

Belief change based on global minimisation

James P. Delgrande
School of Computing Science,
Simon Fraser University
Burnaby, B.C., Canada V5A 1S6
jim@cs.sfu.ca

Jérôme Lang
IRIT-CNRS,
118, route de Narbonne,
F-31062 Toulouse, France
lang@irit.fr

Torsten Schaub*
Universität Potsdam
Institut für Informatik
August-Bebel-Str. 89
D-14482 Potsdam, Germany
torsten@cs.uni-potsdam.de

Abstract

A general framework for minimisation-based belief change is presented. A problem instance is made up of an undirected graph, where a formula is associated with each vertex. For example, vertices may represent spatial locations, points in time, or some other notion of locality. Information is shared between vertices via a process of minimisation over the graph. We give equivalent semantic and syntactic characterisations of this minimisation. We also show that this approach is general enough to capture existing minimisation-based approaches to belief merging, belief revision, and (temporal) extrapolation operators. While we focus on a set-theoretic notion of minimisation, we also consider other approaches, such as cardinality-based and priority-based minimisation.

1 Introduction and Motivation

Minimisation of change is a crucial element in many approaches to knowledge representation and reasoning, including reasoning about action and time, belief revision, and belief merging. Generally speaking, these approaches take as input a collection of formulas (or sets of formulas) and have as output another set of formulas (or sets of formulas). For instance, in standard approaches to belief revision, we start from an initial belief state represented by a formula ϕ ,¹ a formula α , and come up with a new belief state $\phi * \alpha$ that satisfies α while retaining as much as possible from ϕ . In belief merging, we start from a collection $\langle \phi_1, \dots, \phi_n \rangle$ of belief states (again, single formulas), possibly with integrity constraints, and end up with a single belief state, computed by retaining as much as possible from the initial information. In a number of approaches to reasoning about time and change, the input consists of a series of observations and/or actions over a discrete time scale, and the output is series of belief states, one for each time point, obtained by minimising change over

time. (Note that, for each of these areas, there is obviously not a single nor a best way of minimising change.)

Two questions come to mind. First, can we see (some of) these change-minimising approaches as instances of a more general framework? Second, are there other specific cases of this more general setting that are worth considering? This paper answers both questions. The general setting we consider starts with a set of *points* structured in a *graph*, and with a formula attached to each point. Following a minimisation step, a formula is determined at each point, representing the original information associated with the point along with information gleaned from connected points. This minimisation step is driven by disagreements among the formulas attached to connected points. We show that a number of approaches to belief merging fit within this general setting, as do a number of approaches to reasoning about time and change, as well as an approach to revision.

The approach is intended to model any domain that can be modelled by a set of connected points, with data associated with each point. Here are some typical examples:

Multisource merging The graph is a network of interconnected agents (or knowledge bases), where each agent consistently incorporates information from other agents.

Temporal minimisation of change The graph is a set of time points, and we minimise the set of changes between successive time points.

Spatial minimisation of change The points are places or regions in space, vertices express adjacency or proximity between places/regions, and the principle of minimising change is justified by the assumption of spatial persistence between connected places or points.

Thus for the last example, if it is raining in some place, e.g. Amsterdam, and we don't know whether it is raining in a neighbouring place, e.g. Brussels, we might want to infer by default that it is also raining there. However, if we also know that it is not raining in Luxembourg, we might want to conclude that we don't know whether it is raining in Brussels.

The next section introduces basic notation and reviews related work. Section 3 presents the general approach, while Section 4 discusses specific instances of the approach. This is followed by a section on extensions to the approach, including adding weights to edges. The final section gives concluding remarks and directions for future work.

*Affiliated with the School of Computing Science at Simon Fraser University, Canada, and IIS at Griffith University, Australia.

¹More generally, initial beliefs are expressed by belief sets, i.e. closed propositional theories; it is well-known that both representations are equivalent when the set of propositional symbols is finite.

2 Related Work

The approach described herein falls into a broad category of approaches to belief change based on minimisation of change. In our approach, we employ a graph $G = \langle V, E \rangle$ to model various phenomena. Some earlier work can be regarded as a specific case of this general approach (and in Section 4 we show that this is indeed the case). Thus *belief extrapolation* [Dupin de Saint-Cyr and Lang, 2002] involve minimisation of change in a chain (with intended interpretation of the chain being a series of time points). Similarly, the approaches to *merging* of [Delgrande and Schaub, 2006] can be regarded as minimisation over complete graph or star graphs. *Chronological minimisation* corresponds to a version of the approach on a chain with priorities on edges. The approach of [Zhang *et al.*, 2004] to negotiation essentially reduces negotiation to a form of mutual belief revision.

On the other hand, the general approach is distinct from belief merging as usually understood, in that most merging operators begin with a set of belief states but yield a single belief state (see for example [Liberatore and Schaerf, 1995; Konieczny and Pino Pérez, 2002]). The *conciliation* operators of [Gauwin *et al.*, 2005] however are closer to the approach at hand: in this case there is a global merging of the agent’s knowledge after which each agent revises its beliefs by the “consensual” knowledge. Similar remarks apply to [Booth, 2002] which, for our purposes, can be regarded as employing multi-agent contraction. [Liberatore and Schaerf, 2000] describes a specific approach combining merging of multiple sources, revision and update; there the process of propagating information temporally is somewhat similar to extrapolation. One approach that conforms very well to our underlying intuitions is the REV!GIS project [Würbel *et al.*, 2000] which addresses revision in geographical information systems: flooding information (for example) is encoded for different regions, and persistence of water levels can be used for revising information in adjacent regions.

3 Global minimisation of change

Let $\mathcal{L}_{\mathcal{P}}$ be a propositional language over Alphabet $\mathcal{P} = \{p, q, \dots\}$, using the logical constants \top, \perp and connectives $\neg, \vee, \wedge, \supset$ and \leftrightarrow . Formulas are denoted by Greek letters α, φ, \dots . The set of interpretations of $\mathcal{L}_{\mathcal{P}}$ is denoted by $\Omega_{\mathcal{P}}$. (Whenever clear from the context, we omit the underlying alphabet \mathcal{P} .) We denote interpretations explicitly as lists of literals, the falsity of an atom p being indicated by \bar{p} . For instance, for $\mathcal{P} = \{p, q, r\}$, the interpretation where only q is true is denoted by $\bar{p}q\bar{r}$. $\text{Mod}(\alpha)$ is the set of models of α ; $\text{Cn}(\alpha)$ is the deductive closure of α .

For each $i > 0$, define alphabet $\mathcal{P}^i = \{p^i \mid p \in \mathcal{P}\}$; that is, \mathcal{P}^i is a copy of \mathcal{P} where atomic symbols carry the superscript i . We rely on the fact that \mathcal{P}^i and \mathcal{P}^j are disjoint for $i \neq j$. Similarly, \mathcal{L}^i is a shorthand for $\mathcal{L}_{\mathcal{P}^i}$. We associate \mathcal{P}^0 with \mathcal{P} and \mathcal{L}^0 with $\mathcal{L}_{\mathcal{P}}$. For $\alpha \in \mathcal{L}^i$, the function *ren* is defined so that $\text{ren}(\alpha, j)$ is the formula α where every occurrence of $p^i \in \mathcal{P}^i$ is replaced by the corresponding proposition $p^j \in \mathcal{P}^j$. We can thus write α^j for $\alpha \in \mathcal{L}$ as an abbreviation for $\text{ren}(\alpha, j)$. For instance, if $\alpha = (\neg p \wedge (q \vee r))$ then $\text{ren}(\alpha, 3) = (\neg p^3 \wedge (q^3 \vee r^3)) = \alpha^3$.

3.1 Completion and consensus

In this paper we make use of finite undirected graphs only. Thus for graph G , $G = \langle V, E \rangle$ where E is a set of unordered pairs of elements of V . For the sake of simplicity we associate the set of vertices V with an initial sequence of the natural numbers; thus for $|V| = n$ we have $V = \{1, \dots, n\}$. Conventionally we write an edge as (x, y) rather than $\{x, y\}$.

In the entire paper, let $G = \langle V, E \rangle$ be a graph where $|V| = n$, and \mathcal{P} an alphabet.

Definition 3.1 (Scenarios and Interpretations)

- A G -scenario Σ_G is a list $\langle \phi_1, \dots, \phi_n \rangle$, where $\phi_i \in \mathcal{L}_{\mathcal{P}}$ for each $i \in V$.
 Σ_G is consistent iff ϕ_i is consistent for all $i \in V$.
 $\Sigma_G[i]$ denotes ϕ_i , the i th component of Σ_G .
- A G -interpretation \mathcal{M}_G is a list $\langle M_1, \dots, M_n \rangle$, where each M_i is an interpretation from $\Omega_{\mathcal{P}}$.
 \mathcal{M}_G is said to satisfy Σ_G , denoted $\mathcal{M}_G \models \Sigma_G$, iff $M_i \models \phi_i$ holds for all $i \in V$.
 Ω^V is the set of all G -interpretations and $\text{Mod}(\Sigma_G) = \{\mathcal{M}_G \mid \mathcal{M}_G \models \Sigma_G\}$ is the set of all G -models of Σ_G .
- The shared or invariant knowledge of a G -scenario Σ_G , $\text{SK}(\Sigma_G)$, is the formula $\bigvee_{i=1}^n \phi_i$.

When clear from the context we drop the subscript G , and simply write Σ and \mathcal{M} rather than Σ_G and \mathcal{M}_G . Sometimes, to make notation more readable, we make use of this notation for scenarios: $\langle 1 : \phi_1, \dots, n : \phi_n \rangle$ instead of $\langle \phi_1, \dots, \phi_n \rangle$. Similarly, a G -interpretation $\langle M_1, \dots, M_n \rangle$ is sometimes denoted as $\langle 1 : M_1, \dots, n : M_n \rangle$. Slightly abusing notation, we denote by $M_i(p)$ the truth value of $p \in \mathcal{P}$ given by the interpretation M_i .

Definition 3.2 (Change sets)

- A change set for \mathcal{P} and G is a subset Δ of $E \times \mathcal{P}$.
- Let \mathcal{M} be a G -interpretation. The change set associated with \mathcal{M} , $\Delta(\mathcal{M})$, is defined as

$$\{\langle (x, y), p \rangle \mid (x, y) \in E \text{ and } M_x(p) \neq M_y(p)\}.$$
- A change set Δ is a change set for a G -scenario Σ iff there exists a G -interpretation \mathcal{M} such that $\mathcal{M} \models \Sigma$ and $\Delta = \Delta(\mathcal{M})$.
- Δ is a minimal change set for Σ if Δ is a change set for Σ and there is no change set Δ' for Σ such that $\Delta' \subset \Delta$.

Minimal change sets pick out those G -interpretations that in a certain (viz. set-containment) sense maximise “closeness” between individual interpretations of members of Σ . To be sure, there are other measures of “closeness” that can be used, a point that we return to later.

Definition 3.3 (Preferred G -interpretations) Given two G -interpretations \mathcal{M} and \mathcal{M}' , define $\mathcal{M} \succeq \mathcal{M}'$ iff $\Delta(\mathcal{M}) \subseteq \Delta(\mathcal{M}')$, and $\mathcal{M} \succ \mathcal{M}'$ iff $\Delta(\mathcal{M}) \subset \Delta(\mathcal{M}')$.

Given a G -scenario Σ and a G -interpretation \mathcal{M} , we say that \mathcal{M} is a preferred G -interpretation of Σ iff

1. $\mathcal{M} \models \Sigma$;
2. there is no \mathcal{M}' such that $\mathcal{M}' \succ \mathcal{M}$ and $\mathcal{M}' \models \Sigma$.

The set of preferred G -interpretations of Σ is denoted $\text{Pref}(\text{Mod}(\Sigma), \succeq)$.

Thus \mathcal{M} is preferred to \mathcal{M}' iff the change set of \mathcal{M} is included in the change set of \mathcal{M}' .

Definition 3.4 Let Σ be a G -scenario.

- A completion $\Theta(\Sigma)$ of Σ is a G -scenario $\Sigma' = \langle \phi'_1, \dots, \phi'_n \rangle$ such that for every $i \in \{1, \dots, n\}$,
 $\text{Mod}(\phi'_i) = \{M_i \mid \langle M_1, \dots, M_n \rangle \in \text{Pref}(\text{Mod}(\Sigma), \succeq)\}$.
- The consensus of Σ is the invariant of its completion, that is, $\text{Consensus}(\Sigma) = \bigvee_{i=1}^n \phi'_i$

We note that a completion is unique up to logical equivalence. Consequently we henceforth refer to *the* completion, understanding some canonical representative for a set of formulas; see also Proposition 3.3.

Example 1 Let $G = \langle V, E \rangle$ be the graph with $V = \{1, 2, 3, 4\}$ and $E = \{(1, 2), (2, 3), (2, 4), (3, 4)\}$. Let $\Sigma = \langle 1 : \neg p, 2 : (p \vee q) \wedge r, 3 : \neg q, 4 : \neg r \rangle$. The preferred models of Σ are $\langle \bar{p}qr, \bar{p}qr, \bar{p}\bar{q}r, \bar{p}\bar{q}\bar{r} \rangle$, $\langle \bar{p}qr, \bar{p}qr, \bar{p}\bar{q}r, \bar{p}\bar{q}\bar{r} \rangle$, $\langle \bar{p}qr, \bar{p}qr, \bar{p}\bar{q}\bar{r}, \bar{p}\bar{q}\bar{r} \rangle$ and $\langle \bar{p}\bar{q}r, \bar{p}\bar{q}r, \bar{p}\bar{q}\bar{r}, \bar{p}\bar{q}\bar{r} \rangle$.

Therefore, the completion of Σ is $\Theta(\Sigma) = \langle 1 : \neg p \wedge r, 2 : (p \leftrightarrow \neg q) \wedge r, 3 : \neg q, 4 : \neg r \wedge (\neg p \vee \neg q) \rangle$.

Although Vertices 2 and 4 disagree about r , the completion of Σ contains r at Vertex 1 (because 1 is connected to 2 only, and so r being true at 2 “blocks” the potential propagation of $\neg r$ at 3 to Vertex 1). In contrast, the completion contains neither r nor $\neg r$ at 3. The consensus of Σ is $\text{Consensus}(\Sigma) \equiv \neg p \vee \neg q$. As can be verified, this is equivalent to the disjunction of the constituent formulas, viz $(\neg p \wedge r) \vee ((p \leftrightarrow \neg q) \wedge r) \vee \neg q \vee (\neg r \wedge (\neg p \vee \neg q))$.

3.2 Syntactic characterisation

Definition 3.5 (Equivalence sets, Fits, Maximal fits)

- An equivalence set is a subset of
 $\{p^i \leftrightarrow p^j \mid \langle (i, j), p \rangle \in E \times \mathcal{P}\}$.
- The equivalence set induced by a change set Δ is
 $EQ(\Delta) = \{p^i \leftrightarrow p^j \mid \langle (i, j), p \rangle \in (E \times \mathcal{P}) \setminus \Delta\}$.
- An equivalence set EQ is a fit for a G -scenario $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ iff $EQ \cup \bigcup_{1 \leq i \leq n} \text{ren}(\phi_i, i)$ is consistent. A fit EQ for Σ is maximal iff for every $EQ' \supset EQ$ we have $EQ' \cup \bigcup_{1 \leq i \leq n} \text{ren}(\phi_i, i)$ is inconsistent.

Proposition 3.1 Δ is a minimal change set for Σ iff $EQ(\Delta)$ is a maximal fit for Σ .

Hence, maximising fit between propositional symbols or minimising change between models is equivalent. This in turn leads to a syntactic characterisation of completion:

Proposition 3.2 Let $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ be a G -scenario, and let \mathbf{F} be the set of maximal fits of Σ .

Then, $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$ is the completion of Σ iff for every $i \in \{1, \dots, n\}$,

$$Cn(\phi'_i) = \bigcap_{EQ \in \mathbf{F}} \{ \alpha \mid \alpha^i \in Cn(EQ \cup \bigcup_{j=1}^n \text{ren}(\phi_j, j)) \}$$

Here is an equivalent formulation of Proposition 3.2:

Corollary 3.1 Let $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ be a G -Scenario and let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$. Then, for any $\beta \in \mathcal{L}_{\mathcal{P}}$, $\phi'_i \models \beta$ iff for every \subseteq -maximal set $EQ \subseteq \{p^i \leftrightarrow p^j \mid \langle (i, j), p \rangle \in E \times \mathcal{P}\}$ consistent with $\bigcup_{j=1}^n \text{ren}(\phi_j, j)$ we have

$$EQ \cup \bigcup_{j=1}^n \text{ren}(\phi_j, j) \models \text{ren}(\beta, i).$$

This result shows that determining whether some formula holds at some vertex of the completion of Σ can be expressed as skeptical inference in normal default theories without prerequisites [Reiter, 1980]. An immediate consequence of this is that inference from the completion of a scenario is in Π_2^P . (We will see shortly that it is also a lower bound.)

Next we describe a syntactic characterisation that determines a specific formula for the completion at each vertex. We begin by adapting the standard notion of a *substitution* to our requirements: Given alphabets $\mathcal{P}^i, \mathcal{P}^j$ a *substitution* is a set $\{p^i/l(p^j) \mid p^i \in \mathcal{P}^i, p^j \in \mathcal{P}^j\}$ where $l(p^j)$ is p^j or $\neg p^j$. That is, a substitution is simply a set of pairs, where each pair is made up of an atom in one language and that atom or its negation in the other language. For binary relation R , let R^* be its transitive closure.

Definition 3.6 Let EQ be an equivalence set. For $i, j \in V$ define $\sigma_{i,j}^{EQ} = (\sigma_{i,j}^{EQ})^+ \cup (\sigma_{i,j}^{EQ})^-$ where

$$\begin{aligned} (\sigma_{i,j}^{EQ})^+ &= \{p^i/p^j \mid EQ \models p^i \leftrightarrow p^j\} \\ (\sigma_{i,j}^{EQ})^- &= \{p^i/\neg p^j \mid (i, j) \in E^*, EQ \not\models p^i \leftrightarrow p^j\} \end{aligned}$$

A substitution instance of $\alpha^i \in \mathcal{P}^i$ with respect to $\sigma_{i,j}^{EQ}$, written $\alpha^i \sigma_{i,j}^{EQ}$, is as expected: every atom in α^i is replaced by its unique counterpart according to $\sigma_{i,j}^{EQ}$.

We obtain the following finite characterisation:

Proposition 3.3 Let $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ a scenario of G . Let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$ be the completion of Σ , and let \mathbf{F} be the set of maximal fits of Σ . Then for $1 \leq j \leq n$ we have that

$$\phi'_j \leftrightarrow \bigvee_{EQ \in \mathbf{F}} \left(\bigwedge_{(j,i) \in E^*} (\phi_i^i \sigma_{i,j}^{EQ})^0 \right)$$

Thus the above proposition gives a recipe whereby one can compute a formula for each vertex that gives the completion.

Example 2 Consider Example 1 again, and Vertex $j = 3$. For preferred model \mathcal{M}_1 and Vertex $i = 2$ we have that ϕ_2 is $(p \vee q) \wedge r$; thus ϕ_2^2 is $(p^2 \vee q^2) \wedge r^2$. The relevant part of EQ is $p^2 \equiv p^3$ and $r^2 \equiv r^3$. Consequently $\phi_2^2 \sigma_{2,3}^{EQ}$ is $(p^3 \vee \neg q^3) \wedge r^3$. Continuing in this fashion, for the maximal fit based on \mathcal{M}_1 we obtain:

$$\begin{aligned} \phi_1^1 \sigma_{1,3}^{EQ} &= \neg p^3, & \phi_2^2 \sigma_{2,3}^{EQ} &= (p^3 \vee \neg q^3) \wedge r^3, \\ \phi_3^3 \sigma_{3,3}^{EQ} &= \neg q^3, & \phi_4^4 \sigma_{4,3}^{EQ} &= r^3. \end{aligned}$$

The conjunction of these terms is equivalent to $\neg p^3 \wedge \neg q^3 \wedge r^3$; re-expressed in our original language this is $\neg p \wedge \neg q \wedge r$. The other five preferred models similarly yield formulas for Vertex $j = 3$, in this example, each being a conjunction of literals. It can be confirmed that their disjunction is equivalent to $\neg q$.

3.3 Properties of completion and consensus

Here are some properties of completion and consensus.

Proposition 3.4 Let $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ be a G -scenario, and let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$.

1. $\Theta(\Sigma) = \Theta(\Theta(\Sigma))$.
2. If $\exists \phi'_i \in \Theta(\Sigma)$, $\phi'_i \models \perp$ then $\forall \phi'_j \in \Theta(\Sigma)$, $\phi'_j \models \perp$.
3. If G is connected and $\bigwedge_{1 \leq i \leq n} \phi_i$ is consistent then $\phi'_j \leftrightarrow \bigwedge_{1 \leq i \leq n} \phi_i$ for every $j \in \{1, \dots, n\}$.
4. Σ is consistent iff $\Theta(\Sigma)$ is consistent.
5. $\phi'_i \models \phi_i$ for every $\phi_i \in \Sigma$.
6. $\text{Consensus}(\Sigma) \models SK(\Sigma)$.

Note that Property 3 above refers to the structure of the graph. Clearly, the information at a given vertex i can propagate to vertex j only if there is a path from i to j . This brings us to the following (easy but useful) decomposability result:

Proposition 3.5 Let G_1, \dots, G_q be the connected components of G and let $\{V_1, \dots, V_q\}$ be the corresponding partition of V . For each $i \in \{1, \dots, n\}$, let $k(i)$ be the unique integer such that $i \in V_{k(i)}$. Given a G -scenario Σ , $\Theta(\Sigma) = \langle \phi_1, \dots, \phi_n \rangle$, and $k \in \{1, \dots, q\}$, let $\Sigma_k = \langle \phi_i \mid i \in V_k \rangle$, and let $\Theta(\Sigma_k) = \langle \phi'_i \mid i \in V_k \rangle$. Then:

- $\phi'_i = \phi''_i$ for each $i \in \{1, \dots, n\}$;
- $\text{Consensus}(\Sigma) \equiv \bigvee_{1 \leq k \leq q} \text{Consensus}(\Sigma_k)$.

This result is a bit tedious to state but is very intuitive: if the graph G is not connected, then it is sufficient to consider its connected components separately, and to compute the completion of Σ for each one of these.

Example 3 Let $G = \langle V, E \rangle$ with $V = \{1, 2, 3, 4\}$ and $E = \{(1, 2), (3, 4)\}$. Let $\Sigma = \langle 1 : a \vee b, 2 : \neg a \vee \neg b, 3 : a, 4 : \neg a \wedge \neg b \rangle$. The connected components of G are $G_1 = \langle \{1, 2\}, \{(1, 2)\} \rangle$ and $G_2 = \langle \{3, 4\}, \{(3, 4)\} \rangle$, hence,

$\Sigma_1 = \langle 1 : a \vee b, 2 : \neg a \vee \neg b \rangle$ and $\Sigma_2 = \langle 3 : a, 4 : \neg a \wedge \neg b \rangle$.

We obtain $\Theta(\Sigma_1) = \langle 1 : a \leftrightarrow b, 2 : a \leftrightarrow b \rangle$ and $\Theta(\Sigma_2) = \langle 3 : a \wedge \neg b, 4 : \neg a \wedge \neg b \rangle$. Therefore,

$\Theta(\Sigma) = \langle 1 : a \leftrightarrow b, 2 : a \leftrightarrow b, 3 : a \wedge \neg b, 4 : \neg a \wedge \neg b \rangle$.

Thus, without loss of generality, we may now assume that G is connected.

Another decomposition result concerns the propositional symbols the logical theories at different vertices are about. We say that ϕ_i is *irrelevant* to $\mathcal{Q} \subseteq \mathcal{P}$ iff ϕ_i is logically equivalent to some formula ψ_i in which no symbol in \mathcal{Q} appears [Lang *et al.*, 2003].

Proposition 3.6 Let Σ be a G -scenario. Suppose there exist a partition $\{V_1, V_2\}$ of V and a partition $\{\mathcal{Q}_1, \mathcal{Q}_2\}$ of \mathcal{P} such that for each $i \in V_1$ (resp. $i \in V_2$), ϕ_i is irrelevant to \mathcal{Q}_2 (resp. to \mathcal{Q}_1). Let Σ_1 and Σ_2 be the restrictions of Σ to V_1 and V_2 , defined by $\Sigma_1[i] = \Sigma[i]$ if $i \in V_1$ and $\Sigma_1[i] = \top$ otherwise (and similarly for Σ_2). Let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$. Then, for every $i \in V$,

- $\phi'_i = \Theta(\Sigma_1)[i] \wedge \Theta(\Sigma_2)[i]$
- $\text{Consensus}(\Sigma) = \text{Consensus}(\Sigma_1) \wedge \text{Consensus}(\Sigma_2)$.

Example 4 Let $V = \{1, 2, 3, 4, 5\}$, $E = \{(1, 2), (2, 3), (3, 4), (4, 5), (5, 1)\}$ and $\Sigma = \langle 1 : p \vee q, 2 : \neg p, 3 : a \wedge b, 4 : \neg p \wedge q, 5 : a \wedge \neg b \rangle$. Taking $\mathcal{Q}_1 = \{a, b\}$, $\mathcal{Q}_2 = \{p, q\}$, $V_1 = \{3, 5\}$ and $V_2 = \{1, 2, 4\}$, we get $\Theta(\Sigma_1) = \langle 1 : a, 2 : a, 3 : a \wedge b, 4 : a, 5 : a \wedge \neg b \rangle$ and $\Theta(\Sigma_2) = \langle 1 : \neg p \wedge q, 2 : \neg p \wedge q, 3 : \neg p \wedge q, 4 : \neg p \wedge q, 5 : \neg p \wedge q \rangle$. Hence, $\Theta(\Sigma) = \langle 1 : a \wedge \neg p \wedge q, 2 : a \wedge \neg p \wedge q, 3 : a \wedge b \wedge \neg p \wedge q, 4 : a \wedge \neg p \wedge q, 5 : a \wedge \neg b \wedge \neg p \wedge q \rangle$.

Let us now consider *conjunctive* formulas: a formula ϕ_i is conjunctive iff it is a consistent conjunction of literals.

Proposition 3.7 Let $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ be a G -scenario such that each ϕ_i is a conjunctive formula. Let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$. Then, $\phi'_i \models p$ iff the two following holds:

1. $\phi_j \models p$ for some $j \in \{1, \dots, n\}$;
2. If $\phi_l \models \neg p$, then for every path π from l to i , there is a m in π such that $\phi_m \models p$.

Proposition 3.7 gives a simple method for computing $\Theta(\Sigma)$ when Σ consists of conjunctive belief bases. For any $p \in \mathcal{P}$:

1. let G_p (resp. $G_{\neg p}$) be the graph obtained from G by deleting all vertices i such that $\phi_i \models \neg p$ (resp. $\phi_i \models p$) and all edges to or from i ;
2. let X_p (resp. $X_{\neg p}$) be the set of all i such that there is a path in G_p (resp. in $G_{\neg p}$) from a vertex j such that $\phi_j \models p$ (resp. $\phi_j \models \neg p$).

Proposition 3.8 Let $\Sigma = \langle \phi_1, \dots, \phi_n \rangle$ be a G -scenario such that each ϕ_i is a conjunctive formula and let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$. Then for every $i \in \{1, \dots, n\}$,

$$\phi'_i \leftrightarrow \bigwedge_{\{p \mid i \in X_p \cap \overline{X_{\neg p}}\}} p \wedge \bigwedge_{\{p \mid i \in X_{\neg p} \wedge \overline{X_p}\}} \neg p$$

As a corollary, the completion and the consensus of a conjunctive scenario can be computed in polynomial time.

Example 5 Let $V = \{1, 2, 3, 4, 5, 6\}$, $E = \{(1, 2), (2, 3), (3, 4), (4, 5), (5, 6), (1, 6)\}$ and $\Sigma = \langle 1 : p \wedge q, 2 : \neg q, 3 : \neg p, 4 : p, 5 : \top, 6 : \top \rangle$. G_p is the graph whose set of edges is $E_p = \{(1, 2), (4, 5), (5, 6), (1, 6)\}$, therefore $X_p = \{1, 2, 4, 5, 6\}$. $G_{\neg p}$ is the graph whose set of edges is $E_{\neg p} = \{(2, 3), (5, 6)\}$, therefore $X_{\neg p} = \{2, 3\}$. Similarly, $E_{\neg q} = \{(2, 3), (3, 4), (4, 5), (5, 6)\}$, $E_q = \{(3, 4), (4, 5), (5, 6), (1, 6)\}$, $X_q = \{1, 3, 4, 5, 6\}$ and $X_{\neg q} = \{2, 3, 4, 5, 6\}$. Therefore, $\Theta(\Sigma) = \langle 1 : p \wedge q, 2 : \neg q, 3 : \neg p, 4 : p, 5 : p, 6 : p \rangle$.

4 Particular Instances of the Approach

Our approach falls into the broad category of approaches to belief change based on minimisation. We consider initial formulas believed at vertices of G , and select those models that minimise differences between models of the individual formulas. So far, we haven't said much about specific interpretations of G . In this section, we give several examples of such interpretations, sometimes corresponding to particular classes of graphs; and we show that our setting generalises frameworks that have been explored previously.

We first consider three specific types of graphs: *chain graphs*, *star graphs* and *complete graphs*, and show that each of these three cases leads to a specific belief change operator studied previously.

Definition 4.1 (Extrapolation) Let $G = \langle V, E \rangle$ be a chain graph, where $E = \{(i, i+1) \mid i \in V, 1 \leq i < n\}$, and let Σ be a G -scenario. The extrapolation operator $Ext(\Sigma)$ induced by \succeq is unchanged from Definition 3.4: $Ext(\Sigma) = \Theta(\Sigma)$.

Ext is an extrapolation operator in the sense of [Dupin de Saint-Cyr and Lang, 2002], in which V is considered as a set of time points: if $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$, then ϕ'_t is the completed scenario at time t , obtained from the initial observations and a principle of minimal change. Other extrapolation operators can be obtained by varying the preference relation (see Section 5). As to $Consensus(\Sigma)$, it simply expresses the information that remains invariant over time in the scenario.

Definition 4.2 (Projection) Let $G = \langle V, E \rangle$ be a star graph, where $V = \{0, 1, \dots, n\}$ and $E = \{(0, i) \mid i \in V \setminus \{0\}\}$, and let Σ be a G -scenario. The projection operator $Proj$ induced by \preceq is the function $Proj(\Sigma) = \Theta(\Sigma)[0]$

$Consensus(\Sigma)$ here coincides with the projection operator given in [Delgrande and Schaub, 2006]: $\{1, \dots, n\}$ are the labels of the knowledge bases to be merged; ϕ_i represents the beliefs of source/agent i ; and 0 is the label for the result of the merging (with $\phi_0 = IS$ for set of integrity constraints IS).

Definition 4.3 (Consensus) Let G be a complete graph and let Σ be a G -scenario. The consensus operator of G is $Con(\Sigma) = Consensus(\Sigma)$.

That is, the consensus operator contains the common knowledge implicit in the minimal models of Σ . This operator corresponds to the consensus merging operator of [Delgrande and Schaub, 2006].

A crucial difference between the preceding two operators is illustrated in the following example. Let G be the complete graph on 3 vertices, and assume we have $\phi_1 = p \wedge q \wedge r$, $\phi_2 = \neg p \wedge \neg q \wedge r$, $\phi_3 = \top$. The projection on ϕ_3 is r while the consensus is $(p \leftrightarrow q) \wedge r$. See [Delgrande and Schaub, 2006] for a detailed examination of these operators.

Definition 4.4 (Revision) Let $G = \langle \{1, 2\}, \{(1, 2)\} \rangle$ and let $\langle \phi_1, \phi_2 \rangle$ be a G -scenario. The revision operator $\dot{+}$ is the function $\phi_1 \dot{+} \phi_2 = \Theta(\Sigma)[2]$

This notion of revision corresponds to the basic notion of revision given in [Sato, 1988; Delgrande and Schaub, 2003].

Deciding whether $\beta \in \phi_1 \dot{+} \phi_2$ is Π_2^p -complete [Delgrande et al., 2004]. Thus (given the remark in Section 3.2) the general problem of deciding, given a G -scenario Σ , a formula β , and $i \in V$, whether $\beta \in \Theta(\Sigma)[i]$, is Π_2^p -complete (and so no more complex than all the particular cases considered above). A similar result holds for consensus.

There are of course other graph structures that could be investigated, for example trees or cycles. Further, one can consider other interpretations. Thus a graph may be considered under a spatial interpretation, in which nodes correspond to regions of space and edges correspond to connections between adjacent regions. Completion is then based on the principle that, normally, what holds in a region also holds in the adjacent region. Another possible interpretation of the graph would be a mixture of merging and extrapolation, in the same vein as [Liberatore and Schaerf, 2000], where we have several time points, and possibly several independent information sources at each time point.

5 Alternative Definitions of Completion and Consensus

So far, we have focussed on a very simple criterion for selecting preferred interpretations, namely, \mathcal{M} being preferred to \mathcal{M}' iff $\Delta(\mathcal{M}) \subset \Delta(\mathcal{M}')$. Clearly, this notion of preference between change sets (and between interpretations) frequently gives many incomparable interpretations, hence a large number of preferred interpretations, and then completion and consensus may be rather weak. Here we give a more general definition without changing much of our notions and results.

A preference relation on change sets \succeq is a reflexive and transitive relation over $E \times \mathcal{P}$. Any such preference relation induces a preference relation \succeq on interpretations, by $\mathcal{M} \succeq \mathcal{M}'$ iff $\Delta(\mathcal{M}) \succeq \Delta(\mathcal{M}')$. We recover the preference relation considered previously if we define $\Delta \succeq_I \Delta'$ iff $\Delta \subseteq \Delta'$. Another intuitive preference relation is cardinality: $\Delta \succeq_C \Delta'$ iff $|\Delta| \leq |\Delta'|$. The definition of a preferred G -interpretation (w.r.t. \succeq) as well as those of completion and consensus are straightforward generalisations of Definitions 3.3 and 3.4. Note that for a star graph and \succeq_C one obtains majority merging as in [Konieczny and Pino Pérez, 2002].

A second generalisation of the approach begins from the observation that edges, representing connected points, may have differing strengths or reliability. Hence, more sophisticated preference relations may incorporate ranking functions, or numerical weights, on edges. For graph $G = \langle V, E \rangle$, a *priority assignment* P is a mapping from E to the natural numbers \mathbf{N} (any totally ordered set serves equally well).

Definition 5.1 (P-preferred G-interpretations) Let Σ be a scenario of graph $G = \langle V, E \rangle$, and let P be a priority assignment to E . Let \mathcal{M} and \mathcal{M}' be interpretations of Σ , and let $\Delta(\mathcal{M}, e)$ (resp. $\Delta(\mathcal{M}', e)$) be the change set associated with edge $e \in E$ in \mathcal{M} (resp. \mathcal{M}').

Define $\mathcal{M} \succ_P \mathcal{M}'$ iff

- $\mathcal{M} \models \Sigma$ and $\mathcal{M}' \models \Sigma$;
- there is $i \in \mathbf{N}$ such that:
 - for every $e \in E$ such that $P(e) = i$ we have $\Delta(\mathcal{M}, e) \subset \Delta(\mathcal{M}', e)$, and
 - for every $e' \in E$ such that $P(e') > i$ we have $\Delta(\mathcal{M}, e') = \Delta(\mathcal{M}', e')$.

Example 6 Consider Example 1 where $P((1, 2)) = 2$ and other edges are assigned value 1. The preferred models are now $\mathcal{M}_1, \mathcal{M}_2, \mathcal{M}_3$, and \mathcal{M}_4 . Consequently $\neg p$ is true at each point in $\Theta(\Sigma)$.

The completion of Σ under P , $\Theta(\Sigma, P)$, is defined in an obvious extension to Definition 3.4, over the P -preferred interpretations; we omit details given space constraints.

Proposition 5.1 Let Σ be a scenario of graph $G = \langle V, E \rangle$ and let P be a priority assignment to E . Let $\Theta(\Sigma) = \langle \phi'_1, \dots, \phi'_n \rangle$ be the completion of Σ and let $\Theta(\Sigma, P) = \langle \phi''_1, \dots, \phi''_n \rangle$ be the completion of Σ under P . Then $\phi''_i \models \phi'_i$ for every $1 \leq i \leq n$.

The notion of P -preferred G -interpretations captures the notion of *chronological minimisation*: let G be a chain graph $\langle V, E \rangle$ with $V = \{1, \dots, n\}$, $E = \{(t, t+1) \mid t =$

$0, \dots, n-1\}$, and priority assignment $P(t, t+1) = t$. Then the completion of Σ corresponds to chronological ignorance [Shoham, 1988] in which later changes are preferred to earlier changes. Further, when G is star graph (for merging), the priority of the edge connecting a source i to the centre point 0 can encode the reliability of i . Of course other alternatives are possible.

In a further direction, and generalising cardinality-based preference, one can assign weights to propositions, reflecting that some propositions are more likely to persist than others. For instance, when crossing the border between two countries in Europe, the official language is more likely to change than the main religion. These weights would then be aggregated to assign an overall weight to a change set.

6 Discussion

We have described a general approach to belief change based on minimisation of change. In the approach, an undirected graph $G = \langle V, E \rangle$ has associated with each vertex some information represented in a formula. The graph structure can model various phenomena including spatially adjacent points, temporally adjacent points in time, connected agents, etc. Consequently it is quite expressive, and generalises similar frameworks that have been explored previously. The intuition is that change between neighbours is exceptional; this assumption justifies *global* minimisation of change. Following the minimisation process, the *completion* of the graph provides new information at each vertex, representing that information that may be consistently gleaned from other connected vertices. The approach generalises previous frameworks involving change minimisation, including extrapolation and related belief change operators, and some merging and revision operators. However, the general approach is of independent interest. In particular, it proves to be suitable to *spatial minimisation of change*. Along this line, an interesting practical issue would be to apply the approach to a domain such as is addressed in [Würbel *et al.*, 2000]. Lastly, the general approach is computationally no more complex than each of these particular cases.

An interpretation of the graph that warrants more investigation is when it represents a network of agents. Global minimisation of change with respect to a network of agents can be understood as follows: agents have some initial beliefs that they don't want to give up, and after some communication process, the agents come to an equilibrium where changes between the belief bases of neighbouring agents are globally minimal. (Hence, our approach, based on global minimisation, can also be characterised as *consistent expansion* in a network of agents.) One interesting issue for future research is that of *local* minimisation of change, wherein minimisation at a vertex is with respect to that vertex's adjacent neighbours only. Clearly in this case the process admits nontrivial iteration. Our approach also departs from approaches such as [Gauwin *et al.*, 2005], where agents iteratively revise or merge their current beliefs by/with the beliefs of other agents. Another promising issue for further research is to allow information to be similarly revised (and not just strengthened) at a vertex (and so where Theorem 3.4, Part 5 does not necessarily

hold). In this case, an agent may in fact give up beliefs, perhaps as a result of contradicting beliefs at an adjacent vertex with higher reliability.

Acknowledgements

We wish to thank Florence Dupin de Saint-Cyr for many helpful discussions.

References

- [Booth, 2002] R. Booth. Social contraction and belief negotiation. *Proceedings of KR'02*, pages 375–384.
- [Delgrande and Schaub, 2003] J. Delgrande and T. Schaub. A consistency-based approach for belief change. *Art. Intelligence*, 151(1-2):1–41, 2003.
- [Delgrande and Schaub, 2006] J. Delgrande and T. Schaub. A consistency-based framework for merging knowledge bases. *J. Applied Logic*, 2006. To appear.
- [Delgrande *et al.*, 2004] J. Delgrande, T. Schaub, H. Tompits, and S. Woltran. On computing belief change operations using quantified boolean formulas. *J. Logic and Computation*, 14(6):801–826, 2004.
- [Dupin de Saint-Cyr and Lang, 2002] F. Dupin de Saint-Cyr and J. Lang. Belief extrapolation (or how to reason about observations and unpredicted change). *Proceedings of KR'02*, pages 497–508, 2002.
- [Gauwin *et al.*, 2005] O. Gauwin, S. Konieczny, and P. Marquis. Iterated belief merging as conciliation operators. *Proceedings of Commonsense'05*, pages 85–92.
- [Konieczny and Pino Pérez, 2002] S. Konieczny and R. Pino Pérez. Merging information under constraints: A logical framework. *J. Logic and Comput.*, 12(5):773–808, 2002.
- [Lang *et al.*, 2003] J. Lang, P. Liberatore, and P. Marquis. Propositional independence : Formula-variable independence and forgetting. *JAIR*, 18:391–443, 2003.
- [Liberatore and Schaerf, 1995] P. Liberatore and M. Schaerf. Arbitration: A commutative operator for belief revision. *Proceedings of WOCFAI'95*, 217–228.
- [Liberatore and Schaerf, 2000] P. Liberatore and M. Schaerf. Brels: A system for the integration of knowledge bases. *Proceedings of KR'00*, pages 145–152.
- [Reiter, 1980] R. Reiter. A logic for default reasoning. *Art. Intelligence*, 13(1-2):81–132, 1980.
- [Satoh, 1988] K. Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of 5th Generation Computer Systems*, pages 455–462, 1988.
- [Shoham, 1988] Y. Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence*. The MIT Press, 1988.
- [Würbel *et al.*, 2000] E. Würbel, R. Jeansoulin, and O. Papini. Revision: An application in the framework of GIS. *Proceedings of KR'00*, pages 505–515.
- [Zhang *et al.*, 2004] D. Zhang, N. Foo, T. Meyer, and R. Kwok. Negotiation as mutual belief revision. In *Proceedings of AAI'04*, pages 317–322.