A Unit Resolution-Based Approach to Tractable and Paraconsistent Reasoning

Sylvie Coste-Marquis and **Pierre Marquis** [1]

Abstract. A family of unit resolution-based paraconsistent inference relations $\vdash^*_\Sigma$ for propositional logic in the clausal case is introduced. Parameter Σ is any consistent set of clauses representing the beliefs which are intended to be exploited through full deduction, i.e., every classical consequence of Σ must be kept w.r.t. $\vdash^*_\Sigma$, whatever the belief base B. Contrariwise to many paraconsistent inference relations, any relation $\vdash^*_\Sigma$ can be decided in time linear in $|B|$ whenever Σ is of bounded size. We show that $\vdash^*_\Sigma$ exhibits several valuable properties, including strong paraconsistency. We also show how a unit propagation algorithm can be simply turned into a decision procedure for $\vdash^*_\Sigma$. We finally show how the $\vdash^*_\Sigma$ family relates to several inference relations proposed so far as approximations of classical entailment.

1 INTRODUCTION

A fundamental limitation of classical inference in the objective of modeling common-sense reasoning is that it trivializes in presence of inconsistency. In order to address such an issue, many approaches to reasoning from contradictory information have been proposed, in various propositional settings; among them are belief revision, belief merging, reasoning from preferred consistent subsets, argumentative logics, paraconsistent logics, etc. In this family, paraconsistent logics are the most basic techniques, in the sense they enable non trivial reasoning from a single contradictory formula, like $a \wedge \neg a \wedge b$. The fact a belief base B is encoded as a single formula coheres with the cognitively plausible assumption that agents typically do not know the reasons of their pieces of beliefs (this assumption is also made in the influential AGM framework for belief revision [1]). Contrastingly, other approaches to inconsistency tolerant reasoning typically deal with several consistent formulas which are jointly inconsistent (e.g., in belief revision, two formulas are considered: the original belief base and the revision formula), hence they are not suited to the case the available data consists of a single inconsistent formula.

Many paraconsistent inference relations have been defined so far. However, most of them are intractable, their complexity lying at the first or the second level of the polynomial hierarchy [3, 4]. This renders them hard to be used when inconsistent belief bases of large size must be dealt with. Furthermore, while tractability can be obtained for some paraconsistent inference relations (e.g., $\models_3$ defined in a three-valued logical framework when CNF formulas are considered [10]), we are not aware of a propositional framework in which strong paraconsistent reasoning can be efficiently achieved from a limited amount of beliefs Σ (exploited through full deduction), and an unlimited amount of other beliefs B. Especially, standard AGM belief revision operators $*$ do not make the job here for three reasons: (1) $B * \Sigma$ is typically underspecified when B is inconsistent, (2) inference from a revised base $B * \Sigma$ is typically intractable [11, 16], and (3) Σ does not represent more entrenched beliefs than B (especially, Σ is not intended to represent knowledge), but pieces of belief we are interested in the whole set of classical consequences.

To fill this gap, we introduce a family of unit resolution-based paraconsistent inference relations $\vdash^*_\Sigma$ for propositional logic in the clausal case. Parameter Σ is any consistent set of clauses representing the beliefs intended to be exploited through full deduction, i.e., every classical consequence of Σ must be kept w.r.t. $\vdash^*_\Sigma$, whatever the belief base B. Contrariwise to many paraconsistent inference relations, any relation $\vdash^*_\Sigma$ can be decided in time linear in $|B|$ whenever Σ is of bounded size. We show that $\vdash^*_\Sigma$ exhibits several valuable properties. In particular, it satisfies strong paraconsistency: $\vdash^*_\Sigma$ never trivializes. We also show how a unit propagation algorithm can be simply turned into a decision procedure for $\vdash^*_\Sigma$. We finally show how the $\vdash^*_\Sigma$ family relates to several inference relations proposed so far as approximations of classical entailment.

The rest of this paper is organized as follows. After some formal preliminaries (Section 2), our family of relations $\vdash^*_\Sigma$ is introduced in Section 3. We also present their main properties and show how to mechanize them, through an adaptation of a unit propagator. Related work are discussed in Section 4, just before the concluding section (Section 5). Proofs are omitted due to space limitations.

2 FORMAL PRELIMINARIES

$PROP_{PS}$ denotes the propositional language built up from a finite set PS of symbols, the connectives $\neg$, $\vee$, $\wedge$, and a propositional constant $\Box$ (denoting the empty clause and also viewed as falsity) in the usual way. $\Box$ is the irreducible contradiction of $PROP_{PS}$. A literal l is a propositional symbol x from PS (positive literal) or a negated one $\neg x$ (negative literal). If $l = x$ is a positive literal, then its complementary literal $\bar{l}$ is $\neg x$; if $l = \neg x$ is a negative literal, then its complementary literal $\bar{l}$ is x. A clause γ is a finite disjunction of literals, also viewed as the set of its literals when it is convenient. It is Horn when it contains at most one positive literal. A unit clause contains at most one literal, while the empty clause $\Box$ contains no literal at all. A CNF formula Σ is a (finite) conjunction of clauses, also viewed as a set of clauses when it is convenient. In the following, Σ, B and Q denote sets of clauses. $Var(\Sigma)$ denotes the set of propositional variables occurring in Σ. The size of a formula Σ from $PROP_{PS}$, noted $|\Sigma|$, is the number of occurrences of symbols and connectives used to write it.

Formulas are interpreted in the classical truth-functional way. Classical interpretations I over PS are defined in the standard way,

[1] CRIL–CNRS/Université d'Artois, Lens, France, email: {coste,marquis}@cril.univ-artois.fr

as total functions from PS to $\{0,1\}$. As usual, $\models$ denotes classical entailment and $\equiv$ denotes classical equivalence.

3 THE $\vdash^*_\Sigma$ FAMILY

3.1 Definitions

There are mainly two ways of specifying a paraconsistent inference relation: (1) by enlarging the set of classical interpretations (as e.g. in multi-valued logics), and (2) by restricting the set of classical proofs. In the following, we adhere to the second approach and define inference relations in a proof-theoretic way.

Let us first give a brief refresher about resolution. Given two clauses γ_1 and γ_2 from $PROP_{PS}$ s.t. γ_1 contains literal l and γ_2 the complementary literal $\bar{l}$, the resolvent of γ_1 and γ_2 over the pair of complementary literals l, $\bar{l}$ is the clause γ containing every literal of γ_1 except l and every literal of γ_2 except $\bar{l}$.[2] γ is obtained using the resolution rule over the pair of complementary literals l, $\bar{l}$ from its parent clauses γ_1 and γ_2. It is obtained using the unit resolution rule (`ur.` for short) if at least one of its parents is a unit clause. Note that it is not necessary to specify the pair of literals resolved upon since they are necessarily unique in that case.

A (unit) refutation tree from a CNF formula Σ is a finite binary tree whose root is labeled by $\Box$, every leaf node is labeled by a clause from γ and every internal node is labeled by a (unit) resolvent of the clauses labeling its two sons. While (unrestricted) resolution is refutationally complete, meaning that there exists a refutation tree from every inconsistent set of clauses Σ (and no such tree can be generated from consistent set of clauses), it is not the case of unit resolution. For instance, the set of clauses $\Sigma = \{(a \vee b), (\neg a \vee b), (a \vee \neg b), (\neg a \vee \neg b)\}$ is inconsistent but inconsistency cannot be revealed using unit resolution (simply because Σ contains no unit clause). Nevertheless, unit resolution is refutationally complete for significant fragments of clausal propositional logic, like the Horn one and proper supersets of it (like the renamable Horn fragment – the set of all CNF formulas which can be turned into Horn ones through literal renaming, see e.g., [13]).

Let us now present a simple inference relation based on unit resolution:

Definition 3.1 ($\vdash$) *Let B and Q be two sets of clauses from $PROP_{PS}$. We have $B \vdash Q$ iff:*

- *Q contains a single clause $l_1 \vee \cdots \vee l_k$ and there exists a unit refutation tree from $B \cup \{\bar{l}_1, \ldots, \bar{l}_k\}$,*
- *$B \vdash \{\gamma\}$ for every clause γ of Q[3] otherwise.*

For instance, we have $B \vdash Q$ with $B = \{a, (\neg a \vee b)\}$ and $Q = \{b\}$ since there is a unit refutation tree from $B \cup \{\neg b\}$:

$$\cfrac{\cfrac{a, \neg a \vee b}{b}\ \texttt{ur.}, \neg b}{\Box}\ \texttt{ur.}$$

Unit propagation is a linear time algorithm that searches for unit refutations from a set of clauses, and it is complete (a set is found unit-refutable by unit propagation iff a unit refutation from it exists). Hence, $\vdash$ can be decided in time linear in $|B|$. Since unit resolution is a sound inference rule, it is also obvious that $\vdash$ is typically a proper subset of $\models$. Especially, $\vdash$ does not trivialize for every inconsistent B. For instance, it is not the case that every clause is a logical consequence of the inconsistent set $B = \{(a \vee b), (\neg a \vee b), (a \vee \neg b), (\neg a \vee \neg b)\}$: we have $B \vdash \{a, \neg a, b, \neg b\}$ but we do not have $B \vdash \Box$ and we do not have $B \vdash c$. Accordingly, $\vdash$ is weakly paraconsistent: for some inconsistent B, the set of all consequences of B w.r.t. $\vdash$ is not the whole set of CNF formulas.

However, $\vdash$ does not satisfy what we expect here, namely strong paraconsistency: there are some B from which $\vdash$ trivializes. For instance, every clause, like $\neg b$, c, etc. is a consequence w.r.t. $\vdash$ of $B = \{a, \neg a, b\}$, but this is not expected; intuitively, none of $\neg b$, c, etc. is supported by the available evidence $a \wedge \neg a \wedge b$, which just states that a is over-determined (i.e., there is some evidence about it to be true, and some evidence about it to be false) and that there is some evidence for b to be true.

Let us now explain how to turn the basic relation $\vdash$ into our relations $\vdash^*_\Sigma$. Let us do it in a gentle way, through the introduction of two "intermediate" inference relations $\vdash^*$ and $\vdash_\Sigma$.

The first relation $\vdash^*$ is obtained by restricting the set of admissible unit refutations. Only those proofs for which the query is relevant (in the sense that it actually participates at least to one refutation) are kept:

Definition 3.2 ($\vdash^*$) *Let B and Q be two sets of clauses from $PROP_{PS}$. We have $B \vdash^* Q$ iff:*

- *Q contains a single clause $l_1 \vee \ldots \vee l_k$ and there exists a unit refutation tree from $B \cup \{\bar{l}_1, \ldots, \bar{l}_k\}$, s.t. at least one leaf node of it is labeled with a literal from $\{\bar{l}_1, \ldots, \bar{l}_k\}$.*
- *$B \vdash^* \gamma$ for every clause γ of Q otherwise.*

For instance, we have $B = \{a, \neg a, b\} \vdash^* \{a, \neg a, b\}$, and $B \not\vdash^* \neg b$, while we have $B \vdash \neg b$. Indeed, the unique unit refutation tree from $B \cup \{b\}$ is (up to the order of the clauses):

$$\cfrac{a, \neg a}{\Box}\ \texttt{ur.}$$

and b does not participate to the proof.

Now, the family of relations $\vdash_\Sigma$ is obtained by enlarging the set of unit refutations so as to guarantee that each clausal consequence of Σ has a proof w.r.t. $\vdash_\Sigma$ whenever Σ is consistent:

Definition 3.3 ($\vdash_\Sigma$) *Let $\Sigma = \{\gamma_1, \ldots, \gamma_n\}$, B and Q be three set of clauses from $PROP_{PS}$. We have $B \vdash_\Sigma Q$ iff (Σ is empty and $B \vdash Q$), or:*

- *Q contains a single clause γ and for every l_1 of γ_1, ..., for every l_n of γ_n, we have $B \cup \{l_1, \ldots, l_n\} \vdash \gamma$.*
- *$B \vdash_\Sigma \gamma$ for every clause γ of Q otherwise.*

For instance, with $\Sigma = \{(a \vee b \vee c), (\neg a \vee b), (a \vee \neg b), (\neg a \vee \neg b)\}$ and $B = \{(\neg c \vee d)\}$, we have $B \vdash_\Sigma \{\neg a, \neg b, c, d\}$, while we have neither $\Sigma \cup B \vdash c$ nor $\Sigma \cup B \vdash d$. We also have $B \not\vdash_\Sigma a$, $B \not\vdash_\Sigma b$. With $\Sigma = \{(a \vee b \vee c)\}$ and $B = \{(\neg a \vee b), (a \vee \neg b), (\neg a \vee \neg b), (\neg c \vee d)\}$, we have $B \vdash_\Sigma \{\neg a, \neg b, c, d\}$. We also have $B \not\vdash_\Sigma a$ and $B \not\vdash_\Sigma b$. For those two examples, $\Sigma \cup B$ is consistent and $\vdash_\Sigma$ has the same clausal consequences w.r.t. B as $\models$ w.r.t. $\Sigma \cup B$.

Since $\vdash$ can be recovered as a specific case of $\vdash_\Sigma$, it cannot be guaranteed that $\vdash_\Sigma$ is strongly consistent. For instance, with $\Sigma = \emptyset$ and $B = \{a, \neg a, b\}$, every clause γ (including the empty one) is a consequence of B w.r.t. $\vdash_\Sigma$. Note also that $\vdash_\Sigma$ trivializes whenever Σ is an inconsistent set of clauses (just because every $\{l_1, \ldots, l_n\}$ contains a pair of complementary literals, or Σ contains the empty clause).

Finally, our family of inference relations $\vdash^*_\Sigma$ is intended to achieve the right balance:

[2] The set representation of clauses induces an implicit application of the factoring rule: any literal cannot occur more than once in a clause.

[3] To avoid too heavy notations, we will typically replace Q by its element when Q is a singleton.

Definition 3.4 ($\vdash^*_\Sigma$) *Let $\Sigma = \{\gamma_1, \ldots, \gamma_n\}$, B and Q be three set of clauses from $PROP_{PS}$. We have $B \vdash^*_\Sigma Q$ iff (Σ is empty and $B \vdash^* Q$), or:*

- *Q contains a single clause γ and for every l_1 of γ_1, ..., for every l_n of γ_n, we have $B \cup \{l_1, \ldots, l_n\} \vdash \gamma$ and there exist l_1 of γ_1, ..., l_n of γ_n s.t. $B \cup \{l_1, \ldots, l_n\} \vdash^* \gamma$.*
- *$B \vdash^*_\Sigma \gamma$ for every clause γ of Q otherwise.*

Stepping back to the three above examples, we can easily state that B has the same consequences w.r.t. $\vdash^*_\Sigma$ as w.r.t. $\vdash_\Sigma$ regarding the first two examples. However, as it is expected, it does not trivialize for the third example: with $\Sigma = \emptyset$ and $B = \{a, \neg a, b\}$, we have $B \vdash^*_\Sigma \{a, \neg a, b\}$ but $B \not\vdash^*_\Sigma \Box$ and $B \not\vdash^*_\Sigma \neg b$.

3.2 Main properties of $\vdash^*_\Sigma$

Let us first focus on the basic properties one wants satisfied: tractability and strong consistency.

Proposition 3.1 *Let Σ, B, Q be three sets of clauses. Deciding $B \overset{?}{\vdash^*_\Sigma} Q$ is coNP-complete, in the general case, even if the query Q reduces to a literal. However, if $|\Sigma|$ is bounded by a constant, then the inference problem w.r.t. $\vdash^*_\Sigma$ can be decided in time linear in $|B|$.*

To be more precise, when $|\Sigma|$ is bounded by a constant and the query reduces to a single clause γ, determining whether $B \vdash^*_\Sigma \gamma$ can be achieved in time $\mathcal{O}(|B| + |\gamma|)$, hence linear in the input size, while it can be achieved in time $\mathcal{O}(|B| \times ||Q||)$, hence quadratic in the input size in the general case.[4]

Proposition 3.2 *$\vdash^*_\Sigma$ is strongly paraconsistent: for any B, we have $B \not\vdash^*_\Sigma \Box$ (even if Σ is inconsistent).*

In the pathological case Σ contains the irreducible contradiction $\Box$, the set of consequences of any B w.r.t. $\vdash^*_\Sigma$ becomes empty; in all the remaining cases, every tautological clause is a consequence of any B w.r.t. w.r.t. $\vdash^*_\Sigma$.

Our relations also have as consequences all the classical consequences of Σ, provided that Σ is consistent:

Proposition 3.3 *If Σ is consistent and γ is any clause s.t. $\Sigma \models \gamma$, then for any B, we have $B \vdash^*_\Sigma \gamma$. Moreover, under the same consistency assumption, we have $\Sigma \models \gamma$ iff $\emptyset \vdash^*_\Sigma \gamma$.*

The consistency condition for Σ is mandatory (otherwise, Proposition 3.2 would be contradicted); when Σ is inconsistent, the set of consequences of any B contains all the literals belonging to at least one maximal (w.r.t. $\subseteq$) hitting set of Σ; for instance, with $\Sigma = \{a, \neg a, (a \vee b)\}$ and $B = \emptyset$, we have $B \vdash^*_\Sigma \{a, \neg a, b\}$ (the unique maximal hitting set is $\{a, \neg a, b\}$) while we do not have $B \vdash^*_{\Sigma'} b$ with $\Sigma' = \{a, \neg a\}$.

Now, the following proposition details the links between $\vdash^*_\Sigma$ and the other inference relations we have considered:

Proposition 3.4 *Let Σ, B be two sets of clauses s.t. Σ does not contain $\Box$ and let γ be a clause:*

1. *$\vdash \ = \ \vdash_\emptyset$.*
2. *$\vdash^* \ = \ \vdash^*_\emptyset$.*
3. *$\vdash^*_\Sigma \ \subseteq \ \vdash_\Sigma$. The inclusion is proper in the general case.*
4. *If $\Sigma \cup B \vdash^* \gamma$, then $B \vdash^*_\Sigma \gamma$. The converse does not hold.*
5. *If $\Sigma \cup B$ is consistent, then ($B \vdash^*_\Sigma \gamma$ iff $B \vdash_\Sigma \gamma$). The converse does not hold.*
6. *If $B \vdash_\Sigma \gamma$, then $B \cup \Sigma \models \gamma$. The converse does not hold.*

The last item shows in particular that all our relations $\vdash_\Sigma$ (hence, $\vdash^*_\Sigma$) are approximations by below of $\models$. Note that it is not guaranteed that $\vdash^*_\Sigma$ (as well as $\vdash_\Sigma$) coincides with $\models$ whenever $B \cup \Sigma$ is consistent, while some paraconsistent inference relations, like $\models_{LP_m}$ [17, 18] and $\models_{QCL}$ [2, 12] ensure it in the clausal case. Nevertheless, none of these relations is tractable, and in fact no tractable inference relation can satisfy this property, unless P = NP. Furthermore, $\vdash^*_\Sigma$ loses no classical consequences whenever $B \cup \Sigma$ is consistent and B is renamable Horn:

Proposition 3.5 *Let Σ, B be two sets of clauses and γ be a clause. If $B \cup \Sigma$ is consistent and B is renamable Horn, then $B \vdash^*_\Sigma \gamma$ iff $B \cup \Sigma \models \gamma$.*

Let us now present other interesting properties that $\vdash^*_\Sigma$ satisfies:

Proposition 3.6 *Let Σ, B, Q, Q' be sets of clauses s.t. none of Σ, B contains $\Box$ and let γ, γ' be two clauses:*

1. *$\vdash^*_\Sigma$ is reflexive and obviously satisfies the "and" rule (in the sense that if $B \vdash^*_\Sigma Q$ and $B \vdash^*_\Sigma Q'$, then $B \vdash^*_\Sigma Q \cup Q'$).*
2. *$\vdash^*_\Sigma$ is monotonic by clause expansion w.r.t. Σ or B: adding a clause (resp. a non empty clause) to B (resp. Σ) never questions the set of consequences already derived.*
3. *$\vdash^*_\Sigma$ satisfies a weak form of right weakening: if $B \vdash^*_\Sigma \gamma$ and $\gamma \models \gamma'$, then $B \vdash^*_\Sigma \gamma'$.*

All those properties are expected; the first one expresses that the explicit statement of a clause in Σ or B is a sufficient evidence to consider it as a consequence; the second property states that whenever a clause γ is considered as a consequence because the available evidence (Σ and B) supports it, this is still the case when some new pieces of evidence are incorporated; finally, the third property means that whenever the available evidence supports a clause γ, it supports as well all the clausal consequences of γ.

Stronger forms of monotonicity and right weakening cannot be satisfied since $\vdash^*_\Sigma$ is a (strong) paraconsistent inference satisfying reflexivity. Thus, while $b \wedge \neg b \models a$ and $\{a\} \vdash^*_\Sigma a$, we do not have $\{b, \neg b\} \vdash^*_\Sigma a$. Furthermore, though $\{a, \neg a\} \vdash^*_\Sigma \{a, \neg a\}$ holds and $a \wedge \neg a \models b$ holds, we do not have $\{a, \neg a\} \vdash^*_\Sigma b$.

Similarly, left logical equivalence cannot be satisfied by a paraconsistent inference relation (while every inconsistent base B is logically equivalent to the irreducible contradiction $\Box$, they typically do not have the same set of expected consequences). However, when Σ is consistent, it can be replaced by any equivalent set of clauses without questioning the set of consequences of any B w.r.t. $\vdash^*_\Sigma$ (especially, every subsumed clause can be removed from Σ).

Finally, neither transitivity[5] nor the cut rule are satisfied by $\vdash^*_\Sigma$ (in the general case); for instance, let $\Sigma = \emptyset$, $B = \{(a \vee b \vee c), (a \vee b \vee \neg c), (a \vee \neg b \vee c), (a \vee \neg b \vee \neg c)\}$, $B' = \{(a \vee b), (a \vee \neg b)\}$ and $B'' = \{a\}$; we have $B \vdash^*_\Sigma B'$, $B' \vdash^*_\Sigma B''$, but $B \not\vdash^*_\Sigma B''$. Since $\vdash^*_\Sigma$ is monotonic by clause expansion w.r.t. B, this example also serves as a counter-example for the cut rule: we have $B \cup B' \vdash^*_\Sigma B''$, $B \vdash^*_\Sigma B'$, but $B \not\vdash^*_\Sigma B''$. While the lack of transitivity would be dramatic for a fully rational reasoner (since it captures the ability of chaining inferences), it is not so much demanded when the purpose is to design a reasoner with limited computational resources.

[4] $||Q||$ denotes the number of clauses of Q.

[5] Note that since $\vdash^*_\Sigma$ is an approximation by below of $\models$, the fact that transitivity does not hold implies that right weakening does not hold.

3.3 Mechanizing $\vdash^*_\Sigma$

It is quite easy to modify a unit propagation algorithm as used in DLL procedures [8] for the satisfiability problem in order to get decision algorithms for our inference relations $\vdash^*_\Sigma$. Note that many DLL procedures have been proposed so far in the literature so that several implementations of unit propagation are available (including open source ones, see `http://sat.inesc-id.pt/OpenSAT/`).

Function UNIT-PROP

Input: a set of clauses B
Output: B once simplified using unit propagation
while B has a unit clause l **do**
  $B \leftarrow \{\gamma \setminus \{\bar{l}\} \mid \gamma \in B\}$
return (B)

First of all, let us recall that unit propagation[6] (cf. the function `UNIT-PROP`) searches for unit refutations that are also directional [9] (this assumption can be done without loss of generality); furthermore, unit propagation can be achieved in linear time [7]: by maintaining a list of clauses containing each literal and a stack containing unit clauses, propagating a unit clause l requires looking at only those clauses which contain $\bar{l}$; associating a counter with each clause that gives the number of literals left in the clause, it is sufficient to decrement it each time a literal is resolved out of the clause; a unit clause is generated whenever the counter value reaches 1, and the empty clause when it reaches 0. `UNIT-PROP` can be used directly to get a decision procedure for $\vdash$: if the query γ is $(l_1 \vee \ldots \vee l_k)$ with $k > 0$, then add $\{\bar{l}_1, \ldots, \bar{l}_k\}$ to B, run `UNIT-PROP` and we have $B \vdash \gamma$ iff the resulting set B contains $\Box$.

Function DECIDE*?

Input: a set of clauses B,
  a clause $\{l_1, \ldots, l_k\}$
Output: $true$ iff $B \vdash^* l_1 \vee \cdots \vee l_k$
mark $\bar{l}_1, \ldots, \bar{l}_k$
$S \leftarrow B \cup \{\bar{l}_1, \ldots, \bar{l}_k\}$
UNIT-PROP*(S)
for each $\gamma \in S$ **do**
  if $\gamma = \emptyset$ and marked(γ) **then**
    return$(true)$
return$(false)$

It is also quite easy to take advantage of `UNIT-PROP` to get an algorithm `DECIDE*` for deciding $\vdash^*$: the idea is to use an additional bit array so as to mark each clause of $B \cup \{\bar{l}_1, \ldots, \bar{l}_k\}$ that is resolved with a marked clause, while initially only the unit clauses from $\{\bar{l}_1, \ldots, \bar{l}_k\}$ are marked. The unit clauses from $\{\bar{l}_1, \ldots, \bar{l}_k\}$ are processed first, and whenever a unit clause is generated, it is pushed on the stack. When the execution of `UNIT-PROP` ends, we have $B \vdash^* \gamma$ iff the counter value of at least one marked clause is 0.

Now, in order to implement $\vdash^*_\Sigma$, it is sufficient to call iteratively `UNIT-PROP*` (which achieves the main treatment of `DECIDE*?`) on $B \cup \{l'_1, \ldots, l'_n\}$ and $\{\bar{l}_1, \ldots, \bar{l}_k\}$ for each conjunction $\{l'_1, \ldots, l'_n\}$ of literals s.t. for every $i \in 1 \ldots n$, l'_i belongs to the i^{th} clause γ_i of $\Sigma = \{\gamma_1, \ldots, \gamma_n\}$. If the resulting set of clauses does not contain $\Box$, then $B \not\vdash^*_\Sigma \gamma$ (one can exit from the n nested loops); otherwise, if the counter value of at least one marked clause is 0, then a flag is raised and one returns to the loops. Once each conjunction $\{l'_1, \ldots, l'_n\}$ has been considered, we have $B \vdash^*_\Sigma \gamma$ iff the flag has been raised.

[6] Also referred to as "boolean constraint propagation" [15].

Function UNIT-PROP*

Input: a set of clauses B
Output: B once simplified using unit propagation and some marks
while B has a unit clause l **do**
  for each $\gamma \in B$ **do**
    if marked(l) and $\bar{l} \in \gamma$ **then**
      mark(γ)
    $\gamma \leftarrow \gamma \setminus \{\bar{l}\}$
return (B)

Function DECIDE*-SIGMA?

Input: a set of clauses B
  a set of clauses $\Sigma = \{\gamma_1, \ldots, \gamma_n\}$
  a clause $\{l_1, \ldots, l_k\}$
Output: $true$ iff $B \vdash^*_\Sigma l_1 \vee \cdots \vee l_n$
if $\Sigma = \emptyset$ **then**
  DECIDE*?$(B, \{l_1, \ldots, l_k\})$
else
  flag $\leftarrow false$
  for each $\langle l'_1, \ldots, l'_n \rangle \in \langle \gamma_1, \ldots, \gamma_n \rangle$ **do**
    $S \leftarrow B \cup \{l'_1, \ldots, l'_n\} \cup \{\bar{l}_1, \ldots, \bar{l}_k\}$
    unmark(S)
    UNIT-PROP*(S)
    if $\Box \notin S$ **then**
      return$(false)$
    else if marked$(\Box)$ **then**
      flag $\leftarrow true$
  return(flag)

Note that the monotonicity property of $\vdash^*$ w.r.t. clause expansion can be exploited to let aside some conjunctions in the body of the loops: if Σ is consistent and if C, C' are two conjunctions to be considered s.t. $C' \models C$, then it is sufficient to keep C only. Accordingly, it is sufficient to keep only the prime implicants of Σ (when Σ is consistent). When Σ does not often change and is of limited size, such a set can be computed and stored during an off-line preprocessing (compilation) phase. Note however that the set of conjunctions to be considered when Σ is consistent does not reduce to any minimal disjunctive normal form of Σ but the whole set of prime implicants must be considered in general: if $\Sigma = \{(a \vee \neg b), (b \vee c)\}$ and $B = \{(b \vee d \vee e), (b \vee \neg d \vee e), (\neg b \vee d \vee e), (\neg b \vee \neg d \vee e)\}$, we do not have $B \vdash^*_\Sigma e$, just because $B \cup \{a, c\} \not\vdash e$. Though Σ is equivalent to $(a \wedge b) \vee (\neg b \wedge c)$, not considering the term $(a \wedge c)$ would lead to accept e as a consequence.

4 OTHER RELATED WORK

Relation $\vdash$ has been introduced and studied by Dalal in [5, 6] as a tractable approximation by below of classical entailment. In order to preserve more classical consequences, Dalal also introduces a whole family of unit resolution-based inference relations $\vdash^k$, parameterized by an integer k. For any k, any set of clauses B and any clauses γ, γ', $\vdash^k$ is defined by:

- If $B \vdash \gamma$, then $B \vdash^k \gamma$, and
- If $B \vdash^k \gamma'$ and $B \cup \{\gamma'\} \vdash^k \gamma$ and γ' contains at most k literals, then $B \vdash^k \gamma$.

The second rule above allows for restoring chaining for clauses of limited sizes. [5, 6] show that for each fixed k, $\vdash^k$ is tractable and the sequence $(\vdash^k)_k$ is monotonic and stationary from some k_{max}; besides, $\vdash^{k_{max}}$ coincides with $\models$ (restricted to clausal formulas). Nevertheless, it is easy to prove that strong consistency is not guaranteed by any relation $\vdash^k$ (just consider $B = \{a, \neg a, b\}$).

Schaerf and Cadoli [19] introduce a family of inference relations $\models_3^S$ which approximate by below classical entailment. Parameter S is a subset of PS, the variables upon which full resolution is allowed. $\models_3^S$ generalizes the entailment relation $\models_3$ of three-valued logic (it corresponds to the case $S = \emptyset$). Every relation $\models_3^S$ is tractable in the clausal case, provided that the size of S is bounded (it can be decided in time linear in the size of B and exponential in the size of S).

Our inference relation $\vdash^*$ typically captures at least all the consequences of $\models_3^\emptyset$, and much more in many cases (for instance, when B is a consistent renamable Horn formula). Furthermore, every relation $\models_3^S$ can be associated in linear time to a $\vdash_\Sigma$ relation including $\models_3^S$:

Proposition 4.1 *Let* B, Q *be sets of clauses and let* $S = \{x_1, \ldots, x_n\}$ *be a subset of* PS*:*

1. *If* $\Box \notin B$ *and* $B \models_3^\emptyset Q$*, then* $B \vdash^* Q$*. The converse does not hold.*
2. *If* $B \models_3^S Q$*, then* $B \vdash_\Sigma Q$*, with* $\Sigma = \{(x_1 \vee \neg x_1), \ldots, (x_n \vee \neg x_n)\}$*. The converse does not hold.*

Clearly, Point 2. above cannot be extended to $\vdash_\Sigma^*$ because the latter is strongly paraconsistent while $\models_3^S$ is not in the general case (consider $B = \{a, \neg a\}$ and $S = \{a\}$: we have $B \models_3^S b$ while $B \not\vdash^*_{\{(a \vee \neg a)\}} b$). Note also that $\models_3^S$ does not offer the right way to capture all the classical clausal consequences of a designated subset Σ of the beliefs; indeed, considering $S = Var(\Sigma)$ achieves the job but may lead to trivialization, while it is avoided using $\vdash_\Sigma^*$. For instance, let $B = \{a, \neg a, b\}$: with $S = \{a\}$, we have $B \models_3^S \gamma$ for every clause γ, while $B \not\vdash_\Sigma^* \Box$, whatever Σ.

Marquis and Porquet [14] refine the family of $\models_3^S$ relations, so as to guarantee strong consistency, while preserving tractability when $|S|$ is bounded. Whenever $\models_3^S$ trivializes, they suggest to weaken S (i.e., to remove variables from S) so as to recover a paraconsistent inference relation. They present several policies for weakening S, closely related to the policies at work in the approach to inconsistency tolerant reasoning based on the selection of preferred subsets of the belief bases. Thus, given a subset S_0 of S, the inclusion preference policy $\mathcal{IP}$ consists in considering every maximal (w.r.t. $\subseteq$) subset S' of S containing S_0 and s.t. $B \not\models_3^{S'} \Box$: we have $B |\!\!\approx_S^{\mathcal{IP},S_0} \gamma$ iff for every such set S', we have $B \models_3^{S'} \gamma$. With $B = \{a, \neg a \vee b, b\}$, $S_0 = \emptyset$, and $S = \{a, b\}$, we have $B |\!\!\approx_S^{\mathcal{IP},S_0} a \wedge \neg b \wedge (\neg a \vee b)$, but we have $B \not|\!\!\approx_S^{\mathcal{IP},S_0} b$ (while b is a consequence of B w.r.t. $\vdash_\emptyset^*$). Because the inference relations given in [14] are subsets of $\models_3^S$, they typically preserve less information from B than $\vdash_\Sigma^*$, but this is not always the case. For instance, with $B = \{(a \vee b \vee c), (\neg a \vee b), (a \vee \neg b), (\neg a \vee \neg b), d, \neg d\}$, $S_0 = \emptyset$ and $S = \{a, b, c, d\}$, we have $B |\!\!\approx_S^{\mathcal{IP},S_0} c$ while we do not have $B \vdash_\emptyset^* c$. Finally, like $\models_3^S$, the inference relations given in [14] do not offer the right way to capture all the classical clausal consequences of a designated subset Σ of the beliefs.

5 CONCLUSION

We have presented a family of unit resolution-based paraconsistent inference relations $\vdash_\Sigma^*$ for propositional logic in the clausal case. $\vdash_\Sigma^*$ exhibits many valuable features, which are typically not jointly shared by alternative approaches: (1) it can be decided in time linear in the size of the belief base provided that $|\Sigma|$ is bounded; (2) it is strongly paraconsistent; (3) it captures all the classical clausal consequences of Σ when Σ is consistent. While (1) prevents from capturing all the classical consequences of B in the case $B \cup \Sigma$ is consistent (under the usual assumptions of complexity theory), $\vdash_\Sigma^*$ may keep much more expected consequences than $\models_3$, even when $\Sigma = \emptyset$; in particular, $\vdash_\emptyset^*$ coincides with $\models$ whenever B is a consistent renamable Horn formula. Last but not least, our inference relations benefit from many algorithmic insights at work in unit propagators.

This work calls for some perspectives. One of them concerns the semantics issue. It would be interesting to investigate how the model-theoretic semantics of $\vdash$ given in [6] could be adapted to our inference relations. Another perspective is to extend further the deductive power of our inference relations by adding more chaining, i.e., considering $\vdash^k$ as the basic inference relation instead of $\vdash$.

ACKNOWLEDGEMENTS

The authors have been partly supported by the IUT de Lens and by the Région Nord/Pas-de-Calais through the IRCICA Consortium and by the European Community FEDER Program.

REFERENCES

[1] C. E. Alchourròn, P. Gärdenfors, and D. Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *J. of Symbolic Logic*, **50**, 510–530, (1985).

[2] P. Besnard and A. Hunter, 'Quasi-classical logic: non-trivializable classical reasoning from inconsistent information', in *Proc. of ECSQARU'95*, volume 946 of *LNAI*, pp. 44–51, Fribourg (Switzerland), (1995). Springer-Verlag.

[3] M. Cadoli and M. Schaerf, 'On the complexity of entailment in propositional multivalued logics', *Ann. of Mathematics and Artificial Intelligence*, **18**, 29–50, (1996).

[4] S. Coste-Marquis and P. Marquis, 'Complexity results for paraconsistent inference relations', in *Proc. of KR'02*, pp. 61–72, Toulouse (France), (2002).

[5] M. Dalal, 'Anytime families of tractable propositional reasoners', in *Proc. of AI&Maths'96*, pp. 42–45, Fort Lauderdale (FL), (1996).

[6] M. Dalal, 'Semantics of an anytime family of reasoners', in *Proc. of ECAI'96*, pp. 360–364, Budapest (Hungary), (1996).

[7] M. Dalal and D. Etherington, 'A hierarchy of ractable satisfiability problems', *Information Processing Letters*, **44**(4), 173–180, (1992).

[8] M. Davis, G. Logemann, and D. Loveland, 'A machine program for theorem proving', *Comm. of the ACM*, **5**, 394–397, (1962).

[9] R. Dechter and I. Rish, 'Directional resolution: the Davis-Putnam procedure, revisited', in *Proc. of KR'94*, pp. 134–145, Bonn, (1994).

[10] I.M.L. D'Ottaviano and N.C.A. da Costa, 'Sur un problème de Jaśkowski', Technical report, Comptes Rendus de l'Académie des Sciences de Paris, (1970).

[11] T. Eiter and G. Gottlob, 'On the complexity of propositional knowledge base revision, updates, and counterfactuals', *Artificial Intelligence*, **57**, 227–270, (1992).

[12] A. Hunter, 'Reasoning with contradictory information using quasi-classical logic', *J. of Logic and Computation*, **10**(5), 677–703, (2000).

[13] H.R. Lewis, 'Renaming a set of clauses as a Horn set', *J. of the ACM*, **25**, 134–135, (1978).

[14] P. Marquis and N. Porquet, 'Resource-bounded inference from inconsistent belief bases', in *Proc. of IJCAI'01*, pp. 103–108, Seattle (WA), (2001).

[15] D. McAllester, 'Truth maintenance', in *Proc. of AAAI'90*, pp. 1109–1116, Boston (MA), (1990).

[16] B. Nebel, 'How hard is it to revise a belief base?', in *Belief Change*, eds., D. Gabbay and P. Smets, volume 3 of *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, 77–145, Kluwer Academic, (1998).

[17] G. Priest, 'Reasoning about truth', *Artificial Intelligence*, **39**, 231–244, (1989).

[18] G. Priest, 'Minimally inconsistent LP', *Studia Logica*, **50**, 321–331, (1991).

[19] M. Schaerf and M. Cadoli, 'Tractable reasoning via approximation', *Artificial Intelligence*, **74**, 249–310, (1995).