Knowledge Compilation for Belief Change

Maurice Pagnucco

ARC Centre of Excellence for Autonomous Systems and National ICT Australia, School of Computer Science and Engineering, The University of New South Wales, Sydney, NSW 2052, Australia
morri@cse.unsw.edu.au
http://www.cse.unsw.edu.au/~morri/

Abstract. Techniques for knowledge compilation like prime implicates and binary decision diagrams (BDDs) are effective methods for improving the practical efficiency of reasoning tasks. In this paper we provide a construction for a belief contraction operator using prime implicates. We also briefly indicate how this technique can be used for belief expansion, belief revision and also iterated belief change. This simple yet novel technique has two significant features: (a) the contraction operator constructed satisfies all the AGM postulates for belief contraction; (b) when compilation has been effected only syntactic manipulation is required in order to contract the reasoner's belief state.

Keywords: knowledge representation and reasoning, belief revision and update, common-sense reasoning.

Transforming the sentences of a knowledge base KB into an alternate, regular form KB' has proven to be an effective technique in improving the efficiency of reasoning tasks. This idea is referred to as *knowledge compilation* and has been widely applied to problems in truth maintenance, constraint satisfaction, satisfiability, planning and scheduling, etc. Some of the more commonly used transformations for knowledge compilation include prime implicates and implicants, binary decision diagrams (BDDs) and decomposable negation normal form (see [2] for an introduction and analysis of these methods).

We shall concern ourselves with prime implicates and apply this compilation method to constructing an operator for belief change where the idea is to maintain the beliefs of a reasoner as new information is acquired. In this paper we show how to construct a belief contraction operator using prime implicates and also indicate how this can be used for belief expansion and belief revision as well. The elegance of the technique that we describe has two significant advantages: (i) once compilation has been effected only syntactic manipulation is required in order to contract the reasoner's belief state, and; (ii) we do not sacrifice theory for the sake of implementation like many syntactic methods but remain loyal to the AGM postulates for belief contraction. Furthermore, we briefly indicate how this technique can be used to implement iterated belief change.

Knowledge compilation techniques, including prime implicates have been used in implementing belief maintenance systems in the past. The closest work to our own is that of Gorogiannis and Ryan [8] who implement AGM belief change using BDDs, analysing the complexity of these operators in specific instances. However, their method

A. Sattar and B.H. Kang (Eds.): AI 2006, LNAI 4304, pp. 90–99, 2006.

does not give a natural way to achieve iterated belief change. Prime implicates have been suggested for truth maintenance in the work of Reiter and de Kleer [13] as well as that of Kean [11]. Due to the lack of preferential information, it is difficult to implement general forms of belief change and iterated belief change in the frameworks they develop. Implementations of AGM belief change such as those by Dixon and Wobcke [4] and Williams [18] rely on theorem proving which we try to do away with once the reasoner's epistemic stated has been compiled.

The rest of this paper is set out as follows. The next section covers the requisite logical background for prime implicates and belief change. Section 2 explains how prime implicates are utilised to represent the reasoner's epistemic state. The construction of belief contraction operators is described in Sections 3 and 4. Our conclusions together with a brief discussion of how this construction can be used for belief expansion, belief revision and iterated belief change are presented in Section 5.

1 Background

We assume a fixed finite object language $\mathcal{L}$ with standard connectives $\neg$, $\wedge$, $\vee$, $\rightarrow$, $\equiv$ as well as the logical constants $\top$ (truth) and $\bot$ (falsum). Cn represents the classical consequence operator while $\vdash$ represents the classical consequence relation (i.e., $\Gamma \vdash \phi$ iff $\phi \in Cn(\Gamma)$). We also adopt the following linguistic conventions to simplify the presentation. Lower case Greek letters ϕ, ψ, χ, ... denote sentences of $\mathcal{L}$. Upper case Greek letters Δ, Γ, ... denote sets of formulas. Theories (closed under logical consequence) are used to represent reasoners' belief states in AGM belief change and will be denoted by upper case Roman characters H, K, ... (i.e., $K = Cn(K)$). The inconsistent belief set is denoted $K_\bot$. Lower case Roman characters l, k, ... denote literals (both positive and negative). Upper case Roman characters C, D, ... denote clauses.

1.1 Prime Implicates

Prime implicates are minimal length clauses (in terms of set inclusion) implied by a knowledge base. Transforming a knowledge base into a set of prime implicates gives a uniform way of equivalently expressing the sentences in the knowledge base that can be exploited to enhance computational efficiency.

Definition 1. *A clause C is an* implicate *of a set of clauses Γ iff $\Gamma \vdash C$. A non-tautologous clause C is a* prime implicate *of a set of clauses Γ iff C is an implicate of Γ, and $\Gamma \nvdash D$ for any $D \vdash C$ and $D \neq C$.*

When Γ is a single clause, representing clauses as sets of literals allows us to simplify this definition by replacing $\vdash$ by $\subset$ (where $\subset$ denotes proper set inclusion). This will simplify the exposition below. We shall denote the set of prime implicates of an arbitrary set of sentences Γ by $\Pi(\Gamma)$ and note that these sets are logically equivalent.

Proposition 1. *Given a set of clauses Γ, $\Gamma \vdash \phi$ iff $\Pi(\Gamma) \vdash \phi$.*

When Γ is a set of clauses, prime implicates are easily computed by repeated application of resolution and removing subsumed (i.e., implied) clauses [17]. This basic strategy can be improved upon by using incremental methods [12,10] that do not require the implicates to be re-computed when new sentences are added to the knowledge base and through the use of efficient data structures [13] for subsumption checking. Compiling a knowledge base into prime implicate form can lead to an exponential number of implicates in the number of atoms (see [15]).

1.2 AGM Belief Change

Belief change takes the following model as its departure point. A reasoner is in some state of belief when it acquires new information. Belief change studies the manner in which this newly acquired information can be assimilated into its current state of belief in a rational way. We base our account here on the popular framework developed by Alchourrón, Gärdenfors and Makinson (AGM) [1,5,7]. In the AGM the reasoner's belief state is modelled as a deductively closed set of sentences referred to as a *belief set*. That is, for a belief set K, $K = Cn(K)$. In this way, the reasoner has one of three *attitudes* towards a sentence $\phi \in \mathcal{L}$: ϕ is *believed* when $\phi \in K$; ϕ is *disbelieved* when $\neg\phi \in K$; or the reasoner is *indifferent* to ϕ when $\phi, \neg\phi \notin K$. A belief set can be viewed as representing the sentences the reasoner is committed to believing regardless of whether that commitment can be feasibly realised. Shortly, when we discuss epistemic entrenchment, we will consider belief sets with additional preferential information about the reasoner's beliefs which we will refer to as the reasoner's *epistemic state*. These preferences over beliefs can be seen as offering a greater set of epistemic attitudes to the three outlined above.

The AGM considers three types of transformations on beliefs sets as new information is acquired: *belief expansion* by ϕ in which this new information is added to the initial belief set K without removal of any existing beliefs ($K + \phi$); *belief contraction* by ϕ where belief in ϕ is suspended ($K \dot{-} \phi$); and, *belief revision* by ϕ where ϕ is assimilated into K while existing beliefs may need to be suspended in order to maintain consistency ($K * \phi$). Here we are only concerned with belief contraction and so we will present and motivate the AGM rationality postulates for this operation.

(K$\dot{-}$1) For any sentence ϕ and any belief set K,
$K \dot{-} \phi$ is a belief set
(K$\dot{-}$2) $K \dot{-} \phi \subseteq K$
(K$\dot{-}$3) If $\phi \notin K$, then $K \dot{-} \phi = K$
(K$\dot{-}$4) If $\not\vdash \phi$ then $\phi \notin K \dot{-} \phi$
(K$\dot{-}$5) If $\phi \in K$, $K \subseteq (K \dot{-} \phi) + \phi$
(K$\dot{-}$6) If $\vdash \phi \equiv \psi$, then $K \dot{-} \phi = K \dot{-} \psi$
(K$\dot{-}$7) $K \dot{-} \phi \cap K \dot{-} \psi \subseteq K \dot{-} (\phi \wedge \psi)$
(K$\dot{-}$8) If $\phi \notin K \dot{-} (\phi \wedge \psi)$, then $K \dot{-} (\phi \wedge \psi) \subseteq K \dot{-} \phi$

Postulate (K$\dot{-}$1) stipulates that contraction of a belief set should return a new belief set. (K$\dot{-}$2) states that no beliefs should be added during contraction while (K$\dot{-}$3) says that the belief state is unaltered when there is no need to remove anything. (K$\dot{-}$4) ensures that the new information is removed where possible and (K$\dot{-}$5) that removal followed by addition of the same information results in no change. Postulate (K$\dot{-}$6) says that the syntax of the information is irrelevant. (K$\dot{-}$7) states that any beliefs removed in the contraction of $\phi \wedge \psi$ should be removed when contracting by ϕ alone or contracting by ψ

alone (or possibly both) while (K$\dot{-}$8) that if ϕ is removed when given a choice between removing ϕ or ψ, then whatever beliefs are removed along with ϕ are also removed along with $\phi \wedge \psi$. These fairly intuitive conditions are capable of characterising quite powerful forms of contraction.

1.3 Epistemic Entrenchment

It can be easily verified that the rationality postulates for AGM belief contraction describe a class of functions rather than prescribing a particular method. If we wish to focus on a particular function within this class of rational contractions we need to apply extralogical information (which is the whole point of the study of belief change). The two main constructions are a semantic one based on possible worlds due to Grove [9] and a syntactic one providing preferences over beliefs known as *epistemic entrenchment* due to Gärdenfors and Makinson [6]. Intuitively, epistemic entrenchment orders the reasoner's beliefs with tautologies being most strongly believed as these cannot be given up. An entrenchment relation $\phi \leq \psi$ says that ψ is at least as entrenched as ϕ and expresses the notion that when faced with the choice between giving up ϕ and giving up ψ, the reasoner would prefer to relinquish belief in ϕ. The formal properties of this ordering are:

(EE1) If $\phi \leq \psi$ and $\psi \leq \gamma$ then $\phi \leq \gamma$
(EE2) If $\{\phi\} \vdash \psi$ then $\phi \leq \psi$
(EE3) For any ϕ and ψ, $\phi \leq \phi \wedge \psi$ or $\psi \leq \phi \wedge \psi$
(EE4) When $K \neq K_\perp$, $\phi \notin K$ iff $\phi \leq \psi$ for all ψ
(EE5) If $\phi \leq \psi$ for all ϕ then $\vdash \psi$

(EE1) tells us that the relation is transitive. (EE2) states that consequences of belief are at least as entrenched since it is not necessary to give up (all) consequences in order to suspend judgement in the belief itself. Taken together with (EE1) and (EE2), property (EE3) enforces that a conjunction be entrenched at the same level as the least preferred of its conjuncts (i.e., $\phi \wedge \psi = min(\phi, \psi)$). These properties also allow us to state an important property regarding disjunctions—these are at least as entrenched as the most preferred of the two disjuncts (i.e., $\phi \vee \psi \geq max(\phi, \psi)$). This second property will be very important in what follows. Finally, (EE4) says that non-beliefs are least preferred and hence minimally entrenched while (EE5) says that tautologies are most preferred and therefore maximally entrenched.

Putting all of these properties together we see that an entrenchment ordering is a total pre-order—or ranking—of beliefs in which tautologies are most preferred and non-beliefs least preferred. Since the non-beliefs are clumped together as the minimal elements of an epistemic entrenchment ordering, this gives a convenient way of obtaining the belief state of the reasoner. We can simply obtain the belief state as the set of all non-minimal elements of the epistemic entrenchment: $K_{\leq} = \{\phi : \phi > \perp\}$. It is therefore common, particularly in iterated belief change [3], to refer to an epistemic entrenchment relation $\leq$ as the reasoner's *epistemic state* and we shall adopt this convention here. To be more specific, the technique we introduce in this paper is centred around compiling the reasoner's epistemic state (epistemic entrenchment relation) using prime implicates.

It remains to describe how we can effect belief contraction given an epistemic entrenchment ordering and vice versa. The first part is given by the condition [6]

$$(C\dot{-})\quad \psi \in K\dot{-}\phi \text{ iff } \psi \in K \text{ and either } \phi < \phi \vee \psi \text{ or } \vdash \phi$$

It states that a belief is retained after contraction when there is independent evidence for retaining it or it is not possible to effect the contraction[1]. There is independent evidence for retaining a belief ψ whenever the disjunction $\phi \vee \psi$ is strictly more entrenched than ϕ (the sentence we are trying to remove)[2]. Moving in the opposite direction and determining an epistemic entrenchment relation from a contraction function is specified by the following condition [6]:

$$(C \leq)\quad \phi \leq \psi \text{ iff } \phi \notin K\dot{-}(\phi \wedge \psi) \text{ or } \vdash \phi \wedge \psi$$

That is, we prefer ψ to ϕ whenever given the choice between giving up one or the other, we choose to give up ϕ.

The correctness of the foregoing properties and conditions is given by the following two results due to Gärdenfors and Makinson.

Theorem 1. *[6]*
If an ordering $\leq$ satisfies (EE1) – (EE5), then the contraction function which is uniquely determined by ($C\dot{-}$) satisfies ($K\dot{-}1$) – ($K\dot{-}8$) as well as condition ($C \leq$).

Theorem 2. *[6]*
If a contraction function $\dot{-}$ satisfies ($K\dot{-}1$) – ($K\dot{-}8$), then the ordering that is uniquely determined by ($C \leq$) satisfies (EE1) – (EE8) as well as condition ($C\dot{-}$).

2 Compiling Epistemic Entrenchment

We now begin to describe our simple yet elegant technique for compiling an epistemic entrenchment ordering and using it for belief contraction. We proceed as follows. Firstly we describe how to compile the epistemic entrenchment relation using prime implicates. In the next section we use this compiled representation to effect belief contraction using single clauses. We then extend these results to arbitrary formulas by converting them into Conjunctive Normal Form (CNF).

In specifying how to compile the reasoner's epistemic state, represented by an epistemic entrenchment relation, we take as our point of departure an important observation on epistemic entrenchment relations due to Rott.

Proposition 2. *[14] Given an arbitrary sentence $\psi \in \mathcal{L}$, $\{\phi : \psi \leq \phi\} = Cn(\{\phi : \psi \leq \phi\})$.*

Put simply, if we were to "cut" the epistemic entrenchment relation at any level, the beliefs that are at least this entrenched would form a set that is deductively closed (i.e., a belief set)[3]. This gives us yet another way to view the epistemic entrenchment relation. Starting from the tautologies which are maximally entrenched we can see that the cuts form ever expanding theories nested one within the other. This should not be surprising

[1] When the new information ϕ is a tautology.
[2] Recall from above that $\phi \leq (\phi \vee \psi)$.
[3] The term "cut" is due to Rott [14] and we shall adopt it here.

given the relationship established by Gärdenfors and Makinson [6] between epistemic entrenchment and Grove's [9] possible worlds semantics for AGM belief change.

It is this fact which is the key to our representation. Ignoring the tautologies which are always maximally entrenched, we take each successive cut of the epistemic entrenchment relation and compile the resulting theory. In so doing we end up with ordered sets of prime implicates. To be a little more precise, we start with the sentences in the most entrenched cut less than the tautologies and convert this theory into its equivalent set of prime implicates. We then proceed to the next most entrenched cut and use an incremental prime implicate algorithm to compute the additional prime implicates to be added at this level of entrenchment. These prime implicates will be one of two types: (i) prime implicates that subsume more entrenched prime implicates; and, (ii) "new" prime implicates. In either case, we add these additional prime implicates to an epistemic entrenchment relation over prime implicates only.

In this way we take an epistemic entrenchment relation $\leq$ and transform it into a new epistemic entrenchment relation $\leq_{\Pi}$ equivalent to the first but that orders clauses (i.e., prime implicates) only. Clauses occur in the ordering $\leq_{\Pi}$ only when they are prime implicates of some cut of the original epistemic entrenchment relation $\leq$ (i.e., $C \in \Pi(\{\phi : \phi \leq \psi\})$ for some $\psi \in \mathcal{L}$). This new epistemic entrenchment relation is formally specified as follows.

Definition 2. *($\leq_{\Pi}$)*
Given an epistemic entrenchment ordering $\leq$ satisfying properties (EE1)–(EE5) we define a compiled epistemic entrenchment ordering $\leq_{\Pi}$ as follows. For any two clauses C, D, $C \leq_{\Pi} D$ iff all of the following hold:

1. *$C \leq D$;*
2. *$C \in \Pi(\{\phi : \phi \leq \psi\})$ for some $\psi \in \mathcal{L}$; and,*
3. *$D \in \Pi(\{\phi : \phi \leq \chi\})$ for some $\chi \in \mathcal{L}$*

Note that the empty clause is less entrenched than all clauses which we denote $\perp \leq C$ (for all clauses C).

This definition simply specifies that one clause is at least as epistemically entrenched as another precisely when it is at least as epistemically entrenched as the other in the original entrenchment ordering $\leq$ and both clauses are prime implicates of some cut of $\leq$. Importantly, clauses that are not prime implicates for any cut do not appear in the ordering at all and are not required for effecting belief contraction. In this way, the epistemic entrenchment relation only orders a minimal set of required sentences to effect belief contraction rather than ordering all sentences in the language. Furthermore, we can see whether a sentence ϕ is believed ($\phi \in K_{\leq_{\Pi}}$) with respect to the compiled ordering $\leq_{\Pi}$ by converting it into prime implicate form and ensuring that each of the implicates in $\Pi(\phi)$ is subsumed by some non-minimal clause in $\leq_{\Pi}$. That this ordering is correct as far as our intended purpose is concerned can be seen from the following result.

Lemma 1. *Let $\phi \in \mathcal{L}$, $\leq$ satisfy properties (EE1)–(EE5) and $\leq_{\Pi}$ be obtained from $\leq$ via Definition 2. Then for each clause $C \in \Pi(\phi)$ there is some clause D such that $D \subseteq C$ and $\perp <_{\Pi} D$ iff $\phi \in K_{\leq_{\Pi}}$*

3 Belief Change Using the Compiled Epistemic Entrenchment: Contraction by a Single Clause

Using the (C$\dot{-}$) condition we now provide a way of using the compiled entrenchment relation to effect belief contraction. Since this condition is specified in terms of disjunction, prime implicates are an ideal candidate for knowledge compilation and we start by considering contraction of single clauses only. Simply put, our technique works by applying the (C$\dot{-}$) condition directly to the prime implicates for each cut, with one twist. Instead of simply removing clauses as determined using (C$\dot{-}$) we consider whether they should be replaced by weaker clauses otherwise we risk removing too many clauses and not remaining faithful to the AGM postulates (K$\dot{-}$1)–(K$\dot{-}$8).

One concept we require is the level of entrenchment of a clause.

Definition 3. *Let C be a clause, $max_{\leq_\Pi}(C) = D$ such that*

1. *D is a clause where $D \subseteq C$, and*
2. *there is no other clause $D' \subset C$ and $D \leq_\Pi D'$.*

A clause's level of entrenchment is at the same level as the maximally entrenched clause that subsumes it under $\leq_\Pi$ (which is obviously the same level at which it would appear in $\leq$). Using this definition and condition (C$\dot{-}$), when contracting by a clause C we automatically retain all clauses D strictly more entrenched than C (i.e., $max_{\leq_\Pi}(C) < D$). The remaining clauses—those less or equally as entrenched as C—need to be considered in turn.

Again using condition (C$\dot{-}$) for any clauses D, where $D \leq_\Pi max_{\leq_\Pi}(C)$ and $max_{\leq_\Pi}(C) < max_{\leq_\Pi}(C \cup D)$ we can retain D. However, clauses D failing to meet these conditions will be removed from the ordering. As indicated at the start of this section, simply removing these clauses will result in too many beliefs being given up to satisfy all the AGM postulates. In some cases we need to consider how to replace these clauses with weaker versions. For these clauses, $max_{\leq_\Pi}(C) = max_{\leq_\Pi}(C \cup D)$ and we replace D in $\leq_\Pi$ by all $(D \cup E) \setminus C$ (here $\setminus$ is set substraction) where $max_{\leq_\Pi}(C) < max_{\leq_\Pi}(C \cup D \cup E)$.[4] The contracted entrenchment $\leq_\Pi^C$ is now formally defined.

Definition 4. *Given an epistemic entrenchment relation $\leq$ and a clause C we define the contraction $\leq_\Pi^C$ of a compiled epistemic entrenchment $\leq_\Pi$ by C as*
$D \leq_\Pi^C E$ iff either

1. *$D \leq E$ and $max_{\leq_\Pi}(C) < D, E$, or*
2. *$max_{\leq_\Pi}(D) \leq max_{\leq_\Pi}(C)$, $max_{\leq_\Pi}(C) < E$, $D = C \cup F \cup G > max_{\leq_\Pi}(C)$ and $F \leq max_{\leq_\Pi}(C)$ for clauses F, G, or*
3. *$max_{\leq_\Pi}(D) \leq max_{\leq_\Pi}(C)$, $D = C \cup F \cup G > max_{\leq_\Pi}(C)$, $F \leq max_{\leq_\Pi}(C)$ for clauses F, G and $max_{\leq_\Pi}(E) \leq max_{\leq_\Pi}(C)$, $E = C \cup A \cup B > max_{\leq_\Pi}(C)$, $A \leq max_{\leq_\Pi}(C)$ for clauses A, B*

And we can see that this definition is correct as far as single clauses are concerned.

Theorem 3. *Let K be a belief set, $\dot{-}$ an AGM contraction function satisfying (K$\dot{-}$1)–(K$\dot{-}$8) and $\leq$ an epistemic entrenchment relation defined from $\dot{-}$ by condition (C$\leq$). Furthermore, let C be a clause. Then $K \dot{-} C = Cn(\{D : \bot <_\Pi^C D\})$.*

[4] Including where $C \cup D \cup E$ may be a tautology.

4 Belief Change Using the Compiled Epistemic Entrenchment: Belief Contraction by an Arbitrary Sentence

The results of the previous section can now be easily extended to arbitrary sentences ϕ by converting these sentences into CNF and considering some results relating to contraction of conjunctions and epistemic entrenchment. As indicated earlier, belief contraction by an arbitrary sentence ϕ is achieved by contracting by $CNF(\phi)$.

The first result we give states that in contraction by conjunctions we only need to consider the removal of the least entrenched conjuncts.

Proposition 3. *Given a belief set K, an AGM contraction function $\dot{-}$ and epistemic entrenchment relation $\leq$ related by conditions (C$\dot{-}$) and (C$\leq$) and $\phi, \psi \in \mathcal{L}$ such that $\phi < \psi$, then*

1. $K \dot{-} (\phi \wedge \psi) = K \dot{-} \phi$
2. $\psi \in K \dot{-} (\phi \wedge \psi)$

Furthermore, for equally entrenched conjuncts, we can simply take what is common to the contractions of each of the conjuncts.

Proposition 4. *Given a belief set K, an AGM contraction function $\dot{-}$ and epistemic entrenchment relation $\leq$ related by conditions (C$\dot{-}$) and (C$\leq$) and $\phi, \psi \in \mathcal{L}$, then*

If $\phi = \psi$, then $K \dot{-} (\phi \wedge \psi) = K \dot{-} \phi \cap K \dot{-} \psi$

These two results together with (K$\dot{-}$6) mean that contracting by an arbitrary sentence ϕ can be achieved by considering only the minimally entrenched clauses in $CNF(\phi)$, contracting by each of these clauses individually and taking what is common to all. We can specify this formally as follows.

Definition 5. *Let $\phi \in \mathcal{L}$, the set of minimal conjuncts of ϕ is defined as follows: $min_{\leq_\Pi}(CNF(\phi)) = \{C_i \in CNF(\phi) : max_{\leq_\Pi}(C_i) \leq_\Pi max_{\leq_\Pi}(C_j)$ for all $C_j \in CNF(\phi)\}$.*

And extend our previous result to show that this is indeed correct.

Theorem 4. *Let K be a belief set, $\dot{-}$ an AGM contraction function satisfying (K$\dot{-}$1)–(K$\dot{-}$8) and $\leq$ an epistemic entrenchment relation defined from $\dot{-}$ by condition (C$\leq$). Furthermore, let ϕ be a sentence. Then $K \dot{-} \phi = \bigcap_{C_i \in min_{\leq_\Pi}(CNF(\phi))} Cn(\{D : \bot <_\Pi^{C_i} D\})$.*

5 Conclusions

In this paper we have introduced a simple yet elegant technique for compiling an epistemic entrenchment ordering into prime implicates and showing how it can be used to effect AGM belief contraction. Our technique satisfies all the standard AGM postulates and, furthermore, once compilation has been performed only syntactic manipulation and subsumption checking is required.

Belief expansion by a sentence ϕ with respect to a compiled entrenchment ordering is not difficult to achieve but requires additional information to be specified in a similar way to Spohn's [16] framework for belief change. This additional piece of information takes the form of a level of entrenchment where each of the clauses in $CNF(\phi)$ will be entrenched after expansion. If any of these clauses is more entrenched, they can be ignored. Also, prime implicates for each cut lower in the entrenchment ordering may need to be recalculated. Using the Levi Identity: $K * \phi = (K \dot{-} \neg\phi) + \phi$ [5] we can achieve belief revision. Again, we need to specify a level of entrenchment for the clauses in $CNF(\phi)$ that will be required during the expansion phase. It may however be more profitable to attempt to implement belief revision more directly through a modification of the technique presented here. Iterated belief contraction on the other hand is already within our grasp. The definitions above indicate how to modify the compiled entrenchment ordering after contraction by ϕ. The resulting, contracted epistemic entrenchment ordering, can be used for subsequent contractions.

Note that in the worst case [15] since there are potentially an exponential number of prime implicates in the number of atoms we may need to consider exponentially many clauses when contracting by highly entrenched clauses. In practice, knowledge compilation has proven to be more effective in most cases.

Future work will concentrate on the various possibilities for iterated belief change as well as alternate forms of belief contraction. Extending these results to first-order belief change in non-finite domains is also of interest.

Acknowledgments

We would like to thank the anonymous referees whose comments were highly useful in improving this paper.

References

1. Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
2. Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17:229–264, 2002.
3. Adnan Darwiche and Judea Pearl. The logic of iterated belief revision. *Artificial Intelligence*, 89:1–29, 1997.
4. Simon Edmund Dixon and Wayne Wobcke. The implementation of a first-order logic AGM belief revision system. In *Proc. of the Fifth IEEE Int. Conf. on Tools in Artificial Intelligence*, 1993.
5. Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, MIT Press, Cambridge Massachusetts, 1988.
6. Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspect of Reasoning About Knowledge*, pages 83–96, 1988.
7. Peter Gärdenfors and Hans Rott. Belief revision. In Dov M. Gabbay, Christopher J. Hogger, and John A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming Volume IV: Epistemic and Temporal Reasoning*, pages 35–132. Oxford University Press, 1995.

8. Nikos Gorogiannis and Mark D. Ryan. Implementation of belief change operators using BDDs. *Studia Logica*, 70(1):131–156, 2002.
9. Adam Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17:157–170, 1988.
10. Peter Jackson. Computing prime implicates incrementally. In *Proceedings of the Eleventh Conference on Automated Deduction*, June 1992.
11. Alex Kean. A formal characterisation of a domain independent abductive reasoning system. Technical Report HKUST-CS93-4, Dept. of Computer Science, The Hong Kong University of Science and Technology, 1993.
12. Alex Kean and George Tsiknis. An incremental method for generating prime implicants/implicates. *Journal of Symbolic Computation*, 9:185–206, 1990.
13. Raymond Reiter and Johan de Kleer. Foundations of assumption-based truth maintenance systems: Preliminary report. In *Proc. of the Nat. Conf. in Artificial Intelligence*, pages 183–188, 1987.
14. Hans Rott. Preferential belief change using generalized epistemic entrenchment. *Journal of Logic, Language and Information*, 1:45–78, 1992.
15. Robert Schrag and James M. Crawford. Implicates and prime implicates in random 3-SAT. *Artificial Intelligence*, 81(1-2):199–222, 1996.
16. Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William L. Harper and Brian Skryms, editors, *Causation in Decision, Belief Change, and Statistics, II*, pages 105–134. Kluwer Academic Publishers, 1988.
17. Pierre Tison. Generalization of consensus theory and application to the minimization of boolean functions. *IEEE Trans. on Elec. Computers*, 4:446–456, 1967.
18. Mary-Anne Williams. Iterated theory change: A computational model. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1541–1550, San Francisco, 1995.