A Modal Logic Framework for Multi-agent Belief Fusion

CHURN-JUNG LIAU

Institute of Information Science, Academia Sinica, Taipei 115, Taiwan

This paper provides a modal logic framework for reasoning about multi-agent belief and its fusion. We propose logics for reasoning about cautiously merged agent beliefs that have different degrees of reliability. These logics are obtained by combining the multi-agent epistemic logic and multi-source reasoning systems. The fusion is cautious in the sense that if an agent's belief is in conflict with those of higher priorities, then his belief is completely discarded from the merged result. We consider two strategies for the cautious merging of beliefs. In the first, called level cutting fusion, if inconsistency occurs at some level, then all beliefs at the lower levels are discarded simultaneously. In the second, called level skipping fusion, only the level at which the inconsistency occurs is skipped. We present the formal semantics and axiomatic systems for these two strategies and discuss some applications of the proposed logical systems. We also develop a tableau proof system for the logics and prove the complexity result for the satisfiability and validity problems of these logics.

Categories and Subject Descriptors: I.2.4 [**Artificial Intelligence**]: Knowledge Representation Formalisms and Methods—*Modal Logic*; F.4.1 [**Mathematical Logic**]: Modal Logic

General Terms: Design, Languages, Theory

Additional Key Words and Phrases: Epistemic logic, multi-sources reasoning, database merging, belief fusion, belief revision, multi-agent systems

1. INTRODUCTION

Recently, attention has been focussed on the infoglut problem in information retrieval research due to the rapid growth of internet information. If a keyword is input to a commonly-used search engine, it is not unusual to receive a list of thousands of web pages; the real difficulty is not how to find information, but how to locate useful information. Many software agents have been designed to solve the infoglut problem while conducting searches. The agents search through the web and try to find information matching the user's need. However, not all internet information sources are reliable. Some web sites are out-of-date, some news sites provide incorrect information, while others even intentionally spread rumor or hazardous information. Therefore, an important task for information search agents is to merge information coming from different sources according to its degree of

Author's address: Churn-Jung Liau, Institute of Information Science, Academia Sinica, Taipei, 115, Taiwan, E-mail: liaucj@iis.sinica.edu.tw

A preliminary version of the paper has appeared in [Liau 2000].

reliability.

In [Shoham 1993], an agent is characterized with mental attitudes, such as knowledge, belief, obligation, and commitment. This view, in accordance with the intentional stance proposed in [Dennett 1987], has been widely accepted as convenient for the analysis and description of complex systems[Wooldridge and Jennings 1995]. From this viewpoint, each information provider can be considered as an agent and the information provided by the agent corresponds to his belief. The information fusion problem, therefore, is also to merge beliefs of different agents.

The philosophical analysis of these mental attitudes has motivated the development of many non-classical logic systems[Gabbay and Guenthner 1984]. In particular, the analysis of informational attitudes, such as knowledge and belief, has been a traditional concern of epistemology, an important and central branch of philosophy. We needed a formalism more rigorous than natural language to answer key epistemological questions such as "What is knowledge?" "What can we know?" and "What are the characteristic properties of knowledge?". The result was the development of the so-called epistemic logic[Hintikka 1962]. This kind of logic has attracted much attention of researchers from diverse fields such as artificial intelligence(AI), economics, linguistics, and theoretical computer science. In particular, AI researchers and computer scientist have further developed some technically sophisticated formalisms and applied them to the analysis of distributed and multi-agent systems[Fagin et al. 1996; Meyer and van der Hoek 1995].

Though the original epistemic logic in philosophy is mainly concerned with the single-agent case, its application to AI and computer science places emphasis on agent interaction, and this requires a multi-agent epistemic logic. One example of such a logic is proposed by Fagin et al.[Fagin et al. 1996]. In their logic, the knowledge of each agent is represented by a normal modal operator[Chellas 1980], so if no interactions between agents occur, this is merely a multi-modal logic. However, the most novel feature of their logic is the consideration of common knowledge and distributed knowledge among a group of agents. While common knowledge is defined as information that everyone knows, everyone knows that everyone knows, everyone knows that everyone knows that everyone knows, and so on, distributed knowledge is information that can be deduced by pooling together everyone's knowledge. Consequently, it is distributed knowledge that is the main concern in the fusion of knowledge among agents. However, "knowledge" is used in a broad sense in [Fagin et al. 1996] and covers belief and information.[1] Though proper knowledge must be true, the belief of an agent may be wrong and may cause conflicts in the beliefs that are to be merged. In this case, everything can be deduced from the distributed beliefs due to the notorious omniscience property of epistemic logic; the merged result will then be useless for further reasoning.

Instead of directly merging all beliefs of the agents together, there are many sophisticated techniques for knowledge base merging[Baral et al. 1992; 1994; Baral et al. 1991; Cholvy 1994; Cholvy and Hunter 1997; Konieczny 2000; Konieczny and Pérez 1998; 1999; Lin 1994; 1996; Lin and Mendelzon 1999; Nebel 1994; Pradhan

[1] More precisely, the logic for belief is called doxastic logic. However, here we will use the three terms knowledge, belief, and information interchangeably, so epistemic logic is assumed to cover all these notions.

et al. 1995; Subrahmanian 1994]. Most approaches treat belief fusion operators as meta-level constructs, so for a given set of knowledge bases, fusion operators will return the merged results. Some of the approaches propose concrete operators that can be used directly in the fusion process, while others stipulate the desirable properties of reasonable belief fusion operators by postulates. However, few of the approaches have the capability to reason about the fusion process[2]. One of a few exceptions is multi-source reasoning[Cholvy 1994].

Multi-source reasoning models the fusion process of multiple databases in a modal logic. Its goal is to merge a set of databases according to a total ordering of the set that is to be merged. Each database is a finite and satisfiable set of literals. Two attitudes for merging are considered. According to the suspicious attitude, if a database contains a literal that is inconsistent with those in the databases with higher reliability, then the database is completely discarded from the merged result. On the other hand, according to the trusting attitude, if a literal in a database is inconsistent with those in the databases with higher reliability, only the literal is discarded, and other literals in the database are still considered if they are consistent with those in the databases of higher reliability.

Since multi-source reasoning is modelled in a modal logic framework, it is very useful when integrated with epistemic logic. Its restriction is that each database must be a set of literals; however, in the multi-agent epistemic logic, it is expected that more complex compound formulas will be believed by agents. Therefore, we have to extend the multi-sources reasoning to the more general case. To achieve the purpose, the distributed knowledge operators in multi-agent epistemic logic can be used. What we have to do is to adapt the multi-agent epistemic logic so that the distributed knowledge among a group of agents with reliability ordering can also be defined. However, since the set of facts believed by an agent is closed under classical logical equivalence, whereas the trusting attitude of multi-source reasoning adopts a syntax-dependent fusion, it cannot be modelled directly. For example, if p and q are both believed by an agent and $\neg p \vee \neg q$ is believed by another agent with higher reliability, then using the trusting attitude, either p or q should be in the merged result (assuming no other conflicts exist), however, it is obvious that the belief of the first agent is equivalent to $p \wedge q$ and if it is expressed in this way, then no belief of the first agent (except the obvious tautology) should be included in the merged belief. Therefore, we only consider the merging of beliefs according to the suspicious attitude; this approach is very cautious from the viewpoint of belief fusion. However, we show that the syntax-dependent fusion can also be simulated in our logic.

We consider two strategies for the cautious merging of beliefs. In the first, called level cutting fusion, if inconsistency occurs at some level, then all beliefs at the lower levels are discarded simultaneously. In the second, called level skipping fusion, only the level at which the inconsistency occurs is skipped.

On one hand, level cutting fusion is inspired by threshold reasoning with uncertain information. In uncertainty reasoning, it is usual that information with certainty below a certain threshold is discarded. The threshold may be static or dynamic. The inconsistency level is considered to be a dynamic threshold for such

[2]A more detail comparison of these approaches and ours is given in section 5.

reasoning. Therefore, in situations in which the ordering between agents represents a kind of certainty levels, once the threshold is set up due to the occurrence of the inconsistency, all pieces of information with lower certainty should be discarded. We provide examples to show the use of level cutting fusion in such situations.

On the other hand, level skipping fusion is motivated by the maximal consistent subset combination. This kind of combination is one of the most straightforward knowledge base merging techniques. Indeed, it is quite natural to select a maximal consistent subset from an inconsistent set of logical formulas for restoring the consistency of the merged knowledge base. However, as previously mentioned, most maximal consistent subset combination approaches are syntax-dependent, whereas our belief fusion logics are semantic-based. Therefore, we select a maximal consistent subset of agents instead of formulas, though we can simulate the syntax-based approach by assigning to each logical formula a particular agent.

We propose two logics for level cutting and level skipping strategies respectively. Their syntax, semantics, complete axiomatic systems, and tableau proof systems are all presented. Some examples in different realistic domains are used to illustrate the potential applications of these logics. The complexity analysis of the tableau proof systems shows that the logics have the same complexity with the normal modal logic system KD and multi-agent epistemic logic K_n. Therefore, our logics extend the expressive power of the multi-agent epistemic logic without increasing computational complexity.

We integrate multi-source reasoning to enhance the reasoning capability of multi-agent epistemic logic. However, we would also like to consider the extension of logics using some of the more sophisticated fusion operators proposed in the literature. We show a generic extension of our logics to accommodate these belief fusion operators. This means that the belief fusion operators as a standard add-on of multi-agent epistemic logic can be expected.

The rest of the paper is organized as follows. The logics for level cutting fusion and level skipping fusion are introduced in section 2 and 3 respectively. The syntax, semantics, and axiomatic systems of the logics are presented. Some possible applications and realistic examples are also given to illustrate the use of the logics. A tableau proof system for the proposed logics is presented in section 4 and the complexity results are proven via an alternative formulation of the tableau calculus. In section 5, we show the possibility of extending the basic framework to accommodate more sophisticated belief fusion operators and compare the proposed logics with some existing works. Finally, we conclude and discuss some further research directions in section 6.

1.1 Notational preliminary

In the following presentation, we extensively use the notions of ordering relations. Let X be a set, then a binary relation $\geq$ over X is:

(1) *reflexive* if $x \geq x$ for all $x \in X$;
(2) *transitive* if $x \geq y \wedge y \geq z \Rightarrow x \geq z$ for all $x, y, z \in X$;
(3) *anti-symmetric* if $x \geq y \wedge y \geq x \Rightarrow x = y$ for all $x, y \in X$;
(4) a *pre-order* if it is reflexive and transitive;
(5) a *partial order* if it is a pre-order satisfying anti-symmetry;

(6) a *total pre-order* (or connected order) if it is a pre-order and $x \geq y \vee y \geq x$ for all $x, y \in X$;
(7) a *total order* if it is a total pre-order satisfying anti-symmetry.

We write $x > y$ as the abbreviation of $x \geq y$ and $y \not\geq x$, and the binary relation "$>$" is the strict version of $\geq$. For example, if "$\geq$" is a total order, then "$>$" is a strict total order. An element y is said to be an *immediate* $>$*-successor* of another element x if $x > y$, and there does not exist any element $z \in X$ such that $x > z$ and $z > y$. A total order $(X, \geq)$ contains a partial order $(X, \geq')$ if $x \geq' y \Rightarrow x \geq y$ for all $x, y \in X$.

To encode the degrees of reliability of n agents, we use ordering relations over any subset of $\{1, \ldots, n\}$. Let $\mathcal{TO}_n$ denote the set of all possible total orders over any non-empty subset of $\{1, \ldots, n\}$, then we can associate with each total order in $\mathcal{TO}_n$ a unique syntactic notation. Let $X = \{i_1, i_2, \ldots, i_m\}$ be a non-empty subset of $\{1, \ldots, n\}$ and $\geq$ be a total order such that $i_j \geq i_k$ iff $j \leq k$ for all $1 \leq, j, k \leq m$, then the syntactic notation for $(X, \geq)$ is the string

$$i_1 > i_2 > \cdots > i_m.$$

In this paper, the capital letter O is used as meta-variables ranging over such notations. Let O be the string $i_1 > i_2 > \cdots > i_m$, then the set $\{i_1, i_2, \ldots, i_m\}$ is called the domain of O and is denoted by $\delta(O)$. In this case, $O > i_{m+1}$ denotes $i_1 > i_2 > \cdots > i_m > i_{m+1}$ if $i_{m+1} \notin \delta(O)$. As the syntactic notation is unique for each total order, we can also identify the notation with the total order itself, so we can write $O \in \mathcal{TO}_n$. Furthermore, the upper-case Greek letter Ω is used as meta-variables ranging over nonempty subsets of $\mathcal{TO}_n$.

Analogously, if $(X, \geq)$ is a partial order, then its syntactic notation is

$$\{x > y \mid x, y \in X, \ \ y \text{ is an immediate } >\!-\text{successor of } x\}.$$

We use the capital letter Q as meta-variables ranging over such notations and also identify it with the corresponding partial order. Let Q be a partial order over X, then $\mathcal{TO}_Q$ denotes the set of all total orders over X containing Q.

We use some standard notations for binary relations in the following presentation. If $\mathcal{R} \subseteq X \times Y$ is a binary relation between X and Y, we write $\mathcal{R}(x, y)$ for $(x, y) \in \mathcal{R}$ and $\mathcal{R}(x)$ for the subset $\{y \in Y \mid \mathcal{R}(x, y)\}$. A binary relation $\mathcal{R} \subseteq X \times Y$ is *serial* if $\forall x \exists y R(x, y)$.

2. LOGIC FOR LEVEL CUTTING FUSION

In this section, we introduce a logic DBF_n^c for distributed belief fusion using the level cutting strategy, where n is the number of agents.

2.1 Syntax

Let $\mathcal{L}_c$ denote the language of DBF_n^c. We first formally present the syntax of $\mathcal{L}_c$.

Definition 1. The alphabet of $\mathcal{L}_c$ contains the following symbols:

(1) A countable set $\Phi_0 = \{p, q, r, \ldots\}$ of atomic propositions;
(2) The propositional constants $\perp$ (falsum or falsity constant) and $\top$ (verum or truth constant);

(3) The binary Boolean connective $\vee$ (or), and unary Boolean operator $\neg$ (not);

(4) A set $Ag = \{1, 2, \ldots, n\}$ of agents;

(5) The modal operator-forming symbols "[" and "]", set construction symbols "{" and "}", and the ordering symbol ">";

(6) The left and right parentheses "(" and ")", and the punctuation symbol ",".

Definition 2. The well-formed formulas(wffs) of $\mathcal{L}_c$ are defined by the following rules:

(1) if $p \in \Phi_0$, then p is a wff;

(2) $\bot$ and $\top$ are wffs;

(3) if φ is a wff, then $\neg\varphi$ is also a wff;

(4) if φ and ψ are wffs, then so too is $\varphi \vee \psi$;

(5) if φ is a wff, then $[G]\varphi$ and $[O]\varphi$ are wffs for any nonempty $G \subseteq Ag$ and $O \in \mathcal{TO}_n$.

As usual, other classical Boolean connectives $\wedge$ (and), $\supset$ (implication), and $\equiv$ (equivalence) can be defined as abbreviations. Also, we write $\langle G\rangle\varphi$ and $\langle O\rangle\varphi$ as abbreviations of $\neg[G]\neg\varphi$ and $\neg[O]\neg\varphi$, respectively.

The intuitive meaning of $[G]\varphi$ is "The group of agents G has distributed belief φ", whereas $[O]\varphi$ means that φ is derivable from the merged beliefs of agents in $\delta(O)$ according to the specific order O.

Let Q be a partial order on a subset of agents, then define $[Q]\varphi$ as the abbreviation of $\bigwedge_{O\in\mathcal{TO}_Q}[O]\varphi$. Therefore, the restriction of the modalities to total orders is not essential since a partial order can be replaced by the set of total orders compatible with it. The rationale behind the definition of $[Q]\varphi$ is based on the consideration of a partial order as a partial description of a total order[3]. In other words, two agents are not comparable in the partial order simply because it is not known if one is more reliable than the other.

2.2 Semantics

The semantics for DBF_n^c is based on the possible world semantics for multi-agent epistemic logic[Fagin et al. 1996].

Definition 3. A DBF_n^c model for the language $\mathcal{L}_c$ is a structure

$$(W, (\mathcal{R}_i)_{1\le i\le n}, V),$$

where

—W is a set of possible worlds,

—$\mathcal{R}_i \subseteq W \times W$ is a serial binary relation over W for $1 \le i \le n$,

—$V : \Phi_0 \to 2^W$ is a truth assignment mapping each atomic proposition to the set of worlds in which it is true.

[3]Note that the definition corresponds to skeptical reasoning in inheritance systems[Brewka 1991]. If we adopt credulous reasoning, i.e., define $[Q]\varphi$ as $\vee_{O\in\mathcal{TO}_Q}[O]\varphi$, then it is possible that for some φ, both $[Q]\varphi$ and $[Q]\neg\varphi$ are true.

From the binary relations $\mathcal{R}_i$'s, we define two sets of derived relations. First, for each nonempty $G \subseteq Ag$, we define

$$\mathcal{R}_G = \cap_{i \in G} \mathcal{R}_i.$$

Second, for each $O \in \mathcal{TO}_n$, the relation $\mathcal{R}^c_O$ is defined in an inductive way as:

$$\mathcal{R}^c_{O>i}(w) = \begin{cases} \mathcal{R}^c_O(w) & \text{if } \bigcap_{j \in \delta(O>i)} \mathcal{R}_j(w) = \emptyset, \\ \mathcal{R}^c_O(w) \cap \mathcal{R}_i(w) & \text{otherwise,} \end{cases}$$

for any $w \in W$. The superscript c denotes level cutting fusion and can usually be omitted when the context is clear.

Let $O = i_1 > i_2 > \cdots > i_m$ and define $G_j = \{i_1, i_2 \ldots, i_j\}$ for $1 \leq j \leq m$ and assume k is the largest j such that $\bigcap_{i \in G_j} \mathcal{R}_i(w) \neq \emptyset$, then we have

$$\mathcal{R}_O(w) = \bigcap_{i \in G_k} \mathcal{R}_i(w).$$

In other words, beliefs from agents after level k are completely discarded from the merged result. Our rationale behind this is that if belief in level $k+1$ is not acceptable, then any belief in a less reliable level is also not acceptable. Therefore, k plays the role of a dynamic threshold in uncertainty reasoning.

Informally, $\mathcal{R}_i(w)$ is the set of worlds that agent i considers possible under w according to his belief, so $\mathcal{R}_G(w)$ and $\mathcal{R}_O(w)$ are the set of worlds which are considered possible under w, respectively, according to the direct fusion and the ordered fusion of agents' beliefs. The informal intuition is reflected in the definition of the satisfaction relation.

Definition 4. Let $M = (W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$ be a DBF^c_n model and Φ be the set of wffs of $\mathcal{L}_c$, then the satisfaction relation $\models_M \subseteq W \times \Phi$ is defined by the following inductive rules(we use the infix notation for the relation and omit the subscript M for convenience):

(1) $w \models p$ iff $w \in V(p)$ for any $p \in \Phi_0$,
(2) $w \not\models \bot$ and $w \models \top$,
(3) $w \models \neg\varphi$ iff $w \not\models \varphi$,
(4) $w \models \varphi \vee \psi$ iff $w \models \varphi$ or $w \models \psi$,
(5) $w \models [G]\varphi$ iff for all $u \in \mathcal{R}_G(w)$, $u \models \varphi$,
(6) $w \models [O]\varphi$ iff for all $u \in \mathcal{R}_O(w)$, $u \models \varphi$.

A set of wffs Σ is satisfied in a world w, written as $w \models \Sigma$ if $w \models \varphi$ for all $\varphi \in \Sigma$. We write $\Sigma \models_M \varphi$ if for each possible world w in M, $w \models \Sigma$ implies $w \models \varphi$ and $\Sigma \models_{\mathrm{DBF}^c_n} \varphi$ if $\Sigma \models_M \varphi$ for each DBF^c_n model M. Σ can be omitted when it is an empty set, so a wff φ is valid in M, denoted by $\models_M \varphi$, if $\emptyset \models_M \varphi$ and $\models_{\mathrm{DBF}^c_n} \varphi$ denotes $\emptyset \models_{\mathrm{DBF}^c_n} \varphi$. The subscript is also usually omitted if there is no possible confusion.

2.3 Axiomatic system

In [Fagin et al. 1996], some variants of epistemic logic systems are presented. Using the naming convention in [Chellas 1980], the most basic system with distributed

(1) Axioms:
- . P: all tautologies of the propositional calculus
- . G1: $([G]\varphi \wedge [G](\varphi \supset \psi)) \supset [G]\psi$
- . G2: $\neg[\{i\}]\bot$
- . G3: $[G_1]\varphi \supset [G_2]\varphi$ if $G_1 \subset G_2$
- . O1: $\neg[\delta(O > i)]\bot \supset ([O > i]\varphi \equiv [\delta(O > i)]\varphi)$
- . O2: $[\delta(O > i)]\bot \supset ([O > i]\varphi \equiv [O]\varphi)$

(2) Rules of Inference:
- . R1(Modus ponens, MP):

$$\frac{\varphi \quad \varphi \supset \psi}{\psi}$$

- . R2(Generalization, Gen):

$$\frac{\varphi}{[G]\varphi}$$

Fig. 1. The axiomatic system for DBF_n^c

belief is called K_n^D, with n being the number of agents and D denoting the distributed belief operators. In this system, no properties except logical omniscience are imposed on the agents' beliefs. Nevertheless, we assume that the belief of each individual agent is consistent even though the collective beliefs of several agents may not be consistent. Our axiomatic system, then, is based on KD_n^D where the additional axiom D is added to K_n^D to ensure the consistency of each agent's belief.[4] The axiomatic system for DBF_n^c is presented in figure 1.

The axioms G1-G3 and rule R2 are those for KD_n^D. G1 and rule R2 are properties of knowledge for perfect reasoners. They are also the causes of the notorious logical omniscience problem. However, it is appropriate to describe implicit information in this way. G2 is the requirement that the belief of each individual agent is consistent. G3 is a characteristic property of distributed knowledge. The larger the subgroup, the more knowledge it possesses. In [Fagin et al. 1996], another axiom relating distributed knowledge and individual ones is presented. That is,

$$D_{\{i\}}\varphi \equiv B_i\varphi,$$

where $D_{\{i\}}\varphi$ and $B_i\varphi$ correspond, respectively, to wffs $[\{i\}]\varphi$ and $[i]\varphi$ in our logic. However, we do not need this axiom because $[i]\varphi \equiv [\{i\}]\varphi$ can be derived from G2, O1, and MP. We can therefore write $[i]$ for both modalities $[i]$ and $[\{i\}]$. The two axioms O1 and O2 recursively define the merged belief in terms of distributed belief. O1 is the case when $\bigcap_{j\in\delta(O>i)} \mathcal{R}_i(w) \neq \emptyset$, whereas O2 is the opposite case.

The derivability in the system is defined as follows: A wff φ is derivable from the system DBF_n^c, or put simply, φ is a *theorem* of DBF_n^c, if there is a finite sequence

[4]Though it is well accepted that $\mathrm{KD45}_n^D$ is more appropriate for modelling of belief with positive and negative introspection (axioms 4 and 5), we adopt the KD_n^D system for emphasizing that the agents may represent databases and their beliefs may be just the facts stored in the databases.

$\varphi_1, \ldots, \varphi_m$ such that $\varphi = \varphi_m$ and every φ_i is an instance of an axiom schema or obtained from earlier φ_j's by the application of an inference rule. It is written as $\vdash_{\mathrm{DBF}_n^c} \varphi$ if φ is a theorem of DBF_n^c. Let $\Sigma \cup \{\varphi\}$ be a subset of wffs, then φ is derivable from Σ in the system DBF_n^c , written as $\Sigma \vdash_{\mathrm{DBF}_n^c} \varphi$, if there is a finite subset Σ' of Σ such that $\vdash_{\mathrm{DBF}_n^c} \bigwedge \Sigma' \supset \varphi$. We will drop the subscript when no confusion occurs.

Some basic theorems can be derived from the system DBF_n^c.

PROPOSITION 1. *For any* $O = i_1 > i_2 > \cdots > i_m$ *and* $G_j = \{i_1, i_2 \ldots, i_j\}(1 \leq j \leq m)$, *we have:*

(1) $\vdash (\neg[G_j]\bot \wedge [G_{j+1}]\bot) \supset ([O]\varphi \equiv [G_j]\varphi)$, *where the wff* $[G_{j+1}]\bot$ *is deleted from the antecedent when* $j = m$.

(2) $\vdash ([O]\varphi \wedge [O](\varphi \supset \psi)) \supset [O]\psi$,

(3) $\vdash \neg[O]\bot$,

(4) $\dfrac{\varphi}{[O]\varphi}$.

PROOF. The proof of all propositions and theorems can be found in the appendix. □

Proposition 1.1 shows that any total order can be separated into a head and a tail according to some consistency level, and the merged belief according to the ordering is just the distributed belief of the agents from the head part. Proposition 1.2 and 1.4 show that merged belief inherits the properties of the distributed belief since the former is equivalent to the latter for the head part of the ordering. Furthermore, Proposition 1.3 shows that belief fusion keeps consistency.

We have the soundness and completeness results for the system DBF_n^c.

THEOREM 2.1. *For any wff of* DBF_n^c, $\models \varphi$ *iff* $\vdash \varphi$.

2.4 Applications and Examples

The level cutting strategy is especially useful in situations where levels are of crucial importance to reasoning. In particular, if we would like to accept only statements above some threshold level and the threshold is determined by the consistency of statement, then a level cutting strategy can be used. In this subsection, we show several examples for the applications of DBF_n^c in such situations.

2.4.1 *Inconsistency handling in possibilistic logic.* In the first example, we consider knowledge bases consisting of possibilistic logic (PL) formulas. Each formula of the knowledge base is attached with a certainty degree. When the knowledge base is partially inconsistency, we determine its consistency level using the PL mechanism. A deduction from the knowledge base is nontrivial if it uses only formulas above the consistency level. We show that the level cutting strategy can formally model this kind of nontrivial deduction in PL.

PL was originally proposed by Dubois and Prade for uncertainty reasoning[Dubois et al. 1991; 1994; Dubois and Prade 1988]. The semantic basis of PL is the possibility theory developed by Zadeh from fuzzy set theory[Zadeh 1978]. Given a universe W, a *possibility distribution* on W is a function $\pi : W \to [0, 1]$. Obviously, π is a characteristic function of a fuzzy subset of W. Two measures on W can be derived from π. They are called possibility and necessity measures and denoted by Π and N, respectively. Formally, $\Pi, N : 2^W \to [0,1]$ are defined as

$$\Pi(A) = \sup_{w\in A} \pi(w),$$

$$N(A) = 1 - \Pi(\overline{A}),$$

where $\overline{A}$ is the complement of A with respect to W.

A fragment for necessity-valued formulas in PL, called PL1, is introduced in [Dubois et al. 1994]. Let $\mathcal{L}(\Phi_0)$ denote the classical propositional language formed from the set of atomic propositions Φ_0, then each wff of PL1 is of the form (φ, α), where $\varphi \in \mathcal{L}(\Phi_0)$ and $\alpha \in (0,1]$ is a real number. The number α is called the *valuation* or *weight* of the formula. The notation (φ, α) expresses that φ is certain at least to degree α. Formally, a model for PL1 is given by a possibility distribution π on the set W of classical truth assignments for $\mathcal{L}(\Phi_0)$. For any $\varphi \in \mathcal{L}(\Phi_0)$, we can define the *truth set* of φ as the set of truth assignments satisfying φ. By identifying φ and its truth set, a PL1 model π satisfies (φ, α), denoted by $\pi \models (\varphi, \alpha)$, if $N(\varphi) \geq \alpha$. Let $S = \{(\varphi_i, \alpha_i) : 1 \leq i \leq m\}$ be a finite set of PL1 wffs, then $S \models_{\mathrm{PL1}} (\varphi, \alpha)$ if for each π, $\pi \models (\varphi_i, \alpha_i)$ for all $1 \leq i \leq m$ implies $\pi \models (\varphi, \alpha)$. It is shown that the consequence relation in PL1 can be determined completely by the least specific model satisfying S. That is, if $\pi_S : W \to [0,1]$ is defined by

$$\pi_S(w) = \min\{1 - \alpha_i \mid w \models \neg\varphi, 1 \leq i \leq m\},$$

where $\min \emptyset = 1$, then $S \models_{\mathrm{PL1}} (\varphi, \alpha)$ iff $\pi_S \models (\varphi, \alpha)$.

A special feature of PL1 is its capability to cope with partial inconsistency. For S defined as above, let S^* denote the set of classical formulas $\{\varphi \mid 1 \leq i \leq m\}$. The set S is then said to be partially inconsistent when S^* is classically inconsistent. It can be easily shown that S is partially inconsistent iff $\sup_{w\in W} \pi_S(w) < 1$. Thus $\sup_{w\in W} \pi_S(w)$ is called the consistency degree of S, denoted by $Cons(S)$, and $1 - Cons(S)$ is called the inconsistency degree of S, denoted by $Incons(S)$. When S is partially inconsistent, it can be shown that $S \models_{\mathrm{PL1}} (\bot, Incons(S))$, so for any classical wff φ, $(\varphi, Incons(S))$ is a trivial logical consequence of S. On the contrary, if $S \models_{\mathrm{PL1}} (\varphi, \alpha)$ for some $\alpha > Incons(S)$, then φ is called a *nontrivial consequence* of S.

To model the nontrivial deduction of PL1, we assume that the weights of PL1 wffs in a set S are drawn from a finite subset $\mathcal{V} = \{\alpha_1, \ldots, \alpha_n\}$ of $(0,1]$. Without loss of generality, we can assume $\alpha_1 > \cdots > \alpha_n$. For $1 \leq i \leq n$, let us define S_i as:

$$S_i = \{\varphi \mid (\varphi, \alpha_i) \in S\}.$$

Assume that G is a subset of $\{1, 2, \ldots, n\}$, then G is consistent for S if $\bigcup_{i\in G} S_i$ is classically consistent, otherwise, it is inconsistent for S. A subset G is a *maximal consistent agent group* for S if G is consistent for S and for any $i \notin G$, $G \cup \{i\}$ is inconsistent for S. Let $MCAG_S$ denote the class of all maximal consistent agent groups for S, then we can associate with S a set Σ_S of wffs in $\mathcal{L}_c$ as:

$$\Sigma_S = \bigcup_{i=1}^{n} \{[i]\varphi \mid \varphi \in S_i\} \cup \{\neg[G]\bot \mid G \in MCAG_S\}.$$

The nontrivial deduction can then modelled in the logic DBF_n^c in the following way.

PROPOSITION 2. *Let S be a set of PL1 wff and each S_i be classically consistent, then for any $\varphi \in \mathcal{L}(\Phi_0)$, φ is a nontrivial consequence of S iff* $\Sigma_S \vdash_{\mathrm{DBF}_n^c} [1 > 2 > \cdots > n]\varphi$.

We illustrate this application with an example from [Dubois and Prade 1991].

Example 1. [Dubois and Prade 1991] Let us consider a knowledge base S consisting of the following PL1 facts:

. if Léa is a student, then it is rather certain that she is young,

. if Léa is young, then it is fairly certain that she is single,

. if Léa is a student and is a mother, then it is almost certain that she is not single,

. Léa is a student.

Let p_1="Léa is a student", p_2="Léa is young", p_3="Léa is single", and p_4="Léa is a mother", then our premise Σ_S is the following set:

$$\{[1]p_1, [2](p_1 \wedge p_4 \supset \neg p_3), [3](p_2 \supset p_3), [4](p_1 \supset p_2), \neg[\{1,2,3,4\}]\bot\}.$$

Therefore, we can derive

$$\Sigma_S \vdash_{\mathrm{DBF}_4^c} [1 > 2 > 3 > 4](p_2 \wedge p_3).$$

However, if we add to S another fact "Léa is a mother", then Σ_S becomes the set

$$\{[1]p_1, [1]p_4, [2](p_1 \wedge p_4 \supset \neg p_3), [3](p_2 \supset p_3), [4](p_1 \supset p_2),$$
$$\neg[\{1,2,3\}]\bot, \neg[\{1,2,4\}]\bot, \neg[\{1,3,4\}]\bot, \neg[\{2,3,4\}]\bot\},$$

so we can now derive

$$\vdash_{\mathrm{DBF}_4^c} \psi_\Sigma \supset [1 > 2 > 3 > 4](\neg p_2 \wedge \neg p_3).$$

According to [Dubois and Prade 1991](p. 233), this means that:

> *In the terminology of epistemic entrenchment, less entrenched pieces of information are inhibited, only a consistent subbase of strongly entrenched pieces of information remains undisputable.*

□

2.4.2 *Inductive Acceptance.* Scientific induction is the process of generating, evaluating, and accepting hypotheses based on evidence. In [Hempel 1965], three phases in the scientific test of a given hypothesis are described. The first phase is to conduct experiments or observations and obtain the observation reports. The second phase consists in ascertaining whether the observation reports confirm, disconfirm, or are irrelevant to the hypothesis. The final phase consists of either acceptance or rejection of the hypothesis based on the strength of the confirming or disconfirming evidence. In particular, it is pointed that([Hempel 1965], p. 41):

> *The third phase ··· would require the establishment of general "rule of acceptance". Roughly speaking, these rules would state how well a given hypothesis has to be confirmed by the accepted observation reports to be scientifically acceptable itself; i.e., the rules would formulate criteria for the acceptance or rejection of a hypothesis by reference to the kind and amount of confirming or disconfirming evidence for it embodied in the totality of accepted observation reports.*

A direct rule of acceptance is to give a threshold for the strength of confirmation and any hypothesis with a degree of confirmation more than the threshold can be accepted[Kyburg 1961]. However, the rule cannot ensure the set of accepted sentences is consistent due to the notorious "lottery paradox"[Kyburg 1961]. The main aspect of lottery paradox is "a set of statements, each of which is as probable as you please, that are jointly inconsistent"[Kyburg 1961]. While there have been a lot of works in the discussion of this aspect, we emphasize another aspect of the lottery paradox—the acceptance of statements is a kind of threshold reasoning. In other words, a statement under the level of acceptance is rejected even if it is consistent with those at higher levels. This kind of reasoning is, exactly, deduction based on level cutting strategy.

Example 2. Suppose that there are a large number (say N) of people buying the lottery tickets; then according to experience, the following statements are confirmed with different strengths:

. There are some people who will win.
. A particular people will lose.
. A particular people will win.

Let p_i denotes "the i-th people will win", then the degrees of confirmation of the three statements above are listed in table I.

Table I. Degrees of confirmation for statements in the lottery paradox

statement	degree of confirmation
$\bigvee_{i=1}^{N} p_i$	very high
$\neg p_i (1 \leq i \leq N)$	fairly high
$p_i (1 \leq i \leq N)$	very low

According to this table, we can construct a set of wffs Σ as follows:

$$\{[1] \bigvee_{i=1}^{N} p_i\} \cup \{[2]\neg p_i \mid 1 \leq i \leq N\} \cup \{[3]p_i \mid 1 \leq i \leq N\} \cup \{\neg[\{1,3\}]\bot\}$$

Note that $\neg[\{1,3\}]\bot$ is added into Σ because the beliefs of agent 1 and 3 (i.e., $\bigvee_{i=1}^{N} p_i$ and $p_i (1 \leq i \leq N)$) are consistent. Consequently, we have $\Sigma \vdash [1 > 2 > 3] \bigvee_{i=1}^{N} p_i$; we do not however have $\Sigma \vdash [1 > 2 > 3]p_i$ for any p_i, though $\Sigma \vdash \neg[\{1,3\}]\bot$. This is due to the fact that levels play the role of thresholds in the reasoning. When inconsistency occurs at level 2, the threshold of acceptance is set to level 1, so even though the statements at level 3 are completely consistent with those at level 1, they cannot be accepted. Indeed, the result is intuitively reasonable. Though we accept the fact that someone will win, it is very unlikely that all persons will win. □

2.4.3 *Information security in multilevel systems.* Information security is concerned with whether the disclosure of information is authorized, whereas the objective of multilevel information security is to protect information that is classified with respect to a multilevel hierarchy. Two main tasks are essential for the information security in multilevel systems: authorization and authentication. With authorization, an agent can obtain permission to access information, and by authentication, the agent safe-guarding the information can decide whether an agent requesting access right has permission. In a distributed computing environment, the authorization center is, in general, independent of the information center, so it is necessary for the information center to decide whether the request of an agent can be granted. In the next example, we show that level cutting fusion can help reduce the probability of security breach.

Example 3. In the following simplification of a realistic scenario, we assume the existence of n agents with different levels, an authorization center, and an information center in which multilevel secrecy is resident. Authorization is achieved by the distribution of a key to the agents who are permitted to access the information. As the pieces of information change with time, the authorization center may decide to change and redistribute the key. Depending on the classification levels of the information at the time when the key is distributed, the authorization center will dynamically decide to which level of agents information access is permitted. However, since the authorization center is independent of the information center, the latter does not know the decision of the former, so all the information center can do is to decide whether grant the access request of an agent by checking his/her key. The scenario is illustrated by figure 2.

Let q_{key} denote "the key is **key**", where **key** is the correct key value, and p_i denote the access request of the agent i. Assume that the levels of the n agent is ordered by $1 > 2 > \cdots > n$, and at some time, the authorization center distributes the correct key value to level l, then normally, the information center should receive the following set of wffs for an access request:

$$\Sigma = \{[i](q_{key} \wedge p_i) \mid 1 \leq i \leq l\} \cup \{[i]\neg q_{key} \mid l+1 \leq i \leq n\}.$$

The security breach may occur when some agent from level $l+1$ to n is malicious and obtains accidently the key. On the existence of malicious unauthorized agents, the set Σ will become

$$\Sigma = \{[i](q_{key} \wedge p_i) \mid 1 \leq i \leq l\} \cup \{[i]\neg q_{key} \mid i \in G^+\} \cup \{[i](q_{key} \wedge p_i) \mid i \in G^-\},$$

where G^- is the set of agents that are malicious and obtains the key accidently and $G^+ = \{l+1, \cdots, n\} - G^-$. Since the form of wffs received by the information center is rather limited, we can assume that the information center can perform the following closed-world reasoning:

$$\text{if} \quad \not\vdash [G]\bot, \quad \text{then add} \quad \vdash \neg[G]\bot;$$

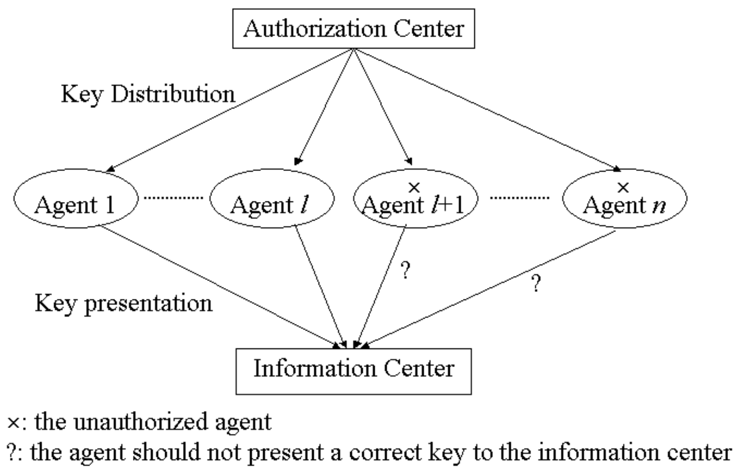

Fig. 2. A scenario for information security

so we do not bother to add the kind of wffs $\neg[G]\bot$ into the knowledge base for any subset G of agents.

Let us now consider what the information center can do when there is a possibility of security breach. On one hand, it can process the access request of each agent independently. That is, it will grant the access request of an agent i if $\Sigma \vdash [i](q_{key} \wedge p_i)$. If the probability that an unauthorized agent is malicious and obtains the key accidently is α, then the probability of security breach is

$$1 - (1 - \alpha)^{n-l}.$$

On the other hand, using level cutting fusion, the information center can do better. In other words, it will grant the access request of an agent i only if $\Sigma \vdash_{\mathrm{DBF}_n^c} [1 > 2 > \cdots > n]p_i$. Since $\Sigma \vdash_{\mathrm{DBF}_n^c} [1 > 2 > \cdots > n]p_i$ for some $i > l$ only when $l + 1 \in G^-$, the probability of security breach reduces from $1 - (1 - \alpha)^{n-l}$ to α. The point is that the agent at level $l + 1$ will help block the security intrusion of other agents if it is a verity agent. □

3. LOGIC FOR LEVEL SKIPPING FUSION

Though the level cutting strategy is useful in practice, it is sometimes too cautious from the perspective of information fusion. A less cautious strategy is to skip only the agent causing inconsistency and continue to consider the next level. This strategy corresponds to the suspicious attitude of multi-source reasoning and has been used in belief revision by Nebel[Nebel 1994]. To model the belief fusion using this strategy, we introduce the logic DBF_n^s for level skipping distributed belief fusion in this section.

3.1 Syntax

Let $\mathcal{L}_s$ denote the language of DBF_n^s with syntax formally defined as:

Definition 5. The alphabet of $\mathcal{L}_s$ is same as that of $\mathcal{L}_c$, and the wffs of $\mathcal{L}_s$ are defined by rules 1-4 for that of $\mathcal{L}_c$ and the following fifth rule.

. 5. if φ is a wff, so are $[\Omega]\varphi$ for any nonempty $\Omega \subseteq \mathcal{TO}_n$.

When Ω is a singleton $\{O\}$, we write $[O]\varphi$ instead of $[\{O\}]\varphi$. If $\Omega = \{O_1, \ldots, O_m\}$ is such that $|\delta(O_i)| = 1$ for all i's, then $[\Omega]$ is the distributed belief operator among ordinary agents. Therefore, the language is more general than that of DBF_n^c.

3.2 Semantics

For the semantics, a DBF_n^s model is same as a DBF_n^c model, so given a model $(W, (\mathcal{R}_i)_{1 \le i \le n}, V)$, we can define $\mathcal{R}_O^s$ inductively.

$$\mathcal{R}_{O>i}^s(w) = \begin{cases} \mathcal{R}_O^s(w) & \text{if } \mathcal{R}_O^s(w) \cap \mathcal{R}_i(w) = \emptyset, \\ \mathcal{R}_O^s(w) \cap \mathcal{R}_i(w) & \text{otherwise,} \end{cases}$$

for any $w \in W$. As in the case of $\mathcal{R}_O^c$, the superscript s denotes the level skipping strategy and can be omitted when the context is clear. We can also define

$$\mathcal{R}_\Omega = \bigcap_{O \in \Omega} \mathcal{R}_O.$$

Then, the satisfaction relation between the possible worlds and wffs of $\mathcal{L}_s$ must satisfy clauses 1-4 of definition 4 and the following fifth clause.

. 5. $w \models [\Omega]\varphi$ iff for all $u \in \mathcal{R}_\Omega(w), u \models \varphi$.

According to the definition, $[O > i]\varphi$ will be equivalent to the distributed fusion of $[O]\varphi$ and $[i]\varphi$ when the belief of i is consistent with the merged belief of O. To axiomatize reasoning under the strategy, we can view O as a virtual agent and consider the distributed belief between O and i. However, to get a bit more general, we also consider the distributed belief among a group of arbitrary virtual agents. This is why we include wffs of the form $[\Omega]\varphi$ in $\mathcal{L}_s$.

3.3 Axiomatic system

Given the language and semantics, the valid wffs of DBF_n^s are captured by the axiomatic system in figure 3.

The axioms V1-V3 and rule R2$'$ correspond to G1-G3 and R2 for distributed belief for virtual agents instead of ordinary agents. Nevertheless, since an ordinary agent is a special case of the virtual agent, these axioms in fact also cover G1-G3 and R2. O1$'$ and O2$'$ are axioms for describing the level skipping strategy and correspond exactly to the inductive definition of $\mathcal{R}_{O>i}$, where Ω in these two axioms denote any subset (empty or not) of $\mathcal{TO}_n$. We can still have the soundness and completeness theorem.

THEOREM 3.1. *For any wff of* DBF_n^s, $\models \varphi$ *iff* $\vdash \varphi$.

Since operator $[O]$ is a special case of $[\Omega]$, propositions 1.2 and 1.4 hold trivially for DBF_n^s. The proposition 1.3 can be easily proved using V2, O1$'$ and O2$'$. However, it is unclear whether a counterpart of proposition 1.1 can be given.

(1) Axioms:
- . P: all tautologies of the propositional calculus
- . V1: $([\Omega]\varphi \wedge [\Omega](\varphi \supset \psi)) \supset [\Omega]\psi$
- . V2: $\neg[i]\bot$
- . V3: $[\Omega_1]\varphi \supset [\Omega_2]\varphi$ if $\Omega_1 \subset \Omega_2$
- . O1′: $\neg[\{O, i\}]\bot \supset ([\Omega \cup \{O > i\}]\varphi \equiv [\Omega \cup \{O, i\}]\varphi)$
- . O2′: $[\{O, i\}]\bot \supset ([\Omega \cup \{O > i\}]\varphi \equiv [\Omega \cup \{O\}]\varphi)$

(2) Rules of Inference:
- . R1(Modus ponens, MP):

$$\frac{\varphi \quad \varphi \supset \psi}{\psi}$$

- . R2′(Generalization, Gen):

$$\frac{\varphi}{[\Omega]\varphi}$$

Fig. 3. The axiomatic system for DBF_n^s

3.4 Applications and Examples

In this section, we discuss some applications of the logic DBF_n^s.

3.4.1 *Multi-agent epistemic reasoning.* In [Fagin et al. 1996], it is argued that multi-agent epistemic logic has many applications in such diverse fields as economics, linguistics, AI, and theoretical computer science. However, the distributed knowledge operator in multi-agent epistemic logic suffers from an inconsistency problem. That is, the distributed belief of a group of agents with conflicting beliefs may crash. The fusion operators proposed in our logic are used to circumvent this problem in the original epistemic logic framework. Let us now look at some examples of integrated reasoning about the multi-agent beliefs and their fusion.

Example 4. If a set of premises $\Sigma = \{\neg[\{1,2\}]\bot \vee \neg[\{1,3\}]\bot, [1](p \supset q), [2]p, [3]\neg q\}$ is given for three agents, then it can be derived that

$$\Sigma \vdash_{\mathrm{DBF}_3^s} [1 > 2 > 3]((p \wedge q) \vee (\neg p \wedge \neg q))$$

and

$$\Sigma \vdash_{\mathrm{DBF}_3^s} [1 > 2 > 3](p \supset q).$$

The wff $\neg[\{1,2\}]\bot \vee \neg[\{1,3\}]\bot$ says that if the beliefs of agents 1 and 2 are incompatible, then those of 1 and 3 are compatible, so the level skipping strategy will either accept the belief of agent 2 or ignore it and consequently accept that of agent 3. This example shows that we can reason with the compatibility of the agents' beliefs in the uniform framework of epistemic reasoning and information fusion. □

The next example shows that the belief about belief may play a role in the fusion process.

Example 5. Assume there are two agents whose beliefs are described by the following set:

$$\Sigma = \{[1]\neg[\{1,2\}]\bot, [1]p, [1][1]p, [1][2]q \\ [2][\{1,2\}]\bot, [2]q, [2][2]q, [2][1]p\}$$

It can then be shown that $[1>2]p \wedge [2>1]q$, $[1][1>2](p\wedge q) \wedge [1][2>1](p \wedge q)$, and $[2][1>2]p \wedge [2][2>1]q$ are derivable from Σ in both DBF_2^s and DBF_2^c. Thus the belief of agent 1 is incorrect because he wrongly believes that he is consistent with agent 2, while agent 2 in fact disagrees with him on the consistency between them. □

Sometimes, it is possible to infer individual agent beliefs from their merged beliefs. The next example shows a very simple case.

Example 6. Assume it is known that two premises $[1>2]p$ and $[2>1]\neg p$ hold, then we have the following derivation in DBF_2^s:

1.$\neg[\{1,2\}]\bot \supset ([1>2]p \supset [\{1,2\}]p)$	$O1'$
2.$\neg[\{1,2\}]\bot \supset ([2>1]\neg p \supset [\{1,2\}]\neg p)$	$O1'$
3.$([1>2]p \wedge [2>1]\neg p) \supset (\neg[\{1,2\}]\bot \supset ([\{1,2\}]p \wedge [\{1,2\}]\neg p))$	$1, 2, P, MP$
4.$([1>2]p \wedge [2>1]\neg p) \supset (\neg[\{1,2\}]\bot \supset [\{1,2\}]\bot)$	$3, P, G1, MP, Gen$
5.$([1>2]p \wedge [2>1]\neg p) \supset [\{1,2\}]\bot$	$4, P$
6.$[\{1,2\}]\bot \supset ([1>2]p \supset [1]p)$	$O2', MP$
7.$[\{1,2\}]\bot \supset ([2>1]\neg p \supset [2]\neg p)$	$O2', MP$
8.$([1>2]p \wedge [2>1]\neg p) \supset ([1]p \wedge [2]\neg p)$	$6, 7, P, MP$

When there are more than two agents, the situation becomes more complicated. However, it is still possible to derive some individual or partially merged beliefs from the totally merged ones. □

3.4.2 *Preferred subtheory.* In the last two decades, non-monotonic reasoning has become an important research area in artificial intelligence and knowledge representation. Many mechanisms for non-monotonic reasoning have been developed and prove very useful in some practical applications. In this section, we show that the logic for level skipping fusion can simulate a kind of non-monotonic reasoning mechanism proposed in [Brewka 1991].

Definition 6 *[Brewka 1991].*

(1) A level default theory T is a tuple $(T_1, \cdots, T_k)$, where each T_i is a set of classical logic formulas.

(2) Let $T = (T_1, \cdots, T_k)$ be a level default theory, then $S = S_1 \cup \cdots \cup S_k$ is a preferred subtheory of T iff for all $i (1 \leq i \leq k)$, $S_1 \cup \cdots \cup S_i$ is a maximal consistent subset of $T_1 \cup \cdots \cup T_i$.

Note that preferred subtheories of a level default theory is syntax-dependent. For example, if $T_1 = \{p, q\}$ and $T_2 = \{\neg p \vee \neg q\}$, then $\{p, \neg p \vee \neg q\}$ and $\{q, \neg p \vee \neg q\}$ are both preferred subtheories of (T_1, T_2). However, if $T_1' = \{p \wedge q\}$, then neither is the preferred subtheory of (T_1', T_2), even though T_1 and T_1' are semantically equivalent. On the other hand, our logical framework is semantics-based, so it is essentially incompatible with the preferred subtheory approach. However, it turns out that reasoning with a preferred subtheory can be simulated in the level skipping fusion logic if the syntax-dependence is preserved in the translation of a level default theory into our logic language. In the following formulation, we assume that each formula $\varphi \in T_1 \cup \cdots \cup T_k$ is classically consistent.

The simulation steps are as follows.

(1) Let $n_i = |T_i|$ for $1 \leq i \leq k$ and let $n = \sum_{i=1}^{k} n_i$. Also let $n_0 = 0$.

(2) Let φ_{ij} denote the jth formula in T_i for $1 \leq i \leq k$ and $1 \leq j \leq n_i$.

(3) Define a two-place function f such that for $1 \leq i \leq k$ and $1 \leq j \leq n_i$

$$f(i,j) = n_1 + \cdots + n_{i-1} + j.$$

(4) Note that f is 1-1 and onto for the range $\{1, 2, \cdots, n\}$, so we can define two inverse functions g_1 and g_2 by

$$g_1(m) = i; \;\; g_2(m) = j$$

if $f(i,j) = m$.

(5) A subset $G \subseteq \{1, 2, \cdots, n\}$ is a maximal consistent agent group if $\{\varphi_{ij} \mid f(i,j) \in G\}$ is consistent and for any $G' \supset G$, $\{\varphi_{ij} \mid f(i,j) \in G'\}$ is inconsistent. Let $MCAG_n$ denote the set of all maximal consistent agent groups.

(6) Let $\Sigma_T = \{[m]\varphi_{ij} \mid 1 \leq i \leq k, 1 \leq j \leq n_i, f(i,j) = m\} \cup \{\neg[G]\bot \mid G \in MCAG_n\}$.

(7) Define a strict partial ordering Q over $\{1, 2, \cdots, n\}$ by

$$x > y \Leftrightarrow g_1(x) < g_1(y).$$

Note that the translation preserves the syntax structure of the original theory. For example, if $T_1 = \{p, q\}$, then it is translated into two modal formulas $\{[1]p, [2]q\}$; however, if $T_1 = \{p \wedge q\}$, then it is simply translated into $\{[1](p \wedge q)\}$. This results in the following proposition.

Proposition 3.

(*1*) *Let S be a preferred subtheory of T, then there exists a total ordering $O \in \mathcal{TO}_Q$ such that $S \models \varphi$ iff $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [O]\varphi$ for any classical logic formula φ.*

(*2*) *Let $O \in \mathcal{TO}_Q$ be a total ordering over $\{1, 2, \cdots, n\}$, then there exists a preferred subtheory S of T such that $S \models \varphi$ iff $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [O]\varphi$ for any classical logic formula φ.*

Example 7. Let us consider a slightly modified version of the meeting example presented in [Brewka 1991]. In that example, the following level default theory is given:

$$\begin{aligned} T_1 &= \{\textsc{Cold}, \textsc{Vacation} \supset \neg R_1, \textsc{Cold} \supset \neg R_2, \textsc{Cold} \supset \textsc{Sick}\} \\ T_2 &= \{R_2, R_2 \supset (\textsc{Sick} \supset \neg R_1)\} \\ T_3 &= \{R_1, R_1 \supset \textsc{Meeting}\}. \end{aligned}$$

Then $S = T_1 \cup T_3 \cup \{R_2 \supset (\textsc{Sick} \supset \neg R_1)\}$ is a preferred subtheory of (T_1, T_2, T_3) and it can be seen that $S \models \textsc{Meeting}$. By our simulation steps, there are 8 agents and the translated theory Σ_T is

$$\{[1]\textsc{Cold}, [2](\textsc{Vacation} \supset \neg R_1), \cdots, [8](R_1 \supset \textsc{Meeting})\} \cup \{\neg[G]\bot \mid G \in MCAG_8\}$$

Since $\{1,2,3,4,6,7,8\} \in MCAG_8$, it can be shown that $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [1 > \cdots > 8]\varphi \equiv [\{1,2,3,4,6,7,8\}]\varphi$ for all wff φ. using the definition of Σ_T, we have $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [\{7,8\}]$MEETING, so $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [1 > \cdots > 8]$MEETING holds. □

One of the earliest approaches to knowledge merging is to manipulate the maximal consistent subsets of the union of component databases. In [Baral et al. 1992; 1994; Baral et al. 1991], knowledge bases with integrity constraints are combined by a meta-level combination operator to form a new knowledge base. While in [Baral et al. 1994; Baral et al. 1991], logic programs and default logic theories that have different semantics than the classical logic are considered, the basic mechanism for combining first-order theories in [Baral et al. 1992] can be considered a special case of reasoning with a preferred subtheory. In [Baral et al. 1992], a combination operator C maps a set of knowledge bases $\{T_1, \cdots, T_k\}$ and a set of integrity constraints IC into a new knowledge base $C(T_1, \cdots, T_k, IC)$ that can be roughly considered as the disjunction of maximally consistent subsets of $T_1 \cup T_2 \cup \cdots \cup T_k$ with respect to IC. More precisely, let $T = (IC, T_1 \cup \cdots \cup T_k)$ be a level default theory and $\mathcal{PF}_T$ is the set of preferred subtheories of T, then

$$C(T_1, \cdots, T_k, IC) = \bigvee_{S \in \mathcal{PF}_T} \bigwedge S.$$

Therefore, by proposition 3, we have the following result.

COROLLARY 1. *Let $T = (IC, T_1 \cup \cdots \cup T_k)$ be a level default theory and the partial ordering Q be defined as in the simulation step 7, then for all classical logic formula φ,*

$$\Sigma_T \vdash_{\mathrm{DBF}_n^s} [Q]\varphi \Leftrightarrow C(T_1, \cdots, T_k, IC) \models \varphi.$$

3.4.3 *Diagnostic reasoning.* Diagnosis is one of the most important areas in which knowledge-based systems are applied. It is a process of finding faults or defects in the observed data and the background knowledge of the diagnosed system. Two familiar examples of diagnosis are the medical diagnosis of a patient's disease from his symptoms and the fault detection of an electrical device from its abnormal behaviors. A review of diagnostic systems based on symbolic reasoning methodology is presented in [Lucas 1997]. There are several formal theories of diagnosis. One of the popular diagnostic reasoning methods is the consistency-based diagnosis proposed in [Reiter 1987]. According to [Lucas 1997], this is a kind of diagnosis based on the knowledge of deviation from normal structure and behavior(DNSB diagnosis). The logical specification of such knowledge is a triple $\mathcal{S} = (SD, COMPS, OB)$, where

(1) SD is a finite set of classical logic formulas, called the system description, which specifies the normal structure and behavior of the system,
(2) $COMPS$ is a set of individual constants, denoting the components of the system, and
(3) OB is another finite set of classical logic formulas, denoting observations.

Each formula in SD is typically in the form:

$$\neg ab(c) \supset \varphi(c)$$

which means that if a component c is not abnormal (i.e., not faulty), then it should satisfy some behavior specification φ. For each $D \subseteq COMPS$, we can form a *hypothesis*

$$H_D = \{ab(c) \mid c \in D\} \cup \{\neg ab(c) \mid c \in COMPS \backslash D\}.$$

The hypothesis indicates that the components in D are faulty and those not in D are normal. A set $D \subseteq COMPS$ is called a diagnosis of the system $\mathcal{S}$ if $SD \cup H_D \cup OB$ is classically consistent. In Reiter's original formulation[Reiter 1987], the problem of diagnostic reasoning is to find diagnoses minimal with respect to set inclusion.

Though minimality is an important criterion in diagnostic reasoning, it is not unique. Sometimes, the reliability of different components may provide clues in the diagnostic process. In particular, when several diagnoses exist, we may have to do more tests to find the real fault, and the part of the system that will be tested first strongly depends on the reliability of the components. In general, it is reasonable to first test the less reliable components. Furthermore, when a system is complex, it is also helpful to hierarchically partition the system into subsystems and consider each subsystem as a component in the diagnostic reasoning process.

To model such diagnostic reasoning in DBF_n^s logic, we associate agent 1 with the system description SD and the observation OB and each component in $COMPS$ or each subsystem to an agent in $\{2, \cdots, n\}$. Therefore, the system $\mathcal{S}$ can be formulated in DBF_n^s by a set $\Sigma_\mathcal{S}$ including the following wffs:

(1) $[1]\varphi$ if φ is a system description or an observed fact,
(2) $[i]\neg ab(c)$ if agent i is associated with an atomic component c,
(3) $[i](\neg ab(c_1) \wedge \cdots \wedge \neg ab(c_k))$ if i is associated with a subsystem consisting of the components $c_1, \cdots, c_k$, and
(4) $\neg[G]\bot$ if G is a maximal consistent agent group, i.e., $[G]\bot$ is not derivable from the preceding three kinds of wffs.

Let $\sigma : \{2, \cdots, n\} \to \{2, \cdots, n\}$ be a bijective function and O denote the total ordering $1 > \sigma(2) > \cdots > \sigma(n)$, then the set

$$D = \{c \in COMPS \mid \Sigma_\mathcal{S} \not\vdash_{\text{DBF}_n^s} [O]\neg ab(c)\}$$

constructs a diagnosis of the system $\mathcal{S}$. We will say that the ordering O respects the reliability of the components if for any c and $c' \in COMPS$ such that c is more reliable than c', the agent associated with c precedes that associated with c' in O. If we have enough reliability information of the components, the number of possible diagnoses will reduce when only orderings respecting the reliability of the components are considered for the diagnostic reasoning.

Example 8 *Flat diagnosis*. Let us consider the diagnosis of a digital circuit shown in figure 4. This circuit has been extensively used in the literature[Reiter 1987]. The atomic components of the circuit consist of two "exclusive-or" gates X_1 and X_2, two "and" gates A_1 and A_2, and one "or" gate R_1. The system description SD contains the following classical logic formulas[5]:

$$\neg ab(X_1) \supset out(X_1) \equiv (in_1(X_1) \oplus in_2(X_1))$$

[5]The formula $p \oplus q$ denotes the abbreviation of $(p \wedge \neg q) \vee (\neg p \wedge q)$.

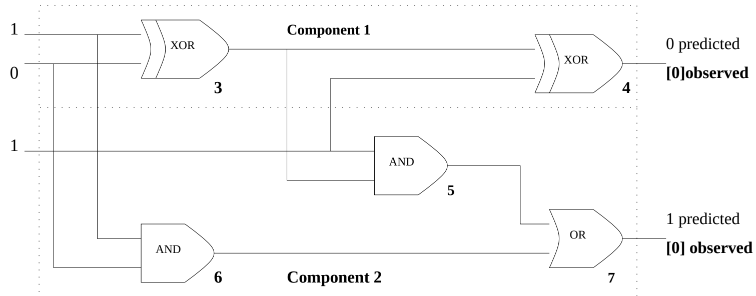

Fig. 4. A full adder

$$\neg ab(X_2) \supset out(X_2) \equiv (in(A_2) \oplus out(X_1))$$

$$\neg ab(A_1) \supset out(A_1) \equiv (in_1(X_1) \wedge in_2(X_1))$$

$$\neg ab(A_2) \supset out(A_2) \equiv (in(A_2) \wedge out(X_1))$$

$$\neg ab(R_1) \supset out(R_1) \equiv (out(A_1) \vee out(A_2))$$

and the observation OB contains

$$in_1(X_1), \neg in_2(X_1), in(A_2), \neg out(X_2), \neg out(R_1).$$

Let $\mathcal{S}$ be the system $(SD, COMPS, OB)$, then we have six agents, where agent 2, 3, 4, 5, 6 corresponds to component X_1, X_2, A_1, A_2, and R_1, respectively. Therefore, the DBF$_n^s$ theory for this system is

$$\Sigma_{\mathcal{S}} = \{[1]\varphi \mid \varphi \in SD \cup OB\} \cup \{[2]\neg ab(X_1), [3]\neg ab(X_2), [4]\neg ab(A_1), [5]\neg ab(A_2), [6]\neg ab(R_1)\} \cup \Sigma_0$$

where $\Sigma_0 = \{\neg[\{1,2,3,4,5\}]\bot, \neg[\{1,2,3,4,6\}]\bot, \neg[\{1,4,5,6\}]\bot, \neg[\{2,3,4,5,6\}]\bot\}$. The set Σ_0 arises from maximal groups of agents whose joint beliefs are consistent. For example, $\neg[\{1,2,3,4,5\}]\bot$ is in Σ_0 because the system description and observed facts are (classically) consistent with the set $\{\neg ab(X_1), \neg ab(X_2), \neg ab(A_1), \neg ab(A_2)\}$. In a practical implementation, we do not need to generate the set Σ_0 for the sake of efficiency. Instead, we adopt a kind of lazy evaluation procedure to test if $[G]\bot$ is derivable from $\Sigma_{\mathcal{S}} \backslash \Sigma_0$ for a group of agents G.

Suppose that the "exclusive-or" gates are in general more reliable than other components, then by considering the ordering $O_1 = 1 > 2 > 3 > 4 > 5 > 6$ and $O_2 = 1 > 2 > 3 > 4 > 6 > 5$, we have

$$\Sigma_{\mathcal{S}} \vdash_{\mathrm{DBF}_n^s} [O_1](\neg ab(X_1) \wedge \neg ab(X_2) \wedge \neg ab(A_1) \wedge \neg ab(A_2))$$

and

$$\Sigma_{\mathcal{S}} \vdash_{\mathrm{DBF}_n^s} [O_2](\neg ab(X_1) \wedge \neg ab(X_2) \wedge \neg ab(A_1) \wedge \neg ab(R_1)).$$

This corresponds to two diagnoses $\{R_1\}$ and $\{A_2\}$ and excludes the diagnosis $\{X_1, X_2\}$ that is possible in Reiter's original formulation. □

Example 9 *Hierarchical diagnosis*. Though the circuit in figure 4 is not too complex, we can still use it to illustrate the notion of hierarchical diagnosis. The dotted lines in figure 4 divide the system into two subsystems C_1 and C_2. Each system is now considered as a compound component, so we can consider three agents with the following set of beliefs:

$$\Sigma_{\mathcal{S}} = \{[1]\varphi \mid \varphi \in SD \cup OB\} \cup \{[2](\neg ab(X_1) \wedge \neg ab(X_2)),$$
$$[3](\neg ab(A_1) \wedge \neg ab(A_2) \wedge \neg ab(R_1))\} \cup \{\neg[\{1,2\}]\bot, \neg[\{1,3\}]\bot, \neg[\{2,3\}]\bot\}.$$

If we assume that the subsystem C_1 is more reliable than C_2, then we have

$$\Sigma_{\mathcal{S}} \vdash_{\mathrm{DBF}_n^s} [1 > 2 > 3](\neg ab(X_1) \wedge \neg ab(X_2)).$$

This corresponds to the diagnosis $\{C_2\}$. Once the diagnosis is given, we can generate further test data directly for subsystem C_2; for example, the second input of A_2 is connected to an external signal source. By observing the output of the test data, we can again carry out the same diagnostic reasoning for the subsystem. □

4. PROOF THEORY

While we have presented semantics and axiomatic systems for modal logics of multi-agent belief fusion, we would like to address some computational issues of the proposed logics. Though axiomatic systems provide an elegant characterization of our logic, it is less practical from a computational viewpoint. It is in general hard to find a proof for a given formula in an automatic way because with the use of the modus ponens rule, we have to look for a proof of ψ and $\psi \supset \varphi$ for the proof of φ. However, since ψ may be an arbitrary formula without any relation to the target formula φ, the search has to be exhaustive.

Other calculi, such as resolution, Gentzen's sequent calculus, and tableau methods, have been proposed for the purpose of automated theorem proving. The subformula property of these methods makes it possible to find a proof by managing only the formula and its subformula. Among them, the tableau method (or its variant, Gentzen's sequent calculus) has received much attention in the automation of modal logic theorem prover recently. While resolution is the commonly-used proof method in classical logic, it is less attractive in the modal logic realm due to the requirement that formulas must be converted into normal form before applying resolution-style rules. The advantage of the tableau method is its strong connection with the semantics as each tableau rule reflects the semantic intuition behind a logical connective or operator.

The pioneering proof methods for modal logic were developed by Fitting[Fitting 1983], in which the systems of prefixed tableaux are presented. A system of prefixed tableaux makes explicit reference to the possible worlds in the Kripke model. We can consider the prefixed tableau method an instance of the general methodology of labelled deductive systems proposed in[Gabbay 1996]. Instead of representing possible worlds as prefixes that describe paths in the model from the initial world as in Fitting's systems, Baldoni et al.[Baldoni et al. 1998] develop a tableau calculus for multimodal logics by giving each new possible world an atomic name and representing explicitly the accessibility relationships among them. Following the method in [Baldoni et al. 1998], we also present the tableau calculus for our logic in that way.

$\alpha-$formula (conjunction)

α	α_1	α_2
$\varphi \wedge \psi$	φ	ψ
$\neg(\varphi \vee \psi)$	$\neg\varphi$	$\neg\psi$
$\neg(\varphi \supset \psi)$	φ	$\neg\psi$

$\beta-$formula (disjunction)

β	β_1	β_2
$\varphi \vee \psi$	φ	ψ
$\neg(\varphi \wedge \psi)$	$\neg\varphi$	$\neg\psi$
$\varphi \supset \psi$	$\neg\varphi$	ψ

$\nu-$formula (necessity)

ν^G	ν_0^G
$[G]\varphi$	φ
$\neg\langle G\rangle\varphi$	$\neg\varphi$

$\pi-$formula (possibility)

π^G	π_0^G
$\langle G\rangle\varphi$	φ
$\neg[G]\varphi$	$\neg\varphi$

Fig. 5. The uniform notation for modal logic formulas

To avoid repetition of some structurally similar rules, a uniform notation invented by Smullyan[Smullyan 1968] and extended in [Fitting 1983] has been adopted extensively in the tableau methods for modal logics. The uniform notation specifies the decomposition of a formula into its components and is shown in figure 5.

The rationale behind the tableau proof method is to find a counter-model for the negation of a wff. If this is impossible, then the original wff is valid. Therefore, to test the validity of a wff φ, let w be a prefix denoting an arbitrary possible world, then a tableau proof tree will be started with the prefixed formula $w : \neg\varphi$ as its root. There are two kinds of formulas in a tableau proof tree. One is the kind of prefixed formulas in the form $w : \varphi$, where w is a prefix and φ is a wff in our logic. The other is the kind of accessibility relation formula $w\rho_G w'$, where w and w' are prefixes and ρ_G is a binary relation symbol for a nonempty subset of agents G[Baldoni et al. 1998]. These two kinds of formulas are called *tableau formulas*.

A tableau rule is applicable when some branch of the proof tree contains the premises of some rule, and the conclusions of the rule are added to the end of the branch after the application of the rule. The set of tableau rules is presented in figure 6. In the presentation, we follow the convention of [Baldoni et al. 1998] by using φ and ψ as metavariables over arbitrary wffs and w, w' as metavariables over prefixes. Also, i, G, and O are respectively taken as metavariables for agents, sets of agents, and total orders over agents. Recall that we identify the singleton $\{i\}$ with i. Also, when the rules are applied, a prefix is *new* if it has never appeared in the branch, otherwise, it is called an *old* prefix.

These rules are divided into three groups. First, there are two kinds of classical rules. The *extension rules* add the conclusions to the end of a branch directly, and the added nodes are organized as a branch and attached to the end of the original branch. The *forking rules* add the conclusions as the children of the end node of the original branch. In the propositional case, all forking rules have exactly two conclusions, so one is the left child and the other is the right child, and the proof tree is binary.

Second, the modal rules can be seen as a special instance of those for inclusion modal logics and incestual modal logics in [Baldoni 1998]. In particular, the rule D is due to the serial accessibility relation for each individual agent, and can be

(1) Classical rules:

(a) Extension rules(α-rule and $\neg\neg$-rule):

$$\frac{w:\alpha}{\begin{array}{c}w:\alpha_1\\ w:\alpha_2\end{array}} \qquad \frac{w:\neg\neg\varphi}{w:\varphi}$$

(b) Forking rule(β-rule):

$$\frac{w:\beta}{w:\beta_1 \mid w:\beta_2}$$

(2) Modal rules:

(a) ν-rule and π-rule:

$$\frac{\begin{array}{c}w:\nu^G\\ w\rho_G w'\end{array}}{w':\nu_0^G} \qquad \frac{w:\pi^G}{\begin{array}{c}w\rho_G w'\\ w':\pi_0^G\end{array}}$$

where in the π-rule, w' is new

(b) Rule D:

$$\frac{w \text{ is old}}{w\rho_i w'}$$

where w' is new.

(c) ρ-rule: if $\emptyset \neq G_2 \subset G_1$, then

$$\frac{w\rho_{G_1}w'}{w\rho_{G_2}w'}$$

(3) DBF_n^c rule:

$$\frac{w:[O>i]\varphi}{\begin{array}{c}w:[\delta(O>i)]\bot\\ w:[O]\varphi\end{array} \Bigg| \begin{array}{c}w:\neg[\delta(O>i)]\bot\\ w:[\delta(O>i)]\varphi\end{array}} \qquad \frac{w:\neg[O>i]\varphi}{\begin{array}{c}w:[\delta(O>i)]\bot\\ w:\neg[O]\varphi\end{array} \Bigg| \begin{array}{c}w:\neg[\delta(O>i)]\bot\\ w:\neg[\delta(O>i)]\varphi\end{array}}$$

(4) DBF_n^s rule:

$$\frac{w:[\Omega\cup\{O>i\}]\varphi}{\begin{array}{c}w:[\{O,i\}]\bot\\ w:[\Omega\cup\{O\}]\varphi\end{array} \Bigg| \begin{array}{c}w:\neg[\{O,i\}]\bot\\ w:[\Omega\cup\{O,i\}]\varphi\end{array}} \qquad \frac{w:\neg[\Omega\cup\{O>i\}]\varphi}{\begin{array}{c}w:[\{O,i\}]\bot\\ w:\neg[\Omega\cup\{O\}]\varphi\end{array} \Bigg| \begin{array}{c}w:\neg[\{O,i\}]\bot\\ w:\neg[\Omega\cup\{O,i\}]\varphi\end{array}}$$

Fig. 6. The tableau rules for DBF_n^C and DBF_n^S

characterized by a special incestual axiom. However, it must be noted that the unrestricted application of rule D may cause non-termination of the proof search process since we can always introduce a new ρ_i-successor for any prefix currently available in the branch. Therefore, we require that the following conditions must be satisfied before the rule D can be applied to an available prefix w:

(1) there are some necessity formulas ν^i such that $w:\nu^i$ has appeared previously,

and

(2) there does not exist any prefix w' and subset of agents G such that $i \in G$ and $w\rho_G w'$ has been in the branch.

In other words, for each available prefix w and agent i, rule D is applied at most once, and only when the prefix appears in the front of some ν formulas.

As for other modal rules, they are only applicable to distributed belief operators. Since a modal operator $[G]$ can be seen as a special kind of $[\Omega]$ operator, it exists in both logics DBF_n^c and DBF_n^s. Therefore, the modal rules, as well as classical rules, are common to DBF_n^c and DBF_n^s.

On the other hand, the DBF_n^c rule and DBF_n^s rule, as the names suggest, are specific to the respective logics. These rules are direct translations of the axioms O1, O2, O1′, and O2′. By repeatedly applying these rules, all modal operators can be ultimately reduced to the distributed belief operators in the form $[G]$, so the modal rules can then be applied.

A branch in a tableau is closed if it contains $w : \bot$, $w : \neg\top$, or both $w : \varphi$ and $w : \neg\varphi$ for some prefix w and wff φ. A branch not closed is called an open branch. A tableau is closed if all its branches are closed. Note that, in ordinary modal tableaux, we have to find an open branch in every newly accessed possible world for the construction of a counter-model. However, by using prefixed wffs, an open branch in our modal tableaux corresponds to a bunch of open branches in ordinary modal tableaux.

A wff φ is tableau provable in DBF_n^c(resp. DBF_n^s), denoted by $\Vdash_{\mathrm{DBF}_n^c}\varphi$(resp. $\Vdash_{\mathrm{DBF}_n^s}\varphi$), if there is a closed tableau with $w : \neg\varphi$ as its root, using the set of classical, modal, and DBF_n^c(resp. DBF_n^s) rules. Such a closed tableau is said to be a proof of φ. Then we have:

THEOREM 4.1. *Let L denote DBF_n^c or DBF_n^s and φ be a wff in the logic L, then $\Vdash_L\varphi$ iff $\models_L \varphi$.*

Though ρ rule is of crucial importance for the completeness of the tableau calculus, it may generate too many redundant accessibility relation formulas not used in the later proof. To circumvent the problem, we can remove the ρ rule if the ν rule is replaced by the following generalized ν rule:

$$\frac{\begin{array}{c} w : \nu^G \\ w\rho_{G'}w' \end{array}}{w' : \nu_0^G}$$

where $G \subseteq G'$. This rule is derived by the successive application of ρ rule and ν rule. The main advantage of this rule is that we do not have to generate accessibility relation formulas $w\rho_G w'$ from $w\rho_{G'}w'$ for all $G \subseteq G'$. An accessibility relation formula $w\rho_G w'$ is used implicitly in the generalized ν rule only when some ν^G-type formulas are available in the premises. Since in the original tableau rules, only the premises of ν rule contain an accessibility relation formula, the generalized ν rule is sufficient for the completeness of the tableau calculus. Therefore, we have the following corollary.

COROLLARY 2. *A wff φ is tableau provable by applications of classical rules, generalized ν rule, π rule, rule D, and DBF_n^c (resp. DBF_n^s) rule iff $\models_{\mathrm{DBF}_n^c} \varphi$ (resp.*

$\models_{\mathrm{DBF}_n^s} \varphi$).

Sometimes, we use a derived rule, called macro DBF_n^s rule, to shorten the height of a tableau tree. The rule is described as follows. Let $O = i_1 > i_2 > \cdots > i_k$, and G be any subset of $\delta(O)$ such that $i_1 \in G$ and $G_j = (G \cap \{i_1, \cdots, i_{j-1}\}) \cup \{i_j\}$ for $1 < j \le k$. Define Σ_G^O as the following set of wffs

$$\{\neg[G]\bot\} \cup \{[G_j]\bot \mid i_j \notin G\}.$$

Then the macro DBF_n^s rule is

$$\frac{w : [\Omega \cup \{O\}]\varphi}{\begin{array}{c|c} w : [\Omega \cup G]\varphi & \cdots\cdots \\ \{w : \psi \mid \psi \in \Sigma_G^O\} & \cdots\cdots \end{array}} \qquad \frac{w : \neg[\Omega \cup \{O\}]\varphi}{\begin{array}{c|c} w : \neg[\Omega \cup G]\varphi & \cdots\cdots \\ \{w : \psi \mid \psi \in \Sigma_G^O\} & \cdots\cdots \end{array}}$$

In other words, when the macro DBF_n^s rule is applied to a branch containing $[\Omega \cup \{O\}]\varphi$ (resp. $\neg[\Omega \cup \{O\}]\varphi$), generate a child for each subset G of $\delta(O)$ that contains the first element of O, and add to the end of the branch the set Σ_G^O and the wff $[\Omega \cup G]\varphi$(resp. $\neg[\Omega \cup G]\varphi$). The following proposition shows the soundness of the macro DBF_n^s rule.

PROPOSITION 4. *The macro DBF_n^s rule is derivable in the tableau calculus for DBF_n^s.*

Example 10. Let us give a tableau proof for example 4. This shows that

$$\varphi = (\neg[\{1,2\}]\bot \vee \neg[\{1,3\}]\bot) \wedge [1](p \supset q) \wedge [2]p \wedge [3]\neg q \supset [1 > 2 > 3]((p \wedge q) \vee (\neg p \wedge \neg q))$$

is tableau provable. The proof is given in figure 7. By using the macro DBF_n^s rule, the tableau tree is no longer binary. Indeed, the proof in figure 7 is split into four branches below the double lines due to the occurrence of $\neg[1 > 2 > 3]((p \wedge q) \vee (\neg p \wedge \neg q))$, where $\delta(1 > 2 > 3)$ has four subsets containing 1. The first three branches are shown in figure 7(a), whereas the last branch is shown in figure 7(b). Before splitting occurs, we apply α rule to the wffs repeatedly. For the first three branches, the π^G and ν^G rules are needed to obtain the desired refutation. However, for the last branch, only β rule is needed. Note that we in fact use the generalized ν^G rule instead of the ρ rule. Therefore, since $w\rho_{\{1,2,3\}}w', w : [1](p \supset q), w : [2]p$, and $w : [3]q$ appear in the first branch, we can add $w' : p \supset q$, $w' : p$, and $w' : q$ to that branch; since the accessibility relation formula appearing on the second branch is $w\rho_{\{1,2\}}w'$, we cannot have $w' : q$ in that branch. □

4.1 Complexity results

An alternative formulation of the tableau method has been used in [Halpern and Moses 1992] to prove that there is a PSPACE algorithm for deciding satisfiability of formulas in multi-agent epistemic logic K_n. The sub-formula property of the tableau method plays a crucial role in the proof. Without using prefixes, the tableau in [Halpern and Moses 1992] is embedded in a pre-tableau: a tree with nodes labelled by sets of wffs, and some edges labelled by agents. Roughly speaking, a label in the nodes of the pre-tableau is a set of wffs with the same prefix in a branch of our prefixed tableau.

$w : \neg((\neg[\{1,2\}]\bot \vee \neg[\{1,3\}]\bot) \wedge [1](p \supset q) \wedge [2]p \wedge [3]\neg q$
$\supset [1 > 2 > 3]((p \wedge q) \vee (\neg p \wedge \neg q)))$
$w : (\neg[\{1,2\}]\bot \vee \neg[\{1,3\}]\bot) \wedge [1](p \supset q) \wedge [2]p \wedge [3]\neg q$
$w : \neg[1 > 2 > 3]((p \wedge q) \vee (\neg p \wedge \neg q))$
$w : \neg[\{1,2\}]\bot \vee \neg[\{1,3\}]\bot$
$w : [1](p \supset q)$
$w : [2]p$
$w : [3]\neg q$

$w : \neg[\{1,2,3\}]\bot$	$w : \neg[\{1,2\}]\bot$
$w : \neg[\{1,2,3\}]((p \wedge q) \vee (\neg p \wedge \neg q))$	$w : [\{1,2,3\}]\bot$
$w\rho_{\{1,2,3\}}w'$	$w : \neg[\{1,2\}]((p \wedge q) \vee (\neg p \wedge \neg q))$
$w' : \neg\bot$	$w\rho_{\{1,2\}}w'$
$w' : p \supset q$	$w' : \neg((p \wedge q) \vee (\neg p \wedge \neg q))$
$w' : p$	$w' : \neg(p \wedge q)$
$w' : \neg q$	$w' : \neg(\neg p \wedge \neg q)$
split: $w' : \neg p$ ($\times$) | $w' : q$ ($\times$)	$w' : p \supset q$
	$w' : p$
	split: $w' : \neg p$ ($\times$) | $w' : \neg q$, then split: $w' : \neg p$ ($\times$) | $w' : q$ ($\times$)

(a) The first two branches of the tableau

⋮

⋮

$w : \neg[\{1,3\}]\bot$	$w : [\{1,2\}]\bot$
$w : [\{1,2\}]\bot$	$w : [\{1,3\}]\bot$
$w : \neg\{1,3\}]((p \wedge q) \vee (\neg p \wedge \neg q))$	$w : \neg[1]\bot$
$w\rho_{\{1,3\}}w'$	$w : \neg[\{1\}]((p \wedge q) \vee (\neg p \wedge \neg q))$
$w' : \neg((p \wedge q) \vee (\neg p \wedge \neg q))$	split: $w : \neg[\{1,2\}]\bot$ ($\times$) | $w : \neg[\{1,3\}]\bot$ ($\times$)
$w' : \neg(p \wedge q)$	
$w' : \neg(\neg p \wedge \neg q)$	
$w' : p \supset q$	
$w' : \neg q$	
split: $w' : \neg p$, then split: ($w' : \neg\neg p$, $w' : p$, $\times$) | ($w' : \neg\neg q$, $w' : q$, $\times$) | $w' : q$ ($\times$)	

(b) The last two branches of the tableau

Fig. 7. A tableau proof for the wff in Example 4.

It is shown that the K_n tableau construction procedure for a wff φ can construct a pre-tableau with a depth polynomial in the size of φ. Furthermore, each formula appearing in the nodes of the pre-tableau is a sub-formula of φ. More precisely, let $Sub(\varphi)$ denote the set of sub-formulas of φ and $Sub^+(\varphi) = Sub(\varphi) \cup \{\neg\psi \mid \psi \in Sub(\varphi)\}$, then each node of the pre-tableau is labelled by a subset of $Sub^+(\varphi)$. Let m be the size of φ, then the label of each node can be represented by a length $2m$ bit-string since it is a subset of $Sub^+(\varphi)$. Therefore, if we search through the pre-tableau in a depth-first way(DFS), then the total space needed is the multiplication of the depth of the pre-tableau and the $O(m)$ space used to store the bit-string and some book-keeping information.

The proof in [Halpern and Moses 1992] can be easily modified to prove the PSPACE upper bound for our logics. The key point is the definition of sub-formula. Due to the DBF_n^c and DBF_n^s rules, the set $Sub^+(\varphi)$ must be expanded to include more wffs. Let $O \in \mathcal{TO}_n$ and $\Omega \subseteq \mathcal{TO}_n$, then define

$$Sub(O) = \{i_1 > i_2 > \cdots > i_k \mid i_1, i_2, \cdots, i_k \in \delta(O)\}$$

as the set of all sub-orderings of O and

$$Sub(\Omega) = \{\Omega' \mid \Omega' \subseteq \bigcup_{O \in \Omega} Sub(O)\}.$$

For a DBF_n^c wff φ, define

$$\begin{aligned} Sub_c(\varphi) \ &= \ Sub(\varphi) \cup \{[O']\psi \mid [O]\psi \in Sub(\varphi), O' \in Sub(O)\} \\ &\quad \cup \{[G]\psi, [G]\bot \mid [O]\psi \in Sub(\varphi), G \subseteq \delta(O)\} \end{aligned}$$

and for a DBF_n^s wff φ, define

$$Sub_s(\varphi) = Sub(\varphi) \cup \{[\Omega']\psi, [\Omega']\bot \mid [\Omega]\psi \in Sub(\varphi), \Omega' \in Sub(\Omega)\}.$$

The sets $Sub_c^+(\varphi)$ and $Sub_s^+(\varphi)$ can then defined in the same way as $Sub_c^+(\varphi) = Sub_c(\varphi) \cup \{\neg\psi \mid \psi \in Sub_c(\varphi)\}$ and $Sub_s^+(\varphi) = Sub_s(\varphi) \cup \{\neg\psi \mid \psi \in Sub_s(\varphi)\}$. Let $c_0 = |\mathcal{TO}_n|$, $c_1 = c_0 + 2^{n+1}$ and $c_2 = 2^{c_0}$, then $|Sub_c^+(\varphi)|$ is bounded by $2(c_1+1)m$ and $|Sub_s^+(\varphi)|$ is bounded by $2(2c_2+1)m$, where $m = |Sub(\varphi)|$ is proportional to the size of φ. Note that though c_0, c_1, and c_2 may be very large numbers, they are essentially constants independent of the size of φ. Furthermore, the bound is quite loose, but it is enough for the present purpose.

The bound on $|Sub_c^+(\varphi)|$ and $|Sub_s^+(\varphi)|$ guarantee that the pre-tableau constructed for φ in DBF_n^c or DBF_n^s has a height at most $O(m^2)$ by using exactly the same proof as that in [Halpern and Moses 1992]. Since the nodes in the pre-tableau for φ are labelled with subsets of $Sub_c^+(\varphi)$ or $Sub_s^+(\varphi)$, the label of each node can be represented by a length $O(m)$ bit-string again. Therefore, by using DFS, the total space needed for deciding the satisfiability of φ is $O(m^3)$. This shows that the satisfiability problem for DBF_n^c and DBF_n^s is in PSPACE.

Since both DBF_n^c and DBF_n^s are extensions of the multi-agent epistemic logic KD_n^D, which in turn is an extension of the normal modal logic KD[Chellas 1980], and it has been shown that the satisfiability problem for KD is PSPACE-complete[Ladner 1977], we can know that the satisfiability problem for DBF_n^c and DBF_n^s is PSPACE-hard. Therefore, we have:

Theorem 4.2. *The satisfiability problem for DBF_n^c and DBF_n^s is PSPACE-complete.*

Using this theorem, the validity problem for DBF_n^c and DBF_n^s is co-PSPACE-complete. However, since PSPACE is deterministic complexity class, co-PSPACE is identical to PSPACE, so the validity problem for DBF_n^c and DBF_n^s is also PSPACE-complete.

5. EXTENSIONS AND RELATED WORK

5.1 A generic extension

While we adopt a modal logic approach to belief fusion, there have also been a lot of work done on knowledge merging using meta-level operators. The main approaches include the following:

(1) Combination based on maximal consistency[Baral et al. 1992; 1994; Baral et al. 1991; Benferhat and Garcia 1998].

(2) Combination by meta-information[Pradhan et al. 1995; Subrahmanian 1994].

(3) Merging by majority[Lin 1994; 1996; Lin and Mendelzon 1999].

(4) Arbitration[Liberatore and Schaerf 1995; Revesz 1993; 1997].

(5) General merging[Konieczny 2000; Konieczny and Pérez 1998; 1999].

(6) Belief revision and update[Alchourrón et al. 1985; Katsuno and Medelzon 1991a; 1991b].

In the meta-level approach, a merging operator is generally used to combine a set of knowledge bases $T_1, T_2, \cdots, T_n$, where each knowledge base is a theory in some logical langauge. The main difference between our approach and the meta-level approach is that the belief fusion operators are incorporated into the object language in our logic, so we can reason with the merged results and also about the fusion process. These meta-level fusion operators can be divided into syntax-based or semantics-based types. A general scheme for semantics-based knowledge merging operators is shown in figure 8.

While we have shown in section 3 that some syntax-dependent combination operators based on maximal consistency[Baral et al. 1992; 1994; Baral et al. 1991] can be simulated in our logic, our logics need to be extended to accommodate other more sophisticated semantics-based knowledge merging operators. We now describe a generic approach for such an extension.

To accommodate such kind of merging operators in the modal logic framework, we use an extended language $\mathcal{L}_e$. The well-formed formulas are defined according to the five rules for $\mathcal{L}_c$ or $\mathcal{L}_s$ as well as the following rule:

. 6. If $G \subseteq \{1, 2, \cdots, n\}$ and φ is a wff, then $[\Gamma(G, \boldsymbol{\kappa})]\varphi$ is also a wff, where $\boldsymbol{\kappa}$ is the sequence of parameters for a merging operator to be included in the language.

For the semantics, a model for $\mathcal{L}_e$ is still a triple $(W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$ as defined in the semantics of DBF_n^c and DBF_n^s except that the seriality requirement of $\mathcal{R}_i$ is dropped out. For any subset $G = \{i_1, \cdots, i_k\}$ and $w \in W$, define

$$\succ_{G,\boldsymbol{\kappa}}^w = \ Ord(\mathcal{R}_{i_1}(w), \cdots, \mathcal{R}_{i_k}(w),\ \boldsymbol{\kappa})$$

Input:. A set of knowledge bases $T_1, T_2, \cdots, T_n$ in some logical language and some extra sequence of parameters, e.g., some weights, $\boldsymbol{\kappa}$.

Output:. A merged knowledge base $\Gamma(T_1, \cdots, T_n, \boldsymbol{\kappa})$ in the same logical language.

Steps:. (1) For $1 \leq i \leq n$, let $\mathcal{M}_i$ denote the set of all interpretations satisfying T_i, that is, $\mathcal{M}_i = Mod(T_i)$.

(2) A ordering $\succ$ over the interpretations of the logical language is determined from all $\mathcal{M}_i$'s and $\boldsymbol{\kappa}$. In other words, there is a function Ord such that

$$\succ = Ord(\mathcal{M}_1, \cdots, \mathcal{M}_n, \boldsymbol{\kappa}).$$

(3) A selection function γ will select a set of interpretations $\mathcal{M}$ from all $\mathcal{M}_i$'s according to the ordering $\succ$, i.e.

$$\mathcal{M} = \gamma(\mathcal{M}_1, \cdots, \mathcal{M}_n, \succ).$$

(4) The resultant merged theory is

$$\Gamma(T_1, \cdots, T_n, \boldsymbol{\kappa}) = \{\varphi \mid \forall I \in \mathcal{M}, I \models \varphi\}.$$

Fig. 8. A general scheme for semantics-based fusion

and

$$\mathcal{R}_{G,\boldsymbol{\kappa}}(w) = \gamma(\mathcal{R}_{i_1}(w), \cdots, \mathcal{R}_{i_k}(w), \succ^{w}_{G,\boldsymbol{\kappa}}).$$

The satisfaction of the wff $[\Gamma(G, \boldsymbol{\kappa})]\varphi$ in a world w is obtained immediately by

$$w \models [\Gamma(G, \boldsymbol{\kappa})]\varphi \text{ iff for all } u \in \mathcal{R}_{G,\boldsymbol{\kappa}}(w), u \models \varphi$$

and the definition of validity is as in the case of DBF_n^c and DBF_n^s.

While the semantics-based fusion operators merge knowledge bases, our extended logic merges beliefs of different agents, so there is some connection between these two formalisms. Let $\mathcal{T} = \{T_1, \cdots, T_n\}$ be a set of knowledge bases in a propositional logic language, then $\Sigma_{\mathcal{T}}$ is defined as the following set of wffs in $\mathcal{L}_e$:

$$\Sigma_{\mathcal{T}} = \{[i]\varphi \mid 1 \leq i \leq n, \varphi \in T_i\}.$$

PROPOSITION 5. *Let $\mathcal{T} = \{T_1, \cdots, T_n\}$ be a set of knowledge bases in a propositional logic language, then for any semantics-based merging operator Γ with extra parameters $\boldsymbol{\kappa}$, there exists an $\mathcal{L}_e$ model $M = (W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$ and $w \in W$ such that $w \models_M \Sigma_{\mathcal{T}}$ and for any propositional formula φ,*

$$w \models_M [\Gamma(\{1, \cdots, n\}, \boldsymbol{\kappa})]\varphi \text{ iff } \Gamma(T_1, \cdots, T_n, \boldsymbol{\kappa}) \models \varphi.$$

A direct corollary of the proposition is:

COROLLARY 3. *If $\Sigma_{\mathcal{T}} \models [\Gamma(G, \boldsymbol{\kappa})]\varphi$, then $\Gamma(T_1, \cdots, T_n, \boldsymbol{\kappa}) \models \varphi$.*

Admittedly, the syntactic and semantic extension of the modal logic framework for the semantics-based merging operators is quite general and preliminary. For each individual merging operator, more details must be worked out in some specific way.

In particular, the axiomatic and proof systems for each specific logic may be quite different. Also, the converse of corollary 3 may not hold for general cases, however, it would be interesting to know whether it indeed holds for some specific merging operators. This will be left for further research and we believe that the general methodology proposed in this paper will be helpful for such development.

5.2 Related works

The most similar previous work is multi-agent epistemic logic[Fagin et al. 1996] and multi-source reasoning[Cholvy 1994]. On the one hand, it is clear that DBF_n^c and DBF_n^s extend the multi-agent epistemic logic KD_n^D. On the other hand, in multi-source reasoning, fusion modalities are only applicable to literals (i.e., atomic formulas or their negations), whereas our fusion modalities may be applied to any complicated wffs. Therefore, our logics add expressive power to both multi-agent epistemic logic and multi-source reasoning.

Modal logic for representing inconsistent beliefs is another research area related to both epistemic logic and belief fusion. In [Meyer and van der Hoek 1998], an epistemic default logic is proposed for the representation of inconsistent beliefs caused by default reasoning. The logic is based on S5P developed in [Meyer and van der Hoek 1991; 1992; 1993] for modelling the monotonic part of default reasoning that deals with plausible assumptions. The basic modalities of S5P consist of an S5 epistemic operator K and a number of K45 belief operators $P_i (1 \leq i \leq n)$. A wff $P_i\varphi$ means that φ is a plausible working belief according to some context or default rules. Since conflict between default rules is not unusual, it is possible that $P_i\varphi \wedge P_j\neg\varphi$ holds. Though P_i corresponds to an application context of some default rules, it can also be seen as the belief operator of some agent, so in this regard, the logic is like a multi-agent epistemic logic with an S5-based epistemic operator as its authority. However, instead of reasoning about the merging of different working beliefs in the logic directly, a downward reflection approach is adopted in [Meyer and van der Hoek 1998]. Since P_i operators are only applied to objective wffs in [Meyer and van der Hoek 1998], the downward reflection function maps a set of S5P wffs (especially wffs of the form $P_i\varphi$) into a set of non-modal formulas. Some downward reflection mechanisms are employed to resolve the inconsistency between working beliefs of different contexts. The mechanism based on the explicit ordering of frames is essentially similar to our logics. The main difference is that we take the orderings as modal operators and reason about the fusion results directly in the object language, while the downward reflection approach considers the fusion in a meta-level.

A dynamic doxastic logic for belief revision is proposed in [Segerberg 1995] and further developed in [Segerberg 2001]. By using the notations of [Segerberg 1995], the doxastic operator B and two kinds of dynamic modal operators $[+\varphi]$ and $[-\varphi]$ for propositional wff φ are taken as the basic constructs of the language. The operators $[+\varphi]$ and $[-\varphi]$ correspond, respectively, to the expansion and contraction operators in AGM theory. Thus the revision operator $[\circ\varphi]$ is defined as $[-(\neg\varphi)][+\varphi]$ according to the so-called Levi's identity[Alchourrón et al. 1985]. However, what can be represented in that logic is the belief revision of a single agent by some new information, whereas in our logic, we are interested in the representation and reasoning of the fusion of multi-agent beliefs.

In our presentation, we assume an agent's belief states are represented as a subset of possible worlds, i.e., $\mathcal{R}_i(w)$ is the belief state of agent i in world w. However, some more fine-grained representations have been proposed, such as total pre-orders over the set of possible worlds [Boutilier 1993; Darwiche and Pearl 1997; Lehmann 1995; Segerberg 1995], ordinal conditional functions [Boutilier 1995; Spohn 1988; Williams 1994], possibility distributions[Benferhat et al. 1997; Dubois and Prade 1992; 2000], belief functions[Smets 2000] and pedigreed belief states[II and Shoham 2001; II and Lehmann 2000]. Perhaps, the most popular representation among them is an ordering of the possible worlds. While a set of possible worlds can be seen as the minimal worlds with respect to a given ordering, it is claimed that the fusion of two orderings is more general than the revision of an ordering by a set of possible worlds[II and Shoham 2001]. To fully utilize the semantic power of an ordering, the syntax of our logics should be further enriched to cover the conditional connectives. Such development of logical systems incorporating the fusion operators based on fine-grained representations of belief states should be a very interesting research direction.

6. CONCLUSIONS AND FURTHER RESEARCH

The main contribution of our work is the integration of belief fusion operators into the multi-agent epistemic logic. We propose two basic logical systems for reasoning about the cautiously merged beliefs of multiple agents. The two systems correspond to two different strategies for discarding information sources. In level cutting fusion, if an information source is discarded, then all those less reliable are also discarded without further examination. On the other hand, in the level skipping strategy, only the level under conflict is skipped, and the next level will be considered independent of those discarded before it.

We present the syntax, semantics, axiomatic systems, and tableau proof systems for our logics. We show that our logics extend the expressive power of the multi-agent epistemic logic without increasing its complexity. Finally, we demonstrate a generic extension of our logics for incorporating more sophisticated belief fusion operators proposed in knowledge base merging. While most of the knowledge base merging research takes the fusion process as a meta-level operator, our approach incorporates it into the object logic directly. Therefore, it is possible to integrate the belief fusion operators into the multi-agent epistemic logic. The advantage of using epistemic logic is its capability to reason with not only the beliefs about the objective world but also the beliefs about beliefs.

In our discussion of the belief fusion logic, we did not distinguish between belief and information. However, in a genuine agent systems, an agent's belief may be different than the information he sends to or receives from other agents. Thus, in general, we should have a set of modal operators $[j]_i$ such that $[j]_i\varphi$ means that agent i receives the information φ from j. In particular, $[i]_i\varphi$ may represent the observation of agent i himself, which should be the most reliable information source for i. Then agent i may form his belief by fusing the information he received from different agents according to the degrees of trust he has on other agents. The fusion may be represented by the operators $[O]_i$. If we consider $[j]_i\varphi$ as the communication of message φ from j to i, then we have a general framework for reasoning about an agent's belief and communication. In such a framework, we can discuss problems like deception of agent. For example, $[O]_i\varphi \wedge [i]_j \neg\varphi$ may mean that agent i deceives agent j by telling j the negation of what i believes. In [Demolombe and Liau 2001], an application of our basic systems to reasoning about beliefs and trusts of multiple agents has been proposed along this direction. However, more work remains to be done for real applications. These applications may also take advantage of some fundamental work on multi-agent belief revision[Dragoni 1992; Dragoni and Giorgini 2001; Dragoni et al. 1997; Galliers 1992; Kfir-Dahav and Tennenholtz 1996].

In this paper, we consider the reliability ordering over agents, so the belief of one agent is either completely discarded or completely accepted. A possible extension is to partially accept the belief of one agent in the fusion process. For example, we can trust agent 1 about topic A and trust agent 2 about topic B, then in the merged belief, we can selectively discard the belief of agent 1 about topic B and the belief of agent 2 about topic A. This kind of extension will require reliability orderings over pairs of agent and information. A logic for partially accepting some of what one agent says has been proposed in [Liau 2003].

In a recent paper, it was shown that multi-source reasoning can be applied to deontic logic under conflicting regulations[Cholvy and Cuppens 1999]. Essentially, conflicting regulations are merged according to their priorities in a way analogous to the fusion of information. However, inherited from the restriction of multi-sources reasoning, it is also required that each regulation to be merged must be a set of deontic literals. Now, by the systems developed here, it is expected that the general forms of regulations can also be merged.

A real difficulty in the application of our logic to model the database merging reasoning is the representational problem of the databases. In some examples of sections 2 and 3, we suggest to find all maximal consistent agent groups in advance and add the wff $\bigwedge_{G\in MCAG} \neg[G]\bot$ to the representation. This is rather time-consuming work. In practice, we can omit this part and check the consistency of some agent groups when it is necessary during the course of proof. Even further, we can consider the implementation of the logic with some non-monotonic reasoning techniques[Antoniou 1997] so that only the explicit information in the knowledge base has to be represented. This will be investigated in the further research.

A. PROOF OF THE PROPOSITIONS AND THEOREMS

A.1 Proof of Proposition 1

(1) By induction on $m = |\delta(O)|$:

. $m = 1$: this is trivial since we identify $[i_1]\varphi$ and $[\{i_1\}]\varphi$ in our language.

. assume this result holds for all $m \leq k$.

. $m = k + 1$: there are two cases:

. $j = m$: $\vdash \neg[G_j]\bot \supset ([O]\varphi \equiv [G_j]\varphi)$ is just an instance of axiom O1,

. $j < m$: let O be written as $O' > i_m$ where $O' = i_1 > \ldots > i_k$, then this proof is as follows:

1.$\neg[G_j]\bot \wedge [G_{j+1}]\bot \supset [G_m]\bot$	$G3, m \geq j+1$
2.$[G_m]\bot \supset ([O]\varphi \equiv [O']\varphi)$	$O2$
3.$\neg[G_j]\bot \wedge [G_{j+1}]\bot \supset ([O]\varphi \equiv [O']\varphi)$	$1, 2, P, MP$
4.$\neg[G_j]\bot \wedge [G_{j+1}]\bot \supset ([O']\varphi \equiv [G_j]\varphi)$	ind. hyp.
5.$\neg[G_j]\bot \wedge [G_{j+1}]\bot \supset ([O]\varphi \equiv [G_j]\varphi)$	$3, 4, P, MP$

(2) By induction on $|\delta(O)|$: if $|\delta(O)| = 1$, this is an instance of G1. Assume this holds for modal operator $[O]$, let us consider the proof for $[O > i]$. Let p and G denote $[\delta(O > i)]\bot$ and $\delta(O > i)$, respectively

1.$[O > i]\varphi \supset (\neg p \supset [G]\varphi)$	$O1$
2.$[O > i](\varphi \supset \psi) \supset (\neg p \supset [G](\varphi \supset \psi))$	$O1$
3.$[O > i]\varphi \supset (p \supset [O]\varphi)$	$O2$
4.$[O > i](\varphi \supset \psi) \supset (p \supset [O](\varphi \supset \psi))$	$O2$
5.$[O > i]\varphi \wedge [O > i](\varphi \supset \psi) \supset (\neg p \supset [G]\psi)$	$1, 2, G1, P, MP$
6.$[O > i]\varphi \wedge [O > i](\varphi \supset \psi) \supset (p \supset [O]\psi)$	$3, 4,$ ind. hyp., P, MP
7.$[O > i]\varphi \wedge [O > i](\varphi \supset \psi) \supset (\neg p \supset [O > i]\psi)$	$5, O1, P, MP$
8.$[O > i]\varphi \wedge [O > i](\varphi \supset \psi) \supset (p \supset [O > i]\psi)$	$6, O2, P, MP$
9.$[O > i]\varphi \wedge [O > i](\varphi \supset \psi) \supset [O > i]\psi$	$7, 8, P, MP$

(3) By induction on $|\delta(O)|$: if $|\delta(O)| = 1$, this is an instance of G2. Assume we have $\vdash \neg[O]\bot$, then the proof of $\vdash \neg[O > i]\bot$ is as follows:

1.$\neg[\delta(O > i)]\bot \supset ([O > i]\bot \supset [\delta(O > i)]\bot)$	$O1$
2.$[\delta(O > i)]\bot \supset ([O > i]\bot \supset [O]\bot)$	$O2$
3.$[O > i]\bot \supset [\delta(O > i)]\bot$	$1, P, MP$
4.$[O > i]\bot \supset ([\delta(O > i)]\bot \supset [O]\bot)$	$2, P, MP$
5.$[O > i]\bot \supset [O]\bot$	$3, 4, P, MP$
6.$\neg[O]\bot$	ind. hyp.
7.$\neg[O > i]\bot$	$5, 6, P, MP$

(4) By induction on $|\delta(O)|$: if $|\delta(O)| = 1$, this is an instance of Gen rule. Assume it is the case for modal operator $[O]$, let us consider the proof for $[O > i]$

1.φ	Assumption
2.$[O]\varphi$	Ind. Hyp.
3.$[\delta(O > i)]\varphi$	*Gen*
4.$[O > i]\varphi \equiv ([\delta(O > i)]\varphi \vee [O]\varphi)$	$O1, O2, P, MP$
5.$[O > i]\varphi$	$2, 3, 4, P, MP$

A.2 Proof of Theorem 2.1

The proof of the theorem is based on that for $S5_n^D$ in [Fagin et al. 1996; Fagin et al. 1992]. As usual, the verification of the soundness part is a routine checking, so we focus on the completeness part. Let $\mathcal{L}$ denote a logical system. A wff φ is $\mathcal{L}$-inconsistent if its negation $\neg\varphi$ can be proved in $\mathcal{L}$. Otherwise, φ is $\mathcal{L}$-consistent. A set Σ of wffs is said to be $\mathcal{L}$-inconsistent if there is a finite subset $\{\varphi_1, \ldots, \varphi_k\} \subseteq \Sigma$ such that the wff $\varphi_1 \wedge \cdots \wedge \varphi_k$ is $\mathcal{L}$-inconsistent; otherwise, Σ is $\mathcal{L}$-consistent. A maximal $\mathcal{L}$-consistent set of wffs ($\mathcal{L}$-MCS) is a consistent set χ of wffs such that whenever ψ is a wff not in χ, then $\chi \cup \{\psi\}$ is $\mathcal{L}$-inconsistent.

On the other hand, φ is $\mathcal{L}$-satisfiable iff there exists a $\mathcal{L}$ model M and a possible world w such that $w \models_M \varphi$, otherwise φ is $\mathcal{L}$-unsatisfiable. Sometimes the prefix $\mathcal{L}$ will be omitted without confusion. To prove the completeness, we will show that every DBF_n^c-consistent wff is DBF_n^c-satisfiable.

Let $\mathcal{I} = \mathcal{TO}_n \cup 2^{\{1,2...,n\}} - \{\emptyset\}$ be the set of all modal operators for the language DBF_n^c. A *pseudo* DBF_n^c structure is a tuple $(W, (\mathcal{R}_I^*)_{I\in\mathcal{I}}, V)$ where W and V are defined as in DBF_n^c models and each $\mathcal{R}_I^*$ is a binary relation on W. Furthermore,it is required that $\mathcal{R}_{\{i\}}^*$ is a serial relation for each $1 \leq i \leq n$. The satisfaction clauses for DBF_n^c wffs in pseudo structures are defined as usual, so for example, we have $w \models [O]\varphi$ iff for $u \in \mathcal{R}_O^*(w)$, $u \models \varphi$. What make difference is that in a pseudo structures, each $\mathcal{R}_I^*$ is considered as an independent relation instead of the intersection of other individual ones. A pseudo structure M^* is called a pseudo model if all wffs provable in DBF_n^c are valid in M^*. A DBF_n^c wff φ is pseudo satisfiable if there exists a pseudo model M^* and a possible world w such that $w \models_{M^*} \varphi$.

The following two results will be proved:

LEMMA 1. *(1) If φ is DBF_n^c-consistent, then φ is pseudo DBF_n^c-satisfiable. (2) If φ is pseudo DBF_n^c-satisfiable, then it is DBF_n^c-satisfiable.*

The first result is proved by a standard canonical model construction procedure. A canonical pseudo structure $M^* = (W, (\mathcal{R}_I)^*_{I\in\mathcal{I}}, V)$ is defined as follows

— $W = \{w_\chi \mid \chi \text{ is a DBF}_n^c\text{-MCS}\}$, in other words, each possible world corresponds precisely to a DBF_n^c-MCS.
— $\mathcal{R}_I^*(w_{\chi_1}, w_{\chi_2})$ iff $\chi_1/I \subseteq \chi_2$ for all $I \in \mathcal{I}$, where $\chi_1/I = \{\varphi \mid [I]\varphi \in \chi_1\}$.
— $V : \Phi_0 \to 2^W$ is defined by $V(p) = \{w_\chi \mid p \in \chi\}$.

The most important result for such construction is the truth lemma.

LEMMA 2 TRUTH LEMMA. *For any wff φ and DBF_n^c-MCS χ, we have $w_\chi \models_{M^*} \varphi$ iff $\varphi \in \chi$.*

PROOF. By induction on the structure of the wff, the only interesting case is the wff of the form $[I]\psi$ for some $I \in \mathcal{I}$. By definition, $w_\chi \models_{M^*} [I]\psi$ iff for all $w_{\chi'} \in \mathcal{R}_I^*(w_\chi)$, $w_{\chi'} \models_{M^*} \psi$ iff for all $\chi/I \subseteq \chi'$, $\psi \in \chi'$ (by induction hypothesis) iff $\chi/I \cup \{\neg\psi\}$ is inconsistent iff $[I]\psi \in \chi$ when $[I]$ is a normal modal operator[Chellas 1980]. However, by the axioms P and G1, rules MP and Gen, and propositions 1.2 and 1.4, both kinds of modal operators $[O]$ and $[G]$ are normal ones. □

Since every DBF_n^c-MCS contains all wffs provable in DBF_n^c, by the truth lemma, all provable wffs are valid in M^*. Furthermore, by axiom G2, each $\mathcal{R}_{\{i\}}^*$ is serial for $1 \leq i \leq n$. Thus M^* is indeed a pseudo model. If φ is DBF_n^c-consistent, then there exists an MCS χ containing φ, so by the truth lemma, $w_\chi \models_{M^*} \varphi$, i.e., φ is pseudo DBF_n^c-satisfiable. This proves the first part of lemma 1.

Note that if $|\mathcal{I}| = m$, then a pseudo model is in fact a model for the multi-agent epistemic logic K_m[Fagin et al. 1996]. The logic K_m has m modal operators corresponding to the knowledge or belief of m independent agents. Admittedly, m may be a very large number, however, it does not matter for the current purpose. What is important is that it can be shown that without loss of generality, we can assume a pseudo model is tree-like. The detail definition of a tree-like model and the proof that each pseudo model can be "unwound" into a tree-like one verifying the same set of valid wffs are rather technical and can be found in ([Fagin et al. 1992],pp.354) and ([Fagin et al. 1996],Exercise 3.30). What is needed here is the property that in a tree-like model, if $I \neq J$, then $\mathcal{R}_I^* \cap \mathcal{R}_J^* = \emptyset$.

Thus, from now on, we can assume that if φ is DBF_n^c-consistent, then φ is pseudo DBF_n^c-satisfiable in a tree-like model $M^* = (W, (\mathcal{R}_I^*)_{I \in \mathcal{I}}, V)$. The next step is to construct a DBF_n^c model $M = (W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$ from M^* by defining $\mathcal{R}_i = \bigcup_{i \in G} \mathcal{R}_G^*$. Note that $\mathcal{R}_i$ is serial since $\mathcal{R}_i \supseteq \mathcal{R}_{\{i\}}^*$ which is serial by the definition of pseudo models. From the definition, we can prove the following lemma.

LEMMA 3. *For any $w \in W$ and wff φ, $w \models_{M^*} \varphi$ iff $w \models_M \varphi$.*

PROOF. By induction on the structure of φ, the basis and classical cases are easy since both models have the same truth assignment function V. For the modal cases, if $\varphi = [G]\psi$, then $w \models_M [G]\psi$ iff for all $u \in \bigcap_{i \in G} \mathcal{R}_i(w)$, $u \models_M \psi$ iff for all $u \in \bigcup_{G' \supseteq G} \mathcal{R}_{G'}^*(w)$, $u \models_{M^*} \psi$ (by the definition of $\mathcal{R}_i$, the tree-likeness of M^*, and the induction hypothesis) iff $w \models_{M^*} \bigwedge_{G' \supseteq G} [G']\psi$ (by definition of satisfaction in pseudo model M^*) iff $w \models_{M^*} [G]\psi$ (since axiom G3 is valid in M^*).

If $\varphi = [O]\psi$, then since proposition 1.1 is valid in both M^* (by definition of pseudo models) and M (by soundness), and by the case for modal operators $[G]$, we can find a j such that w satisfies the wff $\neg[G_j]\bot \wedge [G_{j+1}]\bot$ (or just $\neg[G_j]\bot$ in case of $j = |\delta(O)|$) in both M and M^*, so it can be shown that $w \models_M [O]\psi$ iff $w \models_M [G_j]\psi$ iff $w \models_{M^*} [G_j]\psi$ iff $w \models_{M^*} [O]\psi$. □

This finishes the proof for the second part of lemma 1 and by combining the two parts, we have proved the completeness theorem for DBF_n^c.

A.3 Proof of Proposition 2

Let k denote the largest integer such that $\alpha_k > Incons(\Sigma)$, then $G_k = \{1, 2, \ldots, k\}$ is a subset of some maximal consistent agent group for S and $G_{k+1} = \{1, 2, \ldots, k+1\}$ is inconsistent for S, so by the definition of Σ_S, we have $\Sigma_S \vdash_{\mathrm{DBF}_n^c} \neg[G_k]\bot \wedge [G_{k+1}]\bot$. (In the case of $k = n$, G_{k+1} is omitted.) Since φ is a nontrivial consequence of S, it is a classical logical consequence of $\bigcup_{i \in G_k} S_i$, we have $\Sigma_S \vdash_{\mathrm{DBF}_n^c} [G_k]\varphi$. Therefore, by proposition 1.1, we have $\Sigma_S \vdash_{\mathrm{DBF}_n^c} [1 > 2 > \cdots > n]\varphi$.

A.4 Proof of Theorem 3.1

This proof is analogous to that for theorem 2.1. The difference is that we do not have a counterpart for proposition 1.1 in the system DBF_n^s. First, a pseudo DBF_n^s structure is analogously defined as a tuple $(W, (\mathcal{R}_\Omega)_{\emptyset \neq \Omega \subseteq \mathcal{TO}_n}, V)$ and it is required that $\mathcal{R}_{\{i\}}$ is serial for all $1 \leq i \leq n$. Then a pseudo DBF_n^s model is a pseudo DBF_n^s structure in which all wffs provable in DBF_n^s are valid.

We still have to prove the following lemma.

LEMMA 4.

(*1*) *If φ is DBF_n^s-consistent, then φ is pseudo DBF_n^s-satisfiable.*

(*2*) *If φ is pseudo DBF_n^s-satisfiable, then it is DBF_n^s-satisfiable.*

The first part of the lemma is proved exactly in the same way as in lemma 1. It can also be obtained that if φ is DBF^s_n-consistent, then φ is pseudo DBF^s_n-satisfiable in a tree-like model $M^* = (W, (\mathcal{R}^*_\Omega)_{\emptyset \neq \Omega \subseteq \mathcal{TO}_n}, V)$.

However, the proof of the second part is somewhat different. Let us define the *level* of a modal operator Ω as $l(\Omega) = \max_{O \in \Omega} |\delta(O)|$ and the *length* of Ω as $\sharp(\Omega) =$ the number of elements O in Ω such that $|\delta(O)| = l(\Omega)$. Then we define a function $Ag^* : W \times (2^{\mathcal{TO}_n} - \{\emptyset\}) \to (2^{\{1,2,\ldots,n\}} - \emptyset)$ from the model M^* by

$$Ag^*(w,\Omega) = \begin{cases} \bigcup_{O \in \Omega} Ag^*(w,\{O\}) & \text{if } |\Omega| > 1, \\ Ag^*(w,\{O\}) \cup \{i\} & \text{if } \Omega = \{O > i\} \text{ and } w \models_{M^*} \neg[\{O,i\}]\bot, \\ Ag^*(w,\{O\}) & \text{if } \Omega = \{O > i\} \text{ and } w \models_{M^*} [\{O,i\}]\bot, \\ \{i\} & \text{if } \Omega = \{i\}, \end{cases}$$

Note that since $Ag^*(w,\Omega)$ is a subset of agents, it can also be used as a modal operator for level 1. We can now construct a DBF^s_n model $M = (W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$ from M^* such that $\mathcal{R}_i = \bigcup_{l(\Omega)=1, i \in \Omega} \mathcal{R}^*_\Omega$.

LEMMA 5.

(*1*) *For all $w \in W$, modal operators Ω, and wffs φ, we have $w \models_{M^*} [\Omega]\varphi \equiv [Ag^*(w,\Omega)]\varphi$.*

(*2*) *$\mathcal{R}_\Omega(w) = \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w,\Omega) \subseteq \Omega'\}$ for all $w \in W$ and modal operators Ω, where $\mathcal{R}_\Omega$ is defined in section 4.*

PROOF. (1) By induction on the level of Ω:

. The basis case $l(\Omega) = 1$: then by definition, $Ag^*(w,\Omega) = \Omega$, so the result holds trivially.

. Assume the result holds for all Ω such that $l(\Omega) \leq k$,

. $l(\Omega) = k+1$: by induction on the length of Ω:

. $\sharp(\Omega) = 1$: let $\Omega = \{O > i\} \cup \Omega_1$, where $l(\Omega_1) \leq k$, then $Ag^*(w,\Omega) = Ag^*(w,\{O > i\}) \cup Ag^*(w,\Omega_1) = Ag^*(w,\{O\}) \cup Ag^*(w,\Omega_1) \cup \{i\} = Ag^*(w,\Omega_2)$ if $w \models_{M^*} \neg[\{O,i\}]\bot$ and $= Ag^*(w,\{O\}) \cup Ag^*(w,\Omega_1) = Ag^*(w,\Omega_3)$ if $w \models_{M^*} [\{O,i\}]\bot$, where $\Omega_2 = \{O,i\} \cup \Omega_1$ and $\Omega_3 = \{O\} \cup \Omega_1$. Since $l(\Omega_2) = l(\Omega_3) = k$, then by induction hypothesis, we have $w \models_{M^*} [\Omega_i]\varphi \equiv [Ag^*(w,\Omega_i)]\varphi$ for $i = 2, 3$, so by axioms O1$'$ and O2$'$ (recall that all axioms are valid in a pseudo model), $w \models_{M^*} [\Omega]\varphi \equiv [Ag^*(w,\Omega)]\varphi$ no matter whether $w \models_{M^*} [\{O,i\}]\bot$ or not.

. Assume the result holds for all Ω such that $\sharp(\Omega) \leq t$:

. $\sharp(\Omega) = t+1$: the induction step is completely the same as in the basis case except that $l(\Omega_2) = l(\Omega_3) = k+1$ but $\sharp(\Omega_2) = \sharp(\Omega_3) = t$.

(2) By induction on $l(\Omega)$:

. $l(\Omega) = 1$: then $Ag^*(w,\Omega) = \Omega$ and by definition in section 4

$$\begin{aligned} \mathcal{R}_\Omega(w) &= \bigcap_{i \in \Omega} \mathcal{R}_i(w) \\ &= \bigcap_{i \in \Omega} \bigcup \{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge i \in \Omega'\} \\ &= \bigcup \{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge \Omega \subseteq \Omega'\} \\ &= \bigcup \{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w,\Omega) \subseteq \Omega'\} \end{aligned}$$

. Assume the result holds for $l(\Omega) \leq k$.
. $l(\Omega) = k + 1$: there are two cases
. $|\Omega| = 1$: let $\Omega = \{O > i\}$, then by definition in section 4, we have

$$\mathcal{R}_\Omega(w) = \mathcal{R}_{O>i}(w) = \begin{cases} \mathcal{R}_O(w) & \text{if } \mathcal{R}_O(w) \cap \mathcal{R}_i(w) = \emptyset, \\ \mathcal{R}_O(w) \cap \mathcal{R}_i(w) & \text{otherwise,} \end{cases}$$

where by the induction hypothesis and the definitions of Ag^* and $\mathcal{R}_i(w)$,

$$\mathcal{R}_O(w) = \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w, \{O\}) \subseteq \Omega'\}$$

$$\mathcal{R}_O(w) \cap \mathcal{R}_i(w) = \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w, \{O, i\}) \subseteq \Omega'\}.$$

On the other hand, by the result of first part, let $\Omega_1 = Ag^*(w, \{O, i\})$, then $w \models_{M^*} [\{O, i\}]\bot$ iff $w \models_{M^*} [\Omega_1]\bot$ iff $w \models_{M^*} [\Omega']\bot$ for all Ω' such that $l(\Omega') = 1$ and $\Omega_1 \subseteq \Omega'$ (by axiom V3) iff $\mathcal{R}^*_{\Omega'}(w) = \emptyset$ for all such Ω' iff $\mathcal{R}_O(w) \cap \mathcal{R}_i(w) = \emptyset$. Thus, by the definition of Ag^*,

$$Ag^*(w, \Omega) = \begin{cases} Ag^*(w, \{O\}) & \text{if} \mathcal{R}_O(w) \cap \mathcal{R}_i(w) = \emptyset \\ Ag^*(w, \{O, i\}) & \text{otherwise} \end{cases}$$

and the result follows immediately.
. $|\Omega| > 1$: by definition

$$\begin{aligned} \mathcal{R}_\Omega(w) &= \bigcap_{O \in \Omega} \mathcal{R}_O(w) \\ &= \bigcap_{O \in \Omega} \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w, \{O\}) \subseteq \Omega'\} \\ &= \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge \bigcup_{O \in \Omega} Ag^*(w, \{O\}) \subseteq \Omega'\} \\ &= \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w, \Omega) \subseteq \Omega'\} \end{aligned}$$

□

Finally, we can prove the counterpart of lemma 3 for DBF^s_n

LEMMA 6. *For any $w \in W$ and wff φ, $w \models_{M^*} \varphi$ iff $w \models_M \varphi$.*

PROOF. By induction on the structure of φ, the only interesting case is $\varphi = [\Omega]\psi$,

$$\begin{aligned} w \models_{M^*} [\Omega]\psi &\Leftrightarrow w \models_{M^*} [Ag^*(w, \Omega)]\psi \text{ (lemma 5.1)} \\ &\Leftrightarrow w \models_{M^*} [\Omega']\psi \text{ for all } \Omega' \text{ such that } l(\Omega') = 1 \text{ and } Ag^*(w, \Omega) \subseteq \Omega' \text{ (V3)} \\ &\Leftrightarrow u \models_{M^*} \psi, \forall u \in \bigcup\{\mathcal{R}^*_{\Omega'}(w) \mid l(\Omega') = 1 \wedge Ag^*(w, \Omega) \subseteq \Omega'\} \\ &\Leftrightarrow u \models_M \psi, \forall u \in \mathcal{R}_\Omega(w) \text{ (induction hypothesis and lemma 5.2)} \\ &\Leftrightarrow w \models_M [\Omega]\psi \end{aligned}$$

□

This completes the proof of the second part of lemma 4 and the completeness theorem for DBF^s_n.

A.5 Proof of Proposition 3

(1) Let $S = S_1 \cup \cdots \cup S_k$, then extend the partial ordering Q such that for all $x, y \in \{1, 2, \cdots, n\}$, if $\varphi_{ij_1} \in S_i$ and $\varphi_{ij_2} \notin S_i$, where $i = g_1(x) = g_1(y)$, $g_2(x) = j_1$ and $g_2(y) = j_2$, then $x > y$ in the extended ordering. In other words, agents corresponding to those formulas in S have precedence over those not. Let O be any total ordering containing such extended partial ordering, then $O \in \mathcal{TO}_Q$. By the definition of the preferred subtheory, we have $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [O]\varphi \equiv [G_S]\varphi$, where $G_S = \{f(i,j) \mid \varphi_{ij} \in S\}$. From the definition of Σ_T, it follows that $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [G_S]\varphi$ iff $S \models \varphi$ for any classical formula φ.

(2) Let us denote φ_{ij} by φ_m if $f(i,j) = m$ and without loss of generality, we can assume that $O = 1 > 2 > \cdots > n$, then we can define $S^1 = \{\varphi_1\}$ and

$$S^{i+1} = \begin{cases} S^i \cup \{\varphi_{i+1}\} & \text{if } S^i \cup \{\varphi_{i+1}\} \text{ isconsistent,} \\ S^i & \text{otherwise} \end{cases}$$

Let $S = S^n$ and $G = \{m \mid \varphi_m \in S\}$, then it is obvious that $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [O]\varphi \equiv [G]\varphi$. Since O respects the preference between different T_is, it can be seen that S is a preferred subtheory of T by the construction. By the definition of Σ_T, $\Sigma_T \vdash_{\mathrm{DBF}_n^s} [G]\varphi$ iff $S \models \varphi$ for any classical formula φ.

A.6 Proof of Theorem 4.1

The proof of this theorem is based on that for inclusion modal logic in [Baldoni 1998]. We first recall some notations. Let S denote a set of tableau formulas (i.e., prefix formulas of the form $w : \varphi$ or accessibility relation formulas of the form $w\rho_G w'$) and P be the set of prefixes appearing in S. For a given DBF_n^c or DBF_n^s model $M = (W, (\mathcal{R}_i)_{1\le i\le n}, V)$, a prefix assignment is a mapping $I : P \to W$ from the set of prefixes to the set of possible worlds. The satisfaction relation between prefix assignments and tableau formulas with respect to the model M is defined by

(1) $I \Vdash_M w\rho_G w'$ iff $(I(w), I(w')) \in \cap_{i\in G}\mathcal{R}_i$
(2) $I \Vdash_M w : \varphi$ iff $I(w) \models_M \varphi$.

A set S is satisfiable if there exists some prefix assignment I such that I satisfies all tableau formulas in S. A branch of a tableau is satisfiable if the set of tableau formulas on it is satisfiable and a tableau is satisfiable if some of its branch is satisfiable. The following lemma shows that all the tableau rules preserve the satisfiability of a tableau.

LEMMA 7. *Let T be a satisfiable tableau and T' be the resultant tableau after the application of one tableau rule in figure 6, then T' is also satisfiable.*

PROOF. : The satisfiability of T means that some of its branches are satisfiable. If T' is obtained by applying some tableau rule to branches other than those that are satisfiable, then the satisfiable branch remains unchanged, so T' is satisfiable trivially. Therefore, we can assume that the rule is applied to a satisfiable branch. Let S be the set of tableau formulas on that branch and I be a prefix assignment satisfying the tableau formulas in S with respect to a model M, then the result follows from the following simple facts.

(1) Assume that the applied rule is one of the classical rules, ν rule, ρ rule, DBF_n^c rule, or DBF_n^s rule, then it can be shown that if I satisfies the premises of these rules, then I also satisfies at least one branch of the conclusions. This can be easily proven by the semantics of the logic and the definition of the satisfaction of tableau formulas.

(2) Assume that the applied rule is the π rule. If $I \models_M w : \nu^G$, then by definition, $I(w) \models_M \nu^G$, which by the semantics, implies that there exists some possible world u such that $(I(w), u) \in \cap_{i \in G} \mathcal{R}_i$ and $u \models_M \nu_0^G$. Let I' be agree with I in all prefixed appearing in S and $I'(w') = u$, then I' satisfies all formulas in the extended branch $S \cup \{w \rho_G w'\} \cup \{w' : \nu_0^G\}$ with respect to M, so T' is still satisfiable.

(3) If the rule D is applied, then by the seriality of the model M and the occurrence of w in S, there exists some possible world u such that $(I(w), u) \in \mathcal{R}_i$. Therefore, the extended branch $S \cup \{w \rho_i w'\}$ is satisfied by I' defined as above.

□

The soundness of the tableau calculus then follows immediately from lemma 7.

LEMMA 8. *Let L denote DBF_n^c or DBF_n^s and φ be a wff in L, then $\Vdash_L \varphi$ implies $\models_L \varphi$.*

PROOF. First, it is noted that a closed tableau is not satisfiable. If $\not\models_L \varphi$, then there exists an L model $M = (W, (\mathcal{R}_i)_{1 \le i \le n}, V)$ and $u \in W$ such that $u \models_M \neg\varphi$, so the initial tableau that consists of only the root node $w : \neg\varphi$ is satisfiable by the prefix assignment I with $I(w) = u$. Since φ is tableau provable in L, a closed tableau will result from that initial tableau by repeated application of the tableau rules. However, by the preceding lemma, any tableaux obtained in this way, including the closed one, should also be satisfiable. This is a contradiction, so $\not\models_L \varphi$ is impossible if φ is tableau provable in L. □

To prove the completeness of the tableau calculus, we assume that the tableau starting with the formula $w : \neg\varphi$ is not closed, so there is at least an open branch on which no more rules can be applied. The open branch may be finite or infinite. Let S denote the tableau formulas appearing on this open branch. We first show that the set S contains "enough information" for constructing a counter-model of φ.

DEFINITION A.1. *A set S of tableau formulas is said to be downward saturated if it satisfies the following conditions:*

(1) there does not exists any atomic wff p and prefix w such that $w : p \in S$ and $w : \neg p \in S$;

(2) if $\neg\neg\psi \in S$, then $\psi \in S$;

(3) if $w : \alpha \in S$, then $w : \alpha_1 \in S$ and $w : \alpha_2 \in S$;

(4) if $w : \beta \in S$, then $w : \beta_1 \in S$ or $w : \beta_2 \in S$;

(5) if $w : \nu^G \in S$, then $w' : \nu_0^G \in S$ for all w' such that $w \rho_G w' \in S$;

(6) if $w : \pi^G \in S$, then $w' : \nu_0^G \in S$ for some w' such that $w \rho_G w' \in S$;

(7) *for each w, if there exists ν^i such that $w : \nu^i \in S$, then there exists G and w' such that $i \in G$ and $w\rho_G w' \in S$;*

(8) *if $w\rho_{G_1} w' \in S$, then $w\rho_{G_2} w' \in S$ for all $G_2 \subseteq G_1$.*

DEFINITION A.2. *A downward saturated set S of tableau formulas is DBF_n^{c}-saturated if it in addition satisfies the following two conditions:*

(1) *if $w : [O > i]\psi \in S$, then both $w : [\delta(O > i)]\bot$ and $w : [O]\psi \in S$ or both $w : \neg[\delta(O > i)]\bot$ and $w : [\delta(O > i)]\psi \in S$*

(2) *if $w : \neg[O > i]\psi \in S$, then both $w : [\delta(O > i)]\bot$ and $w : \neg[O]\psi \in S$ or both $w : \neg[\delta(O > i)]\bot$ and $w : \neg[\delta(O > i)]\psi \in S$*

DEFINITION A.3. *A downward saturated set S of tableau formulas is DBF_n^{s}-saturated if it in addition satisfies the following two conditions:*

(1) *if $w : [\Omega \cup \{O > i\}]\psi \in S$, then both $w : [\{O, i\}]\bot$ and $w : [\Omega \cup \{O\}]\psi \in S$ or both $w : \neg[\{O, i\}]\bot$ and $w : [\Omega \cup \{O, i\}]\psi \in S$*

(2) *if $w : \neg[\Omega \cup \{O > i\}]\psi \in S$, then both $w : [\{O, i\}]\bot$ and $w : \neg[\Omega \cup \{O\}]\psi \in S$ or both $w : \neg[\{O, i\}]\bot$ and $w : \neg[\Omega \cup \{O, i\}]\psi \in S$*

From these definitions, we immediately have the following lemma.

LEMMA 9. *Let L denote DBF_n^{c} or DBF_n^{c}, and S be the set of all tableau formulas appearing on an open branch of the L-tableau proof construction for φ, then S is L-saturated.*

The saturated sets of tableau formulas will play the role of the MCSs in the proof of the completeness of the axiomatic systems, so the canonical model construction is based on these saturated sets. In the construction, the prefixes occurring in a saturated set S are essentially the possible worlds of the canonical model. However, since the application of rule D is restricted, some possible worlds may not have $\mathcal{R}_i$-successors. To circumvent the problem, we introduce some dummy worlds. First, a prefix w occurring in S is called an i-dead end if there does not exist any G such that $i \in G$ and $w\rho_G w'$ for some w' occurring in S. For each $1 \leq i \leq n$, if w is an i-dead end in S, then a new dummy world $d_{w,i}$ is introduced. Let $\mathcal{P}$ denote the set of prefixes occurring in S and $\mathcal{D}$ denote the set of dummy worlds introduced in this way, then the canonical model is

$$M_S = (W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$$

where

—$W = \mathcal{P} \cup \mathcal{D}$

—For each i, $\mathcal{R}_i$ is defined by the following three conditions:
 —if $w, w' \in \mathcal{P}$, then $\mathcal{R}_i(w, w')$ iff there is some G such that $i \in G$ and $w\rho_G w' \in S$
 —if $w \in \mathcal{P}$ and $d \in \mathcal{D}$, then $\mathcal{R}_i(w, d)$ iff $d = d_{w,i}$,
 —if $d, d' \in \mathcal{D}$, then $\mathcal{R}_i(d, d')$ iff $d = d'$

—$V : \Phi_0 \to 2^W$ is defined by $V(p) = \{w \in \mathcal{P} \mid w : p \in S\}$

Note that each $\mathcal{R}_i$ is serial in the canonical model. Furthermore, let us define $\mathcal{R}_G = \cap_{i \in G} \mathcal{R}_i$ for any subset of agents G, then the following result can be derived:

LEMMA 10. *For all prefixes $w, w' \in \mathcal{P}$, $w\rho_G w' \in S$ iff $\mathcal{R}_G(w, w')$.*

PROOF. . (⇒): If $w\rho_G w' \in S$, then by the first condition for the construction of accessibility relations, $\mathcal{R}_i(w, w')$ holds for all $i \in G$, so $\mathcal{R}_G(w, w')$ holds, too.

. (⇐): If $\mathcal{R}_i(w, w')$ holds for all $i \in G$, then by construction, there exists G_i such that $i \in G_i$ and $w\rho_{G_i} w' \in S$ for each $i \in G$. There are two possible ways by which the prefix w' is introduced into S, i.e., by applications of the π rule or the D rule. In both cases, there is an accessibility relation formula $w\rho_{G'} w'$ that is introduced into S at the same time. Then by the applicability condition of the ρ rule, $G_i \subseteq G'$ for all $i \in G$, so $G \subseteq G'$. Therefore, by the saturation of S, $w\rho_G w' \in S$ since $w\rho_{G'} w'$ has been introduced into S.

□

The truth lemma for the canonical model construction is as follows:

LEMMA 11. *Let L denote DBF_n^c or DBF_n^c, and S be the set of all tableau formulas appearing on an open branch of the L-tableau proof construction, then for any wff φ in the logic L, $w : \varphi \in S$ implies $w \models_{M_S} \varphi$.*

PROOF. As usual, the lemma can be proven by induction on the complexity of the wff.

The basis cases: if $w : p \in S$, then by the construction of V, $w \models_{M_S} p$ and if $w : \neg p \in S$, then by the saturation of S, $w : p \notin S$, so by the construction of V, $w \models_{M_S} \neg p$.

The inductive step is considered case by case.

(1) The classical cases: if $w : \neg\neg\psi \in S$, $w : \alpha \in S$, or $w : \beta \in S$, then it can be easily proven in the usual way.

(2) The π case: if $w : \pi^G \in S$, then by the saturation of S, there exists a prefix w' such that $w' : \pi_0^G \in S$ and $w\rho_G w' \in S$. By lemma 10, $\mathcal{R}_G(w, w')$, and by the inductive assumption, $w' \models_{M_S} \pi_0^G$, so $w \models_{M_S} \pi^G$ according to the semantics.

(3) The ν case: This is a more complicated case. We first note that if $i \neq j$ and w is both i-dead end and j-dead end, then $d_{w,i} \neq d_{w,j}$. Therefore, by the definition of $\mathcal{R}_G$, if $|G| > 1$, then for all $w \in \mathcal{P}$ and $d \in \mathcal{D}$, $\neg\mathcal{R}_G(w, d)$ holds. Furthermore, if w is not an i-dead end, then $\neg\mathcal{R}_i(w, d)$ holds for all $d \in \mathcal{D}$. Assume $w : \nu^G \in S$ for some G such that $|G| > 1$, then for all w' such that $\mathcal{R}_G(w, w')$, we have $w\rho_G w' \in S$ by lemma 10, so by the saturation of S, $w' : \nu_0^G \in S$. Due to the inductive assumption, this means that for all w', if $\mathcal{R}_G(w, w')$, then $w' \models_{M_S} \nu_0^G$, so $w \models_{M_S} \nu^G$ by the semantics. For the case of $w : \nu^i \in S$, if w is not an i-dead end, then the proof is exactly as above. However, if w is an i-dead end, then by the saturation condition 7, it is impossible that there is any ν^i such that $w : \nu^i \in S$, so the ν case is done.

(4) The O case: there are two subcases depending on $w : [O]\psi \in S$ or $w : \neg[O]\psi \in S$. We will prove the first subcase and the second can be proven analogously. The proof is based on another induction on the cardinality of $\delta(O)$. If $|\delta(O)| = 1$, then the result has been proven in the ν case. Assume the result holds for all O such that $|\delta(O)| \leq k$ and let $O > i$ bea total order such that $|\delta(O)| = k$. Due to the DBF_n^c saturation of S, if $w : [O > i]\psi \in S$, then either $w : [\delta(O >$

$i)]\bot, w : [O]\psi \in S$ or $w : \neg[\delta(O > i)]\bot, w : [\delta(O > i)]\psi \in S$. For the former case, $w \models_{M_S} [\delta(O > i)]\bot$ by the ν case and $w \models_{M_S} [O]\psi$ by the inductive assumption. By the semantics, this means that $\cap_{j\in\delta(O>i)}\mathcal{R}_j(w) = \emptyset$ and for all $w' \in \mathcal{R}_O(w)$, $w' \models_{M_S} \psi$. Therefore, $w \models_{M_S} [O > i]\psi$, since $\mathcal{R}_{O>i}(w) = \mathcal{R}_O(w)$ if $\cap_{j\in\delta(O>i)}\mathcal{R}_j(w) = \emptyset$. For the latter case, $w \models_{M_S} \neg[\delta(O > i)]\bot$ by the π case, so it follows $\cap_{j\in\delta(O>i)}\mathcal{R}_j(w) \neq \emptyset$. This implies $\mathcal{R}_{O>i}(w) = \mathcal{R}_{\delta(O>i)}(w)$, so $w \models_{M_S} [O > i]\psi$ since $w \models_{M_S} [\delta(O > i)]\psi$ by the ν case.

(5) The Ω case: there are also two subcases depending on $w : [\Omega]\psi \in S$ or $w : \neg[\Omega]\psi \in S$. We still prove only the first subcase. The proof is based on an induction on the lexicographical ordering $>_{lex}$ over the pairs $(l(\Omega), \sharp(\Omega))$ of levels and lengths of Ω. For the basis case, if $l(\Omega) = 1$, then Ω is simply a subset of agents, so the result holds by the ν case. Let $\Omega' = \Omega \cup \{O > i\}$ be a subset of total orders such that $|\delta(O > i)| = l(\Omega)$, then it can be shown that

$$(l(\Omega'), \sharp(\Omega')) >_{lex} (l(\Omega \cup \{O, i\}), \sharp(\Omega \cup \{O, i\})),$$

$$(l(\Omega'), \sharp(\Omega')) >_{lex} (l(\Omega \cup \{O\}), \sharp(\Omega \cup \{O\})),$$

$$(l(\Omega'), \sharp(\Omega')) >_{lex} (l(\{O, i\}), \sharp(\{O, i\})).$$

By the inductive assumption and DBF_n^s saturation of S, if $w : [\Omega']\psi \in S$, then either $w \models_{M_S} [\{O, i\}]\bot$ and $w \models_{M_S} [\Omega \cup \{O\}]\psi$ or $w \models_{M_S} \neg[\{O, i\}]\bot$ and $w \models_{M_S} [\Omega \cup \{O, i\}]\psi$. The former case implies $\mathcal{R}_O(w) \cap \mathcal{R}_i(w) = \emptyset$, so $\mathcal{R}_{\Omega'}(w) = \mathcal{R}_\Omega(w) \cap \mathcal{R}_{O>i}(w) = \mathcal{R}_\Omega(w) \cap \mathcal{R}_O(w)$. Therefore, $w \models_{M_S} [\Omega \cup \{O\}]\psi$ implies $w \models_{M_S} [\Omega']\psi$. For the latter case, $w \models_{M_S} \neg[\{O, i\}]\bot$ implies $\mathcal{R}_O(w)\cap\mathcal{R}_i(w) \neq \emptyset$, so $\mathcal{R}_{\Omega'}(w) = \mathcal{R}_\Omega(w)\cap\mathcal{R}_{O>i}(w) = \mathcal{R}_\Omega(w)\cap\mathcal{R}_O(w)\cap\mathcal{R}_i(w)$. Therefore, $w \models_{M_S} [\Omega \cup \{O, i\}]\psi$ implies $w \models_{M_S} [\Omega']\psi$.

□

We can now finish the proof of the completeness theorem by using the truth lemma.

Lemma 12. *Let L denote DBF_n^c or DBF_n^s and φ be a wff in L, then $\not\Vdash_L \varphi$ implies $\not\models_L \varphi$.*

Proof. $\not\Vdash_L \varphi$ means that there is an open branch in the L tableau tree starting with the root $w : \neg\varphi$. Let S be the set of all tableau formulas appearing on the branch, then $w : \neg\varphi \in S$ since it is the root of the tree. By the truth lemma, $w \not\models_{M_S} \varphi$, so $\not\models_L \varphi$ since the canonical model M_S is an L model. □

Theorem 4.1 then follows immediately from the combination of lemma 8 and lemma 12.

A.7 Proof of Proposition 4

The macro DBF_n^s rule can be derived by repeat applications of the DBF_n^s rules. The point is that some intermediate formulas, such as $[\{O, i\}]\bot$ and $\neg[\{O, i\}]\bot$, are omitted in the conclusions of the macro DBF_n^s rule. The formal derivation can be done by induction on the $|\delta(O)|$.

(1) The base case: if $|\delta(O)| = 1$, the result holds trivially.

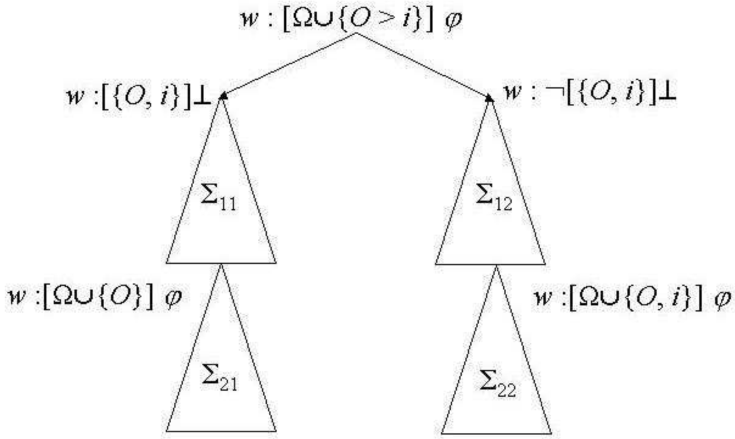

Fig. 9. Repeated application of DBF_n^s rules for the derivation of macro DBF_n^s rules

(2) The inductive assumption: assume the macro DBF_n^s rule is derivable when $|\delta(O)| = k$.

(3) The inductive step: assume $|\delta(O)| = k$ and consider the ordering $O > i$. For each subset G of $\delta(O) \cup \{i\}$ containing the first element of O, we have two cases:

(a) $i \in G$: then $G = G' \cup \{i\}$ where $G' \subseteq \delta(O)$. By the inductive assumption, we can have a branch in the right subtree of figure 9 that contains the results of applying the macro DBF_n^s rule to $w : \neg[\{O, i\}]\bot$ and $w : [\Omega \cup \{O, i\}]\varphi$ for the subset G'. That is,

$$\Sigma_{12} = \{w : \neg[G]\bot\} \cup \{w : \psi \mid \psi \in \Sigma_{G'}^{O}\}$$

$$\Sigma_{22} = \{w : [\Omega \cup G]\varphi\} \cup \{w : \psi \mid \psi \in \Sigma_{G'}^{O}\}.$$

Therefore,

$$\{w : [\Omega \cup G]\varphi\} \cup \{w : \psi \mid \psi \in \Sigma_G^{O>i}\} \subseteq \Sigma_{12} \cup \Sigma_{22}.$$

(b) $i \notin G$, so $G \subseteq \delta(O)$. By inductive assumption (see the left subtree of figure 9), we can show that

$$\Sigma_{11} = \{w : [G \cup \{i\}]\bot\} \cup \{w : \psi \mid \psi \in \Sigma_G^{O}\}$$

$$\Sigma_{21} = \{w : [\Omega \cup G]\varphi\} \cup \{w : \psi \mid \psi \in \Sigma_G^{O}\}$$

and

$$\{w : [\Omega \cup G]\varphi\} \cup \{w : \psi \mid \psi \in \Sigma_G^{O>i}\} \subseteq \Sigma_{11} \cup \Sigma_{21}.$$

A.8 Proof of Theorem 4.2

First, we prove the following lemma.

LEMMA 13. *The satisfiability problem for DBF_n^c and DBF_n^s is PSPACE-hard.*

Proof. The proof can be achieved by reducing the satisfiability problem for propositional modal logic KD to that for DBF_n^c and DBF_n^s. Let Φ_0 be the set of atomic propositions, then the set of well-formed formulas(wff) for KD is the least set containing Φ_0 and closed under Boolean connectives and the unary modal operator $\Box$. For the semantics, a KD model is a triple $(W, \mathcal{R}, V)$, where W is a set of possible worlds, $\mathcal{R}$ is a serial binary relation over W, and V is atruth assignment, and the main clause for the satisfaction are:

—$w \models \Box\varphi$ iff for all $u \in \mathcal{R}(w)$, $u \models \varphi$.

Obviously, we can translate wffs of KD into those of DBF_n^c or DBF_n^s. The translation τ is defined by

(1) $\tau(p) = p$ if $p \in \Phi_0$,
(2) $\tau(\neg\varphi) = \neg\tau(\varphi)$ and $\tau(\varphi \vee \psi) = \tau(\varphi) \vee \tau(\psi)$,
(3) $\tau(\Box\varphi) = [1]\tau(\varphi)$.

By the semantics, we have $\models_{\mathrm{KD}} \varphi$ iff $\models_{\mathrm{DBF}_n^c} \tau(\varphi)$ iff $\models_{\mathrm{DBF}_n^s} \tau(\varphi)$ for any wff φ of KD. Therefore, the PSPACE-hardness of the satisfiability problem for DBF_n^c and DBF_n^s follows immediately from the PSPACE-completeness of KD, which has been shown in [Ladner 1977]. □

Second, we prove that the satisfiability problem for DBF_n^c and DBF_n^s is in PSPACE.

Lemma 14. *There is an algorithm for deciding satisfiability of DBF_n^c and DBF_n^s wffs that needs polynomial space.*

Proof. The algorithm is a slightly modified version of the K_n tableau construction procedure in [Halpern and Moses 1992] and is presented in figure 10. The main modification is to add the applications of DBF_n^c or DBF_n^s rules to step 2(a) and the application of rule D to step 2(c) in that procedure. A set Σ of wff is called *blatantly inconsistent* if for some formula ψ, both ψ and $\neg\psi$ is in Σ or $\bot \in \Sigma$. We assume an arbitrary enumeration of the wffs is given, so in step 2(a), we can find the least witness to the fact that a set of wffs is not closed under all of classical and DBF_n^c (or DBF_n^s) rules. □

Then all proofs that the algorithm for deciding the satisfiability of K_n wffs need polynomial space can be carried out in the cases of DBF_n^c and DBF_n^s by our definitions of $Sub_c^+(\varphi)$ and $Sub_s^+(\varphi)$ (possibly with the difference of a constant factor). The main results are summarized as follows:

Lemma 15.

(1) For all wffs of size m, the DBF_n^c or DBF_n^s pre-tableau construction procedure terminates and the final tree constructed in the procedure has height at most $O(m^2)$.

(2) A wff φ is DBF_n^c(or DBF_n^s) satisfiable iff the DBF_n^c(or DBF_n^s) pre-tableau construction for φ returns "φ is satisfiable".

(3) The DBF_n^c or DBF_n^s pre-tableau construction procedure needs $O(m^3)$ space by using of depth-first search.

(1) Construct a tree consisting of a single node s_0(the "root"), with $L(s_0) = \{\varphi_0\}$.

(2) Repeat until none of (a)-(d) below applies:

(a) If s is a leaf of the tree, $L(s)$ is not blatantly inconsistent, $L(s)$ is not closed under all of α, $\neg\neg$, β, and DBF_n^c (or DBF_n^s) rules, and ψ is the least witness to this fact, then:

i. if ψ is of the form $\neg\neg\psi'$, then create a child s' of s and set $L(s') = L(s) \cup \{\psi'\}$,

ii. if ψ is an α formula, then create a child s' of s and set $L(s') = L(s) \cup \{\alpha_1, \alpha_2\}$,

iii. if ψ is a β formula, then create two children s_1 and s_2 of s and set $L(s_i) = L(s) \cup \{\beta_i\}$, $i = 1, 2$,

iv. if ψ is of the form $[O > i]\psi$ (resp. $\neg[O > i]\psi$) in DBF_n^c, then create two children s_1 and s_2 of s and set

$$L(s_1) = L(s) \cup \{[\delta(O > i)]\bot, [O]\psi \text{ (resp } \neg[O]\psi)\}$$

$$L(s_2) = L(s) \cup \{\neg[\delta(O > i)]\bot, [\delta(O > i)]\psi \text{ (resp } \neg[\delta(O > i)]\psi)\}$$

v. if ψ is of the form $[\Omega \cup \{O > i\}]\psi$ (resp. $\neg[\Omega \cup \{O > i\}]\psi$) in DBF_n^s, then create two children s_1 and s_2 of s and set

$$L(s_1) = L(s) \cup \{[\{O, i\}]\bot, [\Omega \cup \{O\}]\psi \text{ (resp } \neg[\Omega \cup \{O\}]\psi)\}$$

$$L(s_2) = L(s) \cup \{\neg[\{O, i\}]\bot, [\Omega \cup \{O, i\}]\psi \text{ (resp } \neg[\Omega \cup \{O, i\}]\psi)\}$$

(b) If s is a leaf of the tree, $L(s)$ is not blatantly inconsistent, and $L(s)$ is closed under all classical and DBF_n^c (or DBF_n^s) rules, then for each π^G formula in $L(s)$, create a G successor node s' for s and set

$$L(s') = \{\pi_0^G\} \cup \{\psi \mid [G']\psi \in L(s), G' \subseteq G\}$$

(c) If s is a leaf of the tree, $L(s)$ is not blatantly inconsistent, $L(s)$ is closed under all classical and DBF_n^c (or DBF_n^s) rules, and there exists $[i]\psi \in L(s)$ and does not exist any π^G formula in $L(s)$ such that $i \in G$, then create an i successor node s' for s and set

$$L(s') = \{\psi \mid [i]\psi \in L(s)\}$$

(d) If s is not marked "satisfiable", then mark s satisfiable if either

i. $L(s)$ is not closed under all of classical and DBF_n^c (or DBF_n^s) rules and s' is marked "satisfiable" for some child s' of s,

ii. $L(s)$ is closed under all of classical and DBF_n^c (or DBF_n^s) rules, there are no π^G formulas or formulas of the form $[i]\psi \in L(s)$, and $L(s)$ is not blatantly inconsistent, or

iii. $L(s)$ is closed under all of classical and DBF_n^c (or DBF_n^s) rules, s has successors, and all of them are marked "satisfiable".

(3) If the root of the tree is marked "satisfiable", then return "φ_0 is satisfiable"; otherwise return "φ_0 is unsatisfiable".

Fig. 10. The DBF_n^c or DBF_n^s pre-tableau construction for φ_0

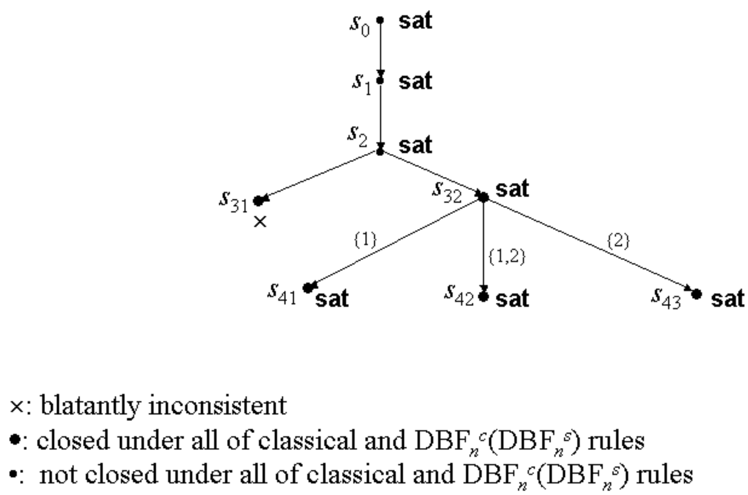

Fig. 11. The pre-tableau for $(\neg[1]p \wedge \neg[2]p) \wedge [1 > 2]p$

Example 11. We present a simple example to illustrate the pre-tableau construction procedure. Let us consider the formula

$$\varphi_0 = (\neg[1]p \wedge \neg[2]p) \wedge [1 > 2]p$$

then our procedure will construct a tree in figure 11. In that tree, $L(s_0) = \{\varphi_0\}$,

$$L(s_1) = L(s_0) \cup \{\neg[1]p \wedge \neg[2]p), [1 > 2]p\}$$

$$L(s_2) = L(s_1) \cup \{\neg[1]p, \neg[2]p\}$$

$$L(s_{31}) = L(s_2) \cup \{[\{1,2\}]\bot, [1]p\}$$

which is blatantly inconsistent,

$$L(s_{32}) = L(s_2) \cup \{\neg[\{1,2\}]\bot, [\{1,2\}]p\}$$

which has three successor nodes labelling respectively by $\{1\}$, $\{1,2\}$, and $\{2\}$,

$$L(s_{41}) = \{\neg p\}, L(s_{42}) = \{\neg\bot, p\}, L(s_{43}) = \{\neg p\}$$

which are all closed under classical and DBF_n^c (or DBF_n^s) rules. According to (d).ii, the three nodes s_{41}, s_{42}, and s_{43} should be marked "satisfiable". Then by (d).iii, s_{32} is marked "satisfiable" and by (d).i, all of nodes from s_0 to s_2 are marked "satisfiable". Therefore, the procedure returns "φ_0 is satisfiable". □

A.9 Proof of Proposition 5

Let $\mathcal{I}$ denote the set of all propositional interpretations over Φ_0 and assign to each interpretation $I \in \mathcal{I}$ a possible world w_I, then $M = (W, (\mathcal{R}_i)_{1 \leq i \leq n}, V)$ is such that

—$W = \{w_0\} \cup \{w_I \mid I \in \mathcal{I}\}$,
—$\mathcal{R}_i(w, w')$ iff $w = w_0$ and $w' \in \{w_I \mid I \in Mod(T_i)\}$,

—$V(p) = \{w_I \mid I(p) = 1, I \in \mathcal{I}\}$ for all $p \in \Phi_0$.

Then by the semantics of $\mathcal{L}_e$, $w \models_M \Sigma_{\mathcal{T}}$ and

$$\mathcal{R}_{\{1,\cdots,n\},\boldsymbol{\kappa}}(w_0) = \{w_I \mid I \in \mathcal{M}\}$$

since $\mathcal{R}_i(w_0) = \{w_I \mid I \in \mathcal{M}_i\}$ for all $1 \leq i \leq n$.

Therefore, for any propositional formula φ,

$$w_0 \models_M [\Gamma(\{1,\cdots,n\},\boldsymbol{\kappa})]\varphi \text{ iff } \Gamma(T_1,\cdots,T_n,\boldsymbol{\kappa}) \models \varphi.$$

ACKNOWLEDGMENTS

We would like to thank three anonymous referees for their helpful comments.

REFERENCES

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

Antoniou, G. 1997. *Nonmonotonic Reasoning.* MIT Press, Cambridge, Massachusetts.

Baldoni, M. 1998. Normal multimodal logics: Automatic deduction and logic programming extension. Ph.D. thesis, Dipartimento di Informatica—Universitá degli Studi di Torino.

Baldoni, M., Giordano, L., and Martelli, A. 1998. “A tableau calculus for multimodal logics and some (un)decidability results”. In *Proc. of the International Conference on Analytic Tableaux and Related Methods*, H. de Swart, Ed. LNAI 1397. Springer-Verlag, 44–59.

Baral, C., Kraus, S., Minker, J., and Subrahmanian, V. S. 1992. “Combining knowledge bases consisting of first-order theories”. *Computational Intelligence 8,* 1, 45–71.

Baral, C., Kraus, S., Minker, J., and Subrahmanian, V. S. 1994. “Combining default logic databases”. *Internationa Journal of Intelligent and Cooperative Information Systems 3,* 3, 319–348.

Baral, C., Minker, J., and Kraus, S. 1991. “Combining multiple knowledge bases”. *IEEE Transactions on Knowledge and Data Engineering 3,* 2, 208–221.

Benferhat, S., Dubois, D., and Prade, H. 1997. “From semantic to syntactic approaches to information combination in possibilistic logic”. In *Aggregation and Fusion of Imperfect Information*, B. Bouchon-Meunier, Ed. Physica-Verlag, 141–161.

Benferhat, S. and Garcia, L. 1998. “A local handling of inconsistent knowledge and default bases”. In *Applications of Uncertainty Formalisms*, A. Hunter and S. Parsons, Eds. LNAI 1455. Springer-Verlag, 325–353.

Boutilier, C. 1993. “Revision sequences and nested conditionals”. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. 519–525.

Boutilier, C. 1995. “Generalized update: Belief change in dynamic settings”. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*. 1550–1556.

Brewka, G. 1991. *Nonmonotonic Reasoning: Logical Foundations of Commonsense.* Cambridge University Press.

Chellas, B. 1980. *Modal Logic: An Introduction*. Cambridge University Press.

Cholvy, L. 1994. “A logiccal approach to multi-souces reasoning”. In *Knowledge Representation and Reasoning under Uncertainty*, M. Masuch and L. Pólos, Eds. LNCS 808. Springer-Verlag, 183–196.

Cholvy, L. and Cuppens, F. 1999. “Reasoning about norms provided by conflicting regulations ”. In *Norms, Logics and Information Systems*, P. McNamara and H. Prakken, Eds. IOS Press, 247–262.

Cholvy, L. and Hunter, A. 1997. “Information fusion in logic: A brief overview”. In *Qualitative and Quantitative Practical Reasoning(ECSQARU'97/FAPR'97).* LNAI 1244. Springer-Verlag, 86–95.

Darwiche, A. and Pearl, J. 1997. “On the logic of iterated belief revision”. *Artificial Intelligence 89,* 1, 1–29.

Demolombe, R. and Liau, C. J. 2001. "A logic of graded trust and belief fusion". In *Proceedings of the 4th Workshop on Deception, Fraud and Trust in Agent Societies.* 13–25.

Dennett, D. 1987. *The Intentional Stance.* MIT Press, Cambridge, MA.

Dragoni, A. 1992. "A model for belief revision in a multi-agent environment". In *Decentralized A.I. 3. Proceedings of the 3rd European Workshop on Modeling Autonomous Agents in a Multi-Agent World*, Y. Demazeau and E. Werner, Eds. Elsevier Science Publisher, 103–112.

Dragoni, A. and Giorgini, P. 2001. "Revisining beliefs received from multiple sources". In *Frontiers in Belief Revision*, H. Roth and M. Williams, Eds. Kluwer Academic Publisher, 431–444.

Dragoni, A., Giorgini, P., and Baffetti, M. 1997. "Distributed belief revision vs. belief revision in a multi-agent environment: First results of a simulation experiment". In *Multi-Agent Rationality: Proceedings of the 8th European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, M. Boman and W. V. de Velde, Eds. LNAI 1237. Springer-Verlag, 45–62.

Dubois, D., Lang, J., and Prade, H. 1991. "Fuzzy sets in approximate reasoning, Part 2: logical approaches". *Fuzzy Sets and Systems 40*, 203–244.

Dubois, D., Lang, J., and Prade, H. 1994. "Possibilistic logic". In *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol 3 : Nonmonotonic Reasoning and Uncertain Reasoning*, D. Gabbay, C. Hogger, and J. Robinson, Eds. Clarendon Press - Oxford, 439–513.

Dubois, D. and Prade, H. 1988. "An introduction to possibilistic and fuzzy logics". In *Non-Standard Logics for Automated Reasoning*, P. Smets, A. Mamdani, D. Dubois, and H. Prade, Eds. Academic Press, 253–286.

Dubois, D. and Prade, H. 1991. "Epsitemic entrechment and possibilistic logic". *Artificial Intelligence 50*, 223–239.

Dubois, D. and Prade, H. 1992. "Belief change and possibility theory". In *Belief Revision*, P. Gärdenfors, Ed. Cambridge University Press, 142–182.

Dubois, D. and Prade, H. 2000. "Possibility theory in information fusion". In *Proc. of the Third International Conference on Information Fusion.* TuA–1.

Fagin, R., Halpern, J., Moses, Y., and Vardi, M. 1996. *Reasoning about Knowledge.* MIT Press.

Fagin, R., Halpern, J., and Vardi, M. 1992. "What can machines know? On the properties of knowledge in distributed systems". *JACM 39,* 2, 328–376.

Fitting, M. 1983. *Proof Methods for Modal and Intuitionistic Logics.* D. Reidel Publishing Company.

Gabbay, D. 1996. *Labelled Deductive Systems.* Oxford UNiversity Press.

Gabbay, D. and Guenthner, F., Eds. 1984. *Handbook of Philosophical Logic, Vol II: Extensions of Classical Logic.* D. Reidel Publishing Co., Dordrecht, The Netherlands.

Galliers, J. 1992. "Autonomous belief revision and communication". In *Belief Revision*, P. Gärdenfors, Ed. Cambridge University Press, 220–246.

Halpern, J. and Moses, Y. 1992. "A guide to completeness and complexity for modal logics of knowledge and belief". *Artificial Intelligence 54*, 311–379.

Hempel, C. 1965. "Studies in the logic of confirmation". In *Aspects of Scienfific Explanation and Other Essays in the Philosophy of Science*, C. Hempel, Ed. Free Press, New York, 53–79.

Hintikka, J. 1962. *Knowledge and Belief.* Cornell University Press.

II, P. M.-R. and Lehmann, D. 2000. "Representing and aggregating conflicting beliefs". In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning.* 153–164.

II, P. M.-R. and Shoham, Y. 2001. "Belief fusion: Aggregating pedigreed belief states". *Journal of Logic, Language and Information 10,* 2, 183–209.

Katsuno, H. and Medelzon, A. 1991a. "On the difference between updating a knowledge base and revising it". In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91).* Morgan Kaufmann Publisher, 387–394.

Katsuno, H. and Medelzon, A. 1991b. "Propositional knowledge base revision and minimal change". *Artificial Intelligence 52*, 263–294.

KFIR-DAHAV, N. AND TENNENHOLTZ, M. 1996. "Multi-agent belief revision". In *Proceedings of the Sixth International Conference on Theoretical Aspects of Reasoning about Knowledge(TARK'96)*. Morgan Kaufmann Publisher, 175–194.

KONIECZNY, S. 2000. "On the difference between merging knowledge bases and combining them". In *Proceedings of the Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR'00)*. Morgan Kaufmann Publisher, 135–144.

KONIECZNY, S. AND PÉREZ, R. P. 1998. "On the logic of merging". In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*. Morgan Kaufmann Publisher, 488–498.

KONIECZNY, S. AND PÉREZ, R. P. 1999. "Merging with integrity constraints". In *Proccedings of the Fifth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'99)*, A. Hunter and S. Parsons, Eds. LNAI 1638. Springer-Verlag, 233–244.

KYBURG, H. 1961. *Probability and the Logic of Rational Belief.* Wesleyan University Press.

LADNER, R. 1977. "The computational complexity of provability in systems of modal propositional logic". *SIAM Journal of Computing 6*, 467–480.

LEHMANN, D. 1995. "Belief revision, revised". In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*. 1534–1540.

LIAU, C. J. 2000. "A conservative approach to distributed belief fusion". In *Proc. of the Third International Conference on Information Fusion*. MoD4–1.

LIAU, C. J. 2003. "Belief, information acquisition, and trust in multi agent systems—A modal logic formulation". *Artificial Intelligence*.

LIBERATORE, P. AND SCHAERF, M. 1995. "Arbitration: A commutative operator for belief revision". In *Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence (WOCFAI '95)*. 217–228.

LIN, J. 1994. "Information sharing and knowledge merging in cooperative information systems". In *Proceedings of the Fourth Workshop on Information Technologies and Systems*. 58–66.

LIN, J. 1996. "Integration of weighted knowledge bases". *Artificial Intelligence 83,* 2, 363–378.

LIN, J. AND MENDELZON, A. 1999. "Knowledge base merging by majority". In *Dynamic Worlds: From the Frame Problem to Knowledge Management*, R. Pareschi and B. Fronhoefer, Eds. Kluwer Academic Publisher.

LUCAS, P. 1997. "Symbolic diagnosis and its formalization". *The Knowledge Engineering Review 12,* 2, 109–146.

MEYER, J.-J. C. AND VAN DER HOEK, W. 1991. "Nonmonotonic reasoning by monotonic means". In *Logics in AI (JELIA'90)*, J. van Eijck, Ed. LNCS 478. Springer-Verlag, 399–411.

MEYER, J.-J. C. AND VAN DER HOEK, W. 1992. "A modal logic for nonmonotonic reasoning". In *Non-monotonic Reasoning and Partial Semantics*, W. van der Hoek, J.-J. C. Meye, Y. Tan, and C. Witteveen, Eds. Ellis Horwood, 37–77.

MEYER, J.-J. C. AND VAN DER HOEK, W. 1993. "A cumulative default logic based on epistemic states". In *Proccedings of the 2nd European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'93)*, M. Clarke, R. Kruse, and S. Moral, Eds. LNCS 747. Springer-Verlag, 265–273.

MEYER, J.-J. C. AND VAN DER HOEK, W. 1995. *Epistemic Logic for AI and Computer Science*. Cambridge University Press.

MEYER, J.-J. C. AND VAN DER HOEK, W. 1998. "Modal logic for representing incoherent knowledge". In *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Vol. 2*, D. Gabbay and P. Smets, Eds. Kluwer Academic Publisher, 37–75.

NEBEL, B. 1994. "Base revision operator and schemes: semantics representation and complexity". In *Proceedings of the 11th European Conference on Artificial Intelligence*. John Wiley & Sons, 341–345.

PRADHAN, S., MINKER, J., AND SUBRAHMANIAN, V. 1995. "Combining databases with prioritized information". *Journal of Intelligent Information Systems 4,* 3, 231–260.

REITER, R. 1987. "A theory of diagnosis from first principles". *Artificial Intelligence 32*, 57–95.

REVESZ, P. Z. 1993. "On the semantics of theory change: Arbitration between old and new information". In *Proceedings of the Twelfth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*. 71–82.

REVESZ, P. Z. 1997. "On the semantics of arbitration". *International Journal of Algebra and Computation 7,* 2, 133–160.

SEGERBERG, K. 1995. "Belief revision from the point of view of doxastic logic". *Bull. of the IGPL 3,* 4, 535–553.

SEGERBERG, K. 2001. "The basic dynamic doxastic logic of AGM". In *Frontiers in Belief Revision*, H. Roth and M. Williams, Eds. Kluwer Academic Publisher, 57–84.

SHOHAM, Y. 1993. "Agent-oriented programming". *Artificial Intelligence 60,* 1, 51–92.

SMETS, P. 2000. " Data fusion in the transferable belief model". In *Proc. of the Third International Conference on Information Fusion*. WeA–1.

SMULLYAN, R. 1968. *First-Order Logic*. Springer-Verlag.

SPOHN, W. 1988. "Ordinal conditional functions: a dynamic theory of epistemic states". In *Causation in Decision, Belief Change, and Statistics, II*, W. Harper and B. Skyrms, Eds. Kluwer Academic Publishers, 105–134.

SUBRAHMANIAN, V. 1994. "Amalgamating knowledge bases". *ACM Transactions on Database Systems 19,* 2, 291–331.

WILLIAMS, M. 1994. "Transmutations of knowledge systems". In *Proceedings of the 4th International Conference on Principle of Knowledge Representation and Reasoning*, J. Doyle, E. Sandewall, and P. Torasso, Eds. Morgan Kaufmann Publishers, 619–629.

WOOLDRIDGE, M. AND JENNINGS, N. 1995. Intelligent agents: theory and practice. *Knowledge Engineering Review 10(2)*, 115–152.

ZADEH, L. 1978. "Fuzzy sets as a basis for a theory of possibility". *Fuzzy Sets and Systems 1,* 1, 3–28.