Paraconsistent Reasoning via Quantified Boolean Formulas, II: Circumscribing Inconsistent Theories*

Philippe Besnard[1], Torsten Schaub[2**], Hans Tompits[3], and Stefan Woltran[3]

[1] IRIT-CNRS,
118, route de Narbonne, F–31062 Toulouse Cedex, France
besnard@irit.fr

[2] Institut für Informatik, Universität Potsdam,
Postfach 90 03 27, D–14439 Potsdam, Germany
torsten@cs.uni-potsdam.de

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

Abstract. Through minimal-model semantics, three-valued logics provide an interesting formalism for capturing reasoning from inconsistent information. However, the resulting paraconsistent logics lack so far a uniform implementation platform. Here, we address this and specifically provide a translation of two such paraconsistent logics into the language of quantified Boolean formulas (QBFs). These formulas can then be evaluated by off-the-shelf QBF solvers. In this way, we benefit from the following advantages: First, our approach allows us to harness the performance of existing QBF solvers. Second, different paraconsistent logics can be compared with in a unified setting via the translations used. We alternatively provide a translation of these two paraconsistent logics into quantified Boolean formulas representing circumscription, the well-known system for logical minimization. All this forms a case study inasmuch as the other existing minimization-based many-valued paraconsistent logics can be dealt with in a similar fashion.

1 Introduction

The capability of reasoning in the presence of inconsistencies constitutes a major challenge for any intelligent system because in practical settings it is common to have contradictory information. In fact, despite its many appealing features for knowledge representation and reasoning, classical logic falls in a trap: A single contradiction may wreck an entire reasoning system, since it may allow for deriving any proposition. This comportment is due to the fact that a contradiction denies any classical two-valued model, since a proposition must be either true or false. We thus aim at providing formal reasoning systems satisfying the *principle of paraconsistency*: $\{\alpha, \neg\alpha\} \not\vdash \beta$ for some α, β. In other words, given a contradictory set of premises, this should not necessarily lead to concluding all formulas.

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

** Affiliated with the School of Computing Science at Simon Fraser University, Burnaby, Canada.

The idea underlying the approaches elaborated upon in this paper is to counterbalance the effect of contradictions by providing a third truth value that accounts for contradictory propositions. As already put forth in [27], this provides us with inconsistency-tolerating three-valued models. However, this approach turns out to be rather weak in that it invalidates certain classical inferences, even if there is no contradiction. Intuitively, this is because there are too many three-valued models, in particular those assigning the inconsistency-tolerating truth-value to propositions that are unaffected by contradictions. For instance, the three-valued logic LP [27] denies inference by disjunctive syllogism. That is, β is not derivable from the (consistent!) premise $(\alpha \vee \beta) \wedge \neg\alpha$. As pointed out in [15], this deficiency applies also to the closely related paraconsistent systems J_3 [17], L [22], and RP [19]. As a consequence, none of the aforementioned systems coincides with classical logic when reasoning from consistent premises.

The pioneering work to overcome this deficiency was done by Priest in [28]. The key idea is to restrict the set of three-valued models by taking advantage of some preference criterion that aims at "minimizing inconsistency". In this way, a "maximum" of a classically inconsistent knowledge base should be recovered. While minimization is understood in Priest's seminal work [28], proposing his logic LP_m, as preferring three-valued models as close as possible to two-valued interpretations, the overall approach leaves room for different preference criteria. Another criterion is put forth in [9] by giving more importance to the given knowledge base. In this approach, one prefers three-valued models that are as similar as possible to two-valued models of the knowledge base in the sense that those models assign *true* to as many items of the knowledge base as possible. Furthermore, [21] considers cardinality-based versions of the last two preference criteria. Even more criteria are conceivable by distinguishing symbols having different importance.

However, up to know, all these advanced approaches lack effectively implementable inference methods. While Priest defines LP_m in purely semantical terms, a Hilbert calculus comprising 26 axiom schemata is proposed by Besnard and Schaub [9] for axiomatizing their approach. Also, inference is not at issue in [21]. This shortcoming is addressed in this paper. To wit, we develop translations for the three-valued paraconsistent logics defined in [28] and [9]. More precisely, our translations allow for mapping the respective entailment problems into the satisfiability problem for *quantified Boolean formulas* (QBFs). These formulas can then be evaluated by off-the-shelf QBF solvers. The motivation of this particular approach to implementing these logics (as opposed to more direct calculizations) stems from its unique uniformity, even beyond the framework of three-valued logics. In fact, we have already developed in a companion paper [11] similar translations for a rather different family of paraconsistent logics, called *signed systems* [10]; a forthcoming paper deals with approaches to paraconsistency based on the selection of maximally consistent subsets [24, 8].

Our general methodology offers several benefits: First, we obtain uniform axiomatizations of rather different approaches. This allows us to compare different paraconsistent logics in a unified setting. Second, once such an axiomatization is available, existing QBF solvers can be used for implementation in a uniform manner. The availability of efficient QBF solvers, like the systems described in [12, 20, 6], makes such a rapid prototyping approach practicably applicable. Third, these axiomatizations provide a direct access to the complexity of the original approach. Conversely, we can exploit existing complexity results for ensuring the adequateness of our axiomatizations. 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.

2 Paraconsistent Three-Valued Logics

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

An interpretation is a function $v : \mathcal{P} \rightarrow \{t, f, o\}$ extending to $\overline{v} : \mathcal{L} \rightarrow \{t, f, o\}$ according to the truth tables below.

$\bot$	
	f

$\top$	
	t

$\neg$	
t	f
f	t
o	o

$\wedge$	t	f	o
t	t	f	o
f	f	f	f
o	o	f	o

$\vee$	t	f	o
t	t	t	t
f	t	f	o
o	t	o	o

$\rightarrow$	t	f	o
t	t	f	o
f	t	t	t
o	t	f	o

(1)

We sometimes leave an interpretation v implicit and write $p : x$ instead of $v(p) = x$, for $x \in \{t, f, o\}$. An interpretation v is said to be *two-valued* whenever $v(p) \in \{t, f\}$ for all $p \in \mathcal{P}$; otherwise, it is *three-valued*. A *three-valued model* of a formula α is an interpretation that assigns either t or o to α. Modelhood extends to sets of formulas in the standard way. As usual, given a set S of formulas and a formula ϕ, we define $S \models \phi$ if each model of S is a model of ϕ. Whenever necessary, we write $\models_3$ and $\models_2$ to distinguish three-valued from two-valued entailment.

Note that the truth value of $\alpha \rightarrow \beta$ differs from that of $\neg\alpha \vee \beta$ only in the case of $v = \{\alpha : o, \beta : f\}$ resulting in $\overline{v}(\alpha \rightarrow \beta) = f$ and $\overline{v}(\neg\alpha \vee \beta) = o$. This difference is prompted by the fact that t and o indicate modelhood, which motivates the assignment of the same truth values to $\alpha \rightarrow \beta$ no matter whether we have $\alpha : t$ or $\alpha : o$. This has actually to do with the difference between *modus ponens* and *disjunctive syllogism*: The latter yields β from $\alpha \wedge \neg\alpha \wedge \neg\beta$ because $\alpha \vee \beta$ follows from α. The overall inference seems wrong because in the presence of $\alpha \wedge \neg\alpha$, $\alpha \vee \beta$ is satisfied (by $\alpha : o$) with no need for β to be t. As pointed out in [21], one may actually view $\rightarrow$ as *"the 'right' generalization of classical implication because $\rightarrow$ is the internal implication connective [5] for the defined inference relation in the sense that a deduction (meta)theorem holds for it: $\Sigma \wedge \alpha \models_3 \beta$ iff $\Sigma \models_3 \alpha \rightarrow \beta$."* On the other hand, a formula composed of the connectives $\neg, \vee$, and $\wedge$ can never be inconsistent; that is, each such formula has at least one three-valued model [13]. Finally, we mention that the entailment problem for $\models_3$ is *coNP*-complete, no matter whether $\rightarrow$ is included or not [26, 13, 15].

As mentioned in the introductory section, Priest's logic LP_m [28] was conceived to overcome the failure of disjunctive syllogism in *LP* [27]. *LP* amounts to the three-valued logic obtained by restricting $\mathcal{L}$ to connectives $\neg, \vee$ and $\wedge$ (and defining $\alpha \rightarrow \beta$ as $\neg\alpha \vee \beta$). In LP_m modelhood is then limited to models containing a minimal number of *propositional variables* being assigned o. This allows for drawing *"all classical inferences except where inconsistency makes them doubtful anyway"* [28]. Formally, the consequence relation of LP_m can be defined as follows. For three-valued interpretations v, w, define the partial ordering

$$v \leq_m w \quad \text{iff} \quad \{p \in \mathcal{P} \mid v(p) = o\} \subseteq \{p \in \mathcal{P} \mid w(p) = o\} \; .$$

Then, $T \models_m \phi$ iff every three-valued model of T that is minimal with respect to $\leq_m$ is a three-valued model of ϕ.

Unlike this, the approach of Besnard and Schaub [9] prefers three-valued models that assign *true* to as many items of the knowledge base T as possible: For three-valued interpretations v, w, define the partial ordering

$$v \leq_n w \quad \text{iff} \quad \{\phi \in T \mid \overline{v}(\phi) = o\} \subseteq \{\phi \in T \mid \overline{w}(\phi) = o\} \; .$$

Then, $T \models_n \phi$ iff each three-valued model of T which is $\leq_n$-minimal is a three-valued model of ϕ.

The major difference between the last two approaches is that the restriction of modelhood in LP_m focuses on models as close as possible to two-valued *interpretations*, while the one in the last approach aims at models next to two-valued *models* of the considered premises. According to [9], the effects of making the formula select its preferred models can be seen by looking at $T = \{p, \neg p, (\neg p \vee q)\}$: While LP_m yields two $\leq_m$-preferred models, $\{p : o, q : t\}$ and $\{p : o, q : f\}$, from which one obtains $p \wedge \neg p$, the second approach yields q as additional conclusion. In fact, $\{p : o, q : t\}$ is the only $\leq_n$-preferred model of the premises $\{p, \neg p, (\neg p \vee q)\}$; it assigns t to $(\neg p \vee q)$, while this premise is attributed o by the second $\leq_m$-preferred model $\{p : o, q : f\}$; hence the latter is not $\leq_n$-preferred. So, while $T \not\models_m q$ and $T \models_n q$, we note that $T \cup \{(p \vee \neg q)\} \not\models_l q$ for $l = m, n$. On the other hand, $\models_n$ is clearly more syntax-dependent than $\models_m$ since the items within the knowledge base are used for distinguishing $\leq_n$-preferred models.

In fact, both inference relations $\models_m$ and $\models_n$ amount to their classical (two-valued) counterpart, whenever the set of premises is classically consistent. Also, it is shown in [15] that deciding entailment for $\models_m$ and $\models_n$ is Π_2^p-complete, no matter whether $\rightarrow$ is included or not. A logical analysis of both relations can be found in [21] and in the original literature [28, 9].

3 Axiomatizing Three-Valued Paraconsistent Logics

In what follows, we provide axiomatizations of the three-valued paraconsistent logics introduced in the last section in terms of QBFs.

Quantified Boolean formulas. As a conservative extension of classical propositional logic, *quantified Boolean formulas* (QBFs) generalize 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. Given that $\mathcal{K}$ is the language of QBFs over a set $\mathcal{P}$ of propositional variables, the semantical meaning of QBFs can be defined as follows: An interpretation is a function $v : \mathcal{P} \rightarrow \{t, f\}$ extending to $\hat{v} : \mathcal{K} \rightarrow \{t, f\}$ according to the truth tables in (1) and the following two conditions, for every $\Phi \in \mathcal{K}$,

$$\hat{v}(\forall p\, \Phi) = \hat{v}(\Phi[p/\top] \wedge \Phi[p/\bot]) \quad \text{and} \quad \hat{v}(\exists p\, \Phi) = \hat{v}(\Phi[p/\top] \vee \Phi[p/\bot]) \; .$$

We write $\Phi[p_1/\phi_1, \ldots, p_n/\phi_n]$ to denote the result of uniformly substituting each free occurrence[4] of a variable p_i in Φ by a formula ϕ_i, for $1 \leq i \leq n$. If Φ contains no free variable occurrences, then Φ is *closed*. Closed QBFs are either true under every interpretation or false under every interpretation. Hence, for closed QBFs there is no need to refer to particular interpretations.

In the sequel, we use the following abbreviations: The set of all atoms occurring in a formula ϕ is denoted by $var(\phi)$. Similarly, for a set S of formulas, $var(S) = \bigcup_{\phi \in S} var(\phi)$. 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 $\bigwedge_{i=1}^{n}(\phi_i \rightarrow \psi_i)$, and $S < T$ stands for $S \leq T \wedge \neg(T \leq S)$.

Encoding three-valued logic. We start with encoding the truth evaluation of the three-valued logic given in Section 2 by means of classical propositional logic.

To this end, we introduce for each atom p a globally new atom p' and define $\mathcal{P}' = \{p' \mid p \in \mathcal{P}\}$ for a given alphabet $\mathcal{P}$.

Let v be a three-valued interpretation over alphabet $\mathcal{P}$. We define the *associated two-valued interpretation* v_2 by setting

$$\begin{aligned} v_2(p) = v_2(p') = t &\quad \text{if } v(p) = t; \\ v_2(p) = v_2(p') = f &\quad \text{if } v(p) = f; \\ v_2(p) = f \text{ and } v_2(p') = t &\quad \text{if } v(p) = o, \end{aligned}$$

for any $p \in \mathcal{P}$ and any $p' \in \mathcal{P}'$. Conversely, for a given two-valued interpretation v over alphabet $\mathcal{P} \cup \mathcal{P}'$ such that $\overline{v}(p \rightarrow p') = t$, we define the *associated three-valued interpretation* v_3 by setting

$$v_3(p) = \begin{cases} v(p) & \text{if } v(p) = v(p') \\ o & \text{if } v(p) = f \text{ and } v(p') = t \end{cases}$$

for any $p \in \mathcal{P}$.

Moreover, we need the following parameterized translation:

Definition 1. *For $p \in \mathcal{P}$ and $\phi, \psi \in \mathcal{L}$, we define*

1. (*a*) $\tau(p, t) = p$;
   (*b*) $\tau(p, f) = \neg p'$;
   (*c*) $\tau(p, o) = \neg p \wedge p'$;
2. (*a*) $\tau(\neg\phi, t) = \tau(\phi, f)$;

[4] An occurrence of a propositional variable p in a QBF Φ is *free* if it does not appear in the scope of a quantifier $\mathsf{Q}p$ ($\mathsf{Q} \in \{\forall, \exists\}$).

   (*b*) $\tau(\neg\phi, f) = \tau(\phi, t)$;
   (*c*) $\tau(\neg\phi, o) = \tau(\phi, o)$;
3. (*a*) $\tau(\phi \wedge \psi, t) = \tau(\phi, t) \wedge \tau(\psi, t)$;
   (*b*) $\tau(\phi \wedge \psi, f) = \tau(\phi, f) \vee \tau(\psi, f)$;
   (*c*) $\tau(\phi \wedge \psi, o) = \neg\tau(\phi \wedge \psi, f) \wedge \neg\tau(\phi \wedge \psi, t)$;
4. (*a*) $\tau(\phi \vee \psi, t) = \tau(\phi, t) \vee \tau(\psi, t)$;
   (*b*) $\tau(\phi \vee \psi, f) = \tau(\phi, f) \wedge \tau(\psi, f)$;
   (*c*) $\tau(\phi \vee \psi, o) = \neg\tau(\phi \vee \psi, t) \wedge \neg\tau(\phi \vee \psi, f)$;
5. (*a*) $\tau(\phi \rightarrow \psi, t) = \tau(\phi, f) \vee \tau(\psi, t)$;
   (*b*) $\tau(\phi \rightarrow \psi, f) = \neg\tau(\phi, f) \wedge \tau(\psi, f)$;
   (*c*) $\tau(\phi \rightarrow \psi, o) = \neg\tau(\phi, f) \wedge \tau(\psi, o)$.

For computing the three-valued models of a set of formulas $T = \{\phi_1, \ldots, \phi_n\}$, we use $\bigwedge_{\phi \in T} \neg\tau(\phi, f)$ and abbreviate the latter by $\neg\tau(T, f)$.[5]

For example, consider $T = \{p, \neg p, (\neg p \vee q)\}$. We get:

$$\begin{aligned}\neg\tau(T, f) &= \neg\tau(p, f) \wedge \neg\tau(\neg p, f) \wedge \neg\tau((\neg p \vee q), f) \\ &= \neg\neg p' \wedge \neg\tau(p, t) \wedge \neg(\tau(\neg p, f) \wedge \tau(q, f)) \\ &= p' \wedge \neg p \wedge \neg(\tau(p, t) \wedge \neg q') \\ &= p' \wedge \neg p \wedge (\neg p \vee \neg\neg q') \\ &= p' \wedge \neg p\end{aligned}$$

The resulting formula possesses four two-valued models, all of which assign $p : f$ and $p' : t$ while varying on q and q'. In order to establish a correspondence among the four two-models of $\neg\tau(T, f)$ and the three three-valued models of T, assigning o to p and varying on q, the relation between the two alphabets $\mathcal{P}$ and $\mathcal{P}'$ must be fixed. In fact, this is accomplished by adding $r \rightarrow r'$ for every $r \in \mathcal{P}$.

In this way, we obtain the following result.

Theorem 1. *Let ϕ be a formula with $P = var(\phi)$, let $P' = \{p' \mid p \in P\}$, and let $x \in \{t, f, o\}$.*

Then, the following conditions hold:

1. *For any three-valued interpretation v over $\mathcal{P}$, if $\overline{v}(\phi) = x$, then $\overline{v}_2((P \leq P') \wedge \tau(\phi, x)) = t$, where v_2 is the associated two-valued interpretation of v.*
2. *For any two-valued interpretation v over $\mathcal{P} \cup \mathcal{P}'$, if $\overline{v}((P \leq P') \wedge \tau(\phi, x)) = t$, then $\overline{v}_3(\phi) = x$, where $\overline{v}_3$ is the associated three-valued interpretation of v.*

Since the formula $\tau(\phi, t) \vee \tau(\phi, f) \vee \tau(\phi, o)$ is clearly a tautology of classical logic, we immediately get the following relation between the three-valued models of a theory and the two-valued models of the corresponding encoding:

Corollary 1. *Let T be a finite set of formulas with $P = var(T)$ and let $P' = \{p' \mid p \in P\}$.*

Then, there is a one-to-one correspondence between the three-valued models of T and the two-valued models of the formula

$$(P \leq P') \wedge \neg\tau(T, f). \tag{2}$$

[5] Note that generally $\bigwedge_{\phi \in T} \tau(\phi, x) \neq \tau(\bigwedge_{\phi \in T} \phi, x)$.

In particular, the three-valued model of T corresponding to a two-valued model v of (2) is given by the associated three-valued interpretation v_3 of v.

For illustration, consider $T = \{p, \neg p, (\neg p \vee q)\}$ along with

$$(\{p, q\} \leq \{p', q'\}) \wedge \neg\tau(T, f) \quad = \quad (p \rightarrow p') \wedge (q \rightarrow q') \wedge (p' \wedge \neg p) \ .$$

Unlike above, we obtain now three two-valued models, $\{p : f, p' : t, q : t, q' : t\}$, $\{p : f, p' : t, q : f, q' : t\}$, and $\{p : f, p' : t, q : f, q' : f\}$, being in a one-to-one correspondence with the three three-valued models, $\{p : o, q : t\}$, $\{p : o, q : o\}$, and $\{p : o, q : f\}$, of T, respectively.

The role of the implications $\{p, q\} \leq \{p', q'\}$ can be further illustrated by looking at the following translation:

$$\begin{aligned}
\tau(T, t) &= \tau(p, t) \wedge \tau(\neg p, t) \wedge \tau((\neg p \vee q), t) \\
&= p \wedge \tau(p, f) \wedge (\tau(\neg p, t) \vee \tau(q, t)) \\
&= p \wedge \neg p' \wedge (\tau(p, f) \vee q) \\
&= p \wedge \neg p' \wedge (\neg p' \vee q) \\
&= p \wedge \neg p'
\end{aligned}$$

While the last formula admits four two-valued models, the formula $(\{p, q\} \leq \{p', q'\}) \wedge \tau(T, t)$ has no two-valued model, which corresponds to the fact that T has no three-valued model assigning t to all members of T.

Consequently, since there are no three-valued models assigning t "to" T, the formulas $\neg\tau(T, f)$ and $\tau(T, o)$ must be equivalent; this can be verified as follows:

$$\begin{aligned}
\tau(T, o) &= \neg\tau(T, f) \wedge \neg\tau(T, t) \\
&= \neg\big(\tau(p, f) \vee \tau(\neg p, f) \vee \tau(\neg p \vee q, f)\big) \wedge \neg(p \wedge \neg p') \\
&= \neg\big(\neg p' \vee p \vee (\tau(\neg p, f) \wedge \tau(q, f))\big) \wedge (\neg p \vee p') \\
&= \neg\big(\neg p' \vee p \vee (p \wedge \neg q')\big) \wedge (\neg p \vee p') \\
&= p' \wedge \neg p \wedge (\neg p \vee q') \wedge (\neg p \vee p') \\
&= p' \wedge \neg p
\end{aligned}$$

Encoding three-valued paraconsistent logics. To begin with, it is instructive to see that the previous elaboration already allows for a straightforward encoding of three-valued entailment, and, in particular, inference in the logic *LP* [27]:

Definition 2. *Let T be a set formulas and ϕ a formula.*

For $P = var(T \cup \{\phi\})$, we define

$$\mathcal{T}_3(T, \phi) = \forall P, P' \Big(\big((P \leq P') \wedge \neg\tau(T, f) \big) \rightarrow \neg\tau(\phi, f) \Big) \ .$$

Then, we have the following result.

Theorem 2. *$T \models_3 \phi$ iff $\mathcal{T}_3(T, \phi)$ is true.*

To be precise, we obtain (original) inference in *LP* [27] when restricting T and ϕ to formulas whose connectives are among $\neg$, $\wedge$, and $\vee$ only.

Let us now turn to Priest's logic LP_m [28]. For this, we must, roughly speaking, enhance the encoding of *LP* in order to account for the principle of "minimizing inconsistency" used in LP_m. This is accomplished in the next definition by means of the QBF named $Min_m(T)$.

Definition 3. *Let T be a set formulas with $P = var(T)$, V an indexed set of globally new atoms corresponding to P, and ϕ a formula. Moreover, let $O_P = \{\tau(p, o) \mid p \in P\}$ and $O_V = \{\tau(v, o) \mid v \in V\}$.*

We define

$$Min_m(T) = (P \leq P') \wedge \neg\exists V, V'\Big((O_V < O_P) \wedge (V \leq V') \wedge \neg\tau(T[P/V], f)\Big)$$

and, for $R = P \cup var(\phi)$,

$$\mathcal{T}_m(T, \phi) = \forall R, R'\Big(\big(Min_m(T) \wedge \neg\tau(T, f)\big) \rightarrow \neg\tau(\phi, f)\Big).$$

For illustration, let us return to $T = \{p, \neg p, (\neg p \vee q)\}$. We have $P = \{p, q\}$ and correspondingly $V = \{u, v\}$. We start our analysis on the subformula

$$(O_V < O_P) \wedge (V \leq V') \wedge \neg\tau(T[P/V], f), \tag{3}$$

having

$$\begin{aligned} O_V < O_P = \quad & (\tau(u, o) \rightarrow \tau(p, o)) \wedge (\tau(v, o) \rightarrow \tau(q, o)) \ \wedge \\ & \neg\big((\tau(p, o) \rightarrow \tau(u, o)) \wedge (\tau(q, o) \rightarrow \tau(v, o))\big). \end{aligned}$$

From the definition of $\leq$ and $<$, one can see that $(O_V < O_P) \wedge (V \leq V')$ is true under a two-valued interpretation v iff

– for any variable from V assigned o under the associated interpretation v_3, the corresponding variable from P is also assigned o under v_3; and
– there exists at least one variable from V which is not assigned o under v_3, although the corresponding variable from P is assigned o under v_3.

Additionally, v has to be a two-valued model of $\neg\tau(f, T[P/V])$. By Corollary 1, v_3 then has to be a three-valued model of $T[P/V]$. From our previous discussion and by renaming, we know that $T[P/V]$ possesses three three-valued models, viz. $\{u : o, v : t\}$, $\{u : o, v : o\}$, and $\{u : o, v : f\}$. In the case of model $\{u : o, v : o\}$, we cannot find an assignment to p, q which has more variables being assigned o. The other two cases extend to two two-valued models, v' and v'', of (3) with their associated three-valued interpretations v_3' and v_3'' given by $\{p : o, q : o, u : o, v : t\}$ and $\{p : o, q : o, u : o, v : f\}$, respectively. Now, one can check that the only three-valued interpretation w such that $Min_m(T)$ is false under w_2 is $\{p : o, q : o\}$. Recalling the three-valued models of T, we have that the two-valued models of $Min_m(T) \wedge \neg\tau(T, f)$ yield two three-valued models, $\{p : o, q : t\}$ and $\{p : o, q : f\}$.

In general, we have the following result.

Theorem 3. *$T \models_m \phi$ iff $\mathcal{T}_m(T, \phi)$ is true.*

To be precise, we obtain (original) inference in *LP*$_m$ [28] when restricting T and ϕ to formulas whose connectives are among $\neg$, $\wedge$, and $\vee$ only.

Analogously, we can now give an axiomatization of Besnard and Schaub's approach [9].

Definition 4. *Let T be a set formulas with $P = var(T)$, Q an indexed set of globally new atoms corresponding to P, and ϕ a formula. Moreover, let $O_T = \{\tau(\phi, o) \mid \phi \in T\}$ and $O_{T[P/Q]} = \{\tau(\phi, o) \mid \phi \in T[P/Q]\}$.*

We define

$$Min_n(T) = (P \leq P') \wedge \neg\exists Q, Q'\Big((O_{T[P/Q]} < O_T) \wedge (Q \leq Q')\Big)$$

and, for $R = P \cup var(\phi)$,

$$\mathcal{T}_n(T, \phi) = \forall R, R'\Big(\big(Min_n(T) \wedge \neg\tau(T, f)\big) \rightarrow \neg\tau(\phi, f)\Big).$$

The salient difference between the previous definition and the one given in Definition 3 manifests itself in the sets O_P, O_V and $O_T, O_{T[P/Q]}$, respectively. Note that the latter take into account the original set of premises T so that the translation formula $\neg\tau(T[P/V], f)$ can be dropped.

Theorem 4. *$T \models_n \phi$ iff $\mathcal{T}_n(T, \phi)$ is true.*

Alternative Encodings. In view of the discussion given below, we may alternatively capture both approaches as follows.

To begin with, concerning *LP*$_m$, we introduce additional new variables $S = \{s_p \mid p \in var(T)\}$ and $S' = \{s'_p \mid p \in var(T)\}$, and define

$$\begin{aligned} Min_{m'}(T) = {} & (P \leq P') \wedge (S \leq O_P) \wedge \\ & \neg\exists S', Q, Q'\Big((S' < S) \wedge (Q \leq Q') \wedge (S' \leq O_Q) \wedge \neg\tau(T[P/Q], f)\Big) \end{aligned}$$

and

$$\mathcal{T}_{m'}(T, \phi) = \forall S, P, P'\Big(\big(Min_{m'}(T) \wedge \neg\tau(T, f)\big) \rightarrow \neg\tau(\phi, f)\Big).$$

For Besnard and Schaub's approach [9], on the other hand, we similarly introduce additional new variables according to the elements of T, viz. $S = \{s_\phi \mid \phi \in T\}$ and $S' = \{s'_\phi \mid \phi \in T\}$, and define

$$\begin{aligned} Min_{n'}(T) = {} & (P \leq P') \wedge (S \leq O_P) \wedge \\ & \neg\exists S', Q, Q'\Big((S' < S) \wedge (Q \leq Q') \wedge (S' \leq O_Q)\Big) \end{aligned}$$

and

$$\mathcal{T}_{n'}(T, \phi) = \forall S, P, P'\Big(\big(Min_{n'}(T) \wedge \neg\tau(T, f)\big) \rightarrow \neg\tau(\phi, f)\Big).$$

In analogy to Theorems 3 and 4, we then obtain the following result:

Theorem 5. *$T \models_\nu \phi$ iff $\mathcal{T}_{\nu'}(T, \phi)$ is true, for both $\nu \in \{m, n\}$.*

Employing Circumscription. In order to shed some more light on the two paraconsistent logics discussed above, let us slightly reformulate their minimization axiom in terms of circumscription [25]: Let T be a propositional theory and (P, Q, Z) a partition of $var(T)$. Assume two (two-valued) models v, v' of T, and define $v \leq_{P;Z} v'$ iff the following conditions are satisfied:

1. $\{q \in Q \mid v(q) = t\} = \{q \in Q \mid v'(q) = t\}$;
2. $\{p \in P \mid v(p) = t\} \subseteq \{p \in P \mid v'(p) = t\}$.

A model v of T is called $(P;Z)$-*minimal* if no model v' of T with $v' \neq v$ satisfies $v' \leq_{P;Z} v$.

Informally, the partition (P, Q, Z) can be interpreted as follows: The set P contains the variables to be minimized, Z are those variables that can vary in minimizing P, and the remaining variables Q are fixed in minimizing P.

Let T be a theory and (P, Q, Z) a partition of $var(T)$, where $P = \{p_1, \ldots, p_n\}$ and $Z = \{z_1, \ldots, z_m\}$. The set of $(P;Z)$-minimal models of T is given by the truth assignments to the QBF

$$Circ(T; P; Z) = T \wedge \neg\exists \tilde{P}, \tilde{Z}\Big((\tilde{P} < P) \wedge T[P/\tilde{P}, Z/\tilde{Z}]\Big),$$

where $\tilde{P} = \{\tilde{p}_1, \ldots, \tilde{p}_n\}$ and $\tilde{Z} = \{\tilde{z}_1, \ldots, \tilde{z}_m\}$ are sets of new variables corresponding to P and Z, respectively.

Then, for $P = var(T)$, we have that $Min_{m'}(T) \wedge \neg\tau(T, f)$ can be written as

$$\mathcal{C}_m(T) = Circ((P \leq P') \wedge (S \leq \{\tau(p, o) \mid p \in P\}) \wedge \neg\tau(T, f); S; P \cup P'),$$

where $S = \{s_p \mid p \in var(T)\}$, and, analogously, $Min_{n'}(T)$ can be written as

$$\mathcal{C}_n(T) = Circ((P \leq P') \wedge (S \leq \{\tau(\phi, o) \mid \phi \in T\}); S; P \cup P')$$

where $S = \{s_\phi \mid \phi \in T\}$.

Summarizing, we have[6]

1. $T \models_m \phi$ iff $\big(\mathcal{C}_m(T) \wedge \neg\tau(T, f)\big) \rightarrow \neg\tau(\phi, f)$ is true;
2. $T \models_n \phi$ iff $\big(\mathcal{C}_n(T) \wedge \neg\tau(T, f)\big) \rightarrow \neg\tau(\phi, f)$ is true.

This demonstrates how the principle of circumscription can be exploited for characterizing the minimization process in the two considered paraconsistent logics.

4 Related work

A whole variety of approaches uses lattices for dealing with inconsistency, e.g., [1, 7, 29]. For instance, [1, 2] describes a system based on four-valued logic that allows for constraining "the most consistent" models in the meta-level by a user-given set of propositions taking classical truth-values only. In fact, in [3] the preference relation $\leq_m$ is generalized to four-valued logics, giving rise to two distinct orderings: Given two four-valued interpretations over truth values $\{t, f, o, o'\}$[7], define

[6] In fact, $(\mathcal{C}_m(T) \wedge \neg\tau(T, f)) \rightarrow \neg\tau(\phi, f)$ is true iff $\mathcal{C}_m(T) \rightarrow \neg\tau(\phi, f)$ is true.

[7] In a four-valued setting, o, o' are usually denoted by $\bot, \top$.

- $v \leq_1 w$ iff $\{p \in \mathcal{P} \mid v(p) = o\} \subseteq \{p \in \mathcal{P} \mid w(p) = o\}$; and
- $v \leq_2 w$ iff $\{p \in \mathcal{P} \mid v(p) \in \{o, o'\}\} \subseteq \{p \in \mathcal{P} \mid w(p) \in \{o, o'\}\}$.

As with $\models_m$, the models minimal with respect to these orderings are then used to define two distinct four-valued consequence relations. Although we do not detail it here, we mention that an appropriate encoding of the underlying four-valued logic (similar to the one given in Definition 1), along with a slightly generalized QBF encoding (similar to the one given in Definition 3), allows for a straightforward encoding of the two four-valued consequence relations by means of QBFs. Interestingly, both four-valued paraconsistent logics have recently been implemented in [4] by appeal to special-purpose circumscription solvers [16]. [14] proposes a translation-based approach to reasoning in the presence of contradictions that translates a logic into a family of other logics, e.g., classical logic into three-valued logics.

Among the existing inference methods for three-valued paraconsistent logics, we mention the following ones. A resolution-based system close to *LP* yet with a stronger disjunction is described in [23]. In fact, there is an indirect way of implementing LP_m because its consequence relation has recently been shown in [15] to be equivalent to a particular relation within the family of signed systems [10], whose inference can also be mapped onto QBFs, as shown in [11]. The resulting encoding is, however, of little interest since it lacks the spirit of "minimizing inconsistency" and thus fails to provide insight into LP_m; also, it is not extendible with the genuine implication $\rightarrow$ or even to alternative approaches such as [9]. We recall from the introductory section that the latter approach was originally axiomatized in [9] by means of a Hilbert system comprising 26 axiom schemata.

5 Conclusion

Considering two paraconsistent logics based on a minimization principle applied to a three-valued logic, we have shown how a translation into the language of quantified Boolean formulas is possible. The translations obtained clearly fall under the same umbrella, giving rise to a uniform setting for the axiomatization of such logics. (In particular, we have provided translations explicitly displaying the connection with circumscription, the classical proof theory for logical minimization.) Moreover, once such an axiomatization is available, existing QBF solvers can be used for implementation without further ado. Having efficient QBF solvers, like the systems described in [12, 20, 6], makes such a rapid prototyping approach practicably applicable. Finally, we remark that what we did allows us, in some sense, to express this kind of paraconsistent reasoning in (higher order) classical propositional logic and so to harness classical reasoning mechanisms from (a conservative extension of) propositional logic.

References

1. O. Arieli and A. Avron. Logical bilattices and inconsistent data. In *Proc. LICS*, pages 468–476, 1994.
2. O. Arieli and A. Avron. Automatic diagnoses for properly stratified knowledge-bases. In *Proc. ICTAI'96*, pages 392–399. IEEE Computer Society Press, 1996.

3. O. Arieli and A. Avron. The value of four values. *Artificial Intelligence*, 102(1):97–141, 1998.
4. O. Arieli and M. Denecker. Modelling paraconsistent reasoning by classical logic. In *Proc. FOIKS'02*, number 2284 in LNCS, pages 1–14. Springer-Verlag, 2002.
5. A. Avron. Simple consequence relations. *Information and Computation*, 92:105–139, 1991.
6. A. Ayari and D. Basin. QUBOS: Deciding quantified Boolean logic using propositional satisfiability solvers. In *Proc. FMCAD 2002*, volume 2517 of *LNCS*, pages 187–201. Springer-Verlag, 2002.
7. N. Belnap. A useful four-valued logic. In J. Dunn and G. Epstein, editors, *Modern Uses of Multiple-Valued Logic*. Reidel, 1977.
8. S. Benferhat, C. Cayrol, D. Dubois, J. Lang, and H. Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. IJCAI'93*, pages 640–647, 1993.
9. P. Besnard and T. Schaub. Circumscribing inconsistency. In *Proc. IJCAI'97*, pages 150–155, 1997.
10. P. Besnard and T. Schaub. Signed systems for paraconsistent reasoning. *Journal of Automated Reasoning*, 20(1-2):191–213, 1998.
11. P. Besnard, T. Schaub, H. Tompits, and S. Woltran. Paraconsistent reasoning via quantified Boolean formulas, I: Axiomatising signed systems. In *Proc. JELIA'02*, volume 2424 of *LNCS*, pages 320–331. Springer-Verlag, 2002.
12. M. Cadoli, A. Giovanardi, and M. Schaerf. An algorithm to evaluate quantified Boolean formulae. In *Proc. AAAI-98*, pages 262–267. AAAI Press, 1998.
13. M. Cadoli and M. Schaerf. On the complexity of entailment in propositional multivalued logics. *Annals of Mathematics and Artificial Intelligence*, 18:29–50, 1996.
14. W. Carnielli, L. Fariñas del Cerro, and M. Lima Marques. Contextual negations and reasoning with contradictions. In *Proc. IJCAI'91*, pages 532–537, 1991.
15. S. Coste-Marquis and P. Marquis. Complexity results for paraconsistent inference relations. In *Proc. KR'02*, pages 61–72, 2002.
16. P. Doherty, W. Łukaszewicz, and A. Szalas. Computing circumscription revisited: A reduction algorithm. *Journal of Automated Reasoning*, 18:297–334, 1997.
17. I. D'Ottaviano and N. da Costa. Sur un problème de Jaśkowski. In *Comptes Rendus de l'Académie des Sciences de Paris*, volume 270, pages 1349–1353, 1970.
18. R. Feldmann, B. Monien, and S. Schamberger. A distributed algorithm to evaluate quantified Boolean formulas. In *Proc. AAAI-00*, pages 285–290. AAAI Press, 2000.
19. A. Frisch. Inference without chaining. In *Proc. IJCAI'87*, pages 515–519, 1987.
20. 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.
21. S. Konieczny and P. Marquis. Three-valued logics for inconsistency handling. In *Proc. JELIA'02*, volume 2424 of *LNCS*, pages 332–344. Springer-Verlag, 2002.
22. H. Levesque. A knowlege-level account of abduction. In *Proc. IJCAI'89*, pages 1061–1067, 1989.
23. F. Lin. Reasoning in the presence of inconsistency. In *Proc. AAAI'87*, pages 139–143, 1987.
24. R. Manor and N. Rescher. On inferences from inconsistent information. *Theory and Decision*, 1:179–219, 1970.
25. J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28:89–116, 1986.
26. D. Mundici. Satisfiability in many-valued sentential logic is NP-complete. *Theoretical Computer Science*, 52(1-2):145–153, 1987.
27. G. Priest. Logic of paradox. *Journal of Philosophical Logic*, 8:219–241, 1979.
28. G. Priest. Reasoning about truth. *Artificial Intelligence*, 39:231–244, 1989.
29. E. Sandewall. A functional approach to non-monotonic logic. *Computational Intelligence*, 1:80–87, 1985.