Confluence Operators:
Negotiation as Pointwise Merging [1]

Sébastien Konieczny

CRIL - CNRS

Université d'Artois, Lens, France

konieczny@cril.fr

Ramón Pino Pérez

Departamento de Matemáticas, Facultad de ciencias

Universidad de Los Andes, Mérida, Venezuela

pino@ula.ve

Abstract

In the logic based framework of knowledge representation and reasoning many operators have been defined in order to capture different kinds of change: revision, update, merging and many others. There are close links between revision, update, and merging. Merging operators can be considered as extensions of revision operators to multiple belief bases. And update operators can be considered as pointwise revision, looking at each model of the base, instead of taking the base as a whole. Thus, a natural question is the following one: Are there natural operators that are pointwise merging, just as update are pointwise revision? The goal of this work is to give a positive answer to this question. In order to do that, we introduce a new class of operators: the confluence operators. These new operators can be useful in modelling negotiation processes.

1 Introduction

Belief change theory has produced a lot of different operators that models the different ways the beliefs of one (or some) agent(s) evolve over time. Among these operators, one can quote revision [1, 5, 10, 6], update [9, 8], merging [23, 14], abduction [20], extrapolation [4], etc.

In this paper we will focus on revision, update and merging. Let us first briefly describe these operators informally:

Revision Belief revision is the process of accomodating a new piece of evidence that is more reliable than the current beliefs of the agent. In belief revision the world is static, it is the beliefs of the agents that evolve.

Update In belief update the new piece of evidence denotes a change in the world. The world is dynamic, and these (observed) changes modify the beliefs of the agent.

Merging Belief merging is the process of defining the beliefs of a group of agents. So the question is: Given a set of agents that have their own beliefs, what can be considered as the beliefs of the group?

Apart from these intuitive differences between these operators, there are also close links between them. This is particularly clear when looking at the technical definitions. There are close relationship between revision [1, 5, 10] and KM update operators [9]. The first ones looking at the beliefs of the agents globally, the second ones looking at them locally (this sentence will be made formally clear later in the paper)[2]. There is also a close connection between revision and merging operators. In fact revision operators can be seen as particular cases of merging operators. From these two facts a very natural question arises: What is the family of operators that are a generalization of update operators in the same way merging operators generalize revision operators? Or, equivalently, what are the operators that can be considered as pointwise merging, just as KM update operators can be considered as pointwise belief revision. This can be outlined in the figure below. The aim of this

---

[1]This paper is a revised version of a paper that will be published in the Proceedings of the 11th European Conference on Logics in Artificial Intelligence (JELIA'08).

[2]See [8, 4, 15] for more discussions on update and its links with revision.

Figure 1: Revision - Update - Merging - Confluence

paper is to introduce and study the operators corresponding to the question mark. We will call these new operators confluence operators.

these new operators are more cautious than merging operators. This suggest that they can be used to define negotiation operators (see [2, 24, 22, 21, 12]), or as a first step of a negotiation process, in order to find all the possible negotiation results.

In order to illustrate the need for these new operators and also the difference of behaviour between merging and confluence we present the following small example.

Example 1 *Mary and Peter are planning to buy a car. Mary does not like a German car nor an expensive car. She likes small cars. Peter hesitates between a German, expensive but small car or a car which is not German, nor expensive and is a big car. Taking three propositional variables* German_car*,* Expensive_car *and* Small_car *in this order, Mary's desires are represented by* $mod(A) = \{001\}$ *and Peter's desires by* $mod(B) = \{111, 000\}$*. Most of the merging operators*[3] *give as solution (in semantical terms) the set* $\{001, 000\}$*. That is the same solution obtained when we suppose that Peter's desires are only a car which is not German nor expensive but a big car* $(mod(B') = \{000\})$*. The confluence operators will take into account the disjunctive nature of Peter's desires in a better manner and they will incorporate also the interpretations that are a trade-off between* 001 *and* 111*. For instance, the worlds* 011 *and* 101 *will be also in the solution if one use the confluence operator* $\diamond^{d_H, Gmax}$ *(defined in Section 7).*

This kind of operators is particularly adequate when the base describes a situation that is not perfectly known, or that can evolve in the future. For instance Peter's desires can either be imperfectly known (he wants one of the two situations but we do not know which one), or can evolve in the future (he will choose later between the two situations). In these situations the solutions proposed by confluence operators will be more adequate than the one proposed by merging operators. The solutions proposed by the confluence operators can be seen as all possible agreements in a negotiation process.

Belief merging is closely related to judgment aggregation as studied in political science and social choice theory (see e.g. [17, 18, 19]). An important difference is that in judgment aggregation there is an agenda on which the agents give their judgments. There is no such agenda in belief merging. The aim is to find the beliefs of the group. So this can be considered as a judgment aggregation problem where the agenda is the full set of formulae of the language (that are consistent with the integrity constraints). So, in a sense, judgment aggregation is an aggregation in a partial (incomplete) information framework (the only available information is about the formulae of the agenda), whereas belief merging is an aggregation in a complete (total/ideal) information framework.

Abstract negotiation processes have been studied both from belief merging [2, 24, 22, 21, 12] and judgment aggregation [16] perspectives. The definition of conciliation operators in this paper can be related to these works.

In the next section we will give the required definitions and notations. In Section 3 we will recall the postulates and representation theorems for revision, update, and merging, and state the links

---

[3]Such as $\triangle^{d_H, \Sigma}$ and $\triangle^{d_H, Gmax}$ [14].

between these operators. In Section 4 we define confluence operators. We provide a representation theorem for these operators in Section 5. In Section 6 we study the links between confluence operators and update and merging. In Section 7 we give examples of confluence operators. And we conclude in Section 8.

2 Preliminaries

We consider a propositional language $\mathcal{L}$ defined from a finite set of propositional variables $\mathcal{P}$ and the standard connectives, including $\top$ and $\bot$.

An interpretation ω is a total function from $\mathcal{P}$ to $\{0, 1\}$. The set of all interpretations is denoted by $\mathcal{W}$. An interpretation ω is a model of a formula $\phi \in \mathcal{L}$ if and only if it makes it true in the usual truth functional way. $mod(\varphi)$ denotes the set of models of the formula φ, i.e., $mod(\varphi) = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. When M is a set of models we denote by φ_M a formula such that $mod(\varphi_M) = M$.

A *base* K is a finite set of propositional formulae. In order to simplify the notations, in this work we will identify the base K with the formula φ which is the conjunction of the formulae of K^4.

A *profile* Ψ is a non-empty multi-set (bag) of bases $\Psi = \{\varphi_1, \ldots, \varphi_n\}$ (hence different agents are allowed to exhibit identical bases), and represents a group of n agents.

We denote by $\bigwedge \Psi$ the conjunction of bases of $\Psi = \{\varphi_1, \ldots, \varphi_n\}$, i.e., $\bigwedge \Psi = \varphi_1 \wedge \ldots \wedge \varphi_n$. A profile Ψ is said to be consistent if and only if $\bigwedge \Psi$ is consistent. The multi-set union is denoted by $\sqcup$.

A formula φ is complete if it has only one model. A profile Ψ is complete if all the bases of Ψ are complete formulae.

If $\leq$ denotes a pre-order on $\mathcal{W}$ (i.e., a reflexive and transitive relation), then $<$ denotes the associated strict order defined by $\omega < \omega'$ if and only if $\omega \leq \omega'$ and $\omega' \not\leq \omega$, and $\simeq$ denotes the associated equivalence relation defined by $\omega \simeq \omega'$ if and only if $\omega \leq \omega'$ and $\omega' \leq \omega$. A pre-order is *total* if $\forall \omega, \omega' \in \mathcal{W}$, $\omega \leq \omega'$ or $\omega' \leq \omega$. A pre-order that is not total is called *partial*. Let $\leq$ be a pre-order on A, and $B \subseteq A$, then $\min(B, \leq) = \{b \in B \mid \not\exists a \in B \; a < b\}$.

3 Revision, Update and Merging

Let us now recall in this section some background on revision, update and merging, and their representation theorems in terms of pre-orders on interpretations. This will allow us to give the relationships between these operators.

3.1 Revision

Definition 1 (Katsuno-Mendelzon [10]) *An operator $\circ$ is an AGM belief revision operator if it satisfies the following properties:*

(R1) $\varphi \circ \mu \vdash \mu$

(R2) *If $\varphi \wedge \mu \not\vdash \bot$ then $\varphi \circ \mu \equiv \varphi \wedge \mu$*

(R3) *If $\mu \not\vdash \bot$ then $\varphi \circ \mu \not\vdash \bot$*

(R4) *If $\varphi_1 \equiv \varphi_2$ and $\mu_1 \equiv \mu_2$ then $\varphi_1 \circ \mu_1 \equiv \varphi_2 \circ \mu_2$*

(R5) $(\varphi \circ \mu) \wedge \phi \vdash \varphi \circ (\mu \wedge \phi)$

(R6) *If $(\varphi \circ \mu) \wedge \phi \not\vdash \bot$ then $\varphi \circ (\mu \wedge \phi) \vdash (\varphi \circ \mu) \wedge \phi$*

---

[4]Some approaches are sensitive to syntactical representation. In that case it is important to distinguish between K and the conjonction of its formulae (see e.g. [13]). But operators of this work are all syntax independant.

When one works with a finite propositional language the previous postulates, proposed by Katsuno and Mendelzon, are equivalent to AGM ones [1, 5]. In [10] Katsuno and Mendelzon give also a representation theorem for revision operators, showing that each revision operator corresponds to a faithful assignment, that associates to each base a plausibility preorder on interpretations (this idea can be traced back to Grove systems of spheres [7]).

Definition 2 *A faithful assignment is a function mapping each base φ to a pre-order $\leq_\varphi$ over interpretations such that:*

1. *If $\omega \models \varphi$ and $\omega' \models \varphi$, then $\omega \simeq_\varphi \omega'$*

2. *If $\omega \models \varphi$ and $\omega' \not\models \varphi$, then $\omega <_\varphi \omega'$*

3. *If $\varphi \equiv \varphi'$, then $\leq_\varphi = \leq_{\varphi'}$*

Theorem 1 (Katsuno-Mendelzon [10]) *An operator $\circ$ is a revision operator (ie. it satisfies (R1)-(R6)) if and only if there exists a faithful assignment that maps each base φ to a total pre-order $\leq_\varphi$ such that*

$$mod(\varphi \circ \mu) = \min(mod(\mu), \leq_\varphi).$$

This representation theorem is important because it provides a way to easily define revision operators by defining faithful assignments. But also because their are similar such theorems for update and merging (we will also show a similar result for confluence), and that these representations in term of assignments allow to more easily find links between these operators.

3.2 Update

Definition 3 (Katsuno-Mendelzon [9, 11]) *An operator $\diamond$ is a (partial) update operator if it satisfies the properties* **(U1)-(U8)**. *It is a total update operator if it satisfies the properties* **(U1)-(U5)**, **(U8)**, **(U9)**.

(U1) $\varphi \diamond \mu \vdash \mu$

(U2) *If $\varphi \vdash \mu$, then $\varphi \diamond \mu \equiv \varphi$*

(U3) *If $\varphi \not\vdash \bot$ and $\mu \not\vdash \bot$ then $\varphi \diamond \mu \not\vdash \bot$*

(U4) *If $\varphi_1 \equiv \varphi_2$ and $\mu_1 \equiv \mu_2$ then $\varphi_1 \diamond \mu_1 \equiv \varphi_2 \diamond \mu_2$*

(U5) $(\varphi \diamond \mu) \wedge \phi \vdash \varphi \diamond (\mu \wedge \phi)$

(U6) *If $\varphi \diamond \mu_1 \vdash \mu_2$ and $\varphi \diamond \mu_2 \vdash \mu_1$, then $\varphi \diamond \mu_1 \equiv \varphi \diamond \mu_2$*

(U7) *If φ is a complete formula, then $(\varphi \diamond \mu_1) \wedge (\varphi \diamond \mu_2) \vdash \varphi \diamond (\mu_1 \vee \mu_2)$*

(U8) $(\varphi_1 \vee \varphi_2) \diamond \mu \equiv (\varphi_1 \diamond \mu) \vee (\varphi_2 \diamond \mu)$

(U9) *If φ is a complete formula and $(\varphi \diamond \mu) \wedge \phi \not\vdash \bot$, then $\varphi \diamond (\mu \wedge \phi) \vdash (\varphi \diamond \mu) \wedge \phi$*

As for revision, there is a representation theorem in terms of faithful assignment.

Definition 4 *A faithful assignment is a function mapping each interpretation ω to a pre-order $\leq_\omega$ over interpretations such that if $\omega \neq \omega'$, then $\omega <_\omega \omega'$.*

One can easily check that this faithful assignment on interpretations is just a special case of the faithful assignment on bases defined in the previous section on the complete base corresponding to the interpretation.

Katsuno and Mendelzon give two representation theorems for update operators. The first representation theorem corresponds to partial pre-orders.

Theorem 2 (Katsuno-Mendelzon [9, 11]) *An update operator $\diamond$ satisfies (U1)-(U8) if and only if there exists a faithful assignment that maps each interpretation ω to a partial pre-order $\leq_\omega$ such that*

$$mod(\varphi \diamond \mu) = \bigcup_{\omega \models \varphi} \min(mod(\mu), \leq_{\varphi_{\{\omega\}}})$$

And the second one corresponds to total pre-orders.

Theorem 3 (Katsuno-Mendelzon [9, 11]) *An update operator $\diamond$ satisfies (U1)-(U5), (U8) and (U9) if and only if there exists a faithful assignment that maps each interpretation ω to a total pre-order $\leq_\omega$ such that*

$$mod(\varphi \diamond \mu) = \bigcup_{\omega \models \varphi} \min(mod(\mu), \leq_{\varphi_{\{\omega\}}})$$

3.3 Merging

Definition 5 (Konieczny-Pino Pérez [14]) *An operator $\triangle$ mapping a pair Ψ, μ (profile, formula) into a formula denoted $\triangle_\mu(\Psi)$ is an IC merging operator if it satisfies the following properties:*

(IC0) $\triangle_\mu(\Psi) \vdash \mu$

(IC1) *If μ is consistent, then $\triangle_\mu(\Psi)$ is consistent*

(IC2) *If $\bigwedge \Psi$ is consistent with μ, then $\triangle_\mu(\Psi) \equiv \bigwedge \Psi \wedge \mu$*

(IC3) *If $\Psi_1 \equiv \Psi_2$ and $\mu_1 \equiv \mu_2$, then $\triangle_{\mu_1}(\Psi_1) \equiv \triangle_{\mu_2}(\Psi_2)$*

(IC4) *If $\varphi_1 \vdash \mu$ and $\varphi_2 \vdash \mu$, then $\triangle_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_1$ is consistent if and only if $\triangle_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_2$ is consistent*

(IC5) $\triangle_\mu(\Psi_1) \wedge \triangle_\mu(\Psi_2) \vdash \triangle_\mu(\Psi_1 \sqcup \Psi_2)$

(IC6) *If $\triangle_\mu(\Psi_1) \wedge \triangle_\mu(\Psi_2)$ is consistent, then $\triangle_\mu(\Psi_1 \sqcup \Psi_2) \vdash \triangle_\mu(\Psi_1) \wedge \triangle_\mu(\Psi_2)$*

(IC7) $\triangle_{\mu_1}(\Psi) \wedge \mu_2 \vdash \triangle_{\mu_1 \wedge \mu_2}(\Psi)$

(IC8) *If $\triangle_{\mu_1}(\Psi) \wedge \mu_2$ is consistent, then $\triangle_{\mu_1 \wedge \mu_2}(\Psi) \vdash \triangle_{\mu_1}(\Psi)$*

There is also a representation theorem for merging operators in terms of pre-orders on interpretations [14].

Definition 6 *A syncretic assignment is a function mapping each profile Ψ to a total pre-order $\leq_\Psi$ over interpretations such that:*

1. *If $\omega \models \Psi$ and $\omega' \models \Psi$, then $\omega \simeq_\Psi \omega'$*

2. *If $\omega \models \Psi$ and $\omega' \not\models \Psi$, then $\omega <_\Psi \omega'$*

3. *If $\Psi_1 \equiv \Psi_2$, then $\leq_{\Psi_1} = \leq_{\Psi_2}$*

4. *$\forall \omega \models \varphi \ \exists \omega' \models \varphi' \ \omega' \leq_{\{\varphi\} \sqcup \{\varphi\}'} \omega$*

339

5. *If* $\omega \leq_{\Psi_1} \omega'$ *and* $\omega \leq_{\Psi_2} \omega'$, *then* $\omega \leq_{\Psi_1 \sqcup \Psi_2} \omega'$

6. *If* $\omega <_{\Psi_1} \omega'$ *and* $\omega \leq_{\Psi_2} \omega'$, *then* $\omega <_{\Psi_1 \sqcup \Psi_2} \omega'$

Theorem 4 (Konieczny-Pino Pérez [14]) *An operator* $\triangle$ *is an IC merging operator if and only if there exists a syncretic assignment that maps each profile* Ψ *to a total pre-order* $\leq_\Psi$ *such that*

$$mod(\triangle_\mu(\Psi)) = \min(mod(\mu), \leq_\Psi)$$

3.4 Revision vs Update

Intuitively revision operators bring a minimal change to the base by selecting the most plausible models among the models of the new information. Whereas update operators bring a minimal change to each possible world (model) of the base in order to take into account the change described by the new infomation whatever the possible world. So, if we look closely to the two representation theorems (propositions 1, 2 and 3), we easily find the following result:

Theorem 5 *If* $\circ$ *is a revision operator (i.e. it satisfies (R1)-(R6)), then the operator* $\diamond$ *defined by:*

$$\varphi \diamond \mu = \bigvee_{\omega \models \varphi} \varphi_{\{\omega\}} \circ \mu$$

is an update operator that satisfies (U1)-(U9).

Moreover, for each update operator $\diamond$, *there exists a revision operator* $\circ$ *such that the previous equation holds.*

As explained above this proposition states that update can be viewed as a kind of pointwise revision.

3.5 Revision vs Merging

Intuitively revision operators select in a formula (the new evidence) the closest information to a ground information (the old base). And, identically, IC merging operators select in a formula (the integrity constraints) the closest information to a ground information (a profile of bases).

So following this idea it is easy to make a correspondence between IC merging operators and belief revision operators [14]:

Theorem 6 (Konieczny-Pino Pérez [14]) *If* $\triangle$ *is an IC merging operator (it satisfies (IC0-IC8)), then the operator* $\circ$, *defined as* $\varphi \circ \mu = \triangle_\mu(\varphi)$, *is an AGM revision operator (it satisfies (R1-R6)).*

See [14] for more links between belief revision and merging.

4 Confluence operators

So now that we have made clear the connections sketched in figure 1 between revision, update and merging, let us turn now to the definition of confluence operators, that aim to be a pointwise merging, similarly as update is a pointwise revision, as explained in Section 3.4. Let us first define p-consistency for profiles.

Definition 7 *A profile* $\Psi = \{\varphi_1, \ldots \varphi_n\}$ *is* p-consistent *if all its bases are consistent, i.e* $\forall \varphi_i \in \Psi$, φ_i *is consistent.*

Note that p-consistency is much weaker than consistency, the former just asks that all the bases of the profile are consistent, while the later asks that the conjunction of all the bases is consistent.

Definition 8 *An operator $\diamond$ is a confluence operator if it satisfies the following properties:*

(UC0) $\diamond_\mu(\Psi) \vdash \mu$

(UC1) *If μ is consistent and Ψ is p-consistent, then $\diamond_\mu(\Psi)$ is consistent*

(UC2) *If Ψ is complete, Ψ is consistent and $\bigwedge \Psi \vdash \mu$, then $\diamond_\mu(\Psi) \equiv \bigwedge \Psi$*

(UC3) *If $\Psi_1 \equiv \Psi_2$ and $\mu_1 \equiv \mu_2$, then $\diamond_{\mu_1}(\Psi_1) \equiv \diamond_{\mu_2}(\Psi_2)$*

(UC4) *If φ_1 and φ_2 are complete formulae and $\varphi_1 \vdash \mu$, $\varphi_2 \vdash \mu$,*
then $\diamond_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_1$ is consistent if and only $\diamond_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_2$ is consistent

(UC5) $\diamond_\mu(\Psi_1) \wedge \diamond_\mu(\Psi_2) \vdash \diamond_\mu(\Psi_1 \sqcup \Psi_2)$

(UC6) *If Ψ_1 and Ψ_2 are complete profiles and $\diamond_\mu(\Psi_1) \wedge \diamond_\mu(\Psi_2)$ is consistent,*
then $\diamond_\mu(\Psi_1 \sqcup \Psi_2) \vdash \diamond_\mu(\Psi_1) \wedge \diamond_\mu(\Psi_2)$

(UC7) $\diamond_{\mu_1}(\Psi) \wedge \mu_2 \vdash \diamond_{\mu_1 \wedge \mu_2}(\Psi)$

(UC8) *If Ψ is a complete profile and if $\diamond_{\mu_1}(\Psi) \wedge \mu_2$ is consistent*
then $\diamond_{\mu_1 \wedge \mu_2}(\Psi) \vdash \diamond_{\mu_1}(\Psi) \wedge \mu_2$

(UC9) $\diamond_\mu(\Psi \sqcup \{\varphi \vee \varphi'\}) \equiv \diamond_\mu(\Psi \sqcup \{\varphi\}) \vee \diamond_\mu(\Psi \sqcup \{\varphi'\})$

Some of the (UC) postulates are exactly the same as (IC) ones, just like some (U) postulates for update are exactly the same as (R) ones for revision.

In fact, (UC0), (UC3), (UC5) and (UC7) are exactly the same as the corresponding (IC) postulates. So the specificity of confluence operators lies in postulates (UC1), (UC2), (UC6), (UC8) and (UC9). (UC2), (UC4), (UC6) and (UC8) are close to the corresponding (IC) postulates, but hold for complete profiles only. The present formulation of (UC2) is quite similar to formulation of (U2) for update. Note that in the case of a complete profile the hypothesis of (UC2) is equivalent to ask coherence with the constraints, *i.e.* the hypothesis of (IC2). Postulates (UC8) and (UC9) are the main difference with merging postulates, and correspond also to the main difference between revision and KM update operators. (UC9) is the most important postulate, that defines confluence operators as pointwise agregation, just like (U8) defines update operators as pointwise revision. This will be expressed more formally in the next Section (Lemma 1).

5 Representation theorem for confluence operators

In order to state the representation theorem for confluence operators, we first have to be able to "localize" the problem. For update this is done by looking to each model of the base, instead of looking at the base (set of models) as a whole. So for "localizing" the aggregation process, we have to find what is the local view of a profile. That is what we call a state.

Definition 9 *A multi-set of interpretations will be called a* state. *We use the letter e, possibly with subscripts, for denoting states. If $\Psi = \{\varphi_1, \ldots, \varphi_n\}$ is a profile and $e = \{\omega_1, \ldots, \omega_n\}$ is a state such that $\omega_i \models \varphi_i$ for each i, we say that e is a state of the profile Ψ, or that the state e models the profile Ψ, that will be denoted by $e \models \Psi$. If $e = \{\omega_1, \ldots, \omega_n\}$ is a state, we define the profile Ψ_e by putting $\Psi_e = \{\varphi_{\{\omega_1\}}, \ldots, \varphi_{\{\omega_n\}}\}$.*

State is an interesting notion. If we consider each base as the current point of view (goals) of the corresponding agent (that can be possibly strengthened in the future) then states are all possible negotiation starting points.

States are the points of interest for confluence operators (like interpretations are for update), as stated in the following Lemma:

Lemma 1 *If $\diamond$ satisfies (UC3) and (UC9) then $\diamond$ satisfies the following*

$$\diamond_\mu(\Psi) \equiv \bigvee_{e \models \Psi} \diamond_\mu(\Psi_e)$$

Defining profile entailment by putting $\Psi \vdash \Psi'$ iff every state of Ψ is a state of Ψ', the previous Lemma has as a corollary the following:

Corollary 1 *If $\diamond$ is a confluence operator then it is monotonic in the profiles, that means that if $\Psi \vdash \Psi'$ then $\diamond_\mu(\Psi) \vdash \diamond_\mu(\Psi')$*

This monotony property, that is not true in the case of merging operators, shows one of the big differences between merging and confluence operators. Remark that there is a corresponding monotony property for update.

Like revision's faithful assignments that have to be "localized" to interpretations for update, merging's syncretic assignments have to be localized to states for confluence.

Definition 10 *A distributed assignment is a function mapping each state e to a total pre-order $\leq_e$ over interpretations such that:*

1. *$\omega <_{\{\omega,\ldots,\omega\}} \omega'$ if $\omega' \neq \omega$*

2. *$\omega \simeq_{\{\omega,\omega'\}} \omega'$*

3. *If $\omega \leq_{e_1} \omega'$ and $\omega \leq_{e_2} \omega'$, then $\omega \leq_{e_1 \sqcup e_2} \omega'$*

4. *If $\omega <_{e_1} \omega'$ and $\omega \leq_{e_2} \omega'$, then $\omega <_{e_1 \sqcup e_2} \omega'$*

Now we can state the main result of this paper, that is the representation theorem for confluence operators.

Theorem 7 *An operator $\diamond$ is a confluence operator if and only if there exists a distributed assignment that maps each state e to a total pre-order $\leq_e$ such that*

$$mod(\diamond_\mu(\Psi)) \quad = \quad \bigcup_{e \models \Psi} \min(mod(\mu), \leq_e) \tag{1}$$

Unfortunately, we have to omit the proof for space reasons. Nevertheless, we indicate the most important ideas therein. As it is usual, the *if* condition is done by checking each property without any major difficulty. In order to verify the *only if* condition we have to define a distributed assigment. This is done in the following way: for each state e we define a total pre-order $\leq_e$ by putting $\forall \omega, \omega' \in \mathcal{W} \ \omega \leq_e \omega'$ if and only if $\omega \models \diamond_{\varphi_{\{\omega,\omega'\}}}(\Psi_e)$. Then, the main difficulties are to prove that this is indeed a distributed assigment and that the equation (1) holds. In particular, Lema 1 is very helpful for proving this last equation.

Note that this theorem is still true if we remove respectively the postulate (UC4) from the required postulates for confluence operators and the condition 2 from distributed assignments.

6 Confluence vs Update and Merging

So now we are able to state the proposition that shows that update is a special case of confluence, just as revision is a special case of merging.

Theorem 8 *If $\diamond$ is a confluence operator (i.e. it satisfies (UC0-UC9)), then the operator $\diamond$, defined as $\varphi \diamond \mu = \diamond_\mu(\varphi)$, is an update operator (i.e. it satisfies (U1-U9)).*

Concerning merging operators, one can see easily that the restriction of a syncretic assignment to a complete profile is a distributed assignment. From that we obtain the following result (the one corresponding to Theorem 5):

Theorem 9 *If $\triangle$ is an IC merging operator (i.e. it satisfies (IC0-IC8)) then the operator $\diamond$ defined by*

$$\diamond_\mu(\Psi) = \bigvee_{e \models \Psi} \triangle_\mu(\Psi_e)$$

is a confluence operator (i.e. it satisfies (UC0-UC9)).
Moreover, for each confluence operator $\diamond$, there exists a merging operator $\triangle$ such that the previous equation holds.

It is interesting to note that this theorem shows that every merging operator can be used to define a confluence operator, and explains why we can consider confluence as a pointwise merging.

Unlike Theorem 5, the second part of the previous theorem doesn't follow straightforwardly from the representation theorems. We need to build a syncretic assignment extending the distributed assignment representing the confluence operator. In order to do that we can use the following construction: Each pre-order $\leq_e$ defines naturally a rank function r_e on natural numbers. Then we put

$$\omega \leq_\Psi \omega' \quad \text{if and only if} \quad \sum_{e \models \Psi} r_e(\omega) \leq \sum_{e \models \Psi} r_e(\omega')$$

As a corollary of the representation theorem we obtain the following

Corollary 2 *If $\diamond$ is a confluence operator then the following property holds:*

If $\bigwedge \Psi \vdash \mu$ and Ψ is consistent then $\bigwedge \Psi \wedge \mu \vdash \diamond_\mu(\Psi)$

But unlike merging operators, we don't have generally $\diamond_\mu(\Psi) \vdash \bigwedge \Psi \wedge \mu$.

Note that this "half of (IC2)" property is similar to the "half of (R2)" satisfied by update operators.

This corollary is interesting since it underlines an important difference between merging and confluence operators. If all the bases agree (i.e. if their conjunction is consistent), then a merging operator gives as result exactly the conjunction, whereas a confluence operator will give this conjunction plus additional results. This is useful if the bases do not represent interpretations that are considered equivalent by the agent, but uncertain information about the agent's current or future state of mind.

7 Example

In this section we will illustrate the behaviour of confluence operators on an example. We can define confluence operators very similarly to merging operators, by using a distance and an aggregation function.

Definition 11 *A pseudo-distance between interpretations is a total function $d : \mathcal{W} \times \mathcal{W} \mapsto \mathbb{R}^+$ s.t. for any $\omega, \omega' \in \mathcal{W}$: $d(\omega, \omega') = d(\omega', \omega)$, and $d(\omega, \omega') = 0$ if and only if $\omega = \omega'$.*

A widely used distance between interpretations is the Dalal distance [3], denoted d_H, that is the Hamming distance between interpretations (the number of propositional atoms on which the two interpretations differ).

Definition 12 *An* aggregation function f *is a total function associating a nonnegative real number to every finite tuple of nonnegative real numbers s.t. for any* $x_1, \ldots, x_n, x, y \in \mathbb{R}^+$:

- *if* $x \leq y$, *then* $f(x_1, \ldots, x, \ldots, x_n) \leq f(x_1, \ldots, y, \ldots, x_n)$ *(non-decreasingness)*

- $f(x_1, \ldots, x_n) = 0$ *if and only if* $x_1 = \ldots = x_n = 0$ *(minimality)*

- $f(x) = x$ *(identity)*

Sensible aggregation functions are for instance $\max$, sum, or leximax $(Gmax)$ [5] [14].

Definition 13 (distance-based confluence operators) *Let* d *be a pseudo-distance between interpretations and* f *be an aggregation function. The result* $\diamondsuit_\mu^{d,f}(\Psi)$ *of the confluence of* Ψ *given the integrity constraints* μ *is defined by:* $mod(\diamondsuit_\mu^{d,f}(\Psi)) = \bigcup_{e \models \Psi} \min(mod(\mu), \leq_e)$, *where the pre-order* $\leq_e$ *on* $\mathcal{W}$ *induced by* e *is defined by:*

- $\omega \leq_e \omega'$ *if and only if* $d(\omega, e) \leq d(\omega', e)$, *where*

- $d(\omega, e) = f(d(\omega, \omega_1) \ldots, d(\omega, \omega_n))$ *with* $e = \{\omega_1, \ldots, \omega_n\}$.

It is easy to check that by using usual aggregation functions we obtain confluence operators.

Proposition 1 *Let* d *be any distance,* $\diamondsuit_\mu^{d,\Sigma}(\Psi)$ *and* $\diamondsuit_\mu^{d,Gmax}(\Psi)$ *are confluence operators (i.e. they satisfy (UC0)-(UC9)).*

Example 2 *Let us consider a profile* $\Psi = \{\varphi_1, \varphi_2, \varphi_3, \varphi_4\}$ *and an integrity constraint* μ *defined on a propositional language built over four symbols, as follows:* $mod(\mu) = \mathcal{W} \setminus \{0110, 1010, 1100, 1110\}$, $mod(\varphi_1) = mod(\varphi_2) = \{1111, 1110\}$, $mod(\varphi_3) = \{0000\}$, *and* $mod(\varphi_4) = \{1110, 0110\}$.

$\mathcal{W}$	1111	1110	0000	0110	e_1		e_2		e_3		e_4		e_5		e_6		$\diamondsuit_\mu^{d,\Sigma}$	$\diamondsuit_\mu^{d,\mathsf{Gmax}}$
					Σ	*Gmax*	Σ	*Gmax*	Σ	*Gmax*	Σ	*Gmax*	Σ	*Gmax*	Σ	*Gmax*		
0000	*4*	*3*	*0*	*2*	*11*	*4430*	*10*	*4420*	*10*	*4330*	*9*	*4320*	*9*	*3330*	*8*	*3320*		
0001	*3*	*4*	*1*	*3*	*11*	*4331*	*10*	*3331*	*12*	*4431*	*11*	*4331*	*13*	*4441*	*12*	*4431*		
0010	*3*	*2*	*1*	*1*	*9*	*3321*	*8*	*3311*	*8*	*3221*	*7*	**3211**	**7**	**2221**	*6*	**2211**	×	×
0011	*2*	*3*	*2*	*2*	*9*	*3222*	*8*	**2222**	*10*	*3322*	*9*	*3222*	*11*	*3332*	*10*	*3322*		×
0100	*3*	*2*	*1*	*1*	*9*	*3321*	*8*	*3311*	*8*	**3221**	*7*	**3211**	**7**	**2221**	*6*	**2211**	×	×
0101	*2*	*3*	*2*	*2*	*9*	*3222*	*8*	**2222**	*10*	*3322*	*9*	*3222*	*11*	*3332*	*10*	*3322*		×
0110	*2*	*1*	*2*	*0*	*7*	*2221*	*6*	*2220*	*6*	*2211*	*5*	*2210*	*5*	*2111*	*4*	*2110*		
0111	*1*	*2*	*3*	*1*	*7*	**3211**	**6**	*3111*	*8*	**3221**	*7*	**3211**	*9*	*3222*	*8*	*3221*	×	×
1000	*3*	*2*	*1*	*3*	*9*	*3321*	*10*	*3331*	*8*	**3221**	*9*	*3321*	**7**	*2221*	*8*	*3221*	×	×
1001	*2*	*3*	*2*	*4*	*9*	*3222*	*10*	*4222*	*10*	*3322*	*11*	*4322*	*11*	*3332*	*12*	*4332*		
1010	*2*	*1*	*2*	*2*	*7*	*2221*	*8*	*2222*	*6*	*2211*	*7*	*2221*	*5*	*2111*	*6*	*2211*		
1011	*1*	*2*	*3*	*3*	*7*	**3211**	*8*	*3311*	*8*	**3221**	*9*	*3321*	*9*	*3222*	*10*	*3322*		×
1100	*2*	*1*	*2*	*2*	*7*	*2221*	*8*	*2222*	*6*	*2211*	*7*	*2221*	*5*	*2111*	*6*	*2211*		
1101	*1*	*2*	*3*	*3*	*7*	**3211**	*8*	*3311*	*8*	**3221**	*9*	*3321*	*9*	*3222*	*10*	*3322*		×
1110	*1*	*0*	*3*	*1*	*5*	*3110*	*6*	*3111*	*4*	*3100*	*5*	*3110*	*3*	*3000*	*4*	*3100*		
1111	*0*	*1*	*4*	*2*	**5**	*4100*	**6**	*4200*	**6**	*4110*	**7**	*4210*	**7**	*4111*	*8*	*4211*	×	

Table 1: Computations of $mod(\diamondsuit_\mu^{d,\Sigma}(\Psi))$ and $mod(\diamondsuit_\mu^{d,Gmax}(\Psi))$

The computations are reported in Table 1. The shadowed lines correspond to the interpretations rejected by the integrity constraints. Thus the result has to be taken among the interpretations that are not shadowed. The states that model the profile are the following ones:
$e_1 = \{1111, 1111, 0000, 1110\}$, $e_2 = \{1111, 1111, 0000, 0110\}$,
$e_3 = \{1111, 1110, 0000, 1110\}$, $e_4 = \{1110, 1111, 0000, 0110\}$,
$e_5 = \{1110, 1110, 0000, 1110\}$, $e_6 = \{1110, 1110, 0000, 0110\}$.

---

[5]leximax $(Gmax)$ is usually defined using lexicographic sequences, but it can be easily represented by reals to fit the above definition (see e.g. [13]).

For each state, the Table gives the distance between the interpretation and this state for the Σ aggregation function, and for the $Gmax$ one. So one can then look at the best interpretations for each state.

So for instance for $\diamondsuit_\mu^{d,\Sigma}(\Psi)$, e_1 selects the interpretation 1111, e_2 selects 0111 and 1111, etc. So, taking the union of the interpretations selected by each state, gives $mod(\diamondsuit_\mu^{d,\Sigma}(\Psi)) = \{0010, 0100, 0111, 1000, 1111\}$.

Similarly we obtain $mod(\diamondsuit_\mu^{d,Gmax}(\Psi)) = \{0100, 0011, 0010, 0101, 0111, 1000, 1011, 1101\}$.

8 Conclusion

We have proposed in this paper a new family of change operators. Confluence operators are pointwise merging, just as update can be seen as a pointwise revision. We provide an axiomatic definition of this family, a representation theorem in terms of pre-orders on interpretations, and provide examples of these operators.

In this paper we define confluence operators as generalization to multiple bases of total update operators (i.e. which semantical counterpart are total pre-orders). A perspective of this work is to try to extend the result to partial update operators.

As Example 1 suggests, these operator can prove meaningful to aggregate the goals of a group of agents. They seem to be less adequate for aggregating beliefs, where the global minimization done by merging operators is more appropriate for finding the most plausible worlds. This distinction between goal and belief aggregation is a very interesting perspective, since, as far as we know, no such axiomatic distinction as been ever discussed.

References

[1] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.

[2] R. Booth. Social contraction and belief negotiation. In *Proceedings of the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, pages 374–384, 2002.

[3] M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. In *Proceedings of the American National Conference on Artificial Intelligence (AAAI'88)*, pages 475–479, 1988.

[4] F. Dupin de Saint-Cyr and J. Lang. Belief extrapolation (or how to reason about observations and unpredicted change). *Proceedings of the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, pages 497–508, 2002.

[5] P. Gärdenfors. *Knowledge in flux*. MIT Press, 1988.

[6] P. Gärdenfors, editor. *Belief Revision*. Cambridge University Press, 1992.

[7] A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(157-180), 1988.

[8] A. Herzig and O. Rifi. Update operations: a review. In *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI'98)*, pages 13–17, 1998.

[9] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 387–394, 1991.

[10] H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.

[11] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Belief revision*, volume 29 of *Cambridge Tracts Theoret. Comput. Sci.*, pages 183–203. Cambridge Univ. Press, Cambridge, 1992.

[12] S. Konieczny. Belief base merging as a game. *Journal of Applied Non-Classical Logics*, 14(3):275–294, 2004.

[13] S. Konieczny, J. Lang, and P. Marquis. DA^2 merging operators. *Artificial Intelligence*, 157(1-2):49–79, 2004.

[14] S. Konieczny and R. Pino Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.

[15] J. Lang. Belief update revisited. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'07)*, pages 2517–2522, 2007.

[16] C. List. Group deliberation and the revision of judgments: An impossibility result.

[17] C. List and P. Pettit. Aggregating sets of judgments: An impossibility result. *Economics and Philosophy*, 18:89–110, 2002.

[18] C. List and P. Pettit. Aggregating sets of judgments: Two impossibility results compared. *Synthese*, 140(1-2):207–235, 2004.

[19] C. List and C. Puppe. Judgment aggregation: a survey. In C. Puppe P. Anand and P. Pattaniak, editors, *Oxford Handbook of Rational and Social Choice*. Oxford University Press, 2009.

[20] J. Lobo and C. Uzcátegui. Abductive change operators. *Fundamenta Informaticae*, 27(4):385–411, 1996.

[21] T. Meyer, N. Foo, D. Zhang, and R. Kwok. Logical foundations of negotiation: Outcome, concession and adaptation. In *Proceedings of the American National Conference on Artificial Intelligence (AAAI'04)*, pages 293–298, 2004.

[22] T. Meyer, N. Foo, D. Zhang, and R. Kwok. Logical foundations of negotiation: Strategies and preferences. In *Proceedings of the Ninth Conference on Principles of Knowledge Representation and Reasoning (KR'04)*, pages 311–318, 2004.

[23] P. Z. Revesz. On the semantics of arbitration. *International Journal of Algebra and Computation*, 7(2):133–160, 1997.

[24] D. Zhang, N. Foo, T. Meyer, and R. Kwok. Negotiation as mutual belief revision. In *Proceedings of the American National Conference on Artificial Intelligence (AAAI'04)*, pages 317–322, 2004.