Modal Nonmonotonic Logics Revisited: Efficient Encodings for the Basic Reasoning Tasks ⋆

Thomas Eiter, Volker Klotz, Hans Tompits, and Stefan Woltran

Institut für Informationssysteme 184/3, Technische Universität Wien,
Favoritenstraße 9–11 , A-1040 Vienna, Austria,
[eiter,volker,tompits,stefan]@kr.tuwien.ac.at

Abstract. Modal nonmonotonic logics constitute a well-known family of knowledge-representation formalisms capturing ideally rational agents reasoning about their own beliefs. Although these formalisms are extensively studied from a theoretical point of view, most of these approaches lack generally available solvers thus far. In this paper, we show how variants of Moore's autoepistemic logic can be axiomatised by means of quantified Boolean formulas (QBFs). More specifically, we provide polynomial reductions of the basic reasoning tasks associated with these logics into the evaluation problem of QBFs. Since there are now efficient QBF-solvers, this reduction technique yields a practicably relevant approach to build prototype reasoning systems for these formalisms. We incorporated our encodings within the system QUIP and tested their performance on a class of benchmark problems using different underlying QBF-solvers.

1 Introduction

Modal nonmonotonic logics constitute one of the basic categories of approaches formalising certain aspects of human common-sense reasoning. In contrast to other well-known nonmonotonic reasoning frameworks, like, e.g., default logic [35], logic programming with negation as failure [16, 34], or circumscription [28], modal nonmonotonic logics employ the language of modal logic to realise nonmonotonicity. More specifically, the aim is to model the behaviour of an ideally rational agent reasoning about its own beliefs, by means of a unary operator L. Informally, $L\phi$ means that the agent believes ϕ. Given a set T of formulas as initial premises, referred to as *theories*, introspection properties generate sets of total beliefs, called *expansions.* Different introspection principles have been proposed in the literature, yielding different classes of modal nonmonotonic logics.

Although modal nonmonotonic logics have been extensively studied from a theoretical point of view (for a recent work discussing proof-theoretical issues, cf., e.g., [1]), besides some early attempts about implementational issues [4, 2, 21, 3], most of these approaches lack generally available solvers thus far. This despite the fact that recent years witnessed an increasing amount of successful implementations for various nonmonotonic formalisms, mostly for the answer-set programming paradigm, as realised,

⋆ This work was partially supported by the Austrian Science Fund Project P15068.

e.g., by the state-of-the-art solvers `dlv` [14] and `smodels` [31] implementing the stable model semantics for logic programs, or the default-logic prover `DeReS` [6].

In this paper, we describe a general method to build a prototype reasoning system for modal nonmonotonic logics, based on a reduction approach. The central idea is to translate a given reasoning task into a quantified Boolean formula (QBF) and then applying some sophisticated QBF-solver to evaluate the translated QBF. The existence of efficient QBF-solvers, like, e.g., the systems developed by Cadoli *et al*. [5], Giunchiglia *et al.* [18], Rintanen [37], Letz [25], or Feldmann *et al.* [15], makes this reduction approach practicably applicable.

Concerning the particular reductions, we provide efficient (polynomial-time) translations of reasoning tasks for the following modal nonmonotonic logics: Moore's autoepistemic logic [29], nonmonotonic logic $\mathbf{N}$ [26] (also called *iterative autoepistemic logic*), reasoning with parsimoniously grounded expansions [13], and Konolige's system of moderately grounded expansions [24].

From a theoretical point of view, the feasibility of the current approach relies on the observation that the evaluation problem of quantified Boolean formulas, QSAT, is PSPACE-complete, so any decision problem in PSPACE can be polynomially reduced to QSAT. In fact, the evaluation problem for QBFs having prenex normal form with $i - 1$ quantifier alternations is complete for the i-th level of the polynomial hierarchy. Since the reasoning tasks considered in this paper belong to the second and third level of the polynomial hierarchy, respectively, efficient translations to QBFs with one or two quantifier alternations must exist.

A similar approach for solving various reasoning tasks belonging to the area of nonmonotonic reasoning has been realised in the system QUIP [11, 10, 12, 7, 32]. This prototypical implementation currently handles the computation of the main reasoning tasks for logic-based abduction, default logic, a consistency-based approach to belief revision, and equilibrium logic, a generalisation of the stable model semantics for logic programs. We implemented the translations for modal nonmonotonic logics by incorporating them into the system QUIP.

Reduction methods to QBFs naturally generalise similar approaches for problems in NP; these latter problems can in turn be solved by translating them (in polynomial time) to SAT, the satisfiability problem of classical propositional logic (see e.g., [22] for such an application in Artificial Intelligence). Besides the implementation of different nonmonotonic reasoning tasks as realised by the system QUIP, successful applications based on reductions to QBFs have also been applied to conditional planning [36].

In order to conduct some experimental evaluation of our translations, we tested their implementation in the system QUIP based on a class of benchmark problems using different underlying QBF-solvers. More specifically, we used the solvers `ssolve` [15], `QuBe` [18], and `semprop` [25] on randomly generated problem instances covering a complete easy-hard-easy pattern.

2 Preliminaries

We deal with propositional languages and use the logical symbols $\top$, $\bot$, $\neg$, $\vee$, $\wedge$, and $\rightarrow$ to construct formulas in the standard way. We write $\mathcal{L}_{\mathcal{P}}$ to denote a language over an alphabet $\mathcal{P}$ of *propositional variables* or *atoms*. Formulas are denoted by Greek lower-case letters (possibly with subscripts).

Given an alphabet $\mathcal{P}$, we define a disjoint alphabet $\mathcal{P}'$ as $\mathcal{P}' = \{p' \mid p \in \mathcal{P}\}$. Accordingly, for a formula $\alpha \in \mathcal{L}_{\mathcal{P}}$, we define α' as the result of replacing in α each atom p from $\mathcal{P}$ by the corresponding atom p' in $\mathcal{P}'$. This is defined analogously for sets of formulas. Observe that this priming mechanism can be applied in an iterative manner, yielding, e.g., formulas of form α'' whose underlying alphabet $\mathcal{P}''$ is disjoint from both $\mathcal{P}$ and $\mathcal{P}'$.

Quantified Boolean formulas (QBFs) generalise ordinary propositional formulas by the admission of quantifications over propositional variables (QBFs are denoted by Greek upper-case letters). Informally, a QBF of form $\forall p\, \exists q\, \Phi$ means that for all truth assignments of p there is a truth assignment of q such that Φ is true. For instance, it is easily seen that the QBF $\exists p_1\, \exists p_2\, ((p_1 \rightarrow p_2) \wedge \forall p_3 (p_3 \rightarrow p_2))$ evaluates to true.

The precise semantical meaning of QBFs is defined as follows. First, some ancillary notation. An occurrence of a variable v in a QBF Φ is *free* iff it does not appear in the scope of a quantifier $\mathsf{Q}v$ ($\mathsf{Q} \in \{\forall, \exists\}$), otherwise the occurrence of v is *bound*. If Φ contains no free variables, then Φ is *closed*, otherwise Φ is *open*. Furthermore, $\Phi[v_1/\psi_1, \ldots, v_n/\psi_n]$ denotes the result of uniformly substituting the free occurrences of variables v_i in Φ by ψ_i ($1 \leq i \leq n$).

By an *interpretation*, I, we understand a set of variables. Informally, a variable v is true under I iff $v \in I$. In general, the truth value, $\nu_I(\Phi)$, of a QBF Φ under an interpretation I is recursively defined as follows:

1. if $\Phi = \top$, then $\nu_I(\Phi) = 1$;
2. if $\Phi = \bot$, then $\nu_I(\Phi) = 0$;
3. if $\Phi = v$ is an atom, then $\nu_I(\Phi) = 1$ if $v \in I$, and $\nu_I(\Phi) = 0$ otherwise;
4. if $\Phi = \neg\Psi$, then $\nu_I(\Phi) = 1 - \nu_I(\Psi)$;
5. if $\Phi = (\Phi_1 \wedge \Phi_2)$, then $\nu_I(\Phi) = min(\{\nu_I(\Phi_1), \nu_I(\Phi_2)\})$;
6. if $\Phi = (\Phi_1 \vee \Phi_2)$, then $\nu_I(\Phi) = max(\{\nu_I(\Phi_1), \nu_I(\Phi_2)\})$;
7. if $\Phi = (\Phi_1 \rightarrow \Phi_2)$, then $\nu_I(\Phi) = 1$ iff $\nu_I(\Phi_1) \leq \nu_I(\Phi_2)$;
8. if $\Phi = \forall v\, \Psi$, then $\nu_I(\Phi) = \nu_I(\Psi[v/\top] \wedge \Psi[v/\bot])$;
9. if $\Phi = \exists v\, \Psi$, then $\nu_I(\Phi) = \nu_I(\Psi[v/\top] \vee \Psi[v/\bot])$.

We say that Φ is *true under* I iff $\nu_I(\Phi) = 1$, otherwise Φ is *false under* I. If $\nu_I(\Phi) = 1$, then I is a *model* (or *satisfying truth assignment*) of Φ. Likewise, for a set S of formulas, if $\nu_I(\Phi) = 1$ for all $\Phi \in S$, then I is a model of S. If Φ has some model, then Φ is said to be *satisfiable*. If Φ is true under any interpretation, then Φ is *valid*. Observe that a closed QBF is either valid or unsatisfiable, because closed QBFs are either true under each interpretation or false under each interpretation. Hence, for closed QBFs, there is no need to refer to particular interpretations.

In the sequel, we employ the following abbreviation in the context of QBFs: For an indexed set $P = \{p_1, \ldots, p_n\}$ of propositional variables and a quantifier $\mathsf{Q} \in \{\forall, \exists\}$, we let $\mathsf{Q}P\, \Phi$ stand for the formula $\mathsf{Q}p_1 \mathsf{Q}p_2 \cdots \mathsf{Q}p_n\, \Phi$.

3 Translations

In this section, we show how several modal nonmonotonic logics can be mapped to QBFs in polynomial time.

The language of each modal nonmonotonic logic contains the unary modal operator L, where $L\phi$ intuitively means that ϕ is believed. By $\mathcal{L}_L$ we denote the language obtained from $\mathcal{L}$ by the adjunction of L. Finite subsets of $\mathcal{L}_L$ are referred to as *autoepistemic theories*, or simply as *theories*, and are identified with the conjuction of their elements.

Formulas of the form $L\phi$ are treated like propositional atoms and are called *modal atoms*.

The following notational convention will be applied: If $S \subseteq \mathcal{L}_L$, then $\neg S = \{\neg\phi \mid \phi \in S\}$. For an autoepistemic theory T, M_T denotes all modal atoms of T, and M_T^0 denotes all modal atoms of T which are not in the scope of another L-operator.

Autoepistemic theories represent an initial set of premises, generating sets of total knowledge, called *expansions*. The specific definition of an expansion depends on the nonmonotonic logic at hand, and incorporates differing notions of *groundedness*.

For all modal nonmonotonic logics introduced in the sequel, we provide reductions of the following reasoning tasks into QBFs:

1. compute all expansions of a given autoepistemic theory;
2. given an autoepistemic theory T and a formula ψ, check whether T possesses some expansion containing ψ ("brave reasoning"); and
3. given an autoepistemic theory T and a formula ψ, check whether ψ is contained in all expansions of T ("skeptical reasoning").

3.1 Autoepistemic Logic

A *stable expansion* [29] of an autoepistemic theory $T \subseteq \mathcal{L}_L$ is a set of formulas $E \subseteq \mathcal{L}_L$ such that

$$E = Cn(T \cup \{L\phi \mid \phi \in E\} \cup \{\neg L\phi \mid \phi \in \mathcal{L}_L \setminus E\}), \tag{1}$$

where $Cn(\cdot)$ is the classical consequence operator with respect to the extended language $\mathcal{L}_L$. A weakening of this concept is that of a *stable set*: $S \subseteq \mathcal{L}_L$ is stable if S satisfies the following three conditions: (i) $S = Cn(S)$, (ii) $\phi \in S$ implies $L\phi \in S$, and (iii) $\phi \notin S$ implies $\neg L\phi \in S$.

Following Niemelä [30], there is a one-to-one correspondence between stable expansions and *full sets*, which are defined as follows: Let $T \subseteq \mathcal{L}_L$ be an autoepistemic theory. Call $\Lambda \subseteq M_T \cup \neg M_T$ *T-full* iff, for all $L\phi \in M_T$, (i) $T \cup \Lambda \models \phi$ iff $L\phi \in \Lambda$, and (ii) $T \cup \Lambda \not\models \phi$ iff $\neg L\phi \in \Lambda$.

Furthermore, the following relation is defined: Given a set $T \subseteq \mathcal{L}_L$ and a modal formula ψ, $T \models_L \psi$ iff $T \cup SB_T(\psi) \models \psi$, where $SB_T(\psi) = \{L\phi \in M_\psi^0 \mid T \cup SB_T(\phi) \models \phi\} \cup \{\neg L\phi \in \neg M_\psi^0 \mid T \cup SB_T(\phi) \not\models \phi\}$.

Proposition 1 ([30]). *Let T be an autoepistemic theory. Then, the function SE_T, defined as*

$$SE_T(\Lambda) = \{\phi \in \mathcal{L}_L \mid T \cup \Lambda \models_L \phi\},$$

gives a bijective mapping from the set of T-full sets to the set of stable expansions of T. Moreover, $SE_T(\Lambda)$ is the unique stable expansion S of T such that

$$\Lambda = (M_T \cup \neg M_T) \cap (\{L\phi \mid \phi \in S\} \cup \{\neg L\phi \mid \phi \notin S\}).$$

Example 1. Let $T = \{Lp \rightarrow p\}$. This theory has two stable expansions, one containing p and Lp, the other containing $\neg Lp$. Concerning the first expansion, the corresponding T-full set is $\Lambda_1 = \{Lp\}$, since $T \cup \{Lp\} \models p$ holds. The second expansion is characterised by the T-full set $\{\neg Lp\}$, since we have that $T \cup \{\neg Lp\} \not\models p$. Observe that $SE_T(\Lambda_2)$ does neither contain p nor $\neg p$.

It has been argued in the literature [24, 26] that the first expansion is in some sense counterintuitive, since the assertion of p is based solely on the assumption that p is believed. The following subsections will deal with variations of autoepistemic logic which circumvent this problem.

For an autoepistemic theory T having propositional variables V and for a set M of new modal atoms, the following QBF will be used as a basic module:

$$\mathcal{F}_{ael}[T, M] = \forall V\Big(T \rightarrow \bigwedge_{L\phi \in M} (L\phi \rightarrow \phi)\Big) \wedge \bigwedge_{L\phi \in M} \Big(\neg L\phi \rightarrow \exists V(T \wedge \neg\phi)\Big).$$

This QBF reflects the conditions for full sets. In fact, we have the following characterisation:

Proposition 2 ([11]). *An autoepistemic theory $T \subseteq \mathcal{L}_L$ has at least one stable expansion iff $\mathcal{F}_{ael}[T, M_T]$ is satisfiable. Moreover, the satisfying truth assignments of $\mathcal{F}_{ael}[T, M_T]$ are in a one-to-one correspondence to the full sets of T.*

Thus, due to Proposition 1, the satisfying truth assignments of $\mathcal{F}_{ael}[T, M_T]$ are in a one-to-one correspondence to the stable expansions of T.

Example 2. We show the functioning of module $\mathcal{F}_{ael}[\cdot, \cdot]$ by using the theory $T = \{Lp \rightarrow p\}$ from Example 1. In this case, we have $M_T = \{Lp\}$ and $V = \{p\}$. So, $\mathcal{F}_{ael}[T, M_T]$ is given by

$$\forall p\Big((Lp \rightarrow p) \rightarrow (Lp \rightarrow p)\Big) \wedge \Big(\neg Lp \rightarrow \exists p((Lp \rightarrow p) \wedge \neg p)\Big). \qquad (2)$$

Since Lp is the single free variable of (2), there are only two interpretations serving as potential models of $\mathcal{F}_{ael}[T, M_T]$, viz. $\emptyset$ and $\{Lp\}$.

Observe that the first conjunct of (2) evaluates trivially to true, since $(Lp \rightarrow p) \rightarrow (Lp \rightarrow p)$ is a tautology. Thus, it remains to analyse the second conjunct. Clearly, $\neg Lp \rightarrow \exists p((Lp \rightarrow p) \wedge \neg p)$ evaluates to true if Lp is set to true. Hence, $\{Lp\}$ is a model of $\mathcal{F}_{ael}[T, M_T]$. On the other hand, setting Lp to false makes $\exists p((Lp \rightarrow p) \wedge \neg p)$ true as well, since $(Lp \rightarrow p) \wedge \neg p$ reduces to $\neg p$ in this case, and there is quite obviously an interpretation which makes $\neg p$ true. Hence, $\emptyset$ is also a model of $\mathcal{F}_{ael}[T, M_T]$. Invoking Proposition 2 yields the expected result that T possesses two stable expansions, one containing Lp, and the other containing $\neg Lp$.

Using Niemelä's relation $\models_L$ as a basis for describing autoepistemic inference tasks, one has to deal with the recursive nature of this relation. However, such recursive definitions are extremely unhandy to be expressed as QBFs. In the following, we give an alternative method to decide containment of a formula in a stable expansion.

Lemma 1. *Let T be an autoepistemic theory, $SE_T(\Lambda)$ a stable expansion of T, and ψ a formula. Then, for any subformula χ of ψ and any $L\phi \in M_T \cap M_\chi^0$, we have that $L\phi \in \Lambda$ iff $L\phi \in SB_{T\cup\Lambda}(\chi)$.*

This lemma gives evidence that any possible element $L\phi \in M_T$ used in the recursive computation of an assertion $T \cup \Lambda \models_L \psi$ has equal polarity in the two sets $SB_{T\cup\Lambda}(\chi)$ and Λ. Like the full-set characterisation of stable models of T, it turns out that the remaining modal atoms of ψ can be utilised analogously to describe membership of formulas in stable models.

Lemma 2. *Let T, $SE_T(\Lambda)$, and ψ be as in Lemma 1. Then, $\psi \in SE_T(\Lambda)$ iff there exists some $\Lambda^+ \subseteq (M_\psi \setminus M_T) \cup \neg(M_\psi \setminus M_T)$ such that $T \cup \Lambda \cup \Lambda^+ \models \psi$, where Λ^+ satisfies the following conditions, for each $L\phi \in M_\psi \setminus M_T$:*

(i) $T \cup \Lambda \cup \Lambda^+ \models \phi$ *iff* $L\phi \in \Lambda^+$; *and*
(ii) $T \cup \Lambda \cup \Lambda^+ \not\models \phi$ *iff* $\neg L\phi \in \Lambda^+$.

We get the following translations into QBFs:

Theorem 1. *Let T be an autoepistemic theory and ψ a formula. Furthermore, let V be the set of propositional atoms occurring in T or ψ, and let $M^+ = M_\psi \setminus M_T$. Define*

$$\mathcal{F}_{ael}^{cred}[T,\psi] = \exists M^+ \Big(\mathcal{F}_{ael}[T, M_T \cup M_\psi] \wedge \forall V(T \to \psi)\Big); \quad \textit{and}$$
$$\mathcal{F}_{ael}^{skept}[T,\psi] = \exists M^+ \Big(\mathcal{F}_{ael}[T, M_T \cup M_\psi] \wedge \neg\forall V(T \to \psi)\Big).$$

Then:

(i) *T has a stable expansion containing ψ iff $\mathcal{F}_{ael}^{cred}[T,\psi]$ is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in $\mathcal{F}_{ael}^{cred}[T,\psi]$ are in a one-to-one correspondence to the stable expansions of T containing ψ.*
(ii) *ψ is contained in all stable expansions of T iff the QBF $\neg\exists M_T(\mathcal{F}_{ael}^{skept}[T,\psi])$ evaluates to true.*

Intuitively, brave reasoning via $\mathcal{F}_{ael}^{cred}[T,\psi]$ works as follows. The first conjunct in the scope of the quantifier $\exists M^+$ determines the T-full set Λ as well as the set Λ^+ satisfying conditions (i) and (ii) of Lemma 2, and the second conjunct, $\forall V(T \to \psi)$, checks, with respect to the selected variables from $M_T \cup M_\psi$ corresponding to $\Lambda \cup \Lambda^+$, whether $T \cup \Lambda \cup \Lambda^+ \models \psi$ holds. As for skeptical reasoning, observe that the models of $\mathcal{F}_{ael}^{skept}[T,\psi]$ correspond to those T-full sets *not* containing ψ. Obviously, if no such set exists (i.e., if $\neg\exists M^+ \mathcal{F}_{ael}^{skept}[T,\psi]$ evaluates to true), ψ is a skeptical consequence of T.

3.2 Nonmonotonic Logic N

As illustrated by Example 1, Moore's autoepistemic logic admits stable expansions which are not sufficiently grounded in the premises. To circumvent this problem, several more restrictive groundedness conditions have been proposed. The nonmonotonic logic $\mathbf{N}$ [26] eliminates unfounded expansions by representing the positive introspection part in the fixed-point condition (1) in terms of a modified consequence operator. More specifically, define the inference relation $\vdash_{\mathbf{N}}$ by adding the *necessitation rule* $\phi/L\phi$ to the postulates of the classical derivability relation $\vdash$. Accordingly, let $Cn_{\mathbf{N}}(T) = \{\phi \mid T \vdash_{\mathbf{N}} \phi\}$. Then, E is an $\mathbf{N}$-*expansion* of T iff

$$E = Cn_{\mathbf{N}}(T \cup \{\neg L\phi \mid \phi \in \mathcal{L}_L \setminus E\}). \tag{3}$$

Example 3. Reconsider Example 1, where we argued that the expansion $SE_T(\{Lp\})$ of the theory $T = \{Lp \rightarrow p\}$ is counterintuitive. This expansion is not an $\mathbf{N}$-expansion. To see this, since $Lp \notin Cn_{\mathbf{N}}(T)$ and $Lp \notin Cn_{\mathbf{N}}(T \cup \{\neg L\phi \mid \phi \in \mathcal{L}_L \setminus E\})$ for any set E of formulas containing Lp, it follows that no set E containing Lp fulfills Condition 3. It is easy to see that the stable expansion $SE_T(\{\neg Lp\})$ is, however, an $\mathbf{N}$-expansion of T.

In order to express $\mathbf{N}$-expansions in terms of QBFs, we need a suitable characterisation of the derivability operator $\vdash_{\mathbf{N}}$.

Proposition 3 ([20]). *Let T be an autoepistemic theory, ψ a formula, and $M = M^0_T \cup M^0_\psi$. Then, $T \vdash_{\mathbf{N}} \psi$ iff there exists no set $K \subseteq M$ such that*

(i) *$T \cup K \not\models \psi$; and*
(ii) *for each $L\phi \in M \setminus K$, $T \cup K \not\models \phi$.*

Proposition 3 allows the construction of a QBF representing the relation $T \vdash_{\mathbf{N}} \psi$ in the following way. Let T be an autoepistemic theory, ψ an autoepistemic formula, V the set of propositional variables occurring in T or ψ, and $M \subseteq (M^0_T \cup M^0_\psi)$. Then, $\Phi_{\mathbf{N}}[T, \psi, M]$ is defined as

$$\neg \exists M' \Big(\exists M \exists V (T \wedge \bigwedge_{L\phi' \in M'} (L\phi' \rightarrow L\phi) \wedge \neg\psi) \wedge$$

$$\bigwedge_{L\phi' \in M'} (\neg L\phi' \rightarrow \exists M \exists V (T \wedge \bigwedge_{L\phi' \in M'} (L\phi' \rightarrow L\phi) \wedge \neg\phi)) \Big).$$

Theorem 2. *For any autoepistemic theory T and any formula ψ, $T \vdash_{\mathbf{N}} \psi$ holds iff $\Phi_{\mathbf{N}}[T, \psi, M^0_T \cup M^0_\psi]$ is true.*

$\mathbf{N}$-expansions can be characterised by means of full sets as follows:

Proposition 4 ([19]). *Let T be an autoepistemic theory T and let Λ be T-full. Then, $SE_T(\Lambda)$ is an $\mathbf{N}$-expansion of T iff*

$$T \cup \{\neg L\phi \,|\, \neg L\phi \in \Lambda\} \vdash_{\mathbf{N}} \{L\phi \,|\, L\phi \in \Lambda\}.$$

Obviously, each $\mathbf{N}$-expansion of a theory T is also a stable expansion of T, but not vice versa. This is reflected by the following characterisations, in which we add an additional conjunct to rule out those models which correspond to stable expansions but not to $\mathbf{N}$-expansions.

Theorem 3. *Let T be an autoepistemic theory and ψ a formula. Furthermore, let $\Lambda^+ = \bigwedge_{L\phi \in M_T}(L\phi \rightarrow L\phi')$ and $\Lambda^- = \bigwedge_{L\phi \in M_T}(\neg L\phi \rightarrow \neg L\phi')$, and consider the following QBFs:*

$$\mathcal{F}_{ael}[T, M_T] \wedge \Phi_{\mathbf{N}}[T' \wedge \Lambda^-, \Lambda^+, M_{T'}]; \tag{4}$$

$$\mathcal{F}_{ael}^{cred}[T, \psi] \wedge \Phi_{\mathbf{N}}[T' \wedge \Lambda^-, \Lambda^+, M_{T'}]; \tag{5}$$

$$\neg \exists M_T \big(\mathcal{F}_{ael}^{skept}[T, \psi] \wedge \Phi_{\mathbf{N}}[T' \wedge \Lambda^-, \Lambda^+, M_{T'}]\big). \tag{6}$$

Then:

(i) *T has an $\mathbf{N}$-expansion iff* (4) *is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in* (4) *are in a one-to-one correspondence to the $\mathbf{N}$-expansions of T.*
(ii) *T has an $\mathbf{N}$-expansion containing ψ iff* (5) *is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in* (5) *are in a one-to-one correspondence to the $\mathbf{N}$-expansions of T containing ψ.*
(iii) *ψ is contained in all $\mathbf{N}$-expansions of T iff* (6) *is true.*

Although the notion of an $\mathbf{N}$-expansion overcomes some difficulties arising with stable expansions (cf. Example 3), there are still some unwanted features of nonmonotonic logic $\mathbf{N}$.

As a case in point, consider the theory $T = \{\neg L \neg Lp \rightarrow p\}$. T has two $\mathbf{N}$-expansions: one containing $\neg L \neg Lp$ (and thus p, by propositional inference, and Lp, by the rule of necessitation), the other containing $L\neg Lp$, but neither p nor $\neg p$. Hence, here we have the situation that the objective (i.e., non-modal) part of one $\mathbf{N}$-expansion (namely the one containing neither p nor $\neg p$) is a proper subset of the objective part of the other $\mathbf{N}$-expansion, which, in some sense, is undesirable.

Note that, as shown in [19], all of the reasoning problems discussed so far lie at the second level of the polynomial hierarchy. It is easily verified that the corresponding encodings (1)–(6) yield QBFs possessing one quantifier alternation. Thus, our transformations reflect the inherent complexity of the expressed tasks. Moreover, since the reductions are constructible in polynomial time, they are, in this sense, *efficient*.

We proceed with two approaches which impose a minimality criterion on expansions. These systems were shown to be located at the third level of the polynomial hierarchy [13]. Consequently, our subsequent encodings possess an additional quantifier alternation.

3.3 Parsimoniously Grounded Expansions

Parsimoniously grounded expansions [13] are a natural strengthening of the concept of stable expansions requiring that the objective (i.e., non-modal) part of a stable expansion must be minimal with respect to set inclusion compared to all other stable expansions.

We define the following partial order: Let $S_1, S_2 \subseteq \mathcal{L}_L$ be stable sets. Then, $S_1 \preceq S_2$ iff $S_1 \cap \mathcal{L} \subseteq S_2 \cap \mathcal{L}$. As usual, we write $S_1 \prec S_2$ if $S_1 \preceq S_2$ and $S_2 \not\preceq S_1$.

A stable expansion E_1 of T is *parsimonious* iff there is no stable expansion E_2 of T such that $E_2 \prec E_1$.

Proposition 5 ([13]). *Let E_1, E_2 be stable expansions of T_1 and T_2, respectively. Then, $E_2 \preceq E_1$ iff $F_{T_1}(E_1) \models F_{T_2}(E_2)$, where $F_T(E)$ is the propositional formula resulting from T by substituting all modal atoms $L\phi$ in T by $\top$ if $L\phi \in E$, and by $\bot$ otherwise.*

We define the following modules: Let S, T be autoepistemic theories, and V the propositional atoms occurring in S or T. Furthermore, let $\hat{S}$ be the result of replacing in S (uniformly) all modal atoms $L\phi$ by $L\phi'$, providing $L\phi \in M_S$. Then:

$$\Phi_{\prec}[S,T] = \forall V(T \rightarrow \hat{S}) \wedge \neg\forall V(\hat{S} \rightarrow T);$$
$$\Phi_{pars}[T] = \neg\exists M_{T'}\Big(\Phi_{\prec}[T,T] \wedge \mathcal{F}_{ael}[T', M_{T'}]\Big).$$

Theorem 4. *Let T be an autoepistemic theory and ψ a formula. Consider the following QBFs:*

$$\mathcal{F}_{ael}[T, M_T] \wedge \Phi_{pars}[T]; \tag{7}$$
$$\mathcal{F}_{ael}^{cred}[T,\psi] \wedge \Phi_{pars}[T]; \tag{8}$$
$$\neg\exists M_T\big(\mathcal{F}_{ael}^{skept}[T,\psi] \wedge \Phi_{pars}[T]\big). \tag{9}$$

Then:

(i) *T has a parsimoniously grounded expansion iff* (7) *is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in* (7) *are in a one-to-one correspondence to the parsimoniously grounded expansions of T.*
(ii) *T has a parsimoniously grounded expansion containing ψ iff* (8) *is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in* (8) *are in a one-to-one correspondence to the parsimoniously grounded expansions of T containing ψ.*
(iii) *ψ is contained in all parsimoniously grounded expansions of T iff* (9) *is true.*

3.4 Moderately Grounded Expansions

Konolige [24] suggested to restrict stable expansions to *moderately grounded expansions*, defined as follows: Let E be a stable expansion of T. Then, E is moderately grounded iff there is no stable set S such that $T \subseteq S$ and $S \prec E$.

The difference between this notion of groundedness and parsimoniously grounded expansions is illustrated by the following example:

Example 4. Consider the autoepistemic theory $T = \{Lp \rightarrow p, p \rightarrow q, Lq\}$. This theory has one parsimoniously grounded expansion, $E = SE_T(\{Lp, Lq\})$. But E is not moderately grounded, since

$$F = \{\phi \mid T \cup \{\neg Lp, Lq\} \models_L \phi\}$$

is a stable set containing T, and the objective part of F is a proper subset of the objective part of E, i.e., $F \prec E$ holds.

We proceed with a characterisation of moderately grounded expansions required for our subsequent QBF encoding.

Proposition 6 ([13]). *A stable expansion E of T is moderately grounded iff there exists no set $\Gamma \subseteq M_T \cup \neg M_T$ such that Γ is T^+-full, where $T^+ = T \cup \{\phi \mid L\phi \in \Gamma\}$, and $SE_{T^+}(\Gamma) \prec E$.*

We define the following QBF module:

$$\Phi_{mod}[T] = \neg \exists M_{T'} \Big(\Phi_{\prec}[S,T] \wedge \mathcal{F}_{ael}[S', M_{S'}] \Big),$$

with $S = T \wedge \bigwedge_{L\phi \in M_T} (L\phi \rightarrow \phi)$. Observe that $M_T = M_S$, and thus $M_{T'} = M_{S'}$.

Theorem 5. *Let T be an autoepistemic theory and ψ a formula. Consider the following QBFs:*

$$\mathcal{F}_{ael}[T, M_T] \wedge \Phi_{mod}[T]; \tag{10}$$

$$\mathcal{F}_{ael}^{cred}[T, \psi] \wedge \Phi_{mod}[T]; \tag{11}$$

$$\neg \exists M_T \big(\mathcal{F}_{ael}^{skept}[T, \psi] \wedge \Phi_{mod}[T] \big). \tag{12}$$

Then:

- (i) *T has a moderately grounded expansion iff* (10) *is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in* (10) *are in a one-to-one correspondence to the moderately grounded expansions of T.*
- (ii) *T has a moderately grounded expansion containing ψ iff* (11) *is satisfiable. Moreover, the satisfying truth assignments of the free variables M_T in* (11) *are in a one-to-one correspondence to the moderately grounded expansions of T containing ψ.*
- (iii) *ψ is contained in all moderately grounded expansions of T iff* (12) *is true.*

4 Implementation

Our methodology for expressing several modal nonmonotonic logics in terms of quantified Boolean formulas is motivated by the availability of several practically efficient QBF-solvers. Among the different tools, there is a propositional theorem prover, `boole`,[1] based on *binary decision diagrams*, a system using a generalised resolution principle [23], several provers implementing an extended Davis-Putnam procedure [5, 15, 18, 25, 37], as well as a distributed algorithm running on a PC-cluster [15]. With the exception of `boole`, these tools do not accept arbitrary QBFs, but require the input formula to be in *prenex conjunctive normal form*. To avoid an exponential increase of formula size, *structure-preserving normal-form translations* [8, 9, 33] can be used to translate a general QBF into the required normal form. In contrast to the usual normal-form translations based on distributivity laws, structure-preserving normal-form translations introduce new labels for subformula occurrences and are polynomial in the length of the input formula.

[1] The system, together with its source code, can be downloaded from the Web at URL http://www.cs.cmu.edu/∼modelcheck/bdd.html.

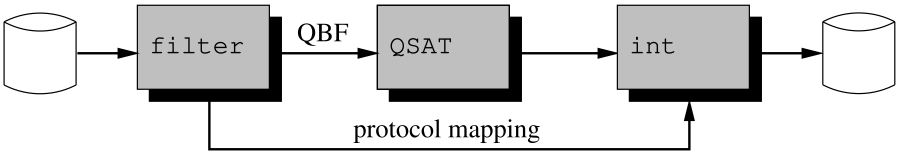

Fig. 1. Architecture to use different QBF-solvers.

The translations discussed in the previous section are implemented as a special module of the reasoning system QUIP [11, 10, 12, 7, 32], which is a prototype tool for solving several nonmonotonic reasoning tasks based on reductions to QBFs. The general architecture of QUIP is depicted in Figure 1. QUIP consists of three parts, namely the `filter` program, a QBF-evaluator, and the interpreter `int`. The input filter translates the given problem description (in our case, a modal nonmonotonic theory and a specified reasoning task) into the corresponding quantified Boolean formula, which is then fed into the QBF-evaluator. The current version of QUIP provides interfaces to most of the sequential QBF-solvers mentioned above. For the solvers requiring prenex normal form, the QBFs are translated into structure-preserving normal form. The result of the QBF-evaluator is interpreted by `int`. Depending on the capabilities of the employed QBF-evaluator, `int` provides an explanation in terms of the underlying problem instance (e.g., listing all stable expansions of a given autoepistemic theory). This task relies on a protocol mapping of internal variables of the generated QBF into concepts of the problem description which is provided by `filter`.

5 Experimental Results

In this section, we report some experimental results conducted on the implementations of the translations from Section 3. We focus here on the encodings for autoepistemic logic (cf. Proposition 2 and Theorem 1), since these translations represent, in some sense, the core parts for the other formalisms we considered. More specifically, we tested the computational behaviour, in terms of running time, of three different QBF solvers using a class of randomly generated benchmark problems, representing skeptical autoepistemic reasoning. The employed solvers are `ssolve` [15], `QuBe` [18], and `semprop` [25]. These solvers have been chosen because they turned out to be the most efficient ones on some preparatory tests.

The problem instances used here are built up as follows: We consider autoepistemic theories of form

$$T = \{\, p_1 \leftrightarrow Lp_1\,,\, \dots\,,\, p_n \leftrightarrow Lp_n\,,\, \neg L\phi \rightarrow u \,\},$$

where ϕ is a randomly generated propositional formula on atoms $p_1, \dots, p_m$ $(m \geq n)$ not containing u, and the task is to check whether u is contained in all stable expansions of T. This particular class of problems is taken from the Π_2^p-hardness proof of skeptical autoepistemic reasoning, following Gottlob [19]. Thus, in some sense, these problems

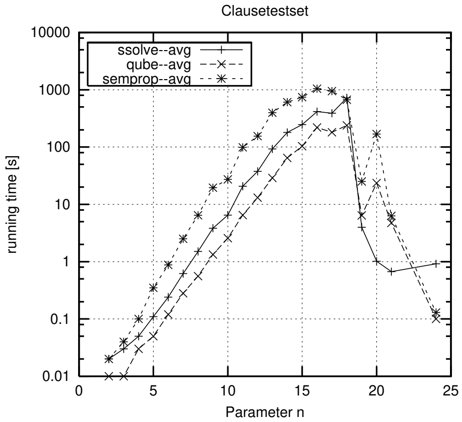

Fig. 2. Running times with varying parameter n and $l = 35$ and $m = 26$ fixed.

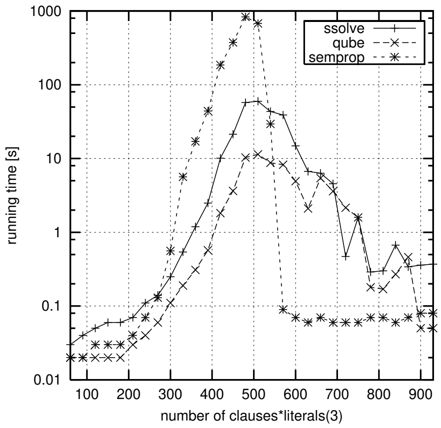

Fig. 3. Running times with varying parameter l and $m = 60$ and $n = 4$ fixed.

are responsible for the inherent worst-case complexity of skeptical autoepistemic reasoning.

For the specific tests, we chose ϕ to be in conjunctive normal form built up by l clauses each of which containing three literals. The distribution of the variables p_i in ϕ is in flavour of other benchmark methodologies [17, 27] and tries to capture computationally hard instances.

All tests have been performed on Pentium II/450 MHz processors with 128MB RAM running the Linux operating system. The running time is measured in seconds (with an upper time limit set to 1000s) and comprises the sum of both user time and system time.

In the first test set, ϕ is build over $m = 26$ variables and $l = 35$ clauses. We set up 21 sample sets, each containing 50 randomly generated formulas, by varying n from 2 to 24. Figure 2 depicts the observed running times. In the second test set, the number of clauses of ϕ is assumed to be varying, where m was set to 60 and n to 4. Each sample set contains again 50 formulas. Figure 3 shows the results. Note that the y-axes in both graphs are scaled logarithmically.

Observe that the figures indicate quite clearly that the considered test sets cover a complete easy-hard-easy pattern. To wit, in Figure 2, the instances with $n < 15$ are easily evaluated to true, whilst instances with $n > 21$ are easily evaluated to false, and the so-called *phase transition*—containing the most difficult problems—is located in the interval $16 \leq n \leq 21$. A similar phase transition occurs in Figure 3 between 400 and 600.

The most significant result of both tests is that, in general, there is no best solver. Each of the three solvers under consideration outperforms the others on certain instances. While `QuBe` seems to be most effective on formulas evaluating to true, `semprop` significantly outperforms the other ones on the instances with $l * 3 \geq 550$ in the second test.

Figure 2 reveals also another interesting behaviour. Here, for both `QuBe` and `semprop`, there is a significant peak at $n = 20$, which is not present for `ssolve`. This peak results from a single instance (out of the 50 tested) which turned out to be hard for `QuBe` and `semprop`, but apparently not for `ssolve`. This, in turn, demonstrates that invoking several solvers in parallel is quite useful in practice.

Clearly, generally speaking, it is a matter of more comprehensive tests in order to obtain a better understanding where the difficulties in solving QBFs of the current kind arise, and thus to be able to provide valuable hints to developers of QBF-solvers.

6 Conclusion

In this paper, we considered the compilation of reasoning tasks into quantified Boolean formulas as an approach to realise prototype reasoning engines for several modal nonmonotonic logics. The need for these logics in Artificial Intelligence has been well recognised in the literature. However, in contrast to other nonmonotonic-reasoning formalisms like default logic or logic programming with negation as failure, none of the considered approaches currently possesses publicly available solvers.

The investigated decision problems belong to the second and third level of the polynomial hierarchy, respectively, and are thus, from a computational point of view, "intractable". The translations described here are polynomial-time constructible and their existence is guaranteed by corresponding complexity results.

The employed framework is a natural generalisation of a similar method successfully applied to problems in NP. In general, the use of QBFs for knowledge representation purposes has been advocated in the literature [5, 37], and, besides the current framework, reductions of other reasoning tasks to QBFs have been discussed in [36, 11].

We reported some experiments using the implementation of our encodings for autoepistemic logic. Due to the absence of available solvers for modal nonmonotonic logics, we focused here on a comparison of different QBF-solvers, constituting core engines for the system QUIP.

Finally, a particular advantage of our modular approach is the straightforward ability to *parallelise* the entire evaluation process. On the one hand, this can be done by using different provers in parallel, and, on the other hand, by using the distributed QBF-solver `PQsolve` [15].

Future issues include a more thorough experimental evaluation, as well as investigating possible optimisations of the current translations.

References

1. P. Bonatti and N. Olivetti. Sequent Calculi for Propositional Nonmonotonic Logics. *ACM Transactions on Computational Logic*, 2002. To appear.
2. G. Brewka. Tweety - Still Flying: Some Remarks on Abnormal Birds, Applicable Rules and a Default Prover. In *Proc. AAAI-86*, pages 8–12, 1986.
3. G. Brewka. *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge University Press, 1991.
4. G. Brewka and K. Wittur. Nichtmonotone Logiken. Universität Bonn, Informatik Berichte 40, 1984.
5. M. Cadoli, A. Giovanardi, and M. Schaerf. An Algorithm to Evaluate Quantified Boolean Formulae. In *Proc. AAAI-98*, pages 262–267, 1998.
6. P. Cholewinski, W. Marek, and M. Truszczyński. Default Reasoning System DeReS. In *Proc. KR-96*, pages 518–528, 1996.
7. J. Delgrande, T. Schaub, H. Tompits, and S. Woltran. On Computing Solutions to Belief Change Scenarios. In *Proc. ECSQARU-01*, pages 510–521, 2001.
8. E. Eder. *Relative Complexities of First-Order Calculi*. Vieweg Verlag, 1992.
9. U. Egly. On Different Structure-preserving Translations to Normal Form. *Journal of Symbolic Computation*, 22:121–142, 1996.
10. U. Egly, T. Eiter, V. Klotz, H. Tompits, and S. Woltran. Computing Stable Models with Quantified Boolean Formulas: Some Experimental Results. In *Proc. AAAI Spring Symposium-01*, pages 53–59, 2001.
11. U. Egly, T. Eiter, H. Tompits, and S. Woltran. Solving Advanced Reasoning Tasks Using Quantified Boolean Formulas. In *Proc. AAAI-00*, pages 417–422, 2000.
12. U. Egly, V. Klotz, H. Tompits, and S. Woltran. A Toolbox for Abduction: Preliminary Report. In *Proc. IJCAR-Workshop on Theory and Applications of Quantified Boolean Formulas*, pages 29–39, 2001.
13. T. Eiter and G. Gottlob. Complexity of Reasoning with Parsimonious and Moderately Grounded Expansions. *Fundamenta Informaticae*, 17(1-2):31–53, 1992.
14. T. Eiter, N. Leone, C. Mateis, G. Pfeifer, and F. Scarcello. A Deductive System for Nonmonotonic Reasoning. In *Proc. LPNMR-97*, pages 363–374, 1997.
15. R. Feldmann, B. Monien, and S. Schamberger. A Distributed Algorithm to Evaluate Quantified Boolean Formulas. In *Proc. AAAI-00*, pages 285–290, 2000.
16. M. Gelfond and V. Lifschitz. The Stable Model Semantics for Logic Programming. In *Proc. ICLP-88*, pages 1070–1080, 1988.
17. I. P. Gent and T. Walsh. Beyond NP: The QSAT Phase Transition. In *Proc. AAAI-99*, pages 648–653, 1999.
18. E. Giunchiglia, M. Narizzano, and A. Tacchella. QUBE: A System for Deciding Quantified Boolean Formulas Satisfiability. In *Proc. IJCAR-01*, pages 364–369, 2001.

19. G. Gottlob. Complexity Results for Nonmonotonic Logics. *Journal of Logic and Computation*, 2(3):397–425, 1992.
20. G. Gottlob. Translating Default Logic into Standard Autoepistemic Logic. *Journal of the ACM*, 42(4):711–740, 1995.
21. U. Junker and K. Konolige. Computing the Extensions of Autoepistemic and Default Logic with a TMS. In *Proc. AAAI-90*, pages 278–283, 1990.
22. H. Kautz and B. Selman. Planning as Satisfiability. In *Proc. ECAI-92*, pages 359–363, 1992.
23. H. Kleine-Büning, M. Karpinski, and A. Flögel. Resolution for Quantified Boolean Formulas. *Information and Computation*, 117(1):12–18, 1995.
24. K. Konolige. On the Relation Between Default and Autoepistemic Logic. *Artificial Intelligence*, 35(3):343–382, 1988.
25. R. Letz. Advances in Decision Procedures for Quantified Boolean Formulas. In *Proc. IJCAR-Workshop on Theory and Applications of Quantified Boolean Formulas*, pages 55–64, 2001.
26. W. Marek and M. Truszczyński. Modal Logic for Default Reasoning. *Annals of Mathematics and Artificial Intelligence*, 1:275–302, 1990.
27. F. Massacci. Design and Results of the Tableaux-99 Non-classical (Modal) Systems Comparison. In *Proc. TABLEAUX-99*, pages 14–18, 1999.
28. J. McCarthy. Circumscription - A Form of Nonmonotonic Reasoning. *Artificial Intelligence*, 13:27–39, 1980.
29. R. C. Moore. Semantical Considerations on Nonmonotonic Logic. *Artificial Intelligence*, 25(1):75–94, 1985.
30. I. Niemelä. On the Decidability and Complexity of Autoepistemic Reasoning. *Fundamenta Informaticae*, 17:117–155, 1992.
31. I. Niemelä and P. Simons. Smodels: An Implementation of the Stable Model and Well-Founded Semantics for Normal Logic Programs. In *Proc. LPNMR-97*, pages 420–429, 1997.
32. D. Pearce, H. Tompits, and S. Woltran. Encodings for Equilibrium Logic and Logic Programs with Nested Expressions. In *Proc. EPIA-01*, pages 306–320. Springer, 2001.
33. D. A. Plaisted and S. Greenbaum. A Structure Preserving Clause Form Translation. *Journal of Symbolic Computation*, 2(3):293–304, 1986.
34. T. Przymusinski. Stable Semantics for Disjunctive Programs. *New Generation Computing Journal*, 9:401–424, 1991.
35. R. Reiter. A Logic for Default Reasoning. *Artificial Intelligence*, 13(1–2):81–132, Apr. 1980.
36. J. Rintanen. Constructing Conditional Plans by a Theorem Prover. *Journal of Artificial Intelligence Research*, 10:323–352, 1999.
37. J. Rintanen. Improvements to the Evaluation of Quantified Boolean Formulae. In *Proc. IJCAI-99*, pages 1192–1197, 1999.