Salem Benferhat
Laurent Garcia

Handling Locally Stratified Inconsistent Knowledge Bases

Abstract. This paper investigates the idea of reasoning, in a local (or contextual) way, under prioritized and possibly inconsistent knowledge bases. Priorities are not supposed to be given globally between all the beliefs in the knowledge base, but locally inside sets of pieces of information responsible for inconsistencies. This local stratification offers more flexibility for representing priorities between beliefs. Given this local ordering, we discuss five basic definitions of influence relations between conflicts. These elementary notions of influence between two conflicts A and B exhaustively explore the situations where solving A leads to solve B. Then we propose natural approaches to restore the coherence of a knowledge base on the basis of influence relations between locally-stratified conflicts.

Keywords: Inconsistency handling, conflicts, local stratification

1. Introduction

The problem of handling conflicting information is important in many areas in Artificial Intelligence. It is particularly present in default reasoning, data fusion, diagnosis and more generally in managing the dynamics of knowledge bases. Recently, two main approaches have been proposed for handling inconsistency: the first one consists in revising the base by restoring consistency and the second one keeps inconsistency but changes the inference relation.

The first approach, called coherence-based approach, can be described into two steps: (1) select one (several) consistent subbase(s) and (2) apply the classical inference on this (these) subbase(s). Examples of systems selecting one consistent subbase are possibilistic logic (Dubois et al., 1994), linear ordering (Nebel, 1994), Adjustment and Maxi-adjustment (Williams, 1994; 1996). Examples of approaches selecting several subbases are acceptable subbases (Rescher, 1976), preferred subbases (Brewka, 1989), Papini's revision function (Papini, 1992) and lexicographical approach (Benferhat et al., 1993; Lehmann, 1995).

The second approach does not use classical logic since the inconsistency is kept. This induces that the handling of inconsistency is made at a higher level (Toulmin, 1956). Examples of such approaches are paraconsistent logics (Da Costa, 1963; Besnard and Hunter, 1995; Damasio and Pereira, 1997) or argumentation (Dung, 1995; Benferhat et al., 1995; Amgoud and Cayrol, 1998; 2000).

Studia Logica **70**: 77–104, 2002.

It is well known that the use of priorities is very important for the purpose of managing the flux of information in databases. Gärdenfors (1988) has proved that any revision process that satisfies natural requirements is implicitly based on priority ordering. Taking preferences into account makes easier the handling of inconsistency but adds a step for the expression of priorities. In most of the existing systems, the priorities expressed between pieces of information are often a total preordering on the beliefs. Three kinds of problems then occur. First, an information whose priority is low can be ignored even if it is not responsible for any inconsistency. This problem is known as the drawning problem (Benferhat et al., 1993). Second, using a total preorder can compare two *unrelated* pieces of information. Lastly, it is not possible to express context-dependent priorities.

In this paper, we will focus on coherence-based approaches, by proposing several ways to selecting consistent subbases based on a local representation of preferences. Priorities are not supposed to be given globally between all the beliefs in the knowledge base, but locally inside sets of pieces of information responsible for inconsistencies. Section 2 gives basic definitions and notations used in the paper and describes the process of locally handling inconsistent information. Section 3 discusses five basic definitions of influence relations between conflicts. From these basic influence relations, we extract a natural order in which conflicts should be solved. Section 4 gives reasonable inference relations to deal with locally stratified inconsistent knowledge bases.

2. Basic definitions and motivations

We consider a finite propositional language denoted by $\mathscr{L}$. The classical consequence relation is denoted by $\vdash$. We denote by Σ a set of beliefs which is not deductively closed. Pieces of information are called 'beliefs' since they are not sure and can be withdrawn in the light of new, more sure, information. Beliefs of Σ are assumed to be provided by different sources. A source can be seen as an entity (which is supposed to be independent of other sources) providing some pieces of information. Let $\mathscr{S} = \{S_1, S_2, \ldots, S_n\}$ be a set of sources. We denote by $\Phi(S_i)$ the logical formula encoding the information provided by the source S_i. Then, our knowledge base Σ is the multi-set $\{\Phi(S_1), \Phi(S_2), \ldots, \Phi(S_n)\}$, which can be inconsistent even if $\Phi(S_i)$'s are here assumed to be consistent. Among the available information, we distinguish the sure knowledge that can not be ignored. This sure knowledge is described by a new source, denoted S_K, and its associated classical formula is denoted by $\Phi(S_K)$. Solving inconsistency leads to decide which sources should be kept and which sources should be ignored. The classical formulas will then be used to determine the presence of inconsistency but its resolution will be made using the sources.

The first principle of local approaches is that, instead of considering the whole set of beliefs, we focus on subsets of the base: the ones which are responsible for an inconsistency, called conflicts. The idea of using the subbases responsible for an inconsistency, to restore the consistency of Σ is not new and has been used by several authors: Reiter (1987) and De Kleer (1976) in the framework of model-based diagnosis, Papini (1992) in belief revision, Cayrol et al. (1993) in consistency check and Yang (1992) in planning. This line of research is still investigated, particularly in the context of multi-agents systems (Tessier et al., 2000) or in the context of diagnosis (Wasserman, 2000; Wurbel et al., 2000).

The second principle of local approaches is to express priorities between beliefs only when they are conflicting. Expressing priorities inside each conflict can be meaningful and easier. In practice, it is impossible for an expert to give priorities upon all the beliefs. Moreover, when two beliefs are unrelated, expressing priorities between them is useless.

2.1. Definition of conflicts

Definition 1 introduces formally the notion of conflicts.

DEFINITION 1. A set $C = \{S_i\}$ of sources (with $C \subseteq \mathscr{S}$) is called a conflict iff it satisfies:

(1) $\bigcup_i \{\Phi(S_i)\} \vdash \bot$ and

(2) $\forall S_j \in C$, $\bigcup_i \{\Phi(S_i)\} - \{\Phi(S_j)\} \nvdash \bot$.

The point (1) means that the set of beliefs provided by the sources is classically inconsistent and the point (2) means that a conflict is minimal with respect to the set inclusion relation. In the following, the set of conflicts induced by a set of sources $\mathscr{S}$ will be denoted by $\mathscr{C}(\mathscr{S})$ and, when no ambiguity occurs, by $\mathscr{C}$. Note that $\Phi(S_K)$ can belong to a conflict but the restoration of consistency must always keep this source.

The set $\mathscr{C}(\mathscr{S})$ can be related to the 'base of nogoods' used in the terminology of the ATMS (Assumptions-based Truth Maintenance Systems) (De Kleer, 1986). A nogood is a minimal set of incompatible assumptions. In the ATMS formalism, the vocabulary used to describe knowledge contains two kinds of symbols: the assumption and the non-assumption symbols. Links between conflicts and nogoods can be established in the following way: let Σ

be a belief base, and let Σ' be a new belief base obtained from Σ by replacing each formula ϕ_i in Σ by $\neg S_i \vee \phi_i$, where S_i is an assumption symbol (all S_i's are different) and can be viewed as the source which provides the formula ϕ_i. Then we can show that the subbase $C = \{S_i \mid i = 1, m\}$ is a conflict of Σ if and only if $\mathscr{H}_A = \{S_i \mid \neg S_i \vee \phi_i \in \Sigma', S_i \in C\}$ is a nogood of Σ'.

Efficient algorithms for computing nogoods (hence conflicts) can be found in (Castell et al., 1996). It is clear that the problem of computing conflicts is NP-complete (Provan, 1988), namely, there are (extreme) situations where expliciting the set of all conflicts needs an exponential time. The algorithm proposed in (Castell et al., 1996) is based on Davis and Putnam procedure (1960). Its efficiency is due to the use of particular heuristics, and has been shown experimentally by comparing it with some existing algorithms for computing nogoods in ATMS.

2.2. Local stratification

Two criteria are used when restoring consistency: the first one is the 'parsimony' or 'minimal change' principle that is to try to keep as much information as possible of initial information (Gärdenfors, 1988). The second one is, when a preference relation or reliability relation between sources is avalaible, to get rid (i.e. ignore) the least reliable sources.

The interesting question is how to mix preferences and minimal change principle. For instance, let us consider three sources S_1, S_2 and S_3 such that S_1 is preferred to S_2 and S_1 is also preferred to S_3. Assume that in order to restore the consistency of the knowledge base, we should either remove S_1 or remove both S_2 and S_3. In most existing systems, and in the approaches given in the paper, we get rid of S_2 and S_3. If one wants to express some 'compensation' principle, we should use a richer representation of preferences, either still qualitatively by means of a preference relation given on subsets of sources (Benferhat et al., 1998), or quantitatively by associating to each source a number (or a degree between [0,1]) and use an operation which allows compensation, like a sum operator in penalty logic (Dupin et al., 1994). In this paper, we will only assume a qualitative preference between sources and, to restore consistency of the knowledge base, we will try to keep as much as possible the most prioritary information without compensation.

Restoring consistency means selecting, in each conflict, one source to ignore. We do not really need a refined stratification of conflicts, but simply need to determine inside each conflict which source(s) is (are) the least preferred one(s). The stratification of a conflict C is hence made of two strata, one denoted by $\underline{C}$ (called non-dominant stratum) containing the least confident sources and the other denoted by $\overline{C}$ (called dominant stratum) containing the other sources. Moreover the source S_K which contains sure knowledge, if it is involved in some conflict C, should necessary belongs to $\overline{C}$.

DEFINITION 2. A local stratification of a conflict C is a partition $(\overline{C}, \underline{C})$ such that:

(1) $\overline{C} \cup \underline{C} = C$, $\overline{C} \cap \underline{C} = \emptyset$,

(2) $S_K \notin \underline{C}$, and

(3) $\underline{C} \neq \emptyset$.

In this definition, (1) simply means that the local stratification is a partition of C, (2) means that sure pieces of information do not belong to the non-dominant stratum and (3) means that there is necessarily at least one non-dominant source in C. Note that $\overline{C}$ can be empty; this can correspond for example to the case where no stratification is available. A source S which is in a dominant stratum of a given conflict is not always guaranteed to be kept, except if it never belongs to the non-dominant stratum of any conflict. For example, a source S which is both in the dominant stratum of a conflict C_1 and in the non-dominant stratum of a conflict C_2 can be ignored if C_2 is solved first.

2.3. Local stratification versus global stratification

The local stratification offers more flexibility for representing priorities between sources with respect to a total pre-ordering, a partial ordering or a multi-ordering. Note that in the existing systems, these orders are generally expressed between beliefs and not between sources which provide the beliefs. Here, for homogeneity of representation of preferences, we assume that preference relations are given between sources.

If we consider a total pre-ordering between the sources of $\mathscr{S}$, as in System Z (Pearl, 1990) or in possibilistic logic (Dubois et al., 1994), the base $\mathscr{S}$ is stratified in the form $\mathscr{S} = \mathscr{S}_1 \cup \cdots \cup \mathscr{S}_n$. Sources in $\mathscr{S}_i$ have the same level of priority and are more prioritary than the ones in $\mathscr{S}_j$ when $j > i$. Hence, $\mathscr{S}_1$ contains the most important sources of $\mathscr{S}$ and $\mathscr{S}_n$ contains the least important ones. Given a total pre-ordering between the sources, the local stratification upon sources in our framework is straightforwardly defined: for each conflict C, we define $\underline{C} = (\mathscr{S}_i \cap C)$ such that $i = \max\{j \mid \mathscr{S}_j \cap C \neq \emptyset\}$. A total ordering is easy to use and to implement but in practice it is not obvious to provide a complete pre-order between the sources of $\mathscr{S}$.

Note that contextual preferences cannot be represented by a total pre-order. Let us consider the following example (inspired from the Condorcet paradox (Moulin, 1988)) which concerns VI Nations rugby tournament. An agent expresses the following preferences: between France and Scotland, France is preferred to Scotland, between England and France, England is preferred to France and between England and Scotland, Scotland is preferred to England. If we want to represent the agent preferences by a total ordering, the only way for respecting these priorities is to put the three teams in the same level. However, when we are dealing with a competition between two of them, we can not recover the contextual preferences given by the agent.

In the case of a partial preordering denoted by $\leq_S$, the set $\underline{C}$ is defined by $\underline{C} = \{S_i \in C \mid \nexists S_j \in C, S_i \leq_S S_j\}$ (the set of minimal elements w.r.t. $\leq_S$). A partial ordering is more flexible than a total pre-ordering but is not easy to use: see Roos (1992) or Brewka (1994) which consider all the total orderings that can be drawn from the partial ordering.

The local stratification can also represent a multi-ordering between beliefs. The multi-ordering comes from the fact that two beliefs can be compared differently according to different features. For example, if we consider the approach based on the topics (Cholvy, 1995), a conflict represents contradictory information about a given topic and the ordering between the beliefs is determined with respect to each topic. An example inspired by Cholvy describes the following beliefs: two witnesses, a taylor and a mechanic, give a description of a suspect. The taylor says that the suspect is a girl wearing a Chanel suit and driving a sport Volkswagen car. For the mechanic, the suspect is a girl wearing a dress and driving a diesel car. The idea of the approach about topics is that the sources of information are ranked according to the different topics. Here, the information given by the taylor has priority over the ones given by the mechanic when speaking about 'clothing' while the priority is the contrary when speaking about 'car'. Such ordering can clearly be represented using the local stratification since several partial preference relations can be simultaneously defined.

2.4. General schemata

First, we can give three natural properties that the revision process should follow. Let E_i be a subbase obtained by a revision process, let $\mathscr{C}$ be the set of initial conflicts.

$$\forall E_i \forall C_j \in \mathscr{C} \exists S_m \in C_j,\ \Phi(S_m) \in \Sigma - E_i\,. \tag{1}$$

This property means that every conflict must be solved.

$$\forall E_i \forall C_j \in \mathscr{C} \forall S_m \notin \bigcup \underline{C_j},\ \Phi(S_m) \in E_i\,. \tag{2}$$

This property means that a source which is not in the non-dominant stratum of any conflict must be in the resulting subbase.

$$\forall E_i \forall E_j,\ \text{we do not have } E_i \subseteq E_j \text{ or } E_j \subseteq E_i. \tag{3}$$

The extensions are incomparable in terms of set inclusion.

The process of handling locally stratified conflicts and restoring the consistency of the knowledge base is described in Figure 1.

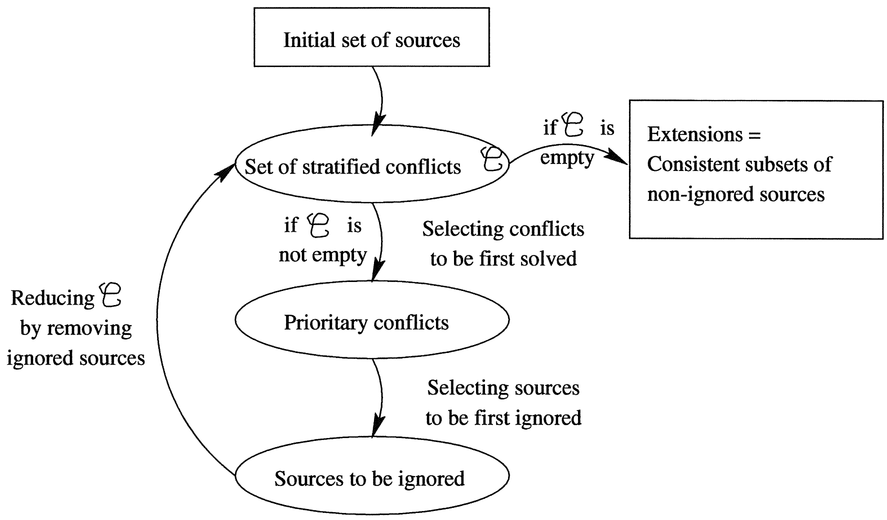

Figure 1. Process of handling locally stratified conflicts

Given an initial set of sources, the conflicts are first computed and locally stratified. To restore consistency by solving conflicts, we need to determine which conflicts must be first handled. This is done by defining natural ordering between conflicts. When the set of prioritary conflicts is determined, strategies for selecting sources to be ignored are applied (where the selection is necessarily done from the non-dominant strata of prioritary conflicts). Once these sources have been selected and ignored, the process is repeated with the remaining conflicts until there is no conflict. Each set of sources corresponds to a consistent subbase, called extension, of the initial base.

The rest of this article develops this process, by first defining a priority relation between conflicts. Then, we study the strategies of selecting sources to ignore and, lastly, the inference relations associated to each strategy.

3. Conflict ordering

When a given source is present in two conflicts, then the resolution of one conflict can have an influence on the resolution of the other conflict. So, it is natural to consider some order in which the conflicts should be solved instead of handling them separately. Example 1, inspired from (Cholvy and Hunter, 1997), illustrates this idea:

EXAMPLE 1. *Three requests for building a house are given: build a fence (denoted by* fe, *provided by* S_1*), build the walls (denoted by* wa, *provided by* S_2*) and build the fundations (denoted by* fu, *provided by* S_3*). Moreover, there are two integrity constraints: it is not possible to build both fences and walls (denoted by* $\neg \mathrm{fe} \vee \neg \mathrm{wa}$*) and it is not possible to build both walls and foundations (denoted by* $\neg \mathrm{wa} \vee \neg \mathrm{fu}$*). These constraints correspond to the source* S_K*. So the belief base is* $\Sigma = \{\Phi(S_1), \Phi(S_2), \Phi(S_3), \Phi(S_K)\}$ *with* $\Phi(S_1) \equiv \mathrm{fe}$, $\Phi(S_2) \equiv \mathrm{wa}$, $\Phi(S_3) \equiv \mathrm{fu}$, $\Phi(S_K) \equiv (\neg\mathrm{fe} \vee \neg\mathrm{wa}) \wedge (\neg\mathrm{wa} \vee \neg\mathrm{fu})$. *We have two conflicts* $A = \{S_1, S_2, S_K\}$ *and* $B = \{S_2, S_3, S_K\}$*. The preferences expressed on these conflicts are: there is no preference between building fence and walls (conflict* A*) and fundations must be built before walls (conflict* B*). So,* $\underline{A} = \{S_1, S_2\}$ *and* $\underline{B} = \{S_2\}$*. If we consider separately the two conflicts, we can ignore* S_1 *(due to conflict* A*) and also* S_2 *(due to conflicts* A *and* B*). The only non-ignored request is to build fundations. However, if we first solve* B *(by ignoring* S_2*),* A *will be also solved, and then we will satisfy more requests (namely, building fundations and building fences).*

Example 1 means that it is natural to relate conflicts to reduce the number of ignored sources. Two conflicts C_1 and C_2 are related if their resolution is not independent, that is if the resolution of C_1 can lead to the resolution of C_2 or conversely.

DEFINITION 3. Two conflicts C_1 and C_2 of $\mathscr{C}(\mathscr{S})$ are *related* if there exists a source $S_i \neq S_K$ such that $S_i \in C_1 \cap C_2$ and $S_i \in \underline{C_1} \cup \underline{C_2}$.

DEFINITION 4. Two conflicts C_1 and C_2 of $\mathscr{C}(\mathscr{S})$ are in the same *group* iff (1) C_1 and C_2 are related or (2) there exists C_3 of $\mathscr{C}(\mathscr{S})$ such that C_1 and C_3 are in the same group and C_2 and C_3 are in the same group.

Conflicts of a given group are necessarily unrelated to conflicts of another group, hence each group can be handled separately from the others. In the following, we focus on solving conflicts of the *same* group.

To define an ordering between conflicts inside each group, we use a notion of positive influence. The positive influence of a conflict A on another conflict B means that the resolution of A has or can have an influence on the resolution of B. In the next subsection, we consider two cases of influence. The first case is a notion of *sure* influence where the resolution of A leads *necessarily* to solve B. The second case of influence is a *possible* influence where the resolution of A *can* only lead to solve B.

3.1. Elementary influence relations between conflicts

Let us start by studying the elementary influences that a conflict A can have on another conflict B. We focus on the sources that belong to $\underline{A}$ since they are the only sources that can be ignored when solving A.

3.1.1. Sure elementary influences

Let us first consider the case where the resolution of A implies the resolution of B whatever the ignored source of $\underline{A}$. This corresponds to the case where all the sources of $\underline{A}$ are in B that is $\underline{A} \subseteq B$. This case can be splitted in three elementary cases.

The first elementary case is when all the sources of $\underline{A}$ are in $\underline{B}$.

DEFINITION 5. A conflict A has an NDS-influence ('ND' for 'Non-Dominant' and 'S' for 'Sure') on a conflict B (denoted $A \leqslant_{\mathrm{NDS}} B$) iff $\underline{A} \subseteq \underline{B}$.

This influence relation is transitive and reflexive. Moreover, when A has a NDS-influence on B, we have $\overline{A} \not\subseteq \overline{B}$. If we again consider Example 1, we can check that B has a NDS-influence on A.

The second elementary sure influence is when the sources of $\underline{A}$ are all in the dominant stratum of B.

DEFINITION 6. A conflict A has a DS-influence ('D' for 'Dominant' and 'S' for 'Sure') on a conflict B (denoted by $A \leqslant_{\mathrm{DS}} B$) iff $\underline{A} \subseteq \overline{B}$.

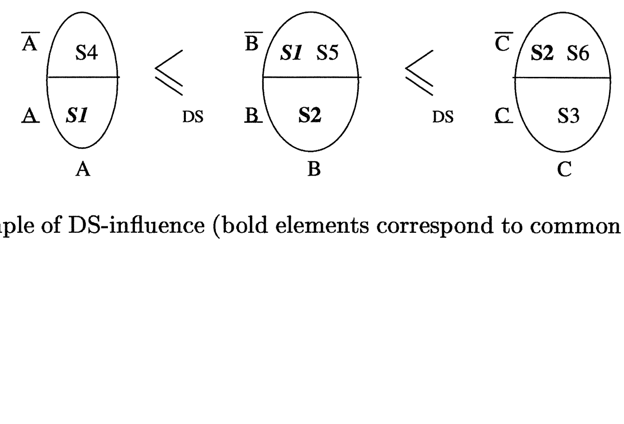

Figure 2. Example of DS-influence (bold elements correspond to common sources between conflicts)

EXAMPLE 2. *An example of DS-influences is given in Figure 2, where,* $\Phi(S_1) \equiv \neg\beta$, $\Phi(S_2) \equiv \beta\vee\gamma\vee\delta$, $\Phi(S_3) \equiv \neg\delta$, $\Phi(S_4) \equiv \alpha\wedge\beta$, $\Phi(S_5) \equiv \neg\gamma\wedge\neg\delta$, $\Phi(S_6) \equiv \neg\beta \wedge \neg\gamma$. *Here, for sake of simplicity, we focus on the conflicts* A, B *and* C. A *has a DS-influence on* B *(since* $\underline{A} \subseteq \overline{B}$ *due to* S_1*) and* B *has a DS-influence on* C *(since* $\underline{B} \subseteq \overline{C}$ *due to* S_2*).*

This relation is neither reflexive nor transitive. Indeed, in this example, we have: A *has a DS-influence on* B, B *has a DS-influence on* C *but* A *has not a DS-influence on* C *(*$\underline{A} \not\subseteq \overline{C}$ *since* $S_1 \notin \overline{C}$*).*

The last elementary case corresponds to the case when $\underline{A} \subseteq B$ but $\underline{A} \not\subseteq \underline{B}$ and $\underline{A} \not\subseteq \overline{B}$.

DEFINITION 7. A conflict A has an OS-influence (with 'O' for 'Other') on a conflict B (denoted by $A \leqslant_{\mathrm{OS}} B$) iff $\underline{A} \subseteq B$, and A has neither a DS-influence nor an NDS-influence on B.

This last elementary case is illustrated by the Example of Figure 3. The OS-influence is neither reflexive neither transitive.

We denote GS-influence the case which recovers all the elementary sure-influences, namely A has a GS-influence on a B iff $\underline{A} \subseteq B$.

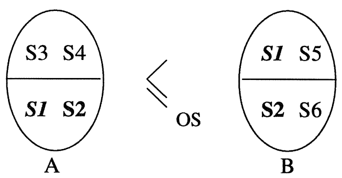

Figure 3. Example of OS-influence

3.1.2. Possible elementary influences

The other case of elementary influences are called Possible influences (denoted by P-influence). The possible influences occur when some source (and not necessarily all) of $\underline{A}$ are in B.

DEFINITION 8. A given conflict A has:

(1) an NDP-influence ('ND' for 'Non-Dominant' and 'P' for 'Possible') on a conflict B (denoted by $A \leqslant_{\mathrm{NDP}} B$) iif $\underline{A} \cap \underline{B} \neq \emptyset$;

(2) a DP-influence on a conflict B (denoted by $A \leqslant_{\mathrm{DP}} B$) iif $\underline{A} \cap \overline{B} \neq \emptyset$.

	Non-Dominant stratum	Dominant stratum	Other	General
Sure influence	$\underline{A} \subseteq \underline{B}$ NDS-influence	$\underline{A} \subseteq \overline{B}$ DS-influence	$\underline{A} \subseteq B$, $\underline{A} \not\subseteq \underline{B}$ and $\underline{A} \not\subseteq \overline{B}$ OS-influence	$\underline{A} \subseteq B$ GS-influence
Possible influence	$\underline{A} \cap \underline{B} \neq \emptyset$ NDP-influence	$\underline{A} \cap \overline{B} \neq \emptyset$ DP-influence		$\underline{A} \cap B \neq \emptyset$ GP-influence

Table 1. Elementary and general influence relations between conflicts

Let us notice that it does not exist a counterpart to OS-influence relation (since it is impossible to have $\underline{A} \cap B \neq \emptyset$, when $\underline{A} \cap \underline{B} = \emptyset$ and $\underline{A} \cap \overline{B} = \emptyset$). The general case, called GP-influence, is: A has a GP-influence on a conflict B iff $\underline{A} \cap B \neq \emptyset$.

The NDP-influence relation is reflexive and symmetric but it is not transitive. Indeed, let A, B, C be three conflicts such that $\underline{A} = \{S_1, S_2\}$, $\underline{B} = \{S_2, S_3\}$, $\underline{C} = \{S_3, S_4\}$. It is easy to check that $A \leqslant_{\mathrm{NDP}} B$, $B \leqslant_{\mathrm{NDP}} C$, but $A \leqslant_{\mathrm{NDP}} C$ is not true.

The DP-influence relation is irreflexive and is not transitive (we can use the same example as Figure 2). The GP-influence relation is reflexive but is neither symmetric nor transitive. It is obvious that NDS-influence (resp. DS) implies NDP-influence (resp. DP).

Table 1 is a recapitulation of elementary influences between conflicts.

3.2. Why are elementary relations not enough?

The aim of this subsection is to provide different counterexamples showing that elementary influences are not enough to determine a natural order in which conflicts should be solved:

- Let us show that $A <_{\mathrm{NDS}} B$ *is not a necessary condition to say that* A *must be solved before* B. Let us take the example of the Figure 4, which concerns a meeting where three persons called Alexandra, Barbara and Claudia are invited. We assume that it is not possible to invite both Alexandra and Barbara and to invite both Barbara and Claudia. The base $\Sigma = \{\Phi(S_1) \equiv \mathrm{Ax}, \Phi(S_2) \equiv \mathrm{Ba}, \Phi(S_3) \equiv \mathrm{Cl}, \Phi(S_K) \equiv (\neg\mathrm{Ax} \vee \neg\mathrm{Ba}) \wedge (\neg\mathrm{Ba} \vee \neg\mathrm{Cl})\}$ is inconsistent. Hence, we need to determine which invitation should be cancelled. The local stratification inside each conflict means that we prefer to invite Alexandra than Barbara (conflict A) and Barbara than Claudia (conflict B). A satisfactory solution is to invite Alexandra and Claudia but not Barbara.

This means that A should be solved before B to make sure that only S_2 is ignored (minimal change principle). However, such natural order in solving conflicts between A and B cannot be obtained if NDS-influence relation is used alone, since there is no NDS-influence of A on B.

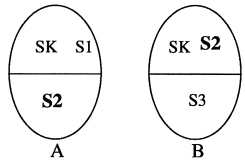

Figure 4. Counterexample for NDS-influence relation

• Let us consider $A <_{\mathrm{DS}} B$ (resp. $A <_{\mathrm{DP}} B$) and the example of Figure 5, which is an extension of Example 1 where ga means 'garden house': $\Sigma = \{\Phi(S_1) \equiv \text{fe}, \Phi(S_2) \equiv \text{wa}, \Phi(S_3) \equiv \text{fu}, \Phi(S_4) \equiv \text{ga}, \Phi(S_K) \equiv (\neg\text{fu} \vee \neg\text{wa}) \wedge (\neg\text{wa} \vee \neg\text{fe} \vee \neg\text{ga})\}$. The stratification is: the foundations must be built before the walls and the fences must be built first.

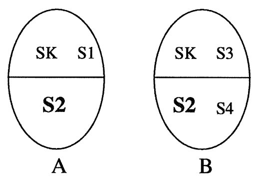

Figure 5. Counterexample for DS-influence relation

If B is first solved, we can ignore S_4 while S_2 will be necessarily ignored when solving A. It is then natural to solve A first and only ignore S_2 which respects the minimal change principle. This priority can not be deduced from the notion of DS-influence. This same counterexample can also be used for the DP-influence. Hence, when A has a sure influence on B, while B has only a possible influence on A, it is more natural to solve A before B.

• Let us consider $A <_{\mathrm{OS}} B$ and the example of Figure 6, which is a variant of meeting example, 'Di' stands for 'Dinah': $\Sigma = \{\Phi(S_1) \equiv \text{Ax}, \Phi(S_2) \equiv \text{Ba}, \Phi(S_3) \equiv \text{Cl}, \Phi(S_4) \equiv \text{Di}, \Phi(S_K) \equiv (\neg\text{Ax} \vee \neg\text{Ba} \vee \neg\text{Cl}) \wedge (\neg\text{Ax} \vee \neg\text{Cl} \vee \neg\text{Di})\}$.

Here, Claudia is the only person that is rejected whatever the considered conflict so it is natural to ignore the source S_3. To make sure that this source is ignored, we have to stand that the conflict A has a priority on the conflict B. This is in contadiction with the consideration of the OS-influence.

• '*$A <_{\mathrm{NDP}} B$ is not a necessary condition to say that A must be solved before B*'. Let us take another example of the meeting example illustrated in

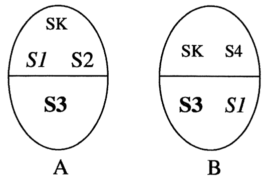

Figure 6. Counterexample for OS-influence relation

Figure 7. The base is $\Sigma = \{\Phi(S_1) \equiv \text{Ax}, \Phi(S_2) \equiv \text{Ba}, \Phi(S_3) \equiv \text{Cl}, \Phi(S_4) \equiv \text{Di}, \Phi(S_K) \equiv (\neg\text{Ax} \vee \neg\text{Ba} \vee \neg\text{Di}) \wedge (\neg\text{Ba} \vee \neg\text{Cl})\}$. Since the suppression of S_2 which solves A leads also to solve B while ignoring any source in $\underline{B}$ can not solve A, it is natural to give priority to A upon B. This priority is not captured by the NDP-influence. This example means that if A has an influence on B while B has no influence on A, it is more natural to first solve A.

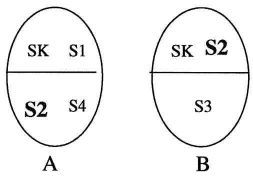

Figure 7. Counterexample for NDP-influence relation

So, it clearly appears from these counterexamples that it is not enough to consider a single elementary influence. Hence, we should define a complex influence relation from a combination of elementary influence relations.

3.3. Determining the positive influence relation between conflicts

In this section, we identify the cases where a conflict A should be solved before B. First, note that some of combinations of elementary influences are *impossible*. For example, this is the case where A has a NDS-influence on B and B has a DS-influence on A (indeed, $\underline{A} \subseteq \underline{B}$ and $\underline{B} \subseteq \overline{A}$ is impossible since it implies that $\underline{A} \subseteq \overline{A}$). Other impossible combinations are (DS, OS), (DS, NDP), (NDS, $\emptyset$), (OS, $\emptyset$), (NDP, $\emptyset$), where (x, y) means A has a x-influence on B and B has a y-influence on A ($\emptyset$ means there is no influence).

The simple case where clearly A should be solved before B is when A has an influence on B and B has not influence on A.

Now consider the case where A has a sure influence on B and B only has a possible influence on A. For this case, whatever the chosen source to ignore in $\underline{A}$ for solving A, it will necessarily belong to B. So the resolution

A on B \ B on A	NDS	DS	OS	NDP	DP	none
NDS	A, B	no	A	A	A	no
DS	no	A, B	no	no	A	A
OS	B	no	A, B	A	A	no
NDP	B	no	B	A, B	A, B	no
DP	B	B	B	A, B	A, B	A
none	no	B	no	no	B	A, B

Table 2. Priority between conflicts given their mutual influences

of A will surely solve B and not conversely. So, it is natural to prefer the sure-influence over the possible influence.

Another interesting case is when A has a NDS-influence on B and B has a OS-influence on A. This means that $\underline{A} \subseteq \underline{B}$ and $\underline{B} \subseteq A$ (but $\underline{B} \not\subseteq \underline{A}$). Then, whatever the chosen source to ignore in $\underline{A}$ (resp. $\underline{B}$), it leads to solve B (resp. A). However each source in $\underline{A}$ is in $\underline{B}$ while a source in $\underline{B}$ may not be in $\underline{A}$. Hence, it is natural to give the priority to solving A that is to prefer the NDS-influence than the OS-influence.

Beside, there are symmetric situations where it is not possible to decide which conflicts should be first solved (for instance, the case where A has only a possible influence on B and B has only a possible influence on A).

Table 2 gives a summary of situations where A should be solved before B given the elementary influences between these two conflicts. A cell of the table is 'A' if the conflict A has a priority on B, 'A, B' if it is not possible to decide between A or B and 'no' if this combination is impossible.

From this table, we can say that A should be solved before B if:

- A has an influence on B and B does not have an influence on A,
- A has a sure influence on B and B does not have a sure influence on A,
- A has a sure influence on B in non-dominant stratum and B does not have a sure influence on A in non-dominant stratum.

Intuitively, A has a positive influence on B when either A should be solved before B or A and B have equal priority. This leads to the following formal definition:

DEFINITION 9. A conflict C_1 has a *positive influence* on a conflict C_2, denoted $C_1 \leqslant_{\mathrm{I}} C_2$ ('I' stands for 'Influence') if and only if

(1) $\underline{C_1} \subseteq \underline{C_2}$, or

(2) $\underline{C_1} \subseteq C_2$ and $\underline{C_2} \nsubseteq \underline{C_1}$, or
(3) $\underline{C_1} \cap C_2 \neq \emptyset$ and $\underline{C_2} \nsubseteq C_1$.

We denote by $\preceq_{\mathrm{I}}$ the transitive closure of positive influence defined by: (1) $A \leqslant_{\mathrm{I}} B \Rightarrow A \preceq_{\mathrm{I}} B$ and (2) $A \preceq_{\mathrm{I}} B$ and $B \preceq_{\mathrm{I}} C \Rightarrow A \preceq_{\mathrm{I}} C$.

We will denote $A \prec_{\mathrm{I}} B$ if $A \preceq_{\mathrm{I}} B$ holds but $B \preceq_{\mathrm{I}} A$ does not hold. Given a set of conflicts $\mathscr{C}$, the relation $\prec_{\mathrm{I}}$ will be used to determine the set of conflicts to be first solved:

DEFINITION 10. A conflict C is in the set of conflicts to be first solved if there is no conflict C' such that $C' \prec_{\mathrm{I}} C$.

We denote by $\min(\mathscr{C})$ the set of conflicts to be first solved.

4. Contextual inference relations

If we consider the process described in Figure 1, the previous Section has shown how to compute prioritary conflicts. In the current section, we study the next step that is how to solve selected conflicts and then determine the corresponding inference relations. An inference relation consists in applying classical inference to one or several consistent subbases. Each subbase corresponds to a solution of conflicts resolution. In the following, we give different possible strategies.

4.1. Ignoring all contestable sources

The first inference relation is based on a simple idea: a source is ignored if it belongs to the non-dominant stratum of a conflict. This leads to compute a single subbase solution.

DEFINITION 11. A source S_i of $\mathscr{S}$ is said to be *contestable* if there exists a conflict of $\mathscr{C}(\mathscr{S})$ such that $S_i \in \underline{C}$. Otherwise, it is said *incontestable*. The set of incontestable sources is denoted $\mathrm{Incont}(\mathscr{S})$.

DEFINITION 12. A formula ψ is an *incontestable conclusion* of a base Σ (denoted $\Sigma \mid\!\sim_{\mathrm{Incont}} \psi$) iff $\bigcup\{\Phi(S_i) \mid S_i \in \mathrm{Incont}(\mathscr{S})\} \vdash \psi$.

$\mathrm{Incont}(\mathscr{S})$ is in some sense an extension of the notion of Kernel (Alchourron et al., 1985). This definition corresponds to the 'safe-contraction' in belief revision (Alchourron et al., 1985), to Non-Defeated inference relation in inconsistent belief bases (Benferhat et al., 1995) and to 'global-adjustment'

(Williams, 1996) in the case of a total pre-ordering. When there is no stratification, this leads to Free-consequence relation based on the notion of free beliefs (Benferhat et al., 1992):

$$\text{Free}(\Sigma) = \{\Phi(S_i) \in \Sigma \mid \forall C \in \mathscr{C}(\mathscr{S}), S_i \notin C\}.$$

The problem with incontestable inference relation is that it is too cautious since too much sources can be ignored:

EXAMPLE 3. *Let us take a variant of the building example. There are three requirements: build the walls* ($\Phi(S_1) \equiv W$)*, build the fence* ($\Phi(S_2) \equiv F$) *and build the garden house* ($\Phi(S_3) \equiv G$)*. Moreover, we assume that it is not possible to satisfy the three requirements (for instance, because of a lack of bricks):* $\Phi(S_K) \equiv \neg W \vee \neg F \vee \neg G$*. We have a single conflict* $A = \{S_1, S_2, S_3, S_K\}$*. The owner wants to build the fence first: the stratification is then* $\underline{A} = \{S_1, S_3\}$*. So the only source that is kept is* S_2 *but it is possible to satisfy another requirement since it is enough to ignore either* S_1*, either* S_3 *to restore consistency.*

4.2. Solving all prioritary conflicts

The idea is to select subsets of sources to ignore such that all prioritary conflicts are solved. We present two ways for solving all the prioritary conflicts: the first way selects all the sources that can be ignored and the second determines minimal sets which are enough to solve conflicts.

We first extend the notion of classical Kernel (Alchourron et al., 1985) to take into account priorities. The Kernel-min of a set of conflicts $\mathscr{A}$ is the union of non-dominant strata of prioritary conflicts in $\mathscr{A}$.

DEFINITION 13. Let $\mathscr{A}$ be a set of conflicts. The Kernel-min of $\mathscr{A}$, is defined by: $\text{Kernel-min}(\mathscr{A}) = \{S_i \mid \exists C \in \min(\mathscr{A}), S_i \in \underline{C}\}$.

4.2.1. Determining a single subbase

The first idea is to ignore all sources belonging to Kernel-min. This leads to solve all the prioritary conflicts. Then, we repeat this step with the remaining conflicts. The construction of the consistent subbases is summarized by the following algorithm where $\text{Sup}\,\mathscr{S}$ denotes the set of sources to ignore and EnsC denotes the set of remaining conflicts at each step.

Input: $\mathscr{C}$ – a set of conflicts
Output: E – the determined single consistent subbase

1. let Sup$\mathscr{S} \leftarrow \emptyset$, let EnsC $\leftarrow \mathscr{C}(\mathscr{S})$
2. repeat until EnsC $= \emptyset$
   2.1. compute Kernel-min(EnsC)
   2.2. for all $S \in$ Kernel-min(EnsC) do
      2.2.1. Sup$\mathscr{S} \leftarrow$ Sup$\mathscr{S} \cup \{S\}$
      2.2.2. EnsC $\leftarrow$ EnsC $-\{C \mid C \in$ EnsC and $S \in C\}$
      /* Remove solved conflicts */
3. return $E = \mathscr{S} -$ Sup$\mathscr{S}$

DEFINITION 14. Let E be the subbase computed in step 3 of the previous algorithm. A formula ψ is a *selective conclusion* of a base Σ (denoted by $\Sigma \mid\!\sim_{Select} \psi$) iff $\bigcup \Phi(S_i) \vdash \psi$, with $S_i \in E$.

PROPOSITION 1. *If ψ is an incontestable consequence of Σ then ψ is a selective consequence of Σ. The converse is false.*

The selective inference relation unfortunately is not completely satisfactory since it can be in some situations too cautious and, in others, too credulous. For the cautious aspect, it is enough to consider Example 3 where the solution of selective inference relation is the same as the incontestable inference relation. For the credulous aspect, let us consider the following example.

EXAMPLE 4. *Let us take an example of cooking where we are interested in melting three ingredients: salt (denoted by* sa*), sugar (denoted by* su*) and pepper (denoted by* pe*). We have the following integrity constraints where it is not possible to melt sugar and salt and we should use salt when using pepper, i.e., $\Phi(S_K) \equiv (\neg\mathrm{su} \vee \neg\mathrm{sa}) \wedge (\neg\mathrm{pe} \vee \mathrm{sa})$. Assume that an agent receives three instructions such that: $\Phi(S_1) \equiv \mathrm{sa}$, $\Phi(S_2) \equiv \mathrm{su}$ and $\Phi(S_3) \equiv \mathrm{pe}$. There are two conflicts $A = \{S_1, S_2, S_K\}$ and $B = \{S_2, S_3, S_K\}$. Let us suppose that $\underline{A} = \{S_1, S_2\}$ (since using salt or using sugar can equally occur) and $\underline{B} = \{S_3\}$ (since pepper is less used than other ingredients). Here, we only have $A \leqslant_{\mathrm{I}} B$ (due to S_1) then min($\mathscr{C}(\mathscr{S})$) $= \{A\}$. To solve A, S_1 and S_2 are selected and the final coherent subbase is then $E = \{S_3, S_K\}$ from which* pe, sa *and* $\neg$su *are computed. This means that using salt and pepper is preferred to using sugar and intuitively contradicts the initial knowledge and beliefs where the use of sugar is never less preferred than the use of other ingredients.*

4.2.2. Using minimal local candidates

In previous inference relations, there are too many ignored sources than what is necessary to restore the coherence of the knowledge base. In order to reduce the number of ignored sources, we first extend the notion of minimal candidates (De Kleer, 1986; Reiter, 1987; Papini, 1992) to the case where the stratification is given locally. We recall that minimal candidates are minimal sets of sources such that they contain at least one source in each conflict. The word 'minimal' is here understood in the sense of cardinality (De Kleer, 1990), namely:

DEFINITION 15. Let $\mathscr{C}$ be a set of conflicts. A set of sources *Cand* is a minimal candidate if (1) $\forall C \in \mathscr{C}$, $C \cap Cand \neq \emptyset$ and (2) $\nexists Cand'$ verifying (1) with $|Cand'| < |Cand|$.

A local minimal candidate (defined below) will be a minimal candidate such that each source belongs to the *non-dominant* stratum of conflicts which are in $\min(\mathscr{C}(\mathscr{S}))$.

DEFINITION 16. Let $\mathscr{S}$ be a set of sources. A subbase $\mathscr{S}_0$ of $\mathscr{S}$ is a local minimal candidate iff

(1) $\forall S \in \mathscr{S}_0 \exists C \in \min(\mathscr{C}(\mathscr{S}))$ such that $S \in \underline{C}$,
(2) $\forall C \in \min(\mathscr{C}(\mathscr{S})), C \cap \mathscr{S}_0 \neq \emptyset$ and
(3) $\nexists \mathscr{S}_1$ verifying (1) and (2) such that $|\mathscr{S}_1| < |\mathscr{S}_0|$.

The point (2) of the definition guarantees that local minimal candidates solve all the prioritary conflicts.

We now indicate how to compute the local minimal candidates. The idea is to start with the Kernel-min($\mathscr{C}(\mathscr{S})$) (see Definition 13). It is clear Kernel-min satisfies conditions (1) and (2) of the definition of a local minimal candidate. Then, we are looking for subsets of Kernel-min($\mathscr{C}$) verifying also condition (3). For this, we start by looking for the subsets of the Kernel-min obtained by removing one source. If the removed source preserves condition (2), then we get better candidates than the Kernel-min and we replace Kernel-min by this subset. This process is repeated until the computed subsets cannot be candidates (hence the current set of solutions is the set of minimal local candidates). This leads to the following algorithm where $\mathscr{D}$ and $\mathscr{D}'$ are sets of candidates, namely set of sets of sources:

Input: $\mathscr{C}$ – a set of conflicts
Output: *Cand* – the set of local minimal candidates

1. $\mathscr{C}urrent \leftarrow \{\text{Kernel-min}(\mathscr{C})\}$

2. repeat until $\mathscr{C}urrent = \emptyset$

   2.1. $\mathscr{S}ol \leftarrow \mathscr{C}urrent$
   2.2. $\mathscr{C}urrent \leftarrow \emptyset$; $\mathscr{I}nterm \leftarrow \emptyset$
   2.3. for all set $\mathscr{X}_i$ of $\mathscr{S}ol$ do
   $\forall S_j \in \mathscr{X}_i$ do
   if $\mathscr{X}_i - \{S_j\} \notin \mathscr{I}nterm$
   then
   2.3.1. $\mathscr{I}nterm \leftarrow \mathscr{I}nterm \cup \mathscr{X}_i - \{S_j\}$
   2.3.2. if $\mathscr{X}_i - \{S_j\}$ satisfies (2) of Definition 16
   and $\mathscr{X}_i - \{S_j\} \notin \mathscr{C}urrent$
   then $\mathscr{C}urrent \leftarrow \mathscr{C}urrent \cup \{\mathscr{X}_i - \{S_j\}\}$

3. $Cand \leftarrow \mathscr{S}ol$

In this algorithm, $\mathscr{C}urrent$ is the set containing the *potential* minimal candidates computed in loop 2.

The algorithm starts with Kernel-min and computes new subsets from it. The subsets contain sources which are in the non-dominant stratum of a prioritary conflict (hence, condition (1) of Definition 16 is always satisfied). Condition (3) is always satisfied since the cardinality of the sets is decreased after each loop. Step 2.4.2 guarantees that point (2) of Definition 16 is satisfied.

The computation of consistent subbases is given by the following algorithm, where Sup$\mathscr{S}$ denotes the set of sources to ignore and EnsC denotes the set of remaining conflicts at each step:

Input : $\mathscr{C}$ – a set of conflicts
Output : E – a set of extensions

1. let Sup$\mathscr{S} \leftarrow \emptyset$, let EnsC $\leftarrow \mathscr{C}(\mathscr{S})$
2. repeat until EnsC $= \emptyset$

   2.1. compute min(EnsC) for determining prioritary conflicts
   2.2. Let $\mathscr{S}_0$ be a local minimal candidate of min(EnsC)
   2.2.1. Sup$\mathscr{S} \leftarrow$ Sup$\mathscr{S} \cup \mathscr{S}_0$
   2.2.2. EnsC $\leftarrow$ EnsC $-\{C' \mid C' \in$ EnsC and $C' \cap \mathscr{S}_0 \neq \emptyset\}$

3. return $E = \mathscr{S} -$ Sup$\mathscr{S}$

As there are several local minimal candidates, there will be several consistent subbases (step 2.2). The set of the final selected consistent subbases is denoted by $\mathscr{E}_{\mathrm{card}}(\mathscr{S})$.

DEFINITION 17. A formula ψ is a *candidate conclusion* of a base Σ (denoted $\Sigma \mathrel{|\!\sim}_{Cand} \psi$) iff $\forall E_j \in \mathscr{E}_{\mathrm{card}}(\mathscr{S})$, $\bigcup \Phi(S_i) \vdash \psi$, with $S_i \in E_j$.

The following Example 5 explains how the candidate inference relation works.

EXAMPLE 5. *Let us consider the base* $\Sigma = \{\Phi(S_1) \equiv a, \Phi(S_2) \equiv b, \Phi(S_3) \equiv c, \Phi(S_4) \equiv d, \Phi(S_5) \equiv \neg d, \Phi(S_6) \equiv \neg a \vee \neg b, \Phi(S_7) \equiv \neg b \vee \neg c, \Phi(S_8) \equiv \neg c \vee \neg d\}$. *Figure 8 describes a stratification of the conflicts corresponding to* Σ *and Figure 9 describes the process of selecting consistent subbases according to candidate inference relation.*

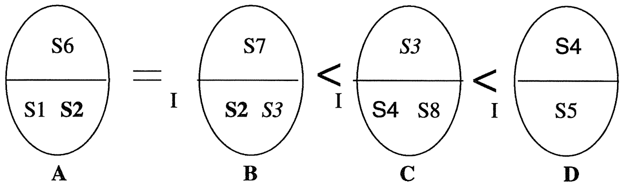

Figure 8. A stratification of conflicts for base of Example 5

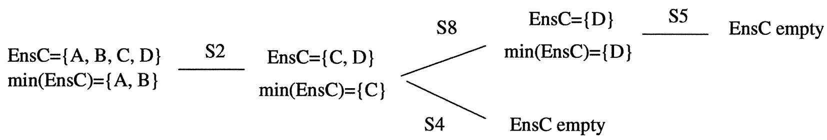

Figure 9. Solving conflicts for base of Example 5

There are two solutions for restoring consistency and then two extensions: $E_1 = \{\Phi(S_1) \equiv a, \Phi(S_3) \equiv c, \Phi(S_5) \equiv \neg d, \Phi(S_6) \equiv \neg a \vee \neg b, \Phi(S_7) \equiv \neg b \vee \neg c, \Phi(S_8) \equiv \neg c \vee \neg d\}$ *and* $E_2 = \{\Phi(S_1) \equiv a, \Phi(S_3) \equiv c, \Phi(S_4) \equiv d, \Phi(S_6) \equiv \neg a \vee \neg b, \Phi(S_7) \equiv \neg b \vee \neg c\}$. *The belief* c *is inferred from both* E_1 *and* E_2 *so* c *is a candidate consequence of* Σ. *The belief* $\neg d$ *is only inferred from* E_1 *while the belief* d *is only inferred from* E_2. *So, neither* $\neg d$ *nor* d *are candidate consequence of* Σ.

PROPOSITION 2. *If* ψ *is an incontestable consequence of* Σ *then* ψ *is a candidate consequence of* Σ.

The following example shows that the selective inference relation and the candidate inference relation are incomparable.

EXAMPLE 6. *Let the base be* $\Sigma = \{\Phi(S_1) \equiv \neg a \vee \neg b, \Phi(S_2) \equiv \neg b \vee \neg c, \Phi(S_3) \equiv a, \Phi(S_4) \equiv b, \Phi(S_5) \equiv c, \Phi(S_6) \equiv \neg c\}$. *There are three conflicts* $A = \{S_1, S_3, S_4\}$, $B = \{S_2, S_4, S_5\}$ *and* $C = \{S_5, S_6\}$ *where we assume that* $\underline{A} = \{S_3, S_4\}$, $\underline{B} = \{S_4, S_5\}$ *and* $\underline{C} = \{S_6\}$. *Here,* c *is a candidate consequence while* $\neg c$ *is a selective consequence.*

The problem with the candidate inference relation is that it depends on the repetition of formulas. If we take the example of the base $\Sigma = \{\Phi(S_1) \equiv a, \Phi(S_2) \equiv \neg a, \Phi(S_3) \equiv a\}$. There are two conflicts $A = \{S_1, S_2\}$ and $B = \{S_2, S_3\}$. Let us suppose that $\underline{A} = \{S_1, S_2\}$ and $\underline{B} = \{S_2, S_3\}$. The use of local minimal candidates leads to ignore S_2. So, the conclusion a is deduced only because it is given by two sources instead of one for $\neg a$. This kind of inference can make sense if we are dealing with independent sources.

4.3. Ignoring one source per step

In this subsection, only one source is ignored at each loop of the resolution process.

4.3.1. Universal inference relation

The first idea is to remove a source belonging to a non-dominant stratum of a conflict of min(EnsC) such that it solves at least one prioritary conflict. The computation of consistent subbases is given by the following algorithm (Sup$\mathscr{S}$ is still a set of sources to ignore and EnsC the set of remaining conflicts):

Input: $\mathscr{C}$ – a set of conflicts
Output: E – a set of extensions

1. let Sup$\mathscr{S} \leftarrow \emptyset$, let EnsC $\leftarrow \mathscr{C}(\mathscr{S})$
2. repeat until EnsC $= \emptyset$
   - 2.1. compute min(EnsC)
   - 2.2. Let $S_i \in \underline{C}$ with $C \in$ min(EnsC)
     - 2.2.1. Sup$\mathscr{S} \leftarrow$ Sup$\mathscr{S} \cup \{S_i\}$
     - 2.2.2. EnsC $\leftarrow$ EnsC $-\{C' \mid C' \in$ EnsC and $S_i \in C\}$
3. return $E = \mathscr{S} -$ Sup$\mathscr{S}$

The resulting consistent subsets are the different sets E that can be computed when considering in each step 2.2. all the possible sources to ignore. However, taking into account all the sources does not guarantee that the computed extensions are maximal as illustrated in example 7.

EXAMPLE 7. *Let* $\Sigma = \{\Phi(S_1) \equiv \neg a \vee \neg b, \Phi(S_2) \equiv \neg b \vee \neg c, \Phi(S_3) \equiv a,$ $\Phi(S_4) \equiv b, \Phi(S_5) \equiv c, \Phi(S_6) \equiv d, \Phi(S_7) \equiv \neg c \vee \neg d\}$. *There are three conflicts and we suppose that they are stratified as in Figure 10. The resolution corresponds to the Figure 11. In this tree, the first branch is not minimal and the fourth branch is redundant with the second one. In fact, there are only three maximal consistent subbases.*

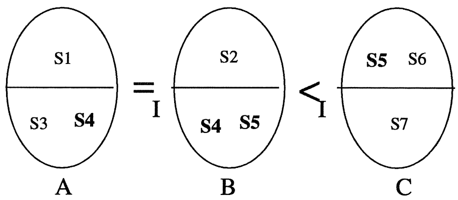

Figure 10. A stratification of conflicts for base of Example 7

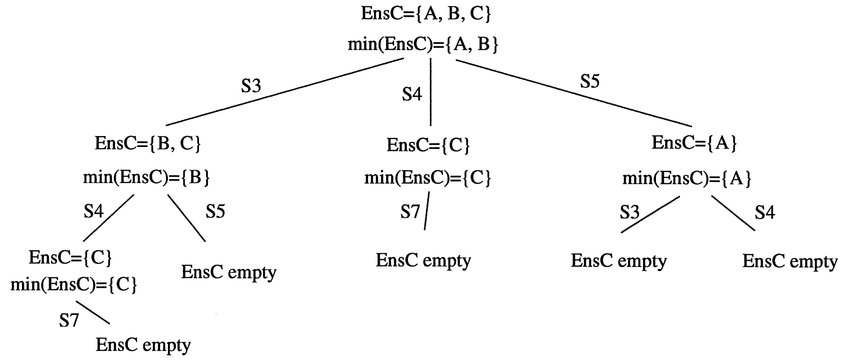

Figure 11. Solving conflicts for base of Example 7

Then, it is natural to avoid developing undesirable branches. For this aim, we develop the tree with a breadth-first search. Subtrees of a given node will be developed only if such node is not yet considered. Let us denote:

EnsC-N the remaining conflicts to solve associated with the node N and Sup S-N the selected sources to ignore associated with the node N and let $\mathscr{N}$ be the set of nodes which have been yet studied.

A node N is a non-minimal node if there exists a node $N' \in \mathscr{N}$ such that: EnsC-N' = EnsC-N and Sup S-$N' \subseteq$ Sup S-N.

A node N is redundant if there exists a node $N' \in \mathscr{N}$ such that Sup S-N = Sup S-N'.

This leads to elaguate the tree by avoiding developing all the nodes. In our example, we find the three correct solutions (see Figure 12).

The set of extensions which are maximal (that is computed after elaguing the tree) is denoted by $\mathscr{E}_{\text{univ}}(\mathscr{S})$. From this set, we can give the Definition 18 of universal inference relation.

DEFINITION 18. A formula ψ is a *universal conclusion* of a base Σ (denoted $\Sigma \mid\!\sim_{\text{Univ}} \psi$) iff for all $E_j \in \mathscr{E}_{\text{univ}}(\mathscr{S})$, $\bigcup \Phi(S_i) \vdash \psi$, with $S_i \in E_j$.

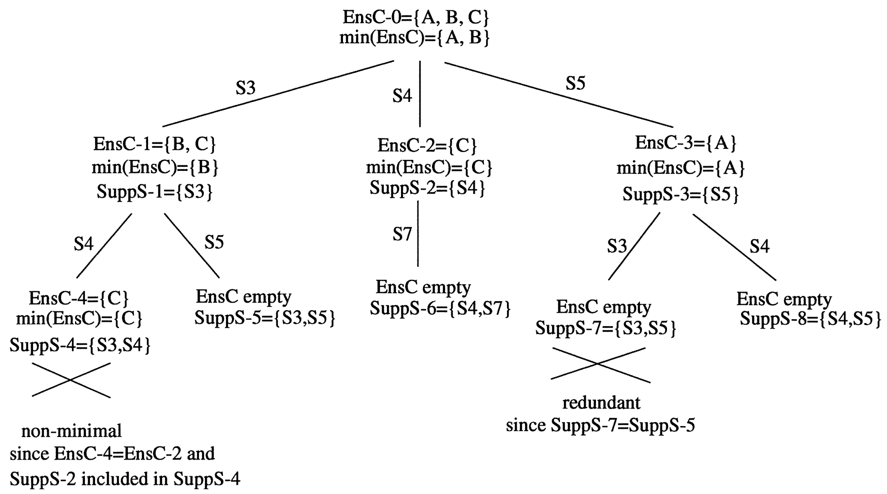

Figure 12. Elaguing tree of Example 7

We can link the universal consequence relation with other consequence relations. This is done in Propositions 3 and 4.

PROPOSITION 3. *If ψ is an incontestable consequence of Σ then ψ is a universal consequence of Σ.*

PROPOSITION 4. *The universal consequence relation is neither comparable with the selective consequence relation nor with the candidate consequence relation.*

The universal inference relation is not too cautious and it avoids to infer undesirable conclusions. However, its drawback is that all the combinations must be computed. The next subsection gives ways to choose particular sources to ignore at each loop.

4.3.2. Ordered inference relation

The idea of this inference relation is that, when there is choice between two sources to ignore, we prefer the one which solves maximal number of conflicts in min($\mathscr{C}$). Therefore, the step 2.2 of the algorithm used to determine the extensions of the universal consequence relation (that is 'Let $S_i \in \underline{C}$ with $C \in$ min(EnsC)') will be modified to select a source $S_i \in \underline{C}$ with $C \in$ min(EnsC) such that the source S_i solves the more possible conflicts.

Inference relation	Priority between conflicts	Selecting several subbases	Solving conflicts one source at the same time	Selecting a particular source
Incontestable	no	no	no	–
Selective	yes	no	no	–
Candidate	yes	yes	no	–
Universal	yes	yes	yes	no
Ordered	yes	yes	yes	yes

Table 3. Inference relations and their corresponding strategies

To each source S, we can associate a set denoted by $\mathrm{Conf}(S)$ representing the set of conflicts in $\min(\mathscr{C})$ solved when S is ignored, namely:

$$\mathrm{Conf}(S) = \{C \in \min(\mathscr{C}) \mid S \in C\}).$$

This set is used to define a priority relation between sources.

DEFINITION 19. Let S_i be a source of $\underline{C}$ with $C \in \min(\mathrm{EnsC})$. S_i is *order-selected* iff $\nexists S_j \in \underline{C'}$ with $C' \in \min(\mathrm{EnsC})$ such that: $|\mathrm{Conf}(S_j)| > |\mathrm{Conf}(S_i)|$.

The computation of consistent subbases is the same as the one given by the algorithm used to determine extensions for the universal consequence relation except that step 2.2 is replaced by 'Let $S_i \in \underline{C}$ with $C \in \min(\mathrm{EnsC})$ such that S_i is order-selected'. The set of extensions is denoted by $\mathscr{E}_{ord}(\mathscr{S})$.

DEFINITION 20. A formula ψ is an *ordered conclusion* of a base Σ (denoted $\Sigma \mid\!\sim_{\mathrm{ord}} \psi$) iff for all $E_j \in \mathscr{E}_{\mathrm{ord}}(\mathscr{S})$, $\bigcup \Phi(S_i) \vdash \psi$, with $S_i \in E_j$.

4.4. Summary and comparative study

Table 3 summarizes local inference relations described in this paper. For each inference relation, we precise if it exploits or not the natural ordering between conflicts, if it selects one or several consistent subbases, and if it selects one or several sources to ignore at each step of restoring the coherence of the knowledge base.

Figure 13 gives the inclusion between sets of conclusions obtained with each relation. For example, an arrow between Incontestable consequence and Universal consequence means that if α is an incontestable consequence of Σ then α is a universal consequence of Σ. Note that all these inference relations collapse with classical logic inference when the set of sources is consistent.

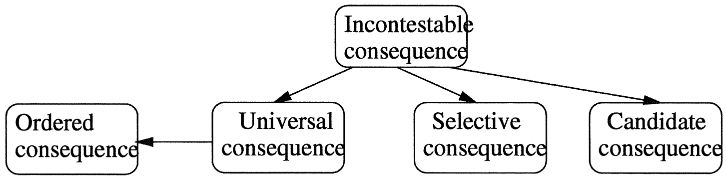

Figure 13. Links between the local inference relations

5. Conclusion

This paper has investigated several local approaches to deal with inconsistent information. An immediate advantage is the use of local stratification which recovers both total preorder, partial preorder and multiple order. Moreover, contextual preferences can also be represented. Several strategies have been used to restore, in a parsimonious way, the consistency of knowledge bases. This is done with the help of defining a natural order in which conflicts should be solved. This natural order guarantees that independent conflicts are handled separately.

In (Garcia, 1998), a comparative study with existing systems, when we restrict to a total preorder, has been done. Moreover, the application of some local inference relations have been studied in the framework of default reasoning (Benferhat and Garcia, 1998). A future work is to apply our local approaches to Geographical Information Systems, as it is done in (Wurbel et al., 2000).

References

Alchourron, C., P. Gärdenfors and D. Makinson. On the logic of theory change: partial meet contraction and revision functions. *J. of Symb. Logic*, 50, 1985.

Amgoud, L., and C. Cayrol. On the acceptability of arguments in preference-based argumentation. *Proceedings of 14th Conference on Uncertainty in Artificial Intelligence (UAI'98)*, 1–7, 1998.

Amgoud, L., and C. Cayrol. A reasoning model based on the production of acceptable arguments. *8th Int. Workshop on Non-Monotonic Reasoning (NMR'2000)*, 2000.

Benferhat, S., C. Cayrol, D. Dubois, J. Lang and H. Prade. Inconsistency management and prioritized syntax-based entailment. *Proceedings of 13th International Joint Conference on Artificial Intelligence (IJCAI'93)*, 640–645, 1993.

Benferhat, S., D. Dubois and H. Prade. Representing default rules in possibilistic logic. *Proceedings of 3rd International Conference of Principles of Knowledge Representation and Reasoning (KR'92)*, 673–684, 1992.

Benferhat, S., D. Dubois and H. Prade. How to infer from inconsistent beliefs without revising? *Proc. of 14th Int. Joint Conf. on Artif. Intelligence (IJCAI'95)*, 1995.

Benferhat, S., D. Dubois, J. Lang, H. Prade, A. Saffiotti and P. Smets. A general approach for inconsistency handling and merging information in prioritized knowledge bases. *Proceedings of 6th International Conference of Principles of Knowledge Representation and Reasoning (KR'98)*, 466–477, 1998.

Benferhat, S., and L. Garcia. A local handling of inconsistent knowledge and default bases. *Applications of Uncertainty Formalisms*, LNAI State-of-the-Art Survey, LNAI 1455, 325–353, 1998.

Besnard, P., and A. Hunter. Quasi-classical logic: non-trivializable classical reasoning from inconsistent information. *Proc. of ECSQARU'95*, LNAI 946, 44–51, 1995.

Brewka, G. Preferred subtheories: an extended logical framework for default reasoning. *Proceedings of 11th International Joint Conference on Artificial Intelligence (IJCAI'89)*, 1043–1048, 1989.

Brewka, G. Reasoning about priorities in default logic. *Proceedings of National Conference on Artificial Intelligence (AAAI'94)*, 940–945, 1994.

Castell, T., C. Cayrol, M. Cayrol and D. Le Berre. Using the Davis and Putnam procedure for an efficient computation of preferred models. *Proceedings of 12th European Conference on Artificial Intelligence (ECAI'96)*, 1996.

Cayrol, C., V. Royer and C. Saurel. Managements of preferences in assumption-based reasoning. *Advanced Methods in Artificial Intelligence*, LNCS 682, 13–22, 1993.

Cholvy, L. Automated reasoning with merged contradictory information whose reliability depends on topics. *Proceedings of ECSQARU'95*, LNAI 946, 1995.

Cholvy, L., and T. Hunter. Information fusion in logic: a brief overview. *Proc. of Qual. and Quant. Practical Reasoning (ECSQARU'97)*, LNAI 1244, 86–95, 1997.

Da Costa, N. C. A. Calcul propositionnel pour les systèmes formels inconsistants. *Compte Rendu Académie des Sciences (Paris)*, 257:3790–3792, 1963.

Damasio, C. V., and L. M. Pereira. A paraconsistent semantics with contradiction support detection. *Proceedings of 4th Conference on Logic Programming and NonMonotonic Reasoning (LPNMR'97)*, LNAI 1265, 224–243, 1997.

Davis, M., and H. Putnam. A computing procedure for quantification theory. *Journal of the Assoc. for Computing Machinery*, 7:201–215. 1960.

De Kleer, J. *Local methods for localizing faults in electronic circuits*, MIT AI Memo 394, 1976.

De Kleer, J. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

De Kleer, J. Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45, 1990.

Dubois, D., J. Lang and H. Prade. Possibilistic logic. *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol3:439–513, 1994.

Dung, P. M. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357. 1995.

Dupin de Saint-Cyr, F., J. Lang and T. Schiex. Penalty logic and its link with Dempster-Shafer theory. *Proceedings of the 10th Conf. on Uncertainty in Artificial Intelligence (UAI'94)*, 204–211, 1994.

Garcia, L. *Raisonnement local dans les bases de connaissances ordonnées : Applications au traitement des incohérences et au raisonnement plausible*, PhD Thesis. Université Paul Sabatier, Toulouse, France, 1998.

Gärdenfors, P. *Knowledge in flux – Modeling the dynamic of epistemic states*, MIT Press, 1988.

Lehmann, D. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15:61–82, 1995.

Moulin, H. *Axioms of cooperative decision making*, Wiley, New York, 1988.

Nebel, B. Base revision operator and schemes: semantics representation and complexity. *Proc. of 11th European Conf. on Artif. Intelligence (ECAI'94)*, 1994.

Papini, O. A complete revision function in propositional calculus. *Proceedings of 10th European Conference on Artificial Intelligence (ECAI'92)*, 1992.

Pearl, J. System Z: a natural ordering of defaults with tractable applications to default reasoning. *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge (TARK'90)*, 121–135, 1990.

Provan, G. M. The computational complexity of assumption-based truth maintenance systems. *Tech. Report*, Univ. of British Columbia, Vancouver, 1988.

Reiter, R. A theory of diagnosis from first principles. *Artif. Intelligence*, 41, 1987.

Rescher, N. *Plausible reasoning: An introduction to the theory and practice of plausibilistic inference*, 1976.

Roos, N. A logic for reasoning with inconsistent knowledge. *Artif. Intell.*, 57, 1992.

Tessier, C., L. Chaudron and H. J. Müller (eds). *Conflicting Agents – Conflict management in multi-agent systems*, Kluwer Academic Publishers, 2000.

Toulmin, S. *The uses of argument*, Cambridge University Press, 1956.

Wassermann, R. An Algorithm for Belief Revision. *Proc. of 7th International Conf. of Principles of Knowledge Representation and Reasoning (KR'2000)*, 2000.

Williams, M.-A. Transmutations of Knowledge Systems. *Proc. of 4th International Conf. of Principles of Knowledge Representation and Reasoning (KR'94)*, 1994.

Williams, M.-A. A practical approach to belief revision: reason-based change. *Proceedings of 5th International Conference of Principles of Knowledge Representation and Reasoning (KR'96)*, 412–421, 1996.

Wurbel, E., O. Papini and R. Jeansoulin. Revision: an application in the framework of GIS. *Proceedings of 7th International Conference of Principles of Knowledge Representation and Reasoning (KR'2000)*, 505–515, 2000.

Yang, Q. A theory of conflict resolution in planning. *Artificial Intelligence*, 58:361–392, 1992.

SALEM BENFERHAT
Institut de Recherche en Informatique de Toulouse (IRIT)
Université Paul Sabatier
118, route de Narbonne
31065 Toulouse Cedex, France
benferhat@irit.fr

LAURENT GARCIA
Laboratoire d'Études et de Recherche
en Informatique d'Angers (LERIA)
Université d'Angers
2, boulevard Lavoisier
49045 Angers Cedex 01, France
garcia@info.univ-angers.fr