A Minimization-Based Approach to Iterated Multi-Agent Belief Change

Paul Vicol[1] and **James Delgrande**[1] and **Torsten Schaub**[2,3]

Abstract. We investigate minimization-based approaches to iterated belief change in multi-agent systems. A network of agents is represented by an undirected graph, where propositional formulas are associated with vertices. Information is shared between vertices via a procedure where each vertex minimizes disagreement with other vertices in the graph. Each iterative approach takes into account the proximity between vertices, with the underlying assumption that information from nearby sources is given higher priority than information from more distant sources. We have identified two main approaches to iteration: in the first approach, a vertex takes into account the information at its immediate neighbours only, and information from more distant vertices is propagated via iteration; in the second approach, a vertex first takes into account information from distance-1 neighbours, then from distance-2 neighbours, and so on, in a prioritized fashion. There prove to be three distinct ways to define the second approach, so in total we have four types of iteration. We define these types formally, find relationships between them, and investigate their basic logical properties. We also implemented the approaches in a software system called *Equibel*.

1 INTRODUCTION

We investigate several approaches for iterated belief change in multi-agent systems, each based on minimizing disagreements between agents in a prioritized manner. A problem instance is an undirected graph with formulas attached to vertices. Information is shared between vertices via a process of minimization over the graph. Previous work [2] dealt with one-shot belief change, where every vertex updates its beliefs through a *global* minimization process, with a weak notion of distance between vertices in a graph. We generalize this work to model iterated approaches, where distance between vertices is explicitly taken into account. One approach is for each vertex to repeatedly update its beliefs by taking into account only the beliefs of its immediate neighbours; another is for a vertex to take into account the beliefs of its distance-1 neighbours, and then take into account the beliefs of its distance-2 neighbours, and so on. We show that this second approach can be defined in three different ways, which lead to different behaviours. We implemented the approaches described in this paper in a software system called *Equibel*, available at `https://github.com/asteroidhouse/equibel`.

To motivate this work, we consider two interpretations of a graph: the first is where the graph represents a network of communicating agents, and the second is where the graph represents some general domain (such as a spatial domain), with local information contained at each vertex. Iterated belief change in these settings can be understood as follows. In the multi-agent setting, each agent *consistently incorporates* information from other agents in a stepwise fashion, where the agent is more inclined to trust close acquaintances compared to more distant ones. In the general setting, the goal is to get an overall picture of the state of the world by combining information from multiple sources.

To illustrate one of the approaches to iterated belief change, consider a specific example in which a graph models a weather-sensing system, where vertices represent weather stations from which observations are made, and edges encode the adjacency of the spatial regions where the stations are located. The goal is to determine what information holds at a region by combining the information that is known to hold at that region (through local observations) with information coming from neighbouring regions. By the assumption of *spatial persistence* between adjacent regions, information from nearby regions should be prioritized over information from more distant regions. Informally, an approach to do this is as follows. We first consider the observations from directly adjacent regions (i.e., vertices at distance 1), and determine which observations minimize disagreement with the observations of the first region; then, we consider the observations from the next-nearest regions (i.e., vertices at distance 2), and we find observations which further minimize disagreement with the observations from distance 1, and so on. This can be seen as prioritized minimization of disagreements with respect to increasingly large *neighbourhoods* around a vertex of interest.

Multi-agent approaches in which an agent considers the beliefs of other agents in increasingly large neighbourhoods can be thought of as modeling a *group conversation* — an agent takes into account the beliefs of other agents, as well as the connections between those agents, in order to decide how to update its beliefs so as to minimize disagreement with the group. An agent tries to satisfy all protagonists to the greatest possible degree, in a prioritized manner (where nearby agents are given higher priority than distant agents).

The next section discusses related work in multi-source merging. Section 3 defines the iterative approaches we examine. Section 4 presents results concerning the relationships between the approaches. It proves to be the case that each of the approaches we propose yields non-comparable results when applied iteratively. Section 5 shows basic logical properties of the approaches. Sections 6 and 7 discuss our work and present the conclusion, respectively.

2 RELATED WORK

Most work on updating knowledge bases given new information stems from the AGM approach to belief revision [1, 7]. Belief merging [8, 5, 10] can be seen as an extension of belief revision to situ-

[1] Simon Fraser University, Burnaby, BC, {pvicol@sfu.ca, jim@cs.sfu.ca}
[2] University of Potsdam, Potsdam, Germany, {torsten@cs.uni-potsdam.de}
[3] INRIA, Rennes, France

ations involving multiple belief bases, where the goal is to combine several, possibly conflicting, bases into a coherent whole.

Here we provide some background behind standard approaches to merging. Given a language $\mathcal{L}_\mathcal{P}$, a belief base K is a finite set of propositional formulas, and a belief profile $\mathcal{K} = \langle K_1, \ldots, K_n \rangle$ is a finite vector consisting of n belief bases which are not necessarily pairwise different. A merging operator Δ is a function $\mathcal{L}_\mathcal{P} \times \mathcal{L}_\mathcal{P}^n \rightarrow \mathcal{L}_\mathcal{P}$ that associates a formula μ and a belief profile $\mathcal{K}$ with a new formula $\Delta_\mu(\mathcal{K})$, which is called the merged belief state. The operator Δ aims at consistently merging the beliefs in $\mathcal{K}$ under the integrity constraint given by μ. A set of nine postulates denoted **(IC0)-(IC8)** have been proposed to capture the notion of rational belief merging. These are called the Integrity Constraint (IC) merging postulates [9], listed below:

(IC0) $\Delta_\mu(\mathcal{K}) \vdash \mu$

(IC1) If $\mu \nvdash \bot$, then $\Delta_\mu(\mathcal{K}) \nvdash \bot$

(IC2) If $\bigwedge_{K \in \mathcal{K}} K \wedge \mu \nvdash \bot$, then $\Delta_\mu(\mathcal{K}) \equiv \bigwedge_{K \in \mathcal{K}} K \wedge \mu$

(IC3) If $\mathcal{K}_1 \equiv \mathcal{K}_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(\mathcal{K}_1) \equiv \Delta_{\mu_2}(\mathcal{K}_2)$

(IC4) If $K_1 \vdash \mu$, $K_2 \vdash \mu$ and $\Delta_\mu(\langle K_1, K_2 \rangle) \wedge K_1 \nvdash \bot$, then $\Delta_\mu(\langle K_1, K_2 \rangle) \wedge K_2 \nvdash \bot$

(IC5) $\Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2) \vdash \Delta_\mu(\mathcal{K}_1 \sqcup \mathcal{K}_2)$

(IC6) If $\Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2)$ is consistent, then $\Delta_\mu(\mathcal{K}_1 \sqcup \mathcal{K}_2) \vdash \Delta_\mu(\mathcal{K}_1) \wedge \Delta_\mu(\mathcal{K}_2)$

(IC7) $\Delta_{\mu_1}(\mathcal{K}) \wedge \mu_2 \vdash \Delta_{\mu_1 \wedge \mu_2}(\mathcal{K})$

(IC8) If $\Delta_{\mu_1}(\mathcal{K}) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(\mathcal{K}) \vdash \Delta_{\mu_1}(\mathcal{K}) \wedge \mu_2$

Any operator Δ that satisfies these postulates is called an *IC merging operator*. Classical approaches to belief merging, such as [9] and [11], begin with a *set* of belief bases and produce a single, merged base. Our approach differs from these in that we deal with updating *multiple belief bases simultaneously*.

Distance-based merging operators $\Delta^{d,f}$ are characterized by a pseudo-distance d (that is, d does not have to satisfy the triangle inequality) between models and an aggregation function $f : \mathbb{R}^+ \times \cdots \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ [8]. Commonly used distances include the *drastic distance* (that is 0 if two models are equal and 1 otherwise), and the Hamming distance (the number of atoms on which two models differ). Our approach to minimization uses a set-theoretic distance between models, which is distinct from any of the standard distances used by IC merging operators.

[6] presents a framework for updating the beliefs of a group of agents via an iterated merge-and-revise procedure. The paper introduces *conciliation operators* which map a belief profile to a new belief profile in each step of the process. These operators are defined in terms of IC merging operators and the revision operators they induce. The authors describe two approaches: in the skeptical one, each agent gives priority to its previous beliefs over the merged beliefs of the group, while in the credulous one, each agent views the merged beliefs of the group as more important than its previous beliefs. The work in [6] focuses on the issue of updating a belief profile; it does not consider graphs. Since each agent considers the beliefs of all other agents simultaneously, the conciliation approach would correspond in our approach to connecting all agents in a complete graph.

More closely related to our approach is [13]. Belief Revision Games (BRGs) are games that model the dynamics of the beliefs of a group of communicating agents. Beliefs are associated with nodes in a directed graph. In each iterative step, each agent updates its beliefs by considering the beliefs of its neighbours. The authors define 18 different *revision policies* based on various IC merging operators that an agent can use to combine its beliefs with those of its neighbours. Each revision policy ascribes a different level of importance to the beliefs of an agent's neighbours compared to the current belief of the agent itself. The revision policies range from one in which an agent completely relinquishes its prior belief and replaces it with the merged beliefs of the group, to one where an agent does not give up its initial beliefs, but strengthens its opinion by incorporating consistent information from its neighbours.

The REV!GIS system [14] deals with belief revision in geographic information systems, using information at one location to revise adjacent locations. This conforms closely to the interpretation of a graph as representing a spatial-domain. BReLS [12] is a framework for integrating information from multiple sources, using an approach that combines merging, revision, and update.

In this paper, we build on the general framework for belief change introduced in [2] and [3]. Our approach to minimizing change between vertices in a graph is similar to that used in *belief extrapolation* [4], which can be seen as minimization of change in a chain graph where vertices represent successive points in time.

3 ITERATIVE APPROACHES

3.1 Preliminaries

We work with a propositional language $\mathcal{L}$ defined over a finite alphabet $\mathcal{P} = \{p, q, r, \ldots\}$ of atoms. We use the constants $\top$ (resp. $\bot$) to represent formulas that are always true (resp. false), and the connectives $\neg, \wedge, \vee, \rightarrow$, and $\leftrightarrow$ to construct formulas in the standard way. An interpretation of $\mathcal{L}$ is an assignment of truth values to the atoms in $\mathcal{P}$. We represent an interpretation by the set of atoms that are *true* in the interpretation. For example, given $\mathcal{P} = \{p, q, r\}$, the interpretation where p is false but q and r are true is expressed by the set $\{q, r\}$. The set of all interpretations of $\mathcal{L}$ is denoted $\mathcal{W}$. Given a formula $\alpha \in \mathcal{L}$ and an interpretation $w \in \mathcal{W}$, we write $w \models \alpha$ iff w makes α true in the usual truth-functional way; then we say that w is a *model* of α. We denote the set of models of α by $Mod(\alpha)$. If ω is a model over the finite alphabet $\mathcal{P}$, let $form(\omega) = \bigwedge_{p \in \omega} p \wedge \bigwedge_{p \in (\mathcal{P} \setminus \omega)} \neg p$. For a *set* of models $\Gamma \subseteq \mathcal{W}$, let $form(\Gamma) = \bigvee_{\omega \in \Gamma} form(\omega)$. For a set A, let $P(A)$ denote its *power set*, i.e., the set of all subsets of A.

3.2 Model Graphs

In this paper, we consider only connected, undirected graphs $G = \langle V, E \rangle$, where the vertices are identified by an initial sequence of natural numbers, e.g. for $|V| = n$, we have $V = \{1, \ldots, n\}$.

Definition 3.1 (G-Scenario)**.** Let $G = \langle V, E \rangle$ be a graph. A *G-scenario* is a function $\sigma : V \rightarrow \mathcal{L}$ that associates a propositional formula with each vertex in the graph. σ is *consistent* iff $\sigma(v)$ is consistent for all $v \in V$.

Next, we define a graph-theoretic representation for a graph G and an associated G-scenario σ, to make explicit the process by which information is shared between vertices. The idea is that a vertex $v \in V$ with formula $\sigma(v)$ is replaced by a *set* of vertices representing the models of $\sigma(v)$. For each edge $(v, w) \in E$, the vertices representing the models of $\sigma(v)$ and those representing the models of $\sigma(w)$ are connected in a complete bipartite graph. Each edge of the complete bipartite graph is given a label representing the level of disagreement between the models it connects. Various approaches for updating the

information at a vertex can be defined in terms of selecting one or more models corresponding to each original vertex, such that the labels of the edges involved are collectively minimal in some way. Next we formally define the *model graph* corresponding to a *base graph* G and an associated G-scenario.

Definition 3.2 (Model Graph)**.** Let $G = \langle V, E\rangle$ be a graph, and let σ be a G-scenario. For a vertex $v \in V$, $\lambda(v) = \{(v, m) \mid m \in Mod(\sigma(v))\}$ is the set of *model vertices* corresponding to v. For an edge $(v, w) \in E$, $\delta(v, w) = \lambda(v) \times \lambda(w)$ is the set of *model edges* corresponding to (v, w). The *model graph of G under σ*, denoted $\mathfrak{J}(G, \sigma)$, is the graph $\mathfrak{J}(G, \sigma) = \langle \bigcup_{v\in V} \lambda(v), \bigcup_{(v,w)\in E} \delta(v, w)\rangle$.

For a model vertex (v, m), let $M((v, m)) = m$, so $M : V \times \mathcal{W} \to \mathcal{W}$ is a function that extracts the *model* from a model vertex. For a set Γ of model vertices, let $M(\Gamma) = \{M((v, m)) \mid (v, m) \in \Gamma\}$.

As an example, given the base graph shown in Figure 1, the corresponding model graph is shown in Figure 2. Here we have $\mathcal{P} = \{p, q, r\}$ and $\sigma(1) = p \wedge \neg q$, so $Mod(\sigma(1)) = \{\{p\}, \{p, r\}\}$. Thus, $\lambda(1) = \{(1, \{p\}), (1, \{p, r\})\}$. This is represented in Figure 2 by the rectangle labelled **1** that contains nodes labelled $\{\mathbf{p}\}$ and $\{\mathbf{p}, \mathbf{r}\}$. Each edge in Figure 2 is labelled by the symmetric difference of the models at its endpoints.

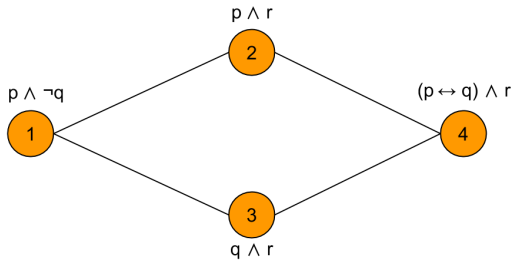

Figure 1. Base graph

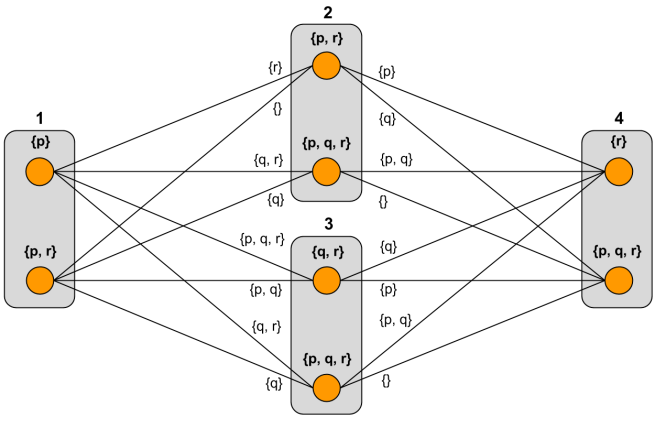

Figure 2. Model graph corresponding to the base graph

3.3 Types of Neighbourhoods

The iterative approaches we define in the next section are based on the notion of *distance prioritization*, which informally states that a vertex gives higher priority to information coming from nearby sources compared to more distant sources. In order to deal with sources of information *within a certain distance* from a vertex of interest, we define *neighbourhoods* about vertices. In this section, we give formal definitions of the neighbourhoods we consider. In the following, let $dist(v, w)$ denote the length of the shortest path between nodes v and w, where we define $dist(v, v) = 0$.

Definition 3.3 (Neighbourhood, Pseudo-Neighbourhood)**.** Let $G = \langle V, E\rangle$ be a graph.

- A *neighbourhood of* $v \in V$ is a connected subgraph $G' = \langle V', E'\rangle$ of G such that $v \in V'$.
- A *pseudo-neighbourhood of* $v \in V$ is a connected graph $G' = \langle V', E'\rangle$ where $v \in V' \subseteq V$ and $E' \subseteq V' \times V'$.

Note that a pseudo-neighbourhood is not a subgraph of G, i.e., there is no relationship between the edges in E and those in E'. Shortest-path trees are a natural starting point to define neighbourhoods consisting of nodes within a certain distance of a given node.

Definition 3.4 (Shortest-Path Tree)**.** Given a graph $G = \langle V, E\rangle$, a *shortest-path tree* rooted at a vertex $v \in V$ is a spanning tree T of G, such that the path between the root v and any other vertex u in T is the shortest path between v and u in G. In graphs with un-weighted edges, shortest-path trees are equivalent to breadth-first search trees. Define a *shortest-path tree of radius r rooted at v* to be a breadth-first search tree of depth r from v.

Definition 3.5 (Cross-Edge)**.** Let $G = \langle V, E\rangle$ be a graph, and let $v \in V$. A *cross-edge for v* is an edge $(s, t) \in E$ such that $dist(v, s) = dist(v, t)$.

Cross-edges are not included in shortest-path trees (because they break the tree property). Using the above definitions, we propose four types of neighbourhoods to be considered for iterative approaches. The following neighbourhood types are defined in alternative ways with respect to a *root* or *central* node v and a *radius* $r \in \mathbb{N}$:

1. Use Shortest-Path Trees (SPTs) to define neighbourhoods. There may be multiple SPTs from a vertex to a specified radius, because there may be multiple shortest paths between two nodes. Let $\{\text{SPT}_j(v, r) \mid j \in \{1, \ldots, k\}\}$ be the set of all distinct SPTs of radius r rooted at v. One type of neighbourhood about v is an *arbitrary SPT of depth r*: $\text{SPT}_j(v, r)$.
2. Use the union of all SPTs to a specified radius from a certain node: $\text{USPT}(v, r) = \bigcup_j \text{SPT}_j(v, r)$. This removes the ambiguity involved in selecting an arbitrary SPT, and makes sense intuitively, since by considering all possible shortest paths from the root node to all other nodes in the neighbourhood, we look at all possible means of propagation of information to the root node. Note that a USPT neighbourhood does not contain cross-edges.
3. Use the union of SPTs together with all cross-edges; we call this a *complete neighbourhood*. Let $D(v, r) = \{w \mid dist(v, w) \leq r\}$. We define the *complete neighbourhood about v of radius r* to be the graph $W(v, r) = \langle D(v, r), \{(i, j) \mid (i, j) \in E, i \in D(v, r), j \in D(v, r)\}\rangle$. That is, the complete neighbourhood consists of all nodes within distance r of v, and all the edges that involve those nodes as endpoints.
4. Use the *ring* of nodes at distance r from the root node, i.e., find nodes at distance r and construct a star graph that connects those nodes directly to the root node. We define the *ring about v of radius r* to be the graph $R(v, r) = \langle T(v, r), S(v, r)\rangle$ where $T(v, r) = \{v\} \cup \{w \mid w \in V, dist(v, w) = r\}$ and $S(v, r) = \{(w, v) \mid w \in V, dist(v, w) = r\}$.

Note that alternatives 1-3 are *neighbourhoods* of v, while alternative 4 is a *pseudo-neighbourhood* of v. In the discussion of our approaches, we use the diameter of a graph and eccentricity of a node, defined as follows:

Definition 3.6 (Eccentricity, Diameter)**.** Let $G = \langle V, E\rangle$ be a graph. The *eccentricity* of a node $v \in V$, denoted $\epsilon(v)$, is the *greatest shortest distance* between v and any other vertex: $\epsilon(v) = \max_{w\in V} dist(v, w)$. The *diameter* of the graph G, denoted $diam(G)$, is the maximum eccentricity of any vertex: $diam(G) = \max_{v\in V} \epsilon(v)$.

For simplicity and uniformity in the definitions of different approaches, for any neighbourhood type $N \in \{\text{SPT}_j, \text{USPT}, W, R\}$ and any vertex v, we *restrict* the radius of the neighbourhood so that it cannot exceed the eccentricity of v:

$$N(v, r) = \begin{cases} N(v, r) & \text{if } r < \epsilon(v) \\ N(v, \epsilon(v)) & \text{if } r \geq \epsilon(v) \end{cases}$$

By definition $diam(G) = \max_{v\in V} \epsilon(v)$, so for every node $v \in V$, $N(v, diam(G)) = N(v, \epsilon(v))$. Note that the complete neighbourhood is the only one with the property that if we look at a neighbourhood with a very large radius about any node, then that is equivalent to considering the original graph G. So $N(v, diam(G)) = G$ when $N = W$ (complete neighbourhood), but in general $N(v, diam(G)) \neq G$ when $N \neq W$ (for example, if N is a union of shortest-path trees).

3.4 Defining the Approaches

In this section, we define four approaches to iterated belief change, which we call the *simple*, *augmenting*, *expanding*, and *ring* approaches. Each approach involves a different procedure to *select model vertices* corresponding to each node in a graph. The following definition formalizes the notion of a *model selection*.

Definition 3.7 (Model Selection)**.** Let $G = \langle V, E\rangle$ be a graph. A *model selection* is a function $s : P(V \times \mathcal{W}) \to V \times \mathcal{W}$ that selects exactly one model vertex from the set $\lambda(v)$ of model vertices for each $v \in V$. Let $S(G)$ be the *set of all possible model selections over* G.

- For a subgraph $G' = \langle V', E'\rangle$ of G and a model selection $s \in S(G)$, let $s(G')$ denote the selection s *restricted to* G'.
- For $N \in \{\text{SPT}_j, \text{USPT}, W, R\}$, let $S_N(v, r)$ be the *set of all model selections restricted to neighbourhood* $N(v, r)$. For a neighbourhood centered about a node v, and a model selection $s \in S_N(v, r)$, let $h(s) = s(\lambda(v))$, so that $h(s)$ is the model vertex selected at the central node v by the selection s. For a set $T \subseteq S_N(v, r)$ of model selections, let $h(T) = \{h(s) \mid s \in T\}$ be the *set* of model vertices selected at v by any of the model selections in T.
- For a model selection $s \in S_N(v, r)$, we define the *subselection of radius* r', denoted $sub(s, r')$ for $r' \leq r$, to be the selection s restricted to the subgraph $N(v, r')$.

Next we define change sets induced by model selections. These change sets measure the *overall disagreement* between the models selected by a model selection at adjacent vertices.

Definition 3.8 (Change Set Induced by a Model Selection)**.** The change set *induced* by a model selection $s \in S(G)$, denoted $\Delta(s)$, is defined as $\Delta(s) = \{\langle (v, w), p\rangle \mid (v, w) \in E, s(\lambda(v)) = (v, a), s(\lambda(w)) = (w, b), p \in a \oplus b\}$, where $a \oplus b$ denotes the symmetric difference between sets a and b.

The following definition uses change sets to define a *preference relation* over model selections.

Definition 3.9 (Preferred Model Selections)**.** Given two model selections $s, s' \in S$, we define $s \succeq s'$ iff $\Delta(s) \subseteq \Delta(s')$ and $s \succ s'$ iff $\Delta(s) \subset \Delta(s')$. We say that s is a *preferred or minimal model selection* iff $\nexists s' \in S : s' \succ s$. We denote the set of all preferred model selections by $Pref(S, \succeq)$.

Thus, s is preferred to s' iff the change set induced by s is included in the change set induced by s'. *Preferred* model selections are those that induce *inclusion-minimal* change sets, and therefore minimize disagreement between the models at adjacent vertices.

It is useful to distinguish between two classes of model selections: those that induce change sets *not induced by other model selections* (which we call *unique model selections*), and those that induce change sets that are *also induced by other model selections* (which we call *duplicated model selections*).

Definition 3.10 (Unique and Duplicated Model Selections)**.** Let Γ be a set of model selections. Then, $q(\Gamma) = \{s \in \Gamma \mid \nexists s' \in \Gamma : s' \neq s \text{ and } \Delta(s') = \Delta(s)\}$ and $d(\Gamma) = \{s \in \Gamma \mid \exists s' \in \Gamma : s' \neq s \text{ and } \Delta(s') = \Delta(s)\}$. The functions q and d define *unique* and *duplicated* model selections, respectively, referring to the existence of other model selections that induce identical change sets. These functions *partition* Γ, so $\Gamma = q(\Gamma) \cup d(\Gamma)$ and $q(\Gamma) \cap d(\Gamma) = \varnothing$.

We are particularly interested in the case where $\Gamma = Pref(S(G), \succeq)$. For a model selection s to be *strictly preferred* to another selection s', it must be the case that for *any* subgraph g of G, $s(g) \succeq s'(g)$ and for *some* subgraph g' of G, $s(g') \succ s'(g')$. A special case of this is when we consider subgraphs of G that are neighbourhoods of different radii about a node v. Then, in order for selection s to be strictly preferred to selection s', we must have: $\forall r \,.\, 1 \leq r \leq diam(G) : s(N(v, r)) \succeq s'(N(v, r))$ and $\exists r \,.\, 1 \leq r \leq diam(G) : s(N(v, r)) \succ s'(N(v, r))$.

The following proposition states that a unique minimal model selection for a graph H must be contained within some minimal model selection for any supergraph I of H.

Proposition 1. *Let $G = \langle V, E\rangle$ be a graph, and let $H = \langle V', E'\rangle$ and $I = \langle V'', E''\rangle$ be subgraphs of G such that H is a subgraph of I. If $s \in q(Pref(S(H), \succeq))$ then $\exists s' \in Pref(S(I), \succeq) : s'(H) = s$.*

Proof. Let $s \in q(Pref(S(H), \succeq))$. Consider the set of model selections $\mathcal{O} = \{s' \in S(I) \mid s'(H) = s\}$; that is, the set of all model selections over I that have subselections over H equal to s. $\mathcal{O}$ is a partially-ordered set with respect to $\succeq$, and thus has minimal elements. Denote the set of minimal elements of $\mathcal{O}$ by $Min(\mathcal{O}, \succeq)$. Take an arbitrary element $z \in Min(\mathcal{O}, \succeq)$. We want to show that $z \in Pref(S(I), \succeq)$. Suppose not. Then there must be a model selection $z' \in Pref(S(I), \succeq)$ such that $z' \succ z$. In order for z' to be strictly preferred to z we must have that for any subgraph g of G, $z'(g) \succeq z(g)$. In particular, we must have $z'(H) \succeq z(H) = s$. But since $s \in q(Pref(S(H), \succeq))$, if $z'(H) \succeq s$, then we must have $z'(H) = s$. Since $z' \in S(I)$ and $z'(H) = s$, we have $z' \in \mathcal{O}$. But $z \in Min(\mathcal{O}, \succeq)$, so $z' \not\succ z$, which is a contradiction. Thus, $z \in Pref(S(I), \succeq)$ and $z(H) = s$, so we have constructed a model selection with the desired properties. This proves that $\exists s' \in Pref(S(I), \succeq) : s'(H) = s$. □

Corollary 1.1. *If $x \in \{s(\lambda(v)) \mid s \in q(Pref(S(H), \succeq))\}$ then $x \in \{s(\lambda(v)) \mid s \in Pref(S(I), \succeq)\}$.*

A consequence of this is that $h(q(Pref(S_N(v,r),\succeq))) \subseteq h(Pref(S_N(v,r+1),\succeq))$. This states that if a model vertex is selected at a node v by a *unique minimal selection* s in a neighbourhood of radius r about v, then that model vertex must be selected at v in every neighbourhood of radius greater than r (i.e., that model vertex will not be eliminated when considering larger neighbourhoods). The next proposition gives an analogous result for *duplicated minimal selections*.

Proposition 2. *Let $G = \langle V, E\rangle$ be a graph, and let $H = \langle V', E'\rangle$ and $I = \langle V'', E''\rangle$ be subgraphs of G such that H is a subgraph of I. If $s \in d(Pref(S(H),\succeq))$ then $\exists s' \in Pref(S(I),\succeq) : s'(H) \in \{t \in Pref(S(H),\succeq) \mid \Delta(t) = \Delta(s)\}$.*

Proof. This proof is nearly identical to the proof of Proposition 1; we omit it due to space constraints. □

Corollary 2.1. *At least one of the model vertices selected at the central node v by a duplicated minimal model selection over H must be selected by a minimal model selection over I: $\{s(\lambda(v)) \mid s \in d(Pref(S(H),\succeq))\} \cap \{s(\lambda(v)) \mid s \in Pref(S(I),\succeq)\} \neq \varnothing$.*

A consequence of this for concentric neighbourhoods is:

$$h(d(Pref(S_N(v,r),\succeq))) \cap h(Pref(S_N(v,r+1),\succeq)) \neq \varnothing$$

Before defining the iterative approaches we investigate, we define the *global completion* operation studied previously [2], which performs *one-shot* (as opposed to iterative) belief change. This allows us to compare the global approach with our iterative approaches in Section 4.

Definition 3.11 (Global Completion)**.** Let $G = \langle V, E\rangle$ be a graph, and let σ be an associated G-scenario. We define $C(v) = h(Pref(S_W(v, diam(G)),\succeq)) = \{s(\lambda(v)) \mid s \in Pref(S(G),\succeq)\}$. The *global completion* of σ, denoted σ_C, is the G-scenario such that $\forall v \in V : \sigma_C(v) = form(M(C(v)))$.

Now we define the *simple*, *augmenting*, and *expanding* approaches to iteration, by defining three functions F, A, and E, respectively, that *select specific model vertices* for each node that minimize disagreement between the beliefs of that node and those of its neighbours.

For a model graph $\mathfrak{J}$ of $G = \langle V, E\rangle$ under σ, the set of model selections over $\mathfrak{J}$ is denoted by $S_{\mathfrak{J}}$. Given an approach $\Omega \in \{F, A, E\}$ and neighbourhood type $N \in \{SPT_j, USPT, W\}$, for any iteration $i > 0$, $\Omega^i_N(v)$ is the *set of model vertices* selected for v by Ω in iteration i using neighbourhood N, and we define

$$\mathfrak{J}^i_{\Omega,N} = \left\langle \bigcup_{v\in V} \Omega^i_N(v), \bigcup_{(v,w)\in E} \Omega^i_N(v) \times \Omega^i_N(w) \right\rangle$$

to be the *model graph* that results from the i^{th} iteration of approach Ω using neighbourhood N. The set of all model selections over $\mathfrak{J}^i_{\Omega,N}$ is $S_{\mathfrak{J}^i_{\Omega,N}}$, which we abbreviate to $S^i_{\Omega,N}$. In this context, we let $S^0_{\Omega,N}$ stand for $S_{\mathfrak{J}}$, so that $S^0_{A,N} = S^0_{E,N} = S^0_{F,N}$. Then, $S^i_{\Omega,N}(v,r)$ is the set of model selections over $\mathfrak{J}^i_{\Omega,N}$, restricted to neighbourhood $N(v,r)$. Next, we define the following iterative approaches:

Definition 3.12 (Simple/Fixed-Radius Iteration)**.** Let $G = \langle V, E\rangle$ be a graph, and let σ be an associated G-scenario. For $i \geq 1$, we define $F^i_N(v) = h(Pref(S^{i-1}_{F,N}(v,1),\succeq))$. The G-scenario that results from the i^{th} iteration of the simple approach is denoted $\sigma^i_{F,N}$, and is defined as the G-scenario such that $\forall v \in V : \sigma^i_{F,N}(v) = form(M(F^i_N(v)))$.

In each iteration of the simple approach, a vertex changes its beliefs by considering only the beliefs of its immediate neighbours.

Definition 3.13 (Augmenting Iteration)**.** Let $G = \langle V, E\rangle$ be a graph, and let σ be an associated G-scenario. Then for $i \geq 1$, we define $A^i_N(v) = h(T^i_N(v))$, where

$$T^i_N(v) = \{s \in Pref(S^{i-1}_{A,N}(v, diam(G)),\succeq) \mid \forall r\,.\,1 \leq r \leq diam(G) : sub(s,r) \in Pref(S^{i-1}_{A,N}(v,r),\succeq)\}$$

The G-scenario that results from the i^{th} iteration of the augmenting approach is denoted $\sigma^i_{A,N}$, and is defined as the G-scenario such that $\forall v \in V : \sigma^i_{A,N}(v) = form(M(A^i_N(v)))$.

A model selection $s \in T^i_N(v)$ can be constructed stepwise by: 1) starting from v and considering preferred model selections over a radius 1 neighbourhood about v; then 2) *augmenting* those preferred selections by considering *minimal model edges* that cross the boundary between the radius 1 and radius 2 neighbourhoods about v (as well as the model edges *between* distance-2 nodes, in the case of complete neighbourhoods), and so on. Once we find preferred model selections over a radius 1 neighbourhood, we take into account *only* the model vertices (and model edges) involved in those selections for future steps of the procedure; at each radius, we examine ways to *augment* preferred model selections from the preceding radius.

Definition 3.14 (Expanding Iteration)**.** Let $G = \langle V, E\rangle$ be a graph, and let σ be an associated G-scenario. For $i \geq 1$, we define $E^i_N(v) = I^{\epsilon(v)}_{i,N}(v)$, where

$$I^1_{i,N}(v) = h(Pref(S^{i-1}_{E,N}(v,1),\succeq))$$

and for $r > 1$,

$$I^r_{i,N}(v) = h(Pref(\{s \in S^{i-1}_{E,N}(v,r) \mid s(\lambda(v)) \in I^{r-1}_{i,N}(v)\},\succeq))$$

Note that $I^{r+1}_{i,N}(v) \subseteq I^r_{i,N}(v)$. The G-scenario that results from the i^{th} iteration of the expanding approach is denoted $\sigma^i_{E,N}$, and is defined as the G-scenario where $\forall v \in V : \sigma^i_{E,N}(v) = form(M(E^i_N(v)))$.

The definition of expanding iteration captures the following intuition: at radius r, a subset of model vertices are selected at v that are associated with preferred model selections in the neighbourhood of radius r about v. When we expand to radius $r+1$, we select a subset of models at v associated with preferred model selections in the radius $r+1$ neighbourhood *from the models that remain after the selection at radius r*.

In addition to the three iterative approaches defined above, that are based on minimal model selections, we introduce another approach called the *ring method*, which is defined in terms of generic merging operators. The ring method involves *iterated merging* of beliefs from nodes in different "layers" about a particular node, where layer r consists of nodes at distance r from that node. In contrast to the other approaches, the ring method does not consider the topology of the graph between the central node and nodes at distance r.

Definition 3.15 (Ring Iteration)**.** Let $G = \langle V, E\rangle$ be a graph, and let σ be an associated G-scenario. The G-scenario that results from the i^{th} iteration of the ring method is denoted σ^i_R, and is defined as follows. Let $C^r_i(v) = \{\sigma^{i-1}_R(w) \mid dist(v,w) = r\}$ be the *context* at distance r from v, in the i^{th} iteration of the ring method. Then, $\sigma^0_R = \sigma$ and for $i \geq 1$, we have $R^0_i(v) = \sigma^{i-1}_R(v)$ and $R^t_i(v) = \Delta_{R^{t-1}_i(v)}(C^t_i(v))$. σ^i_R is defined as the G-scenario where $\forall v \in V : \sigma^i_R(v) = R^{\epsilon(v)}_i(v)$.

In the t^{th} step of the i^{th} ring iteration, we merge the belief profile $C_i^t(v)$ under the integrity constraint $R_i^{t-1}(v)$, which represents the result of the *previous step of the* i^{th} *iteration*. Δ can be any merging operator, but certain properties of the approach can be shown if Δ satisfies certain IC merging properties.

For simplicity, when we discuss the global completion and the ring approach together with the other approaches (simple, augmenting, and expanding), we do not explicitly denote the neighbourhood type, because the global and ring approaches are *not* defined to use different neighbourhoods. In this case, the neighbourhood type can be arbitrarily chosen.

Each of the four iterative approaches reaches a fixpoint, because each monotonically reduces the set of models at each node. This leads to the following observation:

Observation 1 (Monotonicity)**.** *For* $\Omega \in \{A, E, F, R\}$*, we have:* $\forall v \in V : \sigma_\Omega^{i+1} \models \sigma_\Omega^i$.

For an approach $\Omega \in \{A, F, E, R\}$, denote by σ_Ω^* the fixpoint reached by iterating the approach, so that $\sigma_\Omega^* = \sigma_\Omega^t$ where t is the smallest integer for which $\sigma_\Omega^t = \sigma_\Omega^{t+1}$.

We observe that when Δ is the *consistency-based projection merging operator* [3], the ring method can be expressed as expanding iteration performed with respect to *ring neighbourhoods*, so that $\sigma_R^i = \sigma_{E,R}^i$.

4 RELATIONSHIPS BETWEEN APPROACHES

In the previous section we discussed one global approach to belief change (C), and four iterative approaches: simple iteration (F), expanding iteration (E), augmenting iteration (A), and ring iteration (R). Here we show how these approaches relate to one another (i.e., whether two approaches are comparable, and if so, which approach is logically stronger or weaker than the other). The notion of logical strength of G-scenarios is defined below:

Definition 4.1 (Equivalence and Logical Strength of G-Scenarios)**.** Let G be a graph, and σ and σ' be two G-scenarios.

- We say that σ and σ' are *equivalent*, denoted $\sigma \equiv \sigma'$, iff for all $v \in V$ we have $\models \sigma(v) \equiv \sigma'(v)$.
- We say that σ is *at least as strong* as σ', denoted $\sigma \models \sigma'$, iff $\models \sigma(v) \rightarrow \sigma'(v)$ for all $v \in V$.

Given $\Omega, \Omega' \in \{A, E, F, R, C\}$ and $r, t \geq 1$, we say that σ_Ω^r and $\sigma_{\Omega'}^t$ are *non-comparable* iff there exists a graph $G = \langle V, E\rangle$ and scenario σ such that $\sigma_\Omega^r \not\models \sigma_{\Omega'}^t$ and $\sigma_{\Omega'}^t \not\models \sigma_\Omega^r$.

First we give a result concerning the relationship between neighbourhoods. The following proposition states that adding a new edge between two nodes in a graph cannot logically strengthen the results of iteration, i.e., the results are either equivalent or weaker.

Proposition 3. *Let* $G = \langle V, E\rangle$ *be a graph, and* σ *be a* G*-scenario. Let* $G' = \langle V', E'\rangle$ *be a subgraph of* G *such that* $\exists v, w \in V' : [(v, w) \in E$ *and* $(v, w) \notin E']$*. Let* (v, w) *be such an edge, and let* $G'' = \langle V', E' \cup \{(v, w)\}\rangle$ *be a copy of* G' *that contains* (v, w)*. Then,* $Pref(S(G'), \succeq) \subseteq Pref(S(G''), \succeq)$.

Corollary 3.1. *For* $\Omega \in \{A, E, F\}$*, we have* $\forall i > 0 : \sigma_{\Omega,SPT_j}^i \models \sigma_{\Omega,USPT}^i$ *and* $\forall i > 0 : \sigma_{\Omega,USPT}^i \models \sigma_{\Omega,W}^i$.

That is, for any of the approaches A, E, or F, arbitrary SPT neighbourhoods yield the strongest results, followed by union-of-SPT neighbourhoods, and finally by complete neighbourhoods.

Propositions 4-11 compare the *first iterations* of all the approaches. Then, Corollaries 8.1 and 9.1 and Propositions 12-15 compare their *fixpoints*. For Propositions 4-7 and 10, let G be an arbitrary connected, undirected graph, let σ be an arbitrary G-scenario, and let $N \in \{SPT_j, USPT, W\}$ be an arbitrary neighbourhood.

Proposition 4. $\sigma_{A,N}^1 \models \sigma_{F,N}^1$

Proof. By definition, every model selection s found by the augmenting approach is preferred over restricted neighbourhoods of all radii $1 \leq r \leq diam(G)$; in particular, $sub(s, 1) \in Pref(S_{F,N}^0(v, 1), \succeq)$, which implies that $s(\lambda(v)) \in F_N^1(v)$. Thus, $\sigma_{A,N}^1 \models \sigma_{F,N}^1$. □

Proposition 5. $\sigma_{A,N}^1 \models \sigma_{E,N}^1$

Proof. Due to space constraints, we just give an outline of this proof. Let $B_{A,N}(v, r) = \{s \in Pref(S_{A,N}^0(v, r), \succeq) \mid \forall r'.1 \leq r' \leq r : sub(s, r') \in Pref(S_{A,N}^0(v, r'), \succeq)\}$ and let $B_{E,N}(v, 1) = Pref(S_{E,N}^0(v, 1), \succeq)$ and for $r > 1$, $B_{E,N}(v, r) = Pref(\{s \in S_{E,N}^0(v, r) \mid s(\lambda(v)) \in I_{1,N}^{r-1}(v)\}, \succeq)$. We prove by induction that every minimal model selection of radius r about an arbitrary node v found by the augmenting approach is also found by the expanding approach, so $B_{A,N}(v, r) \subseteq B_{E,N}(v, r)$. When $r = 1$, we have $S_{A,N}^0(v, 1) = S_{E,N}^0(v, 1)$, so $B_{A,N}(v, 1) = B_{E,N}(v, 1)$. For the inductive step, assume that $B_{A,N}(v, r) \subseteq B_{E,N}(v, r)$. Take $s \in B_{A,N}(v, r + 1)$ and let $t = sub(s, r)$. By definition, $t \in B_{A,N}(v, r)$, and by the inductive assumption $t \in B_{E,N}(v, r)$. There are two cases: if $t \in q(B_{E,N}(v, r))$, then we use Proposition 1 to show that $s \in B_{E,N}(v, r + 1)$; if $t \in d(B_{E,N}(v, r))$, then we use Proposition 2 to show that $s \in B_{E,N}(v, r + 1)$. It follows by induction that $\sigma_{A,N}^1 \models \sigma_{E,N}^1$. □

Proposition 6. $\forall i \geq 1 : \sigma_{A,N}^i \models \sigma_C$

Proof. If $m \in A_W^1(v)$, then by definition, $m \in h(Pref(S_W(v, diam(G)), \succeq))$, so $m \in C(v)$. Since $\forall v \in V : A_W^{i+1}(v) \subseteq A_W^i(v)$, we have $\forall v \in V, \forall i \geq 1 : A_W^i(v) \subseteq C(v)$. Thus, $\forall i \geq 1 : \sigma_{A,W}^i \models \sigma_C$. Since $\sigma_{A,USPT}^i \models \sigma_{A,W}^i$ and $\sigma_{A,SPT_j}^i \models \sigma_{A,USPT}^i$ we see that for $N \in \{SPT_j, USPT, W\}$, $\forall i \geq 1 : \sigma_{A,N}^i \models \sigma_C$. □

Proposition 7. $\sigma_{E,N}^1 \models \sigma_{F,N}^1$

Proof. By the definition of expanding iteration, $\sigma_{E,N}^1(v) = form(M(E_N^1(v)))$, where $E_N^1(v) = I_{1,N}^{\epsilon(v)}(v)$. We have $I_{1,N}^1(v) = h(Pref(S_{E,N}^0(v, 1), \succeq)) = h(Pref(S_{F,N}^0(v, 1), \succeq)) = F_N^1(v)$, where $S_{E,N}^0 = S_{F,N}^0$ for any neighbourhood type. Since $I_{i,N}^{r+1}(v) \subseteq I_{i,N}^r(v)$, we have $E_N^1(v) = I_{1,N}^{\epsilon(v)}(v) \subseteq I_{1,N}^1(v) = F_N^1(v)$. This implies that $\sigma_{E,N}^1 \models \sigma_{F,N}^1$. □

Proposition 8. $\forall i \geq 1 : \sigma_{E,N}^i$ *and* σ_C *are in general non-comparable.*

Proof. We begin with a counterexample to the statement that $\sigma_{E,N}^1$ and σ_C are always comparable: Let $G = \langle V, E\rangle$, where $V = \{1, 2, 3, 4, 5\}$ and $E = \{(1, 2), (2, 3), (3, 4), (4, 5)\}$, and let σ be a G-scenario such that $\sigma = \langle 1 : q \vee (r \wedge \neg s), 2 : p, 3 : (\neg p \wedge \neg q) \vee s, 4 : \top, 5 : \neg s\rangle$. Then, the first iteration of the expanding approach yields $\sigma_{E,N}^1$, where $\sigma_{E,N}^1(1) = \sigma_{E,N}^1(2) = \sigma_{E,N}^1(3) = p \wedge q \wedge s$ and $\sigma_{E,N}^1(4) = \sigma_{E,N}^1(5) = \neg p \wedge \neg q \wedge \neg s$, while the global completion yields $\sigma_C = \langle 1 : (p \wedge q \wedge s) \vee (p \wedge \neg q \wedge r \wedge \neg s), 2 : (p \wedge q \wedge s) \vee (p \wedge \neg q \wedge r \wedge \neg s), 3 : (p \wedge q \wedge s) \vee (\neg p \wedge \neg q \wedge r \wedge \neg s), 4 : (p \wedge q) \vee (\neg p \wedge \neg q \wedge r \wedge \neg s), 5 : (p \wedge q \wedge \neg s) \vee (\neg p \wedge \neg q \wedge r \wedge \neg s)\rangle$. We

see that $\sigma^1_{E,N}(4) \not\models \sigma_C(4)$ and $\sigma_C(4) \not\models \sigma^1_{E,N}(4)$. For this example, we have $\forall i > 1 : \sigma^i_{E,N} = \sigma^1_{E,N}$, so $\forall i \geq 1 : \sigma^i_{E,N}$ and σ_C are in general non-comparable. □

Corollary 8.1. *$\sigma^*_{E,N}$ and σ_C are non-comparable.*

Proposition 9. *$\forall i \geq 1 : \sigma^i_{F,N}$ and σ_C are in general non-comparable.*

Proof. We begin with a counterexample to the statement that $\sigma^1_{F,N}$ and σ_C are always comparable: Let $G = \langle V, E\rangle$, where $V = \{1,2,3,4\}$ and $E = \{(1,2),(2,3),(3,4)\}$, and let σ be a G-scenario such that $\sigma = \langle 1 : \neg r, 2 : \neg p \vee r, 3 : (p \vee q) \wedge r, 4 : \neg q\rangle$. Then the result of the first simple iteration is $\sigma^1_{F,N} = \langle 1 : \neg p \wedge \neg r, 2 : (p \wedge r) \vee (\neg p \wedge q), 3 : p \wedge \neg q \wedge r, 4 : p \wedge \neg q \wedge r\rangle$ while the result of the global completion is $\sigma_C = \langle 1 : (\neg p \vee \neg q) \wedge \neg r, 2 : (\neg p \wedge \neg r) \vee (p \wedge \neg q \wedge r), 3 : r \wedge (p \vee q) \wedge (\neg p \vee \neg q), 4 : \neg q \wedge r\rangle$. We see that $\sigma^1_{F,N}(2) \not\models \sigma_C(2)$ and $\sigma_C(2) \not\models \sigma^1_{F,N}(2)$. For this example, we have $\sigma^2_{F,N} = \langle 1 : \neg p \wedge q \wedge \neg r, 2 : (\neg p \wedge q) \vee (p \wedge \neg q \wedge r), 3 : p \wedge \neg q \wedge r, 4 : p \wedge \neg q \wedge r\rangle$. We find that $\sigma^2_{F,N}(2) \not\models \sigma_C(2)$ and $\sigma_C(2) \not\models \sigma^2_{F,N}(2)$, so the second iteration of the simple approach is in general non-comparable with the global completion. In addition, $\forall i > 2 : \sigma^i_{F,N} = \sigma^2_{F,N}$, so $\forall i \geq 1 : \sigma^i_{F,N}$ and σ_C are non-comparable. □

Corollary 9.1. *$\sigma^*_{F,N}$ and σ_C are non-comparable.*

Now, we show that the first iteration of the ring method produces logically stronger results than the first iteration of the simple approach.

Proposition 10. *$\sigma^1_R \models \sigma^1_{F,N}$ if the merging operator Δ used in the ring method is the consistency-based projection merging operator.*

Proof. From the definition of the ring method, $R^0_1(v) = \sigma^0_R(v) = \sigma(v)$, and $R^1_1(v) = \Delta_{R^0_1(v)}(C^1_1(v)) = \Delta_{\sigma(v)}(C^1_1(v))$. Note that $\Delta_{\sigma(v)}(C^1_1(v)) = \sigma^1_{F,USPT}(v) = \sigma^1_{F,SPT_j}(v)$ and $\Delta_{\sigma(v)}(C^1_1(v)) \models \sigma^1_{F,W}(v)$, so we can say that $\Delta_{\sigma(v)}(C^1_1(v)) \models \sigma^1_{F,N}(v)$ regardless of the type of neighbourhood used for simple iteration. For all $i > 1$, $R^i_1(v) \models R^1_1(v)$; in particular, $R^{\epsilon(v)}_1(v) \models R^1_1(v)$, and since $R^1_1(v) \models \sigma^1_{F,N}(v)$, we have $\sigma^1_R \models \sigma^1_{F,N}$. □

The following proposition shows that the first iteration of the ring method is in general non-comparable to the first iterations of augmenting iteration, expanding iteration, or the global completion.

Proposition 11. *σ^1_R is in general not comparable with any of $\sigma^1_{A,N}$, $\sigma^1_{E,N}$, or σ_C.*

Proof. All these statements can be proven by a single example for which augmenting iteration, expanding iteration, and the global completion produce the *same result*, which is different from, and not comparable to, the result produced by the ring method. Consider the graph $G = \langle V, E\rangle$ where $V = \{1,2,3,4\}$ and $E = \{(1,2),(2,3),(3,4)\}$ and let σ be a G-scenario such that $\sigma = \langle 1 : p \oplus q, 2 : p, 3 : \top, 4 : \neg p\rangle$. We find that $\sigma^1_{A,N} = \sigma^1_{E,N} = \sigma_C = \langle 1 : p \wedge \neg q, 2 : p \wedge \neg q, 3 : \neg q, 4 : \neg p \wedge \neg q\rangle$, while $\sigma^1_R = \langle 1 : p \wedge \neg q, 2 : p \wedge \neg q, 3 : p \oplus q, 4 : \neg p \wedge q\rangle$. Note that $\sigma^1_{A,N}(4) = \sigma^1_{E,N}(4) = \sigma_C(4) = \neg p \wedge \neg q$, while $\sigma^1_R(4) = \neg p \wedge q$. Since $\neg p \wedge \neg q \not\models \neg p \wedge q$ and $\neg p \wedge q \not\models \neg p \wedge \neg q$, we conclude that the first iterations of expanding iteration, augmenting iteration, and the global completion are in general non-comparable with the first iteration of the ring method. □

Propositions 12-15 compare the *fixpoints* of the augmenting, expanding, simple, and ring approaches.

Proposition 12. *$\sigma^*_{E,N}$ and $\sigma^*_{F,N}$ are in general non-comparable.*

Proof. Counterexample to the statement that $\sigma^*_{E,N}$ and $\sigma^*_{F,N}$ are always comparable: Let $G = \langle V, E\rangle$, where $V = \{1,2,3,4,5,6\}$ and $E = \{(i, i+1) \mid 1 \leq i \leq 5\}$, and let σ be a G-scenario such that $\sigma = \langle 1 : \neg r, 2 : \top, 3 : \neg p \vee r, 4 : (p \vee q) \wedge r, 5 : \top, 6 : \neg q\rangle$. The fixpoint of simple iteration is $\sigma^*_{F,N} = \langle 1 : \neg p \wedge q \wedge \neg r, 2 : \neg p \wedge q \wedge \neg r, 3 : \neg p \wedge q \wedge r, 4 : p \wedge \neg q \wedge r, 5 : p \wedge \neg q \wedge r, 6 : p \wedge \neg q \wedge r\rangle$ The fixpoint of expanding iteration is $\sigma^*_{E,N} = \langle 1 : \neg p \wedge q \wedge \neg r, 2 : \neg p \wedge q \wedge \neg r, 3 : p \wedge \neg q \wedge r, 4 : p \wedge \neg q \wedge r, 5 : p \wedge \neg q \wedge r, 6 : p \wedge \neg q \wedge r\rangle$. We see that $\sigma^*_{F,N}(3) \not\models \sigma^*_{E,N}(3)$ and $\sigma^*_{E,N}(3) \not\models \sigma^*_{F,N}(3)$, so the simple and expanding approaches are not comparable in general. □

Proposition 13. *$\sigma^*_{E,N}$ and $\sigma^*_{A,N}$ are in general non-comparable.*

Proof. Counterexample to the statement that $\sigma^*_{E,N}$ and $\sigma^*_{A,N}$ are always comparable: Let $G = \langle V, E\rangle$, where $V = \{1,2,3,4,5\}$ and $E = \{(1,2),(2,3),(3,4),(4,5)\}$, and let σ be a G-scenario such that $\sigma = \langle 1 : p, 2 : \top, 3 : (p \wedge r) \vee (\neg p \wedge \neg q \wedge \neg r), 4 : \neg r, 5 : q\rangle$. Then the fixpoint of augmenting iteration is $\sigma^*_{A,N} = \langle 1 : p \wedge q \wedge r, 2 : p \wedge q \wedge r, 3 : \neg p \wedge \neg q \wedge \neg r, 4 : \neg p \wedge \neg r, 5 : \neg p \wedge q \wedge \neg r\rangle$, while the fixpoint of expanding iteration is $\sigma^*_{E,N} = \langle 1 : p \wedge \neg q \wedge r, 2 : p \wedge \neg q \wedge r, 3 : \neg p \wedge \neg q \wedge \neg r, 4 : \neg p \wedge \neg r, 5 : \neg p \wedge q \wedge \neg r\rangle$. Note that the beliefs produced by each approach at vertices 1 and 2 are not comparable, i.e., $\sigma^*_{E,N}(1) \not\models \sigma^*_{A,N}(1)$ and $\sigma^*_{A,N}(1) \not\models \sigma^*_{E,N}(1)$. □

Proposition 14. *$\sigma^*_{A,N}$ and $\sigma^*_{F,N}$ are in general non-comparable.*

Proof. Counterexample to the statement that $\sigma^*_{A,N}$ and $\sigma^*_{F,N}$ are always comparable: Let $G = \langle V, E\rangle$, where $V = \{1,2,3,4,5\}$ and $E = \{(1,2),(2,3),(3,4),(4,5)\}$, and let σ be a G-scenario such that $\sigma = \langle 1 : p \wedge r, 2 : p, 3 : s, 4 : \neg r \vee s, 5 : \neg s\rangle$. Then the fixpoint of simple iteration is $\sigma^*_{F,N} = \langle 1 : p \wedge r \wedge s, 2 : p \wedge r \wedge s, 3 : p \wedge r \wedge s, 4 : p \wedge \neg r, 5 : p \wedge \neg r \wedge \neg s\rangle$, while the fixpoint of augmenting iteration is $\sigma^*_{A,N} = \langle 1 : p \wedge r \wedge s, 2 : p \wedge r \wedge s, 3 : p \wedge r \wedge s, 4 : p \wedge (r \leftrightarrow s), 5 : p \wedge \neg r \wedge \neg s\rangle$. Note that the beliefs produced by each approach at vertex 4 are not comparable, i.e., $\sigma^*_{F,N}(4) \not\models \sigma^*_{A,N}(4)$ and $\sigma^*_{A,N}(4) \not\models \sigma^*_{F,N}(4)$. □

We also show that the fixpoint of the ring method is in general non-comparable to the fixpoints of augmenting iteration, expanding iteration, simple iteration, or the global completion.

Proposition 15. *σ^*_R is in general not comparable with any of $\sigma^*_{A,N}$, $\sigma^*_{E,N}$, $\sigma^*_{F,N}$, or σ_C.*

Proof. All these statements can be proven by a single example for which augmenting iteration, expanding iteration, simple iteration and the global completion produce the *same fixpoint*, which is different from, and not comparable to, the fixpoint produced by the ring method. Once again consider the graph and associated scenario shown in the proof of Proposition 11. We find that $\sigma^*_{A,N} = \sigma^*_{E,N} = \sigma^*_{F,N} = \sigma_C = \langle 1 : p \wedge \neg q, 2 : p \wedge \neg q, 3 : \neg q, 4 : \neg p \wedge \neg q\rangle$, while $\sigma^*_R = \langle 1 : p \wedge \neg q, 2 : p \wedge \neg q, 3 : p \wedge \neg q, 4 : \neg p \wedge q\rangle$. Note that $\sigma^*_{A,N}(4) = \sigma^*_{E,N}(4) = \sigma^*_{F,N}(4) = \sigma_C(4) = \neg p \wedge \neg q$, while $\sigma^*_R(4) = \neg p \wedge q$. Since $\neg p \wedge \neg q \not\models \neg p \wedge q$ and $\neg p \wedge q \not\models \neg p \wedge \neg q$, we conclude that the fixpoints of expanding iteration, augmenting iteration, simple iteration, and the global completion are in general non-comparable with the fixpoint of the ring method. □

5 LOGICAL PROPERTIES

In this section we investigate logical properties that are satisfied by the approaches defined in Section 3. Observation 1 already noted the monotonicity of the iterative approaches. Here, we start with the basic property that if an initial scenario is consistent, then it remains consistent throughout every iteration of each approach.

Proposition 16 (Consistency Preservation)**.** *Let G be a graph and σ be a G-scenario. If σ is consistent, then σ^i_Ω is consistent for all i and all $\Omega \in \{A, E, F, R, C\}$.*

It proves to be the case that if the conjunction of formulas at all vertices is consistent, then the result at each vertex will be just this conjunction. We formalize this as follows:

Definition 5.1 (Agreement Preservation (AP))**.** Let $G = \langle V, E\rangle$ be a graph, σ be a G-scenario, and σ' be the G-scenario that results from applying an approach to updating the information at the vertices in V. If $\bigwedge_{v\in V} \sigma(v)$ is consistent, then $\forall v \in V : \sigma'(v) \equiv \bigwedge_{v\in V} \sigma(v)$.

The following proposition states that the augmenting, expanding, simple, and global approaches satisfy Agreement Preservation.

Proposition 17. *For $\Omega \in \{A, E, F, C\}$, σ^*_Ω satisfies (AP).*

Proof. This was shown for $\Omega = C$ in [2]. If $\bigwedge_{v\in V} \sigma(v)$ is consistent, then $\bigcap_{v\in V} Mod(\sigma(v)) \neq \varnothing$. For each model $m \in \bigcap_{v\in V} Mod(\sigma(v))$, we can define a model selection $s \in S(G)$ such that for each $v \in V$, $s(\lambda(v)) = (v, m)$. Then, $\Delta(s) = \varnothing$. Clearly, s is minimal over any connected subgraph of G. So, given a vertex v, s is minimal over a neighbourhood of any radius about v. In particular, s is minimal in the radius-1 neighbourhood about v, so (v, m) is selected by the simple approach. s is minimal over neighbourhoods of all radii $1 \leq r \leq diam(G)$, so (v, m) is selected by the augmenting approach. For the expanding approach, s is minimal over a neighbourhood of radius 1, so (v, m) is selected in the first step. Since all other vertices agree on model m, (v, m) will continue to be part of the minimal selection s over increasingly large neighbourhoods. Thus, (v, m) is selected by the expanding approach. Finally, although additional models $m' \notin \bigcap_{v\in V} Mod(\sigma(v))$ may initially be selected by an approach (e.g. due to local minimization), such models are progressively eliminated when the approach is iterated, so that the fixpoint of each approach contains only the information agreed upon by all vertices. □

We also show that the ring method satisfies the Agreement Preservation property under certain conditions.

Proposition 18. *The ring method satisfies (AP) iff Δ satisfies (IC2).*

Proof. Let $G = \langle V, E\rangle$ be a graph, and let $v \in V$ be an arbitrary vertex. Note that $C^1_1(v) \sqcup C^2_1(v) \sqcup \cdots \sqcup C^{\epsilon(v)}_1(v) = \bigsqcup_{1\leq r\leq \epsilon(v)} C^r_1(v) = \{\sigma(v) \mid v \in V\}$, where $\sqcup$ denotes multiset union. Then, by using the definition of the ring method, we have $R^0_1(v) = \sigma(v)$ and

$$R^1_1(v) = \Delta_{R^0_1(v)}(C^1_1(v)) = \bigwedge C^1_1(v) \wedge \sigma(v)$$

$$\vdots$$

$$R^{\epsilon(v)}_1(v) = \Delta_{R^{\epsilon(v)-1}_1(v)}(C^{\epsilon(v)}_1(v)) = \bigwedge_{v\in V} \sigma(v)$$

So $\sigma^1_R(v) = \bigwedge_{v\in V} \sigma(v)$ for all $v \in V$. Since the beliefs of all nodes are equivalent after the first iteration of the ring method, clearly there will be no further change, so $\sigma^*_R(v) = \bigwedge_{v\in V} \sigma(v)$ for all $v \in V$. □

The number of iterations for any approach to reach a fixpoint has an upper bound, given by the following proposition.

Proposition 19. *Given a graph $G = \langle V, E\rangle$ and a G-scenario σ, an upper bound on the number of iterations of any approach $\Omega \in \{A, E, F, R\}$ is $\left(\sum_{v\in V} |Mod(\sigma(v))|\right) - |V|$.*

Proof. By monotonicity, in each iteration the number of models at a node can either stay the same or decrease. Clearly, the smallest possible change from one iteration to the next is for a single model to be eliminated from the set of models at a specific node. The total number of models over all nodes in the graph is $\sum_{v\in V} |Mod(\sigma(v))|$. At the fixpoint (assuming σ is consistent), each node must have at least one model; the minimum number of models over the graph in any iteration is $|V|$. Thus, in the case that each iteration yields the minimal change, we see that there can be *at most* $\left(\sum_{v\in V} |Mod(\sigma(v))|\right) - |V|$ iterations before the scenario reaches a fixpoint. □

6 DISCUSSION

This work calls for some comparisons with Belief Revision Games (BRGs) [13]. A BRG represents a network of communication agents, where each agent iteratively updates its beliefs by applying a *revision policy* that takes into account its current beliefs and the beliefs of its neighbours. By definition, each revision policy takes as input the current belief of an agent and the belief profile consisting of the beliefs of her neighbours, and returns a new set of beliefs. In this approach, agents do not take into account the topology of the graph beyond their immediate neighbours. The fixed-radius approach (using USPT neighbourhoods) introduced in this paper has a similar structure to BRGs, in that agents repeatedly update their beliefs by taking into account the beliefs of their immediate neighbours. However, our method of updating information by minimizing disagreements in a set-theoretic fashion is distinct from any of the IC merging operator-based revision policies. The expanding and augmenting approaches presented here further diverge from the BRG approach; they take as input an entire graph G, an associated scenario σ, and a node v, and yield a new set of beliefs for v. The expanding and augmenting approaches take into account the beliefs of agents throughout the graph, as well as the connections between those agents, to minimize disagreement within increasingly large neighbourhoods about a vertex of interest, in a prioritized manner.

7 CONCLUSION

In this paper, we have generalized the work presented in [2] by defining minimization-based approaches to iterated belief change that take into account the distance between nodes in a graph. We defined four iterated approaches which we call the simple, augmenting, expanding, and ring approaches. We showed that each approach is distinct, in that each produces results that are non-comparable with the results of the other approaches. We also compared these iterated approaches to the global approach studied previously. We found that the augmenting approach produces logically stronger results in each iteration than global completion, while the other approaches produce results that are non-comparable with global completion. We examined basic logical properties for the approaches, and stated an upper bound on the number of iterations needed to reach a fixpoint. Finally, we compared our approaches to related work.

We have also implemented the approaches to belief change described in this paper, in a software system called *Equibel*. This software is publicly available at `https://github.com/asteroidhouse/equibel`.

REFERENCES

[1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *Journal of Symbolic Logic*, **50**(2), 510–530, (1985).

[2] James P. Delgrande, Jérôme Lang, and Torsten Schaub, 'Belief change based on global minimisation', in *Proc. International Joint Conference on Artificial Intelligence (IJCAI'07)*, pp. 2468–2473, Hyderabad, India, (2007).

[3] James P. Delgrande and Torsten Schaub, 'A consistency-based framework for merging knowledge bases', *Journal of Applied Logic*, **5**(3), 459–477, (2006).

[4] Florence Dupin de Saint-Cyr and Jérôme Lang, 'Belief extrapolation (or how to reason about observations and unpredicted change)', in *Proc. of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning (KR'02)*, pp. 497–508, Toulouse, France, (2002).

[5] Patricia Everaere, Sébastien Konieczny, and Pierre Marquis, 'Quota and Gmin merging operators', in *Proc. International Joint Conference on Artificial Intelligence (IJCAI'05)*, pp. 424–429, Edinburgh, UK, (2005).

[6] Olivier Gauwin, Sébastien Konieczny, and Pierre Marquis, 'Iterated belief merging as conciliation operators', in *Seventh International Symposium on Logical Formalizations of Commonsense Reasoning*, pp. 85–92, Corfu, Greece, (2005).

[7] Hirofumi Katsuno and Alberto O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**(3), 263–294, (1991).

[8] Sébastien Konieczny, Jérôme Lang, and Pierre Marquis, 'Distance-based merging: A general framework and some complexity results', in *Proc. of the Eighth International Conference on the Principles of Knowledge Representation and Reasoning (KR'02)*, pp. 97–108, San Francisco, USA, (2002).

[9] Sébastien Konieczny and Ramón Pino Pérez, 'Merging information under constraints: A logical framework', *Journal of Logic and Computation*, **12**(5), 773–808, (2002).

[10] Sébastien Konieczny and Ramón Pino Pérez, 'Logic based merging', *Journal of Philosophical Logic*, **40**(2), 239–270, (2011).

[11] Paolo Liberatore and Marco Schaerf, 'Arbitration: A commutative operator for belief revision', in *Proc. of the Second World Conference on the Fundamentals of Artificial Intelligence (WOCFAI-95)*, pp. 217–228, Paris, France, (1995).

[12] Paolo Liberatore and Marco Schaerf, 'BReLS: A system for the integration of knowledge bases', in *Proc. of the Seventh International Conference on the Principles of Knowledge Representation and Reasoning (KR'00)*, pp. 145–152, Breckenridge, CO, USA, (2000).

[13] Nicolas Schwind, Katsumi Inoue, Gauvain Bourgne, Sébastien Konieczny, and Pierre Marquis, 'Belief revision games', in *Proc. of the Twenty-Ninth National Conference on Artificial Intelligence (AAAI-15)*, pp. 1590–1596, Austin, TX, USA, (2015).

[14] Eric Würbel, Robert Jeansoulin, and Odile Papini, 'Revision: An application in the framework of GIS', in *Proc. of the Seventh International Conference on the Principles of Knowledge Representation and Reasoning (KR'00)*, pp. 505–515, Breckenridge, CO, USA, (2000).