Paraconsistent Reasoning via Quantified Boolean Formulas, I: Axiomatising Signed Systems*

Philippe Besnard[1], Torsten Schaub[1], Hans Tompits[2], and Stefan Woltran[2]

[1] Institut für Informatik, Universität Potsdam,
Postfach 90 03 27, D–14439 Potsdam, Germany
[besnard,torsten]@cs.uni-potsdam.de
[2] Institut für Informationssysteme 184/3, Technische Universität Wien,
Favoritenstraße 9–11, A–1040 Vienna, Austria
[tompits,stefan]@kr.tuwien.ac.at

Abstract. Signed systems were introduced as a general, syntax-independent framework for paraconsistent reasoning, that is, non-trivialised reasoning from inconsistent information. In this paper, we show how the family of corresponding paraconsistent consequence relations can be axiomatised by means of quantified Boolean formulas. This approach has several benefits. First, it furnishes an axiomatic specification of paraconsistent reasoning within the framework of signed systems. Second, this axiomatisation allows us to identify upper bounds for the complexity of the different signed consequence relations. We strengthen these upper bounds by providing strict complexity results for the considered reasoning tasks. Finally, we obtain an implementation of different forms of paraconsistent reasoning by appeal to the existing system QUIP.

1 Introduction

In view of today's rapidly growing amount and distribution of information, it is inevitable to encounter inconsistent information. This is why methods for reasoning from inconsistent data are becoming increasingly important. Unfortunately, there is no consensus on which information should be derivable in the presence of a contradiction. Nonetheless, there is a broad class of consistency-based approaches that reconstitute information from inconsistent data by appeal to the notion of consistency. Our overall goal is to provide a uniform basis for these approaches that makes them more transparent and easier to compare. To this end, we take advantage of the framework of quantified Boolean formulas (QBFs). To be more precise, we concentrate here on axiomatising the class of so-called *signed systems* [2] for paraconsistent reasoning; a second paper will deal with maximal-consistent sets and related approaches (cf. [4, 5]).

Our general methodology offers several benefits: First, we obtain uniform axiomatisations of rather different approaches. Second, once such an axiomatisation is available, existing QBF solvers can be used for implementation in a uniform setting. The availability of efficient QBF solvers, like the systems described in [3, 10, 9], makes such a rapid prototyping approach practicably applicable. Third, these axiomatisations provide

* The work was partially supported by the Austrian Science Foundation under grant P15068.

a direct access to the complexity of the original approach. Finally, we remark that this approach allows us, in some sense, to express paraconsistent reasoning in (higher order) classical propositional logic and so to harness classical reasoning mechanisms from (a conservative extension of) propositional logic.

Our elaboration of paraconsistent reasoning is part of an encompassing research program, analysing a large spectrum of reasoning mechanisms in Artificial Intelligence, among them nonmonotonic reasoning [7], (nonmonotonic) modal logics [8], logic programming [13], abductive reasoning [15], and belief revision [6].

In order to keep our paper self-contained, we must carefully introduce the respective techniques. Given the current space limitations, we have thus decided to reduce the motivation and rather concentrate on a thorough formal elaboration. This brings us to the following outline: Section 2 lays down the formal foundations of our work, introducing QBFs and Default Logic. Section 3 is devoted to signed systems as introduced in [2]. Apart from reviewing the basic framework, we provide new unifying characterisations that pave the way for the respective encodings in QBFs, which are the subject of Section 4. That section comprises thus our major contribution: a family of basic QBF axiomatisations that can be assembled in different ways in order to accommodate the variety of paraconsistent inference relations within the framework of signed systems. We further elaborate upon these axiomatisations in Section 5 for analysing the complexity of the respective reasoning tasks. Finally, our axiomatisations are also of great practical value since they allow for a direct implementation in terms of existing QBF-solvers. Such an implementation is described in Section 6, by appeal to the system QUIP [7, 13, 8].

2 Foundations

We deal with propositional languages and use the logical symbols $\top$, $\bot$, $\neg$, $\vee$, $\wedge$, $\rightarrow$, and $\equiv$ to construct formulas in the standard way. We write $\mathcal{L}_\Sigma$ to denote a language over an alphabet Σ of *propositional variables* or *atoms*. Formulas are denoted by Greek lowercase letters (possibly with subscripts). Finite sets $T = \{\phi_1, \ldots, \phi_n\}$ of formulas are usually identified with the conjunction $\bigwedge_{i=1}^{n} \phi_i$ of its elements. The set of all atoms occurring in a formula ϕ is denoted by $var(\phi)$. Similarly, for a set S of formulas, $var(S) = \bigcup_{\phi \in S} var(\phi)$. The classical derivability operator, $\vdash$, is defined in the usual way. The *deductive closure* of a set $S \subseteq \mathcal{L}_\Sigma$ of formulas is given by $Cn_\Sigma(S) = \{\phi \in \mathcal{L}_\Sigma \mid S \vdash \phi\}$. We say that S is *deductively closed* iff $S = Cn_\Sigma(S)$. Furthermore, S is *consistent* iff $\bot \notin Cn_\Sigma(S)$. If the language is clear from the context, we usually drop the index "Σ" from $Cn_\Sigma(\cdot)$ and simply write $Cn(\cdot)$ for the deductive closure operator. An occurrence of a formula φ is *positive* (resp., *negative*) in a formula ψ iff the number of implicit or explicit negation signs preceding φ in ψ is even (resp., odd).

Given an alphabet Σ, we define a disjoint alphabet $\Sigma^\pm$ as $\Sigma^\pm = \{p^+, p^- \mid p \in \Sigma\}$. For $\alpha \in \mathcal{L}_\Sigma$, we define $\alpha^\pm$ as the formula obtained from α by replacing each negative occurrence of p by $\neg p^-$ and by replacing each positive occurrence of p by p^+, for each propositional variable p in Σ. For example $(p \wedge (p \rightarrow q))^\pm = p^+ \wedge (\neg p^- \rightarrow q^+)$. This is defined analogously for sets of formulas. Observe that for any set $T \subseteq \mathcal{L}_\Sigma$, $T^\pm$ is consistent, even if T is inconsistent.

Quantified Boolean Formulas. 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. The precise semantical meaning of QBFs is defined as follows.

First, some ancillary notation. An occurrence of a propositional variable p in a QBF Φ is *free* iff it does not appear in the scope of a quantifier $\mathsf{Q}p$ ($\mathsf{Q} \in \{\forall, \exists\}$), otherwise the occurrence of p is *bound*. If Φ contains no free variable occurrences, then Φ is *closed*, otherwise Φ is *open*. Furthermore, we write $\Phi[p_1/\phi_1, \ldots, p_n/\phi_n]$ to denote the result of uniformly substituting each free occurrence of a variable p_i in Φ by a formula ϕ_i, for $1 \leq i \leq n$.

By an *interpretation*, M, we understand a set of atoms. Informally, an atom p is true under M iff $p \in M$. In general, the truth value, $\nu_M(\Phi)$, of a QBF Φ under an interpretation M is recursively defined as follows:

1. if $\Phi = \top$, then $\nu_M(\Phi) = 1$;
2. if $\Phi = p$ is an atom, then $\nu_M(\Phi) = 1$ if $p \in M$, and $\nu_M(\Phi) = 0$ otherwise;
3. if $\Phi = \neg\Psi$, then $\nu_M(\Phi) = 1 - \nu_M(\Psi)$;
4. if $\Phi = (\Phi_1 \wedge \Phi_2)$, then $\nu_M(\Phi) = min(\{\nu_M(\Phi_1), \nu_M(\Phi_2)\})$;
5. if $\Phi = \forall p\, \Psi$, then $\nu_M(\Phi) = \nu_M(\Psi[p/\top] \wedge \Psi[p/\bot])$;
6. if $\Phi = \exists p\, \Psi$, then $\nu_M(\Phi) = \nu_M(\Psi[p/\top] \vee \Psi[p/\bot])$.

The truth conditions for $\bot$, $\vee$, $\rightarrow$, and $\equiv$ follow from the above in the usual way. We say that Φ is *true under* M iff $\nu_M(\Phi) = 1$, otherwise Φ is *false under* M. If $\nu_M(\Phi) = 1$, then M is a *model* of Φ. If Φ has some model, then Φ is said to be *satisfiable*. If Φ is true under any interpretation, then Φ is *valid*. As usual, we write $\models \Phi$ to express that 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. Two sets of QBFs (or ordinary formulas) are *logically equivalent* iff they possess the same models.

In the sequel, we use the following abbreviations in the context of QBFs: For a 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 $\mathsf{Q}p_1 \mathsf{Q}p_2 \cdots \mathsf{Q}p_n\, \Phi$. Furthermore, for indexed sets $S = \{\phi_1, \ldots, \phi_n\}$ and $T = \{\psi_1, \ldots, \psi_n\}$ of formulas, $S \leq T$ abbreviates $\bigwedge_{i=1}^{n}(\phi_i \rightarrow \psi_i)$.

The following result is needed in the sequel:

Proposition 1. *Let $S = \{\phi_1, \ldots, \phi_n\}$ and T be finite sets of formulas, let $P = var(S \cup T)$, and let $G = \{g_1, \ldots, g_n\}$ be a set of new variables. Furthermore, for any $S' \subseteq S$, define the interpretation $M_{S'} \subseteq G$ such that $\phi_i \in S'$ iff $g_i \in M_{S'}$, for $1 \leq i \leq n$. Then,*

1. *$T \cup S'$ is consistent iff $M_{S'}$ is a model of the QBF $\mathcal{C}[T, S] = \exists P(T \wedge (G \leq S))$.*
2. *S' is a maximal subset of S consistent with T iff $M_{S'}$ is a model of the QBF $\mathcal{C}[T, S] \wedge \bigwedge_{i=1}^{n}(\neg g_i \rightarrow \neg\, \mathcal{C}[T \cup \{\phi_i\}, S \setminus \{\phi_i\}])$.*

Default Logic. The primary technical means for dealing with "signed theories" is *default logic* [14], whose central concepts are *default rules* along with their induced *extensions* of an initial set of premises. A default rule (or *default* for short) $\frac{\alpha : \beta}{\gamma}$ has two types of antecedents: a *prerequisite* α which is established if α is derivable and a *justification* β which is established if β is consistent. If both conditions hold, the *consequent* γ is concluded by default. For convenience, we denote the prerequisite of a default δ by $p(\delta)$, its justification by $j(\delta)$, and its consequent by $c(\delta)$. Accordingly, for a set of defaults D, we define $p(D) = \{p(\delta) \mid \delta \in D\}$, $j(D) = \{j(\delta) \mid \delta \in D\}$, and $c(D) = \{c(\delta) \mid \delta \in D\}$.

A *default theory* is a pair (D, W) where D is a set of default rules and W a set of formulas. A set E of formulas is an *extension* of (D, W) iff $E = \bigcup_{n\in\omega} E_n$, where $E_1 = W$ and, for $n \geq 1$, $E_{n+1} = Cn(E_n) \cup \{\gamma \mid \frac{\alpha:\beta}{\gamma} \in D, \alpha \in E_n, \neg\beta \notin E\}$.

3 Signed Systems

The basic idea of signed systems is to transform an inconsistent theory into a consistent one by renaming propositional variables and then to extend the resulting signed theory by equivalences using default logic.

Starting with a possibly inconsistent finite theory $W \subseteq \mathcal{L}_\Sigma$, we consider the default theory obtained from $W^\pm$ and a set of default rules $D_\Sigma = \{\delta_p \mid p \in \Sigma\}$ defined in the following way. For each propositional letter p in Σ, we define

$$\delta_p = \frac{:\ p^+ \equiv \neg p^-}{(p \equiv p^+) \wedge (\neg p \equiv p^-)}\ . \tag{1}$$

Using this definition, we define the first family of paraconsistent consequence relations:

Definition 1. *Let W be a finite set of formulas in $\mathcal{L}_\Sigma$ and let φ be a formula in $\mathcal{L}_\Sigma$. Let* Ext *be the set of all extensions of $(D_\Sigma, W^\pm)$. For each set of formulas $S \subseteq \mathcal{L}_{\Sigma\cup\Sigma^\pm}$, let $\Pi_S = \{c(\delta_p) \mid p \in \Sigma, \neg j(\delta_p) \notin S\}$. Then, we define*

$W \vdash_c \varphi$ *iff* $\varphi \in \bigcup_{E\in\mathsf{Ext}} Cn(W^\pm \cup \Pi_E)$ (*credulous unsigned*[1] *consequence*)
$W \vdash_s \varphi$ *iff* $\varphi \in \bigcap_{E\in\mathsf{Ext}} Cn(W^\pm \cup \Pi_E)$ (*skeptical unsigned consequence*)
$W \vdash_p \varphi$ *iff* $\varphi \in Cn(W^\pm \cup \bigcap_{E\in\mathsf{Ext}} \Pi_E)$ (*prudent unsigned consequence*)

For illustration, consider the inconsistent theory $W = \{p, q, \neg p \vee \neg q\}$. For obtaining the above paraconsistent consequence relations, W is turned into the default theory[2] $(D_\Sigma, W^\pm) = \big(\{\delta_p, \delta_q\}, \{p^+, q^+, p^- \vee q^-\}\big)$. We obtain two extensions, viz. $Cn(W^\pm \cup \{c(\delta_p)\})$ and $Cn(W^\pm \cup \{c(\delta_q)\})$. The following relations show how the different consequence relations behave: $W \vdash_c p,\ W \not\vdash_s p,\ W \not\vdash_p p$, but, for instance, $W \vdash_c p \vee q,\ W \vdash_s p \vee q,\ W \not\vdash_p p \vee q$.

For a complement, the following "signed" counterparts are defined.

Definition 2. *Given the prerequisites of Definition 1, we define*

$W \vdash_c^\pm \varphi$ *iff* $\varphi^\pm \in \bigcup_{E\in\mathsf{Ext}} Cn(W^\pm \cup \Pi_E)$ (*credulous signed consequence*)
$W \vdash_s^\pm \varphi$ *iff* $\varphi^\pm \in \bigcap_{E\in\mathsf{Ext}} Cn(W^\pm \cup \Pi_E)$ (*skeptical signed consequence*)
$W \vdash_p^\pm \varphi$ *iff* $\varphi^\pm \in Cn(W^\pm \cup \bigcap_{E\in\mathsf{Ext}} \Pi_E)$ (*prudent signed consequence*)

[1] The term "unsigned" indicates that only unsigned formulas are taken into account.
[2] For simplicity, we omitted all δ_x for $x \in \Sigma \setminus \{p, q\}$.

As shown in [2], these relations compare to each other in the following way.

Theorem 1. *Let C_i be the operator corresponding to $C_i(W) = \{\varphi \mid W \vdash_i \varphi\}$ where i ranges over $\{p, s, c\}$, and similarly for $C_i^{\pm}$. Then, we have*

1. $C_i(W) \subseteq C_i^{\pm}(W)$*;*
2. $C_p(W) \subseteq C_s(W) \subseteq C_c(W)$ *and* $C_p^{\pm}(W) \subseteq C_s^{\pm}(W) \subseteq C_c^{\pm}(W)$*.*

That is, signed derivability gives more conclusions than unsigned derivability and within each series of consequence relations the strength of the relation is increasing.

Moreover, they enjoy the following logical properties:

Theorem 2. *Let C_i be the operator corresponding to $C_i(W) = \{\varphi \mid W \vdash_i \varphi\}$ where i ranges over $\{p, s, c\}$, and similarly for $C_i^{\pm}$. Then, we have*

3. $W \subseteq C_i^{\pm}(W)$*;*
4. $C_p(W) = Cn(C_p(W))$ *and* $C_s(W) = Cn(C_s(W))$*;*
5. $C_i^{\pm}(W) = C_i^{\pm}(C_i^{\pm}(W))$*;*
6. $Cn(W) \neq \mathcal{L}_\Sigma$ *only if* $Cn(W) = C_i(W) = C_i^{\pm}(W)$*;*
7. $C_i(W) \neq \mathcal{L}_\Sigma$ *and* $C_i^{\pm}(W) \neq \mathcal{L}_\Sigma$*;*
8. $W \subseteq W'$ *does not imply* $C_i(W) \subseteq C_i(W')$*, and* $W \subseteq W'$ *does not imply* $C_i^{\pm}(W) \subseteq C_i^{\pm}(W')$*.*

The last item simply says that all of our consequence relations are nonmonotonic. For instance, we have $C_i(\{A, A \to B\}) = C_i^{\pm}(\{A, A \to B\}) = Cn(\{A, B\})$, while neither $C_i(\{A, \neg A, A \to B\})$ nor $C_i^{\pm}(\{A, \neg A, A \to B\})$ contains B.

Refinements. The previous relations embody a somewhat global approach in restoring semantic links between positive and negative literals. In fact, the application of a rule δ_p re-establishes the semantic link between all occurrences of proposition p and its negation $\neg p$ at once. A more fine-grained approach is to establish the connections between complementary occurrences of an atom individually.

Formally, for a given W and an index set I assigning different indices to all occurrences of all atoms in W, define

$$\delta_p^{i,j} = \frac{:\ (p \equiv p_i^+) \land (\neg p \equiv p_j^-)}{(p \equiv p_i^+) \land (\neg p \equiv p_j^-)} \tag{2}$$

for all $p \in \Sigma$ and all $i, j \in I$, provided that i and j refer to complementary occurrences of p in W, otherwise set $\delta_p^{i,j} = \delta_p$. Denote by D_Σ^1 this set of defaults and by $W_I^{\pm}$ the result of replacing each $p^+ \in W^{\pm}$ (resp., $p^- \in W^{\pm}$) by p_i^+ (resp., p_i^-) where i is the index assigned to the corresponding occurrence, provided that there are complementary occurrences of p in W.

Finally, abandoning the restoration of semantical links and foremost restoring original (unsigned) literals leads to the most adventurous approach to signed inferences. Consider the following set of defaults, defined for all $p \in \Sigma$ and $i, j \in I$,

$$\delta_p^{i+} = \frac{:\ (p \equiv p_i^+)}{(p \equiv p_i^+)} \qquad \delta_p^{j-} = \frac{:\ (\neg p \equiv p_j^-)}{(\neg p \equiv p_j^-)} \tag{3}$$

for all positive and negative occurrences of p, respectively. As above, we use these defaults provided that there are complementary occurrences of p in W, otherwise use δ_p. A set of defaults of form (3) with respect to W is denoted by D^2_Σ.

Thus, further consequence relations are defined when $(D_\Sigma, W^\pm)$ in Definition 1 is replaced by $(D^1_\Sigma, W^\pm_I)$ or by $(D^2_\Sigma, W^\pm_I)$. Similar results to Theorem 1 and 2 can be shown for these families of consequence relations.

In the following, we identify all introduced default theories as follows. Given a finite set $W \subseteq \mathcal{L}_\Sigma$, the class $\mathsf{DT}(W)$ contains (D_Σ, W), as well as $(D^1_\Sigma, W^\pm_I)$ and $(D^2_\Sigma, W^\pm_I)$ for any index set I. Furthermore, $\mathsf{DT} = \bigcup_{W \subseteq \mathcal{L}_\Sigma} \mathsf{DT}(W)$ denotes the class of all possible default theories under consideration.

Whenever a problem instance may give rise to several solutions, it is useful to provide a preference criterion for selecting a subset of preferred solutions. This is accomplished in [2] by means of a *ranking function* $\varrho : \Sigma \to \mathbb{N}$ on the alphabet Σ for inducing a hierarchy on the default rules in D_Σ:

Definition 3. *Let $\varrho : \Sigma \to \mathbb{N}$ be some ranking function on alphabet Σ, and $(D, V) \in \mathsf{DT}$. We define the hierarchy of D with respect to ϱ as the partition $\langle D_n \rangle_{n\in\omega}$ of D such that for each $\delta \in D$ with δ of form $\delta_p, \delta^{i,j}_p, \delta^{i+}_p, \delta^{i-}_p$, for $p \in \Sigma$ and $i, j \in I$, $\delta \in D_n$ iff $\varrho(p) = n$ holds.*

Strictly speaking, $\langle D_n \rangle_{n\in\omega}$ is not always a genuine partition, since D_n may be the empty set for some values of n.

Definition 4. *Let W be a finite set of formulas in $\mathcal{L}_\Sigma$, $(D, V) \in \mathsf{DT}(W)$, and E a set of formulas. Let $\langle D_n \rangle_{n\in\omega}$ be the hierarchy of D with respect to some ranking function ϱ.*

Then, $E = \bigcup_{n\in\omega} E_n$ is a hierarchic extension of (D, V) relative to ϱ *if $E_1 = V$ and E_{n+1} is an extension of (D_n, E_n) for all $n \geq 1$.*

Let $\langle D_n \rangle_{n\in\omega}$ be the hierarchy of D with respect to some ranking function ϱ, and let Ext_h be the set of all hierarchic extensions of a default theory $(D, V) \in \mathsf{DT}$ in Definition 1. Then, we immediately get corresponding consequence relations $\vdash_{ch}$, $\vdash_{sh}$, and $\vdash_{ph}$. Furthermore, applying hierarchic extensions on default theories (D_Σ, W) in accordance to Definition 2 yields new relations $\vdash^\pm_{ch}$, $\vdash^\pm_{sh}$, and $\vdash^\pm_{ph}$.

In concluding this section, let us briefly recapitulate all paraconsistent consequence relations introduced so far. As a basic classification, we have credulous, skeptical and prudent consequence. For each of these relations, we defined unsigned operators, which are invokable on three different classes of default theories (viz. on $(D_\Sigma, W^\pm)$, $(D^1_\Sigma, W^\pm_I)$, and $(D^2_\Sigma, W^\pm_I)$), either on ordinary extensions ($\vdash_i$) or on hierarchic extensions ($\vdash_{ih}$), and, on the other hand, signed operators also relying on ordinary extensions ($\vdash^\pm_i$) or hierarchic extensions ($\vdash^\pm_{ih}$) of the default theory $(D_\Sigma, W^\pm)$. This gives in total 18 unsigned and 6 signed paraconsistent consequence relations, which shall all be considered in the following two sections.

4 Reductions

In this section, we show how the above introduced consequence relations can be mapped into quantified Boolean formulas in polynomial time.

Recall the set $\mathsf{DT}(W)$ for finite $W \subseteq \mathcal{L}_\Sigma$. In what follows, we use finite default theories $\mathsf{DT}^*(W) = \{(D_W, V) \mid (D, V) \in \mathsf{DT}(W)\}$ where $D_W = \{\delta \in D \mid var(\delta) \cap var(W) \neq \emptyset\}$. Hence, D_W contains each default from D having an unsigned atom which also occurs in W.

We first show the adequacy of these default theories, and afterwards we develop our QBF-reductions based on these finite default theories.

Lemma 1. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas and $(D, V) \in \mathsf{DT}(W)$ a default theory. Moreover, let $C \subseteq D$ and $C_W = \{\delta \in C \mid var(\delta) \cap var(W) \neq \emptyset\}$. Then,*

1. *$Cn(V \cup c(C_W)) \cap \mathcal{L}_\Sigma = Cn(V \cup c(C)) \cap \mathcal{L}_\Sigma$; and*
2. *for each $\varphi^\pm \in \mathcal{L}_{\Sigma^\pm}$, $\varphi^\pm \in Cn(V \cup c(C))$ iff $\varphi^\pm \in Cn(V \cup c(C_W) \cup c(D_\varphi))$ where $D_\varphi = \{\delta_p \mid p \in var(\varphi) \setminus var(W)\}$.*

Both results show that having computed a (possibly hierarchic) extension, one has a finite set of generating defaults sufficient for deciding whether a paraconsistent consequence relation holds. The following result shows that these sets are also sufficient to compute the underlying extensions themselves.

Theorem 3. *Let W, (D, V), C, and C_W be as in Lemma 1, and let $D_W = \{\delta \in D \mid var(\delta) \cap var(W) \neq \emptyset\}$.*

Then, there is a one-to-one correspondence between the extensions of (D, V) and the extensions of (D_W, V). In particular, $Cn(V \cup c(C))$ is an extension of (D, V) iff $Cn(V \cup c(C_W))$ is an extension of (D_W, V). Similar relations hold for hierarchic extensions as well.

The next result gives a uniform characterisation for all default theories under consideration. It follows from the fact that, for each δ_p, the consequent $(p \equiv p^+) \wedge (\neg p \equiv p^-)$ is actually equivalent to $(p^+ \equiv \neg p^-) \wedge (p \equiv p^+)$, and, furthermore, that defaults of form (2) and (3) share the property that their justifications and consequents are identical. Hence, given W and I as usual, it holds that $c(\delta) \models j(\delta)$, for each $\delta \in D$, with $(D, V) \in \mathsf{DT}^*(W)$.

Theorem 4. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas, let $(D, V) \in \mathsf{DT}^*(W)$ be a default theory, and let $C \subseteq D$.*

Then, $Cn(V \cup c(C))$ is an extension of (D, V) iff $j(C)$ is a maximal subset of $j(D)$ consistent with V.

Note that the subsequent QBF reductions, obtained on the basis of the above result, represent a more compact axiomatics than the encodings given in [7] for arbitrary default theories.

We derive an analogous characterisation for hierarchic extensions. In fact, each hierarchic extension is also an extension (but not vice versa) [2]. Thus, we can characterise hierarchic extensions of a default theory (D, V) as ordinary extensions, viz. by $Cn(W \cup c(C))$ with $C \subseteq D$ suitably chosen. The following result generalises Theorem 4 with respect to a given partition on the defaults. In particular, if $\langle D_n \rangle_{n \in \omega} = \langle D \rangle$, Theorem 5 corresponds to Theorem 4.

Theorem 5. *Let W, (D,V), and C be given as in Theorem 4.*

Then, $Cn(V \cup c(C))$ is a hierarchic extension of (D,V) with respect to partition $\langle D_n \rangle_{n\in\omega}$ on D iff for each $i \in \omega$, $j(D_i \cap C)$ is a maximal subset of $j(D_i)$ consistent with $V \cup \bigcup_{j<i} c(D_j \cap C)$.

Finally, in order to relate extensions of default theories to paraconsistent consequence operators, we note the following straightforward observations.

Let Π_S be as in Definition 1. Then, for each extension E of $(D,V) \in \mathsf{DT}(W)$, there exists a $C \subseteq D$ such that $c(C) = \Pi_E$. However, since we have to check whether a given formula is contained in some $Cn(V \cup \Pi_E)$, by Lemma 1 it is obviously sufficient to consider just the generating defaults of an extension of the corresponding restricted default theory from $\mathsf{DT}^*(W)$. In view of Theorems 4 and 5, this immediately implies that all paraconsistent consequence relations introduced so far can be characterised by maximal subsets of the consequences $c(D)$ of the corresponding default theory $(D,V) \in \mathsf{DT}^*(W)$. More specifically, credulous and skeptical paraconsistent consequence reduces to checking whether a given formula is contained in at least one or respectively all such maximal subsets. Additionally, prudent consequence enjoys the following property.

Lemma 2. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas, and $(D,V) \in \mathsf{DT}^*(W)$.*

Then, for each $\varphi \in \mathcal{L}_\Sigma$, we have that $W \not\vdash_p \varphi$ (resp., $W \not\vdash_{ph} \varphi$) iff there exists a set $C \subseteq D$ such that $\varphi \notin Cn(V \cup c(C))$ and, for each $\delta \in D \setminus C$, there is some extension (resp., hierarchic extension) E of (D,V) such that $c(\delta) \notin E$. An analogous result holds for relations $\vdash_p^{\pm}$ and $\vdash_{ph}^{\pm}$.

Main Construction. We start with some basic QBF-modules. To this end, recall the schema $\mathcal{C}[\cdot,\cdot]$ from Proposition 1.

Definition 5. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas and $\varphi \in \mathcal{L}_\Sigma$. For each finite default theory $T = (D,V) \in \mathsf{DT}^*(W)$, let $D = \{\delta_1,\dots,\delta_n\}$, and define*

$$\mathcal{E}[T] = \mathcal{C}[V, j(D)] \wedge \bigwedge_{i=1}^{n} \Big(\neg g_i \rightarrow \neg\, \mathcal{C}[V \cup \{j(\delta_i)\}, j(D \setminus \{\delta_i\})]\Big);$$

$$\mathcal{D}[T,\varphi] = \forall P \Big(V \wedge (G \leq c(D)) \rightarrow \varphi\Big),$$

where P denotes the set of atoms occurring in T or φ, and $G = \{g_i \mid \delta_i \in D\}$ is an indexed set of globally new variables corresponding to D.

Lemma 3. *Let W, $T = (D,V)$, and G be as in Definition 5. Furthermore, for any set $C \subseteq D$, define the interpretation $M_C \subseteq G$ such that $g_i \in M_C$ iff $\delta_i \in C$, for $1 \leq i \leq n$.*

Then, the following relations hold:

1. *$Cn(V \cup c(C))$ is an extension of T iff $\mathcal{E}[T]$ is true under M_C; and*
2. *$\varphi \in Cn(V \cup c(C))$ iff $\mathcal{D}[T,\varphi]$ is true under M_C, for any formula φ in $\mathcal{L}_\Sigma$.*

Observe that the correctness of Condition 1 follows directly from Proposition 1(2), since we have that $\mathcal{E}[T]$ is true under M_C iff $j(C)$ is a maximal subset of $j(D)$ consistent with V, and, in view of Theorem 4, the latter holds iff $Cn(V \cup c(C))$ is an extension of T. Moreover, Condition 2 is reducible to Proposition 1(1). Combining these two QBF-modules, we obtain encodings for the basic inference tasks as follows:

Theorem 6. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas, $T = (D, V)$ a default theory from $\mathsf{DT}^*(W)$ with $D = \{\delta_1, \ldots, \delta_n\}$, φ a formula in $\mathcal{L}_\Sigma$, and $G = \{g_1, \ldots, g_n\}$ the indexed set of variables occurring in $\mathcal{E}[T]$ and $\mathcal{D}[T, \varphi]$.*

Then, paraconsistent credulous and skeptical consequence relations can be axiomatised by means of QBFs as follows:

1. *$W \vdash_c \varphi$ iff $\models \exists\, G(\mathcal{E}[T] \wedge \mathcal{D}[T, \varphi])$; and*
2. *$W \vdash_s \varphi$ iff $\models \neg\exists\, G(\mathcal{E}[T] \wedge \neg\, \mathcal{D}[T, \varphi])$.*

Moreover, for prudent consequence, let $G' = \{g'_i \mid g_i \in G\}$ be an additional set of globally new variables and $\Psi = \bigwedge_{i=1}^{n}(\neg g'_i \rightarrow \exists\, G(\mathcal{E}[T] \wedge \neg\, \mathcal{D}[T, c(\delta_i)]))$. Then,

3. *$W \vdash_p \varphi$ iff $\models \neg\exists G'(\neg\mathcal{D}_{G\leftarrow G'}[T, \varphi] \wedge \Psi)$,*

where $\mathcal{D}_{G\leftarrow G'}[T, \varphi]$ denotes the QBF obtained from $\mathcal{D}[T, \varphi]$ by replacing each occurrence of an atom $g \in G$ in $\mathcal{D}[T, \varphi]$ by g'.

In what follows, we discuss the remaining consequence relations under consideration. We start with signed consequence. Here, we just have to adopt the calls to $\mathcal{D}[(D, V), \varphi]$ with respect to Lemma 2, by adding those defaults δ_p to $W^\pm$ such that $p \in var(\varphi) \setminus var(W)$. Observe that in the following theorem this addition is *not* necessary for the module Ψ. Furthermore, recall that signed consequence is applied only to default theories $(D_\Sigma, W^\pm)$.

Theorem 7. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas and φ a formula in $\mathcal{L}_\Sigma$. Moreover, let $D_W = \{\delta_p \mid p \in var(W)\}$ and $D_\varphi = \{\delta_p \mid p \in var(\varphi) \setminus var(W)\}$, with the corresponding default theories $T = (D_W, W^\pm)$ and $T' = (D_W, W^\pm \cup c(D_\varphi))$, and let G, G', and Ψ be as in Theorem 6.*

Then, paraconsistent signed consequence relations can be axiomatised by means of QBFs as follows:

1. *$W \vdash^\pm_c \varphi$ iff $\models \exists\, G(\mathcal{E}[T] \wedge \mathcal{D}[T', \varphi^\pm])$;*
2. *$W \vdash^\pm_s \varphi$ iff $\models \neg\exists\, G(\mathcal{E}[T] \wedge \neg\, \mathcal{D}[T', \varphi^\pm])$; and*
3. *$W \vdash^\pm_p \varphi$ iff $\models \neg\exists G'(\Psi \wedge \neg\mathcal{D}_{G\leftarrow G'}[T', \varphi^\pm])$,*

where, as above, $\mathcal{D}_{G\leftarrow G'}[\cdot,\cdot]$ replaces each g by g'.

It remains to consider the consequence relations based on hierarchical extensions. To this end, we exploit the characterisation of Theorem 5.

Definition 6. *Let $W \subseteq \mathcal{L}_\Sigma$ be a finite set of formulas, $T = (D, V)$ a default theory from $\mathsf{DT}^*(W)$ with $D = \{\delta_1, \ldots, \delta_n\}$, and $P = \langle D_n \rangle_{n\in\omega}$ a partition on D. We define*

$$\mathcal{E}_h[T, P] = \bigwedge_{i\in\omega} \Big(\mathcal{E}[(V \wedge \bigwedge_{\delta_j \in D_1 \cup \ldots \cup D_{i-1}} (g_j \rightarrow c(\delta_j))\,,\ D_i)]\Big),$$

where $G = \{g_i \mid \delta_i \in D\}$ is the same indexed set of globally new variables corresponding to D as above appearing in each $\mathcal{E}[\cdot]$.

Lemma 4. *Let W, (D, V), G, and P be as in Definition 6. Furthermore, for any set $C \subseteq D$, define the interpretation $M_C \subseteq G$ such that $g_i \in M_C$ iff $\delta_i \in C$, for $1 \leq i \leq n$.*

Then, $Cn(V \cup c(C))$ is a hierarchic extension of T with respect to P iff $\mathcal{E}_h[T, P]$ is true under M_C.

Theorem 8. *Paraconsistent consequence relations $\vdash_{ch}$, $\vdash^{\pm}_{ch}$, $\vdash_{sh}$, $\vdash^{\pm}_{sh}$, $\vdash_{ph}$, and $\vdash^{\pm}_{ph}$ are expressible in the same manner as in Theorems 6 and 7 by replacing $\mathcal{E}[T]$ with $\mathcal{E}_h[T, P]$.*

5 Complexity Issues

In the sequel, we derive complexity results for deciding paraconsistent consequence in all variants discussed previously. We show that all considered tasks are located at the second level of the polynomial hierarchy. This is in some sense not surprising, because the current approach relies on deciding whether a given formula is contained in an extension of a suitably constructed default theory. This problem was shown to be Σ_2^P-complete by Gottlob [11], even if normal default theories are considered. However, this completeness result is not directly applicable here because of the specialised default theories in the present setting. Furthermore, for dealing with hierarchic extensions, it turns out that the complexity remains at the second level of the polynomial hierarchy as well. This result is interesting, since the definition of hierarchic extensions is somewhat more elaborate than standard extensions. In any case, this observation mirrors in some sense complexity results derived for cumulative default logic (cf. [12]).

In the same way as the satisfiability problem of classical propositional logic is the "prototypical" problem of NP, i.e., being an NP-complete problem, the satisfiability problem of QBFs in *prenex form* possessing k quantifier alternations is the "prototypical" problem of the k-th level of the polynomial hierarchy, as expressed by the following well-known result:

Proposition 2 ([16]). *Given a propositional formula ϕ whose atoms are partitioned into $i \geq 1$ sets $P_1, \ldots, P_i$, deciding whether $\exists P_1 \forall P_2 \ldots \mathsf{Q}_i P_i \phi$ is true is Σ_i^P-complete, where $\mathsf{Q}_i = \exists$ if i is odd and $\mathsf{Q}_i = \forall$ if i is even, Dually, deciding whether $\forall P_1 \exists P_2 \ldots \mathsf{Q}'_i P_i \phi$ is true is Π_i^P-complete, where $\mathsf{Q}'_i = \forall$ if i is odd and $\mathsf{Q}_i = \exists$ if i is even.*

Given the above characterisations, we can estimate upper complexity bounds for the reasoning problems discussed in Section 3 simply by inspecting the quantifier order of the respective QBF encodings. This can be argued as follows. First of all, by applying quantifier transformation rules similar to ones in first-order logic, each of the above QBF encodings can be transformed in polynomial time into a QBF in prenex form having exactly one quantifier alternation. Then, by invoking Proposition 2 and observing that completeness of a decision problem D for a complexity class C implies membership of D in C, the quantifier order of the resultant QBFs determines in which class of the polynomial hierarchy the corresponding reasoning task belongs to.

Table 1. Complexity results for all paraconsistent consequence relations.

	$T_0 = (D_\Sigma, W^\pm)$	$T_1 = (D^1_\Sigma, W^\pm_I)$	$T_2 = (D^2_\Sigma, W^\pm_I)$
$\vdash_c$	Σ^P_2	Σ^P_2	Σ^P_2
$\vdash_s$	Π^P_2	Π^P_2	Π^P_2
$\vdash_p$	Π^P_2	in Π^P_2	in Π^P_2
$\vdash^\pm_c$	Σ^P_2	-	-
$\vdash^\pm_s$	Π^P_2	-	-
$\vdash^\pm_p$	in Π^P_2	-	-
$\vdash_{ch}$	Σ^P_2	Σ^P_2	Σ^P_2
$\vdash_{sh}$	Π^P_2	Π^P_2	Π^P_2
$\vdash_{ph}$	Π^P_2	in Π^P_2	in Π^P_2
$\vdash^\pm_{ch}$	Σ^P_2	-	-
$\vdash^\pm_{sh}$	Π^P_2	-	-
$\vdash^\pm_{ph}$	in Π^P_2	-	-

Applying this method to our considered tasks, we obtain that credulous paraconsistent reasoning lies in Σ^P_2, whilst skeptical and prudent paraconsistent reasoning are in Π^P_2. Furthermore, note that the QBFs expressing paraconsistent reasoning using the concept of hierarchical extensions share exactly the same quantifier structures as those using ordinary extensions.

Concerning lower complexity bounds, it turns out that most of the above given estimations are *strict*, i.e., the considered decision problems are hard for the respective complexity classes. The results are summarised in Table 1. There, all entries denote completeness results, except where a membership relation is explicitly stated. The following theorem summarises these relations:

Theorem 9. *The complexity results in Table 1 hold both for ordinary as well as for hierarchical extensions of* T_i ($i = 0, 1, 2$) *as underlying inference principle.*

Some of these complexity results have already been shown elsewhere. As pointed out in [2], prudent consequence, $W \vdash_p \varphi$, on the basis of the default theory $(D_\Sigma, W^\pm)$ captures the notion of *free-consequences* as introduced in [1]. This formalism was shown to be Π^P_2-complete in [4].

Finally, [5] considers the complexity of a number of different paraconsistent reasoning principles, among them the completeness results for $\vdash_s$ and $\vdash^\pm_s$. Moreover, that paper extends the intractability results to some restricted subclasses as well.

6 Discussion

We have shown how paraconsistent inference problems within the framework of signed systems can be axiomatised by means of quantified Boolean formulas. This approach has several benefits: First, the given axiomatics provides us with further insight about how paraconsistent reasoning works within the framework of signed systems. Second, this axiomatisation allows us to furnish upper bounds for precise complexity results, going beyond those presented in [5]. Last but not least, we obtain a straightforward implementation technique of paraconsistent reasoning in signed systems by appeal to existing QBF solvers.

For implementing our approach, we rely on the existing system QUIP [7, 13, 8]. The general architecture of 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 signed system and a specified reasoning task) into the corresponding quantified Boolean formula, which is then sent to the QBF-evaluator. The current version of QUIP provides interfaces to most of the currently available QBF-solvers. 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. This task relies on a protocol mapping of internal variables of the generated QBF into concepts of the problem description.

References

1. S. Benferhat, D. Dubois, and H. Prade. Argumentative Inference in Uncertain and Inconsistent Knowledge Bases. In *Proc. UAI-93*, pages 411–419, 1993.
2. P. Besnard and T. Schaub. Signed Systems for Paraconsistent Reasoning. *Journal of Automated Reasoning*, 20:191–213, 1998.
3. M. Cadoli, A. Giovanardi, and M. Schaerf. An Algorithm to Evaluate Quantified Boolean Formulae. In *Proc. AAAI-98*, pages 262–267. AAAI Press, 1998.
4. C. Cayrol, M. Lagasquie-Schiex, and T. Schiex. Nonmonotonic Reasoning: From Complexity to Algorithms. *Ann. of Mathematics and Artificial Intelligence*, 22(3–4):207–236, 1998.
5. S. Coste-Marquis and P. Marquis. Complexity Results for Paraconsistent Inference Relations. In *Proc. KR-02*, pages 61–72, 2002.
6. J. Delgrande, T. Schaub, H. Tompits, and S. Woltran. On Computing Solutions to Belief Change Scenarios. In *Proc. ECSQARU-01*, pages 510–521. Springer Verlag, 2001.
7. U. Egly, T. Eiter, H. Tompits, and S. Woltran. Solving Advanced Reasoning Tasks Using Quantified Boolean Formulas. In *Proc. AAAI-00*, pages 417–422. AAAI Press, 2000.
8. T. Eiter, V. Klotz, H. Tompits, and S. Woltran. Modal Nonmonotonic Logics Revisited: Efficient Encodings for the Basic Reasoning Tasks. In *Proc. TABLEAUX-02*, 2002. To appear.
9. R. Feldmann, B. Monien, and S. Schamberger. A Distributed Algorithm to Evaluate Quantified Boolean Formula. In *Proc. AAAI-00*, pages 285–290. AAAI Press, 2000.
10. E. Giunchiglia, M. Narizzano, and A. Tacchella. QuBE: A System for Deciding Quantified Boolean Formulas Satisfiability. In *Proc. IJCAR-01*, pages 364–369. Springer Verlag, 2001.
11. G. Gottlob. Complexity Results for Nonmonotonic Logics. *Journal of Logic and Computation*, 2(3):397–425, 1992.
12. G. Gottlob and Z. Mingyi. Cumulative Default Logic: Finite Characterization, Algorithms, and Complexity. *Artificial Intelligence*, 69(1–2):329–345, 1994.
13. 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 Verlag, 2001.
14. R. Reiter. A Logic for Default Reasoning. *Artificial Intelligence*, 13(1–2):81–132, 1980.
15. H. Tompits. Expressing Default Abduction Problems as Quantified Boolean Formulas. In *AI Communications*, 2002. To appear.
16. C. Wrathall. Complete Sets and the Polynomial-Time Hierarchy. *Theoretical Computer Science*, 3(1):23–33, 1976.