policy. This constraint says that in the context of a given regulation there is no sentence p such that a confidential user is both permitted to know p and forbidden to know p . The formal representation of the constraint is:

³If $n = 0$, then conjunction should be interpreted to true and disjunction should be interpreted to false.

$$(CR) \quad \vdash R \rightarrow \neg(PKB_{cp} \wedge FKB_{cp})$$

Property 7: From (CR) we can derive that $C(w) \cap S(w) = \emptyset$ for every world $w \in W$.

The proof of property 7 is trivial. Notice also that the converse of this property is generally not true.

3.1 Deductive Channels

The security constraint which has to be satisfied to prevent deductive channels means that for every sentence p there is no world where a confidential user u knows p and u is not permitted to know p . In other words, the set of sentences $\{R, CDB, KB_{up}, \neg PKB_{cp}\}$ should never be satisfied in the Σ logic. Therefore this constraint is represented by:

$$(SCD) \quad \vdash R \wedge CDB \rightarrow (KB_{up} \rightarrow PKB_{cp})$$

The constraint (SCD) is violated for a sentence if we have: $\not\vdash R \wedge CDB \rightarrow (KB_{up} \rightarrow PKB_{cp})$. This constraint corresponds to the so-called restrictive attitude. Notice that from (CR) and (SCD) we can infer:

$$\vdash R \wedge CDB \rightarrow (KB_{up} \rightarrow \neg FKB_{cp})$$

In fact, from (CR) we have $\vdash R \rightarrow (PKB_{cp} \rightarrow \neg FKB_{cp})$. This last constraint corresponds to the so-called permissive attitude. Notice that the converse does not hold. In the case where there exist sentences which are neither confidential nor secret, the (SCD) constraint imposes a stronger constraint. This means that the restrictive attitude formalizes a more conservative policy than the policy represented by the permissive attitude. For instance, let us assume that, in a model M we have: $C(w) = \emptyset$, $S(w) = \{\|a\|\}$ and $db = \{\|a\|\}$ and let us consider a confidential user who ask the query $a \vee b$. According to the restrictive attitude, the multilevel database should not answer this query, because in w we have $\neg PKB_c(a \vee b)$, whereas according to the permissive attitude, this user should be provided with true as an answer, because in w we have $\neg FKB_c(a \vee b)$.

The (SCD) constraint can be enforced in two different ways. The first is to allow user u to access only the sentences that do not lead to a violation of (SCD). That is, to control accesses to cdb to avoid violations. The second is to modify the regulation, in other words, to extend the set of sentences in C .

3.2 Abductive Channels

In this section, we investigate situations where abductive channels may exist. Informally, if sentence q cannot be deduced from p , it is said that it can be inferred by abductive reasoning if there exists a sentence h such that q can be deduced from $\{p, h\}$. However, this definition has to be refined to remove trivial abductions. There are two kinds of trivial abductions. The first is when $\{p, h\}$ is an inconsistent set of sentences. In this case, any sentence can be derived from $\{p, h\}$ and we can derive q even if there is neither a logical link between p and q nor a logical link between h and q . The second kind of trivial abduction

is when h allows us to derive q without using p . Here again, there is no logical link between h and p . In formal terms, we say that q can be derived from p by abductive reasoning, using the assumption h , and this fact is denoted by $ABD(p,h,q)$, if we have:

$$(1) \vdash h \rightarrow (p \rightarrow q), (2) \not\vdash h \rightarrow \neg p \text{ and } (3) \not\vdash h \rightarrow q$$

These properties are represented in the language L' by:

$$ABD(p, h, q) \stackrel{\text{def}}{=} \Box(h \rightarrow (p \rightarrow q)) \wedge \neg\Box(h \rightarrow \neg p) \wedge \neg\Box(h \rightarrow q)$$

There exists an abductive channel for the sentence q if the set of sentences $\{R, CDB, ABD(p, h, q), KB_{up}, \neg PKB_c q\}$ is satisfiable. From a practical point of view, we have to distinguish two cases to prevent the existence of an abductive channel.

In the first case, it is assumed that to know whether sentence h holds, it is not necessary for a user to access the database. In informal terms we say that h is part of common knowledge. In this case, it would be useless to define the regulation in order to have $FKB_c(h)$. Thus, the security constraint to be satisfied is:

$$(SCA) \quad \vdash R \wedge CDB \wedge ABD(p, h, q) \rightarrow (KB_{up} \rightarrow PKB_c q)$$

To enforce this constraint, we have either to restrict cdb in such a way that p cannot be deduced from cdb or to extend set C in order to have $PKB_c(q)$.

In the second case, user u has to access the database to know h . Then the security constraint to be satisfied is:

$$(SCA') \quad \vdash R \wedge CDB \wedge ABD(p, h, q) \rightarrow (KB_{up} \rightarrow (PKB_c q \vee FKB_c h))$$

To enforce the security constraint, we have to control access to cdb in such a way that p cannot be deduced from cdb , to extend C in order to have $PKB_c q$ or to extend S in order to have $FKB_c h$. A practical solution to extend S in this way is to add the sentence $p \rightarrow q$ to S .

For two given sentences p and q , the most general sentence h such that $ABD(p,h,q)$ is $h = p \rightarrow q$, because condition (1) imposes that any h is an implicant of $p \rightarrow q$. Therefore, if $p \rightarrow q$ is in S , we have $FKB_c(p \rightarrow q)$. From properties 2 and 4, for all the implicants h of $p \rightarrow q$ we have $FKB_c h$. This means that if $p \rightarrow q$ is in S , for all the h such that $ABD(p,h,q)$ a confidential user is forbidden to know that the database believes h . In other words, all the abductive channels from p to q are broken.

4 Conclusion

We have presented formal definitions of the concepts of permission and prohibition to know some parts of the content of a database. In the semantics they are interpreted by two accessibility relations D_{cr} and D_{cp} that independently characterize ideal worlds in regard to a restrictive and a permissive attitude. This enables us to consider that regulation is not necessarily complete, in the

sense that some sentences are neither permitted nor prohibited. These relations are respectively combined with accessibility relations R_{xcr} and R_{xcp} in order to characterize mental states of abstract agents xcr and xcp that are compatible with the two different attitudes.

In the axiomatics, the combination of D_{cr} and R_{xcr} is represented by the modality PKB_c and the combination of D_{cp} and R_{xcp} is represented by the modality FKB_c .

It is worth noting that the logic allows us to represent the fact that we have a complete knowledge of the regulation even if the regulation itself is not complete, in the sense that if, in the context of a given regulation, we cannot infer PKB_cp (resp. FKB_cp), then we can infer $\neg PKB_cp$ (resp. $\neg FKB_cp$).

We have also shown practical applications of the logic to define constraints to guarantee the consistency of a regulation or to prevent the existence of inference channels.

Further extensions of this work might be to extend the expressive power of the language to have two independent representations of the deontic concepts and doxastic concepts, for instance to represent facts of the form: $P_c(KB_c(p) \vee KB_c(q))$. Another interesting extension might be to apply deontic modalities to deontic modalities themselves, for instance to represent facts like: "it is permitted to a user cleared at the secret level to know that a user cleared at the confidential level is forbidden to know p ".

Acknowledgement

This work was partially supported by the ESPRIT Basic Research Action MED-LAR 2. We are very grateful to Jose Cramo for his valuable comments.

References

- [1] Thuraisingham B., Ford W., Collins M., and O'Keefe J. Design and implementation of a database inference controller. *Data & Knowledge Engineering*, 11(3), December 1993.
- [2] Hinke T. H. Inference Aggregation Detection in Database Management Systems. In *IEEE Symposium on Security and Privacy*, Oakland, 1988.
- [3] Garvey T. D., Lunt T. F., and Stickel M. E. Abductive and Approximate Reasoning Models for Characterizing Inference Channels. In *Proc. of the computer security foundations workshop*, Franconia, 1991.
- [4] Garvey T. D., Lunt T. F., Qian X., and Stickel M. Toward a Tool to Detect and Eliminate Inference Problems in the Design of Multilevel Databases. In *Proc. of the Sixth IFIP WG 11.3 Working Conference on Database Security*, Vancouver, 1992.
- [5] Chen J. The Generalized Logic of Only Knowing (GOL) that Covers the Notion of Epistemic Specifications. In Z. Ras and M. Zemankova, editors, *Proceedings of the 8th International Symposium, ISMIS'94*, volume 869 of *Lecture Notes in Artificial Intelligence*, Charlotte, North Carolina, 1994. Springer Verlag.

- [6] Cuppens F. A Logical Analysis of Authorized and Prohibited Information Flows. In *IEEE Symposium on Security and Privacy*, Oakland, 1993.
- [7] Chellas B. F. *Modal Logic: An introduction*. Cambridge University Press, 1988.
- [8] Cuppens F. Roles and Deontic Logic. In A. J. I. Jones and M. Sergot, editors, *Second International Workshop on Deontic Logic in Computer Science*, Oslo, Norway, 1994.
- [9] Pörn I. *Action Theory and Social Science; Some Formal Models*, volume 120 of *Synthese Library*. D. Reidel, Dordrecht, 1977.
- [10] Carmo J. and Jones A. J. I. Deontic database constraints and the characterization of recovery. In A. J. I. Jones and M. Sergot, editors, *Second International Workshop on Deontic Logic in Computer Science*, Oslo, Norway, 1994.
- [11] Jones A. J. I. and Porn I. Ideality, Sub-ideality and Deontic Logic. *Synthese*, 65, 1985.
- [12] Brewer D. and Nash M. The Chinese wall security policy. In *IEEE Symposium on Security and Privacy*, Oakland, 1989.