INSTITUT FÜR INFORMATIK
Lehr- und Forschungseinheit für
Programmier- und Modellierungssprachen
Oettingenstraße 67 D–80538 München

Ludwig–
Maximilians–
Universität
München

Minimal Model Generation with Positive Unit Hyper-Resolution Tableaux

François Bry **Adnan Yahya**

http://www.pms.informatik.uni-muenchen.de/publikationen
Forschungsbericht/Research Report PMS-FB-1996-1
Mai 1996

Minimal Model Generation with Positive Unit Hyper-Resolution Tableaux

François Bry Adnan Yahya*

Abstract

Herbrand models for clausal theories are useful in several areas of computer science. In most cases, however, the conventional model generation algorithms are inappropriate because they generate nonminimal Herbrand models and can be inefficient. This article describes a novel approach for generating minimal Herbrand models of clausal theories. The approach builds upon *positive unit hyper-resolution (PUHR) tableaux,* that are in general smaller than conventional tableaux. To generate only minimal Herbrand models, a *complement splitting* expansion rule and a specific search strategy are applied. The proposed procedure is optimal in the sense that each minimal model is generated only once, and nonminimal models are rejected before their complete construction. First measurements on an implementation point to the efficiency of the procedure.

1 Introduction:

Generating Herbrand models of clausal theories is useful in several areas of computer science. In automated theorem proving, models can assist in making conjectures, that can be later checked for provability with conventional provers. In automated theorem proving and program verification, model generation can also be applied to searching for counter-examples to conjectures. Model generation is useful in logic programming and deductive databases for specifying their declarative semantics, in some approaches to query answering [3, 7, 20], and in nonmonotonic reasoning [6].

The conventional tableaux methods [16, 4, 18, 19] are however inappropriate as model generation procedures because they often return redundant models [6, 13, 17]. The a posteriori detection of redundant models is tedious and might be time consuming. Moreover, redundant models are a source of inefficiency because they blow up the search space. This article describes a method for generating minimal Herbrand models of clausal theories. The proposed procedure is optimal in the sense that each minimal model is generated only once, and nonminimal models are rejected as soon as possible,

*Also: Electrical Engineering Department, Birzeit University, Birzeit, Palestine

in general before their complete construction. First measurements on an implementation point to the efficiency of the procedure.

The method is based on *positive unit hyper-resolution tableaux* (short *PUHR tableaux*), a (novel) formalization of an approach first introduced with the theorem prover SATCHMO [10, 11]. PUHR tableaux are ground and positive, more precisely their nodes consist of sets of ground atoms and disjunctions of ground atoms. They are expanded by means of only two rules, the *positive unit hyper-resolution* and the *splitting* (a simple version of β expansion [16, 4]) rules, from *range-restricted* clauses. Range-restriction is a syntactical property required in deductive database languages which is comparable to Skolemization: although requiring an extension of the language, it preserves models in a certain sense. Both the branching factor and the size of the nodes of PUHR tableaux can be significantly smaller than those of conventional tableaux. Positive unit hyper-resolution makes it possible not to blindly instantiate universally quantified variables. Instead, it combines in one step instantiations (or γ expansions [16, 4]) and splittings (or β expansion [16, 4]), thus reducing the depth of PUHR tableaux.

In order to restrict the generation of nonminimal models, a *complement splitting* (or reduction in [14], folding-down in [8]) rule is applied in lieu of the classical splitting rule. Because PUHR tableaux are ground, complement splitting can be nicely and efficiently built into the method. While discarding many nonminimal models, and preventing the generation of duplicate models, complement splitting is not always sufficient to reject all nonminimal models. In order to prune redundant models as soon as possible, a special depth first search strategy is applied. The resulting procedure is sound in the sense that it generates only minimal Herbrand models, and complete in the sense that it returns all minimal Herbrand models of the input theory. A variation, we call MM-SATCHMO, of the SATCHMO program is given, which implements the minimal model generation procedure.

The plan of the rest of this paper is as follows. Section 2 introduces terminology and notations, defines range-restriction and PUHR tableaux, and recalls the SATCHMO implementation of PUHR tableaux. Section 3 is devoted to model generation using PUHR tableaux. Section 4 defines the minimal model generation procedure as a modified PUHR tableaux method and gives the implementation of MM-SATCHMO. The last chapter compares the proposed procedure with other approaches discussed in the literature, draws some conclusions, and points to possible directions for future research.

An abridged version of this report has been published in the Proceedings of the Fifth Workshop on Theorem Proving with Analytic Tableaux and Related Methods [1].

2 Preliminaries

2.1 Terminology and Notation

Throughout the paper usual terminology and notations are used, as in e.g. [16, 4]. When not explicitly otherwise stated, a first-order language $\mathcal{L}$ is implicitly assumed. It is also assumed that two special atoms $\top$ and $\bot$ are available, expressing respectively truth and falsity, i.e. $\top$ is satisfied in every interpretation, no interpretations satisfy $\bot$.

Every clause $C = L_1 \vee ... \vee L_k$ with negative literals $\{\neg A_1, ..., \neg A_n\}$ and positive literals $\{B_1, ..., B_m\}$ can be represented by a *clause in implication form*: $C' = A_1 \wedge .... \wedge A_n \rightarrow B_1 \vee ... \vee B_m$. $A_1 \wedge .... \wedge A_n$ is called the body of C', $B_1 \vee ... \vee B_m$ its head. If C contains no negative literals, $C' = \top \rightarrow B_1 \vee ... \vee B_m$. If C contains no positive literals, $C' = A_1 \wedge .... \wedge A_n \rightarrow \bot$.

A unifier (resp. most general unifier) σ of a conjunction of atoms $(A_1 \wedge .... \wedge A_n)$ and a sequence of atoms $(B_1, ..., B_n)$ (possibly with repeated atoms) is a unifier such that $A_i\sigma = B_i\sigma$, for all $i = 1, ..., n$. If $(A_1 \wedge .... \wedge A_n)$ and $(B_1, ..., B_n)$ have a unifier, they are unifiable. Note that, since repeatitions in the sequence $(B_1, ..., B_n)$ are allowed, a conjunction $(A_1 \wedge .... \wedge A_n)$ might be unifiable with a sequence containig less than n (distinct) atoms.

An atom A is said to *subsume* an atom B (a disjunction of atoms $B_1 \vee ... \vee B_n$, resp.) if there exists a substitution σ such that $A\sigma = B$ ($A\sigma = B_i$ for some $i \in \{1, ..., n\}$, resp.).

An interpretation of $\mathcal{L}$ will be denoted as a pair $(\mathcal{D}, m)$ where the non-empty set $\mathcal{D}$ is the universe (or domain) and m is the mapping interpreting the symbols and expressions of the language.

The *universal closure* of a clause C is $\forall x_1 ... \forall x_n C$, where $x_1, ..., x_n$ are the variables occurring in C. A clause (resp. a set of clauses) is said to be *satisfied* by an interpretation when the universal closure of the clause (resp. the set of the universal closures of the clauses) is satisfied by this interpretation. A clause (resp. a set of clauses) is said to be *satisfiable* if it has at least one interpretation in which it is satisfied. A clause (resp. a set of clauses) is said to be *finitely satisfiable* if it is satisfied by an interpretation with a finite domain.

A term or formula in which no variable occurs is said to be *ground.* If $\mathcal{A}$ is a set of ground atoms, $H(\mathcal{A})$ denotes the Herbrand interpretation which satisfies a ground atom B if and only if $B \in \mathcal{A}$. A Herbrand interpretation $H(\mathcal{A})$ is said to be *finitely representable* if $\mathcal{A}$ is finite. Since confusions can be avoided from the context, a set of formulas having a finitely representable Herbrand model will be said to be *finitely representable.* Note that finite representability (of sets of formulas) and finite satisfiability are two distinct properties.

The subset relationship $\subseteq$ over sets of ground atoms induces an order $\leq$ on Herbrand interpretations: given two sets $\mathcal{A}_1$ and $\mathcal{A}_2$ of ground atoms,

$H(\mathcal{A}_1) \leq H(\mathcal{A}_2)$ if and only if $\mathcal{A}_1 \subseteq \mathcal{A}_2$. If $\mathcal{S}$ is a set of clauses, $\leq$ induces an order on Herbrand models of $\mathcal{S}$. A Herbrand model $H(\mathcal{A})$ of $\mathcal{S}$ is said to be a minimal Herbrand model of $\mathcal{S}$ if it is minimal for $\leq$, i.e. for every Herbrand model $H(\mathcal{A}')$ of $\mathcal{S}$, if $H(\mathcal{A}') \leq H(\mathcal{A})$, then $\mathcal{A}' = \mathcal{A}$.

If $\mathcal{E}$ is a set of formulas, $Units(\mathcal{E})$ denotes the set of unit clauses that are elements of $\mathcal{E}$.

Variables are denoted by x with or without subscripts, constants by a, b, c or d, predicate symbols by P, Q, and R, and function symbols by f.

In this paper a tableau method and a minimal model generation procedure for *clausal* theories are defined, i.e. it is assumed that existential quantifications have been removed through Skolemization.

2.2 Range Restriction

Definition 1 (Range restricted clause) *A clause (resp. a clause in implication form) is said to be range restricted if every variable occurring in a positive (resp. head) literal also appears in a negative (resp. body) literal.*

Clearly, a range restricted clause in implication form is ground if its body is ground, e.g. if it is $\top$. A transformation is first defined, which associates a set $RR(\mathcal{S})$ of range restricted clauses in implication form with every set $\mathcal{S}$ of clauses in implication form.

Definition 2 (Range restriction transformation) *Let $\mathcal{L}'$ be an extension of the language $\mathcal{L}$ with a unary predicate D (not belonging to $\mathcal{L}$).*

For every $\mathcal{L}$-clause $C = A_1 \wedge ... \wedge A_n \rightarrow B_1 \vee ... \vee B_m$, let $RR(C)$ be the following $\mathcal{L}'$-clause:

$$RR(C) := \begin{cases} C \text{ if } C \text{ is range restricted;} \\ D(x_1) \wedge ... \wedge D(x_k) \wedge A_1 \wedge ... \wedge A_n \rightarrow B_1 \vee ... \vee B_m \text{ otherwise,} \\ \text{where } x_1, ..., x_k \text{ are the variables occurring in the } B_i\text{s and in} \\ \text{none of the } A_j\text{s.} \end{cases}$$

Let $\mathcal{S}$ be a set of $\mathcal{L}$-clauses. For a term t distinct from a variable occurring in $\mathcal{S}$, let C_t be the $\mathcal{L}'$-clause:

$$C_t := \begin{cases} D(x_1) \wedge ... \wedge D(x_k) \rightarrow D(t) \text{ if the variables } x_1, ..., x_k \text{ occur in } t; \\ \top \rightarrow D(t) \text{ if no variables occur in } t. \end{cases}$$

Let τ be the set of nonvariable terms occurring in $\mathcal{S}$. Let $\mathcal{S}'$ be the following set of $\mathcal{L}'$-clauses:

$$\mathcal{S}' := \begin{cases} \{C_t \mid t \in \tau\} \text{ if } \tau \text{ contains a constant;} \\ \{C_a\} \cup \{C_t \mid t \in \tau\} \text{ otherwise, for some constant } a. \end{cases}$$

$RR(\mathcal{S}) := \{RR(C) \mid C \in \mathcal{S}\} \cup \mathcal{S}'$ *is the range restriction of $\mathcal{S}$.*

Note that by construction the clauses in $RR(\mathcal{S})$ are range restricted and that $RR(\mathcal{S})$ is finite if $\mathcal{S}$ is finite. Strictly speaking, the range restriction transformation does not preserve models because it extends the language. The following theorem, however, shows that it preserves models in a certain sense, similar to the way Skolemization does.

Theorem 3 *Let $\mathcal{S}$ be a set of clauses in a language $\mathcal{L}$ and $RR(\mathcal{S})$ be the range restriction of $\mathcal{S}$ (in an extension $\mathcal{L}'$ of $\mathcal{L}$ with a unary predicate D).*

1. *If $(\mathcal{D}, m)$ is a model of $\mathcal{S}$ and if m' is the mapping over $\mathcal{L}'$ defined as follows:*

$$m'(s) := \begin{cases} m(s) & \text{if } s \neq D, \\ \mathcal{D} & \text{if } s = D. \end{cases}$$

   then $(\mathcal{D}, m')$ is a model of $RR(\mathcal{S})$.

2. *If $(\mathcal{D}, m')$ is a model of $RR(\mathcal{S})$, then $(\mathcal{D}, m'|_{\mathcal{L}})$ is a model of $\mathcal{S}$, where $m'|_{\mathcal{L}}$ denotes the restriction of m' to $\mathcal{L}$.*

Proof: Point 1 follows immediately from Definition 2. For point 2 the non-emptiness of $\mathcal{S}'$ (cf. Definition 2) is necessary, because the clauses $RR(C)$ are satisfied over any interpretation mapping the added unary predicate D to the empty set. ■

This result means that range restrictedness can be seen as just a special syntactic form rather than a real restriction – from a theoretical point of view. For practical purposes, on the other hand, range restrictedness does make a difference. In the context of PUHR tableaux, the range restriction transformation induces an enumeration of the ground terms, making the γ expansion rule of conventional tableaux [16, 4] superfluous. Thus, if the procedures presented in this paper are applied to a set $RR(\mathcal{S})$ obtained from $\mathcal{S}$ by the transformation above, then the newly introduced atoms with predicate D have basically the same effect as an instantiation rule for the clauses of the original set $\mathcal{S}$.

When applied in a refutation procedure, instantiation is often a source of inefficiency. Note, however, that this is not the case for model generation. In contrast to refutation, model generation requires instantiation anyway, indeed, for Herbrand models are characterized as sets of *ground* atoms.

Definition 4 (Positive unit hyperresolution) *Let $C = A_1 \wedge ... \wedge A_n \rightarrow E_1 \vee ... \vee E_m$ be a clause in implication form, $B_1, ..., B_n$ be n (non necessarily distinct) atoms such that $(A_1 \wedge ... \wedge A_n)$ unifies with $(B_1, ..., B_n)$. If σ is a most general unifier of $(A_1 \wedge ... \wedge A_n)$ and $(B_1, ..., B_n)$, then $(E_1 \vee ... \vee E_m)\sigma$ is a positive unit hyper-resolvent of C and $B_1, ..., B_n$.*

Lemma 5 *The positive unit hyper-resolvent of a range restricted clause in implication form and ground atoms is a ground atom or a disjunction of ground atoms.*

Proof: Immediate. ∎

Note that no occur-checks need to be performed for computing the positive unit hyper-resolvent of a range restricted clause in implication form and ground atoms.

In the next section, positive unit hyper-resolution tableaux are defined for range restricted clauses. This is not a significant restriction, for there is a transformation of any set of clauses into a set of range-restricted clauses which preserves models in the sense of Theorem 3.

2.3 Positive Unit Hyper-Resolution Tableaux

Starting from the set $\{\top\}$, the PUHR tableaux method expands a tree – or positive unit hyper-resolution (PUHR) tableau – for a set $\mathcal{S}$ of range restricted clauses in implication form by applying the following expansion rules that are defined with respect to $\mathcal{S}$. The nodes of a PUHR tableau are sets of ground atoms or disjunctions of ground atoms.

Definition 6 (PUHR tableaux expansion rules) *Let $\mathcal{S}$ be a set of clauses in implication form.*

Positive unit hyper-resolution (PUHR) rule:

$$\frac{\begin{array}{c} B_1 \\ \vdots \\ B_n \end{array}}{E\sigma}$$

where σ is a most general unifier of the body of a clause $(A_1 \wedge \ldots \wedge A_n \rightarrow E) \in \mathcal{S}$ with $(B_1, \ldots, B_n)$.

Splitting rule:

$$\frac{E_1 \vee E_2}{E_1 \quad | \quad E_2}$$

In the following definition, the splitting rule is applied to *ground* disjunctions.

Definition 7 (PUHR tableaux) *Positive unit hyper-resolution (PUHR) tableaux for a set $\mathcal{S}$ of clauses in implication form are trees whose nodes are sets of ground atoms and disjunctions of ground atoms. They are inductively defined as follows:*

1. *$\{\top\}$ is a positive unit hyper-resolution tableau for $\mathcal{S}$.*

2. *If T is a positive unit hyper-resolution tableau for $\mathcal{S}$, if L is a leaf of T such that an application of the PUHR rule (resp. splitting rule) to formulas in L yields a formula E (resp. two formulas E_1 and E_2) not subsumed by an atom in L, then the tree T' obtained from T by adding the node $L \cup \{E\}$ (resp. the two nodes $L \cup \{E_1\}$ and $L \cup \{E_2\}$) as successor(s) to L is a positive unit hyper-resolution tableau for $\mathcal{S}$.*

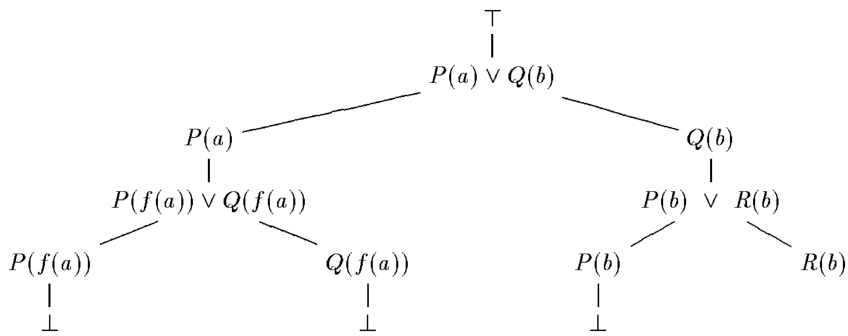

Figure 1: A PUHR tableau for the set of clauses of Example 1.

A branch of a positive unit hyper-resolution tableau is said to be closed, if it includes a node containing the atom $\perp$. *A positive unit hyper-resolution tableau is said to be closed if all its branches are closed. A branch (resp. tableau) which is not closed is said to be open.*

A positive unit hyper-resolution tableau T *for* $\mathcal{S}$ *is said to be satisfiable if the union of* $\mathcal{S}$ *with the nodes of a branch of* T *is satisfiable.*

Note that if $\mathcal{P}$ is a path from the root to a node N of a PUHR tableaux, then by Definition 7, $N = \cup\mathcal{P}$.

Convention. If N_1 and $N2$ are the nodes of a PUHR tableau T containing respectively E_1 and E_2 and resulting from an application of the splitting rule to $E_1 \vee E_2$, it is assumed in the sequel that the PUHR tableau T is ordered such that N_1 is the left sibling of E_2.

Example 1 Figure 1 gives a PUHR tableau for the following set of clauses in implication form:

$$\begin{array}{ll}
\top \rightarrow P(a) \vee Q(b) & P(b) \rightarrow \perp \\
P(x) \rightarrow P(f(x)) \vee Q(f(x)) & P(f(x)) \rightarrow \perp \\
Q(x) \rightarrow P(x) \vee R(x) & P(x) \wedge Q(f(x)) \rightarrow \perp
\end{array}$$

For the sake of readability, the nodes of the tree of Figure 1 are labeled with the ground atoms or disjunctions of ground atoms added at these nodes. We recall that by Definition 7 the nodes of a PUHR tableau are sets of ground atoms and disjunctions of ground atoms.

By Lemma 5 the nodes of a positive unit hyper-resolution tableau for a set of range restricted clauses are sets of ground atoms and disjunctions of ground atoms. Sets of clauses for which PUHR tableaux are defined may be infinite. According to Definition 6 clauses whose heads are $\perp$ only contribute to close branches. Since negative formulas do not explicitly occur in PUHR tableaux, closure is simply detected by the presence of $\bot$, which is simpler than checking for atomic closure [4].

Definition 8 *Let $\mathcal{S}$ be a set of range-restricted clauses in implication form and $\mathcal{A}$ a set of ground atoms and disjunctions of ground atoms. $\mathcal{A}$ is said to be saturated with respect to $\mathcal{S}$ for the positive unit hyper-resolution and splitting expansion rules when the following properties hold:*

1. *if $(A_1 \wedge ... \wedge A_n \rightarrow E) \in \mathcal{S}$, $B_1 \in \mathcal{A}$, ..., and $B_n \in \mathcal{A}$, and $(A_1 \wedge ... \wedge A_n)$ and $(B_1, ..., B_n)$ are unifiable, then $E\sigma \in \mathcal{A}$ for a most general unifier σ of $(A_1 \wedge ... \wedge A_n)$ and $(B_1, ..., B_n)$.*

2. *If $(E_1 \vee E_2) \in \mathcal{A}$, then $E_1 \in \mathcal{A}$ or $E_2 \in \mathcal{A}$.*

Note that if $\mathcal{B}$ is an open or a closed branch of a PUHR tableau, then $\cup\mathcal{B}$ is not necessarily saturated. As well, if $\cup\mathcal{B}$ is saturated, then $\mathcal{B}$ is neither necessarily open, nor necessarily closed.

Lemma 9 *The application of an expansion rule to a satisfiable PUHR tableau results in a satisfiable PUHR tableau.*

Proof: If M is a model of a set $\mathcal{F}$ of clauses, atoms and disjunctions, and if E is a positive unit hyper-resolvent of elements of $\mathcal{F}$, then $M \models E$. If M is a model of $\mathcal{F}$ and $E_1 \vee E_2 \in \mathcal{F}$, then $M \models E_1$ or $M \models E_2$. ∎

Theorem 10 (Refutation soundness) *Let $\mathcal{S}$ be a set of range-restricted clauses in implication form. If there exists a closed PUHR tableau for $\mathcal{S}$, then $\mathcal{S}$ is unsatisfiable.*

Proof: Assume $\mathcal{S}$ is satisfiable. By Lemma 9 there are no closed PUHR tableaux for $\mathcal{S}$. ∎

Definition 11 *A PUHR tableau is said to be fair, if the union of the nodes of each of its open branches is saturated for the expansion rules.*

Informally, a PUHR tableaux is fair if along each of its open branches, each possible application of an expansion rule is performed at least once.

If $\mathcal{B}$ is a branch of a tableau, then $Units(\cup\mathcal{B})$ denotes the set of unit clauses that are elements of some nodes in $\mathcal{B}$. In the sequel, $Units(\mathcal{E})$ will always be refered to in cases where all unit clauses in $\mathcal{E}$ are *ground atoms*. Recall that if $Units(\mathcal{E})$ is a set of ground atoms, it characterizes the Herbrand interpretation $H(Units(\mathcal{E}))$.

Lemma 12 *Let $\mathcal{S}$ be a set of range-restricted clauses in implication form and $\mathcal{E}$ be a set of ground atoms and disjunctions of ground atoms. If $\mathcal{S} \cup \mathcal{E}$ is saturated for the expansion rules with respect to $\mathcal{S}$ and if $\mathcal{E}$ does not contain $\bot$, then $H(Units(\mathcal{E}))$ is a model of $\mathcal{S}$.*

Proof: Immediate. ∎

Theorem 13 (Refutation completeness) *Let $\mathcal{S}$ be a set of range-restricted clauses in implication form. If $\mathcal{S}$ is unsatisfiable, then every fair positive unit hyper-resolution tableau for $\mathcal{S}$ is closed.*

Proof: Let T be an open fair PUHR tableau for $\mathcal{S}$, and $\mathcal{B}$ an open branch of T. Since T is fair, then $\cup\mathcal{B}$ is saturated for the expansion rules. By Lemma 12 $H(Units(\cup\mathcal{B}))$ is a model of $\mathcal{S}$. Hence $\mathcal{S}$ is satisfiable. ∎

PUHR tableaux are defined for sets of range restricted clauses. Combined with the PUHR expansion rule of Definition 6, the range restriction transformation induces an enumeration of the ground terms, as observed in [9].

2.4 Implementation in Prolog

The Prolog program of Figure 2 expands fair PUHR tableaux for sets of range-restricted clauses in implication form under a depth-first search strategy. The tableaux expanded by this program are strict [4] and subsumption-free. Strictness means that no application of an expansion rule is performed more than once to given clauses, atoms, or disjunctions. Subsumption-freeness means that only ground disjunctions that are not subsumed by previously generated atoms or disjunctions can be split.

Backtracking over `satisfiable` returns Herbrand models $H(\mathcal{M})$. The ground atoms of $\mathcal{M}$ are inserted into the Prolog database by the predicate `assume`. On backtracking, they are removed. A clause $A_1 \wedge ... \wedge A_n \rightarrow B_1 \vee ... \vee B_m$ is represented in the Prolog database as

```
A1, ..., An ---> B1 ; ... ; Bm,
```

where `--->` is declared as an infix binary predicate. $\perp$ is represented as `false`, $\top$ as the built-in predicate `true`, which is always satisfied.

Fairness is ensured by the call to the all-solutions built-in predicate `findall`. The predicate `component` on backtracking successively returns the atoms of a disjunction. The predicate `satisfy` on backtracking successively returns the components of a disjunction that are not subsumed by atoms previoulsy inserted into the Prolog database. For each ground instance `_B ---> H` of a clause returned by the the call

```
findall(Clause, violated_instance(Clause), Set)
```

the predicate `satisfy_all` selects an atom in the head `H` and asserts it in the Prolog database. On backtracking, the different ways to satisfy the head `H` of each ground instance `_B ---> H` returned by the call to `findall` are considered.

```
satisfiable :-
     findall(Clause, violated_instance(Clause), Set),
     not (Set = []), !,
     satisfy_all(Set),
     satisfiable.
satisfiable.

violated_instance(B ---> H) :-
     (B ---> H), B, not H.

satisfy_all([]).
satisfy_all([_B ---> H | Tail]) :-
     H, !, satisfy_all(Tail).
satisfy_all([_B ---> H | Tail]) :-
     satisfy(H), satisfy_all(Tail).

satisfy(E) :-
    component(Atom, E), not (Atom = false),
    assume(Atom).

component(Atom, (Atom ; _Rest)).
component(Atom, (_ ; Rest)) :-
     !, component(Atom, Rest).
component(Atom, Atom).

assume(Atom) :-
     asserta(Atom).
assume(Atom) :-
     once(retract(Atom)),
     fail.
```

Figure 2: The SATCHMO program.

The program of Figure 2, called SATCHMO, as well as variations of it have been first published in [10, 11]. In these articles, the programs are explained in more detail and performance on bench mark examples is reported. The PUHR tableaux introduced in Section 2.3 are a formalization of the principle of the SATCHMO programs.

It is worth pointing out that `satisfy_all` is a simple and straightforward implementation which, in some cases, has drawbacks. Consider for example the following Prolog representations $\mathcal{R}_1$ and $\mathcal{R}_2$ of the same set of clauses:

$\mathcal{R}_1$:	$\mathcal{R}_2$:
`true ---> p(a)`	`true ---> p(b) ; p(a)`
`true ---> p(b) ; p(a)`	`true ---> p(a)`

Applied to $\mathcal{R}_1$, the call to

```
findall(Clause, violated_instance(Clause), Set),
```

instantiates `Set` with the list `[(true ---> p(a)),(true ---> p(b);p(a))]`. Then the call to `satisfy_all` first asserts `p(a)` into the Prolog database so as to satisfy the head of `true ---> p(a)`. Since now `p(b) ; p(a)` is satisfied, no further actions are taken, as specified by the second clause of `satisfy_all`. If in contrast $\mathcal{R}_2$ is considered, the call to

```
findall(Clause, violated_instance(Clause), Set)
```

binds `Set` to the list `[(true ---> p(b) ; p(a)) (p(a), true ---> p(a))]`. The call to `satisfy_all` now satisfies first `p(b) ; p(a)`, then `p(a)`. That is `p(b)` is first asserted, then `p(a)`. On backtracking, `p(a)` only is asserted.

Such a behaviour depending on the order of the clauses in Prolog can be avoided with a more sophisticated implementation of `satisfy_all` which satisfies the considered set of heads of ground clauses by a *minimal* set of atoms. Since such a refined implementation of `satisfy_all` is not needed for the purpose of this report, it is not given here.

3 Model Generation with PUHR Tableaux

In the previous section, PUHR tableaux were considered from the angle of refutation. In this section, their properties with respect to model generation are investigated.

Theorem 14 (Model soundness) *Let $\mathcal{S}$ be a satisfiable set of range-restricted clauses in implication form and T a fair PUHR tableau for $\mathcal{S}$. If $\mathcal{B}$ is an open branch of T, then $H(Units(\cup\mathcal{B}))$ is a model of $\mathcal{S}$.*

Proof: Fairness ensures saturation with respect to the expansion rules. Theorem 14 follows from Lemma 12. ∎

Theorem 15 *Let $\mathcal{S}$ be a satisfiable set of range-restricted clauses in implication form, T be a PUHR tableau for $\mathcal{S}$, and $\mathcal{M}$ a set of ground atoms. If $H(\mathcal{M})$ is a model of $\mathcal{S}$, then there is a branch $\mathcal{B}$ of T such that $Units(\cup\mathcal{B}) \subseteq \mathcal{M}$.*

Proof: Let $\mathcal{B}_0, ..., \mathcal{B}_i, ...$ be an enumeration of the branches of T, whose atoms are not in $\mathcal{M}$. For each $i \in \mathbf{N}$ let A_i be an atom of the branch $\mathcal{B}_i$ which is not in $\mathcal{M}$. Let $\mathcal{S}' = \mathcal{S} \cup \{A_i \rightarrow \perp \ : i \in \mathrm{N}\}$. By definition of $\mathcal{S}'$, since no A_i is in $\mathcal{M}$, $H(\mathcal{M})$ is also a model of $\mathcal{S}'$. Furthermore T can be extended into a positive unit hyper-resolution tableau T' of $\mathcal{S}'$ by adding $\perp$ to the successor nodes of those nodes of T that contain some A_i. Let $\mathcal{B}'_i$ denote such an extension of the branch $\mathcal{B}_i$ in T'. By Theorem 10, T' has an open branch, say $\mathcal{B}$. Since $\mathcal{B}$ is open, it is none of the $\mathcal{B}'_i$. Since all clauses of $\mathcal{S}$, whose heads are $\perp$, are also in $\mathcal{S}'$, $\mathcal{B}$ is also an open branch of T. $\mathcal{B}$ is none of the $\mathcal{B}_i$ because otherwise, by definition of T', it would be one of the $\mathcal{B}'_i$. By definition of the $\mathcal{B}_i$s $Units(\cup\mathcal{B}) \subseteq \mathcal{M}$. ∎

Corollary 16 (Minimal model completeness) *Let $\mathcal{S}$ be a satisfiable set of range-restricted clauses in implication form, T be a fair positive unit hyper-resolution tableau for $\mathcal{S}$, and $\mathcal{M}$ a set of ground atoms. If $H(\mathcal{M})$ is a minimal model of $\mathcal{S}$, then there is a branch $\mathcal{B}$ of T such that $Units(\cup\mathcal{B}) = \mathcal{M}$.*

Proof: By Theorem 15, there is a branch $\mathcal{B}$ of T such that $Units(\mathcal{B}) \subseteq \mathcal{M}$. Since T is fair, by Theorem 14 $H(Units(\mathcal{B}))$ is a model of $\mathcal{S}$. Since $H(\mathcal{M})$ is a minimal model of $\mathcal{S}$, $Units(\cup\mathcal{B}) = \mathcal{M}$. ∎

The following example demonstrates that a plain PUHR tableau can generate both, nonminimal and duplicate models.

Example 2 Let $\mathcal{S}$ be the following set of clauses:

$$\begin{array}{ll} \top \rightarrow P(a) \vee P(b) & \quad P(a) \rightarrow P(b) \vee P(d) \\ \top \rightarrow P(a) \vee P(c) & \quad P(b) \rightarrow P(a) \vee P(d) \end{array}$$

Figure 3 is a PUHR tableau for $\mathcal{S}$. The minimal model $H(\{P(a), P(b)\})$ of $\mathcal{S}$ is generated twice, at the leftmost branch and at the third branch from the left of the PUHR tableau. The fourth branch from the left of the PUHR tableau generates the nonminimal model $H(\{P(a), P(b), P(c)\})$. Note that the PUHR tableau returns among others all minimal models of $\mathcal{S}$, i.e. $H(\{P(a), P(b)\})$, $H(\{P(a), P(d)\})$, and $H(\{P(b), P(c), P(d)\})$.

Corollary 16 is established, though in a different context, in [2] and mentioned without proof in [7]. As the following counter-example shows, fairness is necessary in Corollary 16, although not in Theorem 15.

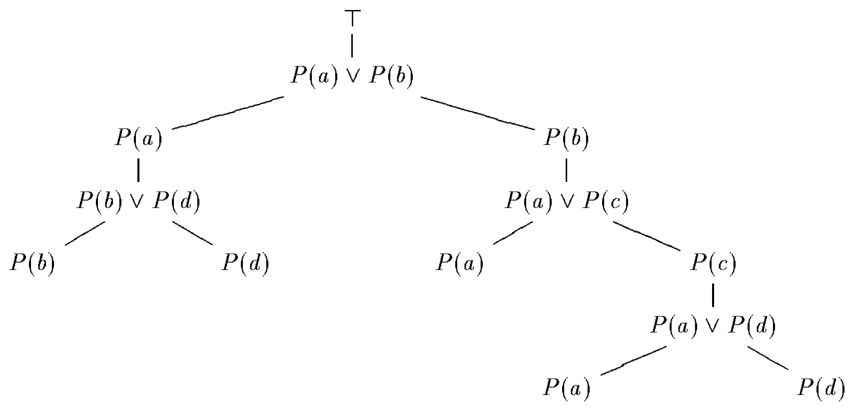

Figure 3: A PUHR tableau for Example 2 with nonminimal and duplicate models.

Example 3 With the theory $\mathcal{S} = \{\top \rightarrow P(a), P(x) \rightarrow P(f(x)) \vee P(b), P(a) \rightarrow P(b)\}$ consistently expanding on the second clause will not allow the generation of the (only) minimal model $H(\{P(a), P(b)\})$ of $\mathcal{S}$.

4 Minimal Model Generation

By Corollary 16 fair PUHR tableaux generate all minimal models. However, they often also generate duplicate and/or nonminimal models, as e.g. in Example 2 above. A naive approach to minimal model generation consists in first expanding fair PUHR tableaux, and later pruning them from redundant branches. In this section a more efficient approach is described which consists in pruning PUHR tableaux from redundant branches as soon as possible. The pruning involves a refinement of the splitting rule, and a specific search strategy based on depth-first search. Under certain finiteness conditions, the proposed procedure is complete.

4.1 Finiteness Properties

Theorem 17 *Let $\mathcal{S}$ be a set of formulas. If $\mathcal{S}$ has a finitely representable Herbrand model it also has a finite model.*

Proof: Let $(\mathcal{D}, m)$ be a finitely representable Herbrand model of $\mathcal{S}$, and $\mathcal{A}$ be the set of ground atoms that are satisfied in $(\mathcal{D}, m)$. A finite model of $\mathcal{S}$ is built by identifying the elements of the universe $\mathcal{D}$ over which no terms occurring in $\mathcal{A}$ are mapped. Formally, let $\sim$ be the equivalence relation over $\mathcal{D}$ defined by: $d_1 \sim d_2$ if and only if $d_1 = d_2$ or for all $R(t_1, ..., t_n) \in \mathcal{A}$ and for all $i = 1, ..., n$, $m(t_i) \neq d_1$ and $m(t_i) \neq d_2$. Let f be the mapping of an element of $\mathcal{D}$ to its equivalence class for $\sim$ in $\mathcal{D}/\sim$. Let $\mathcal{D}' = \mathcal{D}/\sim$ and $m' = f \circ m$. Since $\mathcal{A}$ is finite, $\mathcal{D}/\sim$ is finite. By definition of $\mathcal{D}'$ and m', a ground atom is satisfied in $(\mathcal{D}', m')$ if and only if it is satisfied in $(\mathcal{D}, m)$. Since $(\mathcal{D}, m) \models \mathcal{S}$, it follows that $(\mathcal{D}', m') \models \mathcal{S}$. ∎

The following result relates the finiteness of the set of minimal models to the finite representability of the minimal models. Let us call *finitary* a set of clauses, whose minimal Herbrand models are all finitely representable.

Theorem 18 *Let $\mathcal{S}$ be a set of clauses. If $\mathcal{S}$ is finitary, then $\mathcal{S}$ has finitely many minimal Herbrand models.*

Proof: Let $\mathcal{S}$ be a set of clauses with an infinite number of finitely representable minimal Herbrand models. Let $H(\mathcal{A}_0), ..., H(\mathcal{A}_n), ...$ be an enumeration of all finitely representable minimal Herbrand models of $\mathcal{S}$, such that the $\mathcal{A}_i$s are pairwise distinct. If $\mathcal{A}$ is a finite set of atoms $\{A_1, ..., A_k\}$, let $Neg(\mathcal{A})$ denote the (singleton) set of clauses $\{A_1 \wedge ... \wedge A_k \rightarrow \bot\}$. For every $n \in \mathbf{N}$, let $\mathcal{S}^n = \mathcal{S} \cup Neg(\mathcal{A}_0) \cup ... \cup Neg(\mathcal{A}_n)$. Since all $\mathcal{A}_i$s are finite, the $\mathcal{S}^n$s are also finite. Each $\mathcal{S}^n$ is satisfiable because, by definition, $H(\mathcal{A}_{n+1})$ is a model of $\mathcal{S}^n$. Let $\mathcal{S}^\omega = \cup\{\mathcal{S}^n : n \in \mathbf{N}\}$. Since every $\mathcal{S}^n$ is satisfiable, every finite subset of $\mathcal{S}^\omega$ is satisfiable By the compactness theorem, $\mathcal{S}^\omega$ is therefore satisfiable. Since $\mathcal{S}^\omega$ is a set of clauses, it has a Herbrand model, and therefore also some minimal Herbrand model $H(\mathcal{M})$. By definition of $\mathcal{S}^\omega$, $H(\mathcal{M})$ is none of the finitely representable models $H(\mathcal{A}_n)$. Therefore $\mathcal{M}$ is infinite. ∎

Although finite representability (of a set of formulas) is a stronger property than finite satisfiability, we conjecture that it is semi-decidable like finite satisfiability. We also conjecture that the finitary property is semi-decidable.

Let $\mathcal{S}$ be a set of clauses whose minimal Herbrand models are all finitely representable. By Theorem 18 a PUHR tableau for $\mathcal{S}$ pruned from those branches corresponding to nonminimal models is finite. Note, however, that a finitary theory may have inifinite nonminimal Herbrand models, as is shown by Example 4 below.

In many applications, the finite representability of the minimal Herbrand models is often assumed. This is the case in particular of disjunctive databases and of some forms of nonmonotonic reasoning [6]. Thus, from the viewpoint of applications, Theorem 18 seems to be an important result.

4.2 Complement Splitting

If $C = A_1 \vee ... \vee A_n$ is a disjunction of atoms, let $Neg(C)$ denote the finite set of clauses in implication form $Neg(C) := \{A_1 \rightarrow \bot, ..., A_n \rightarrow \bot\}$.

Definition 19 (Complement splitting rule)

$$\frac{E_1 \vee E_2}{\begin{array}{ll|l} E_1 & & E_2 \\ Neg(E_2) & & \end{array}}$$

Like the splitting rule, the complement splitting rule (already mentioned in [11], called reduction in [14] and folding-down in [8]) is applied in the following definitions to *ground* disjunctions. Tableaux expanded with the positive unit hyper-resolution and the complement splitting rules are defined inductively, similarly as in Definition 7. Let us call such tableaux *PUHR complement tableaux.* Note that nodes of PUHR complement tableaux are sets of ground atoms, disjunctions of ground atoms, and ground implications of the form $A \rightarrow \bot$.

Definition 20 (PUHR complement tableaux)
Positive unit hyper-resolution (PUHR) complement tableaux for a set $\mathcal{S}$ of clauses in implication form are trees whose nodes are sets of ground atoms, disjunctions of ground atoms, and ground implications of the form $A \rightarrow \bot$. They are inductively defined as follows:

1. *$\{\top\}$ is a positive unit hyper-resolution complement tableau for $\mathcal{S}$.*
2. *If T is a positive unit hyper-resolution complement tableau for $\mathcal{S}$, if L is a leaf of T such that an application of the PUHR rule (resp. complement splitting rule) to formulas in L yields a formula E (resp. two sets of formulas $\{E_1, Neg(E_2)\}$ and $\{E_2\}$), then the tree T' obtained from T by adding the node $L \cup \{E\}$ (resp. the two nodes $L \cup \{E_1, Neg(E_2)\}$ and $L \cup \{E_2\}$) as successor(s) to L is a positive unit hyper-resolution complement tableau for $\mathcal{S}$.*

For PUHR complement tableaux, closedness and openness of branches and tableaux are defined like in Definition 7: A branch of a PUHR complement tableau is said to be *closed,* if it includes a node containing the atom $\bot$. A PUHR complement tableau is said to be closed if all its branches are closed. A branch (resp. PUHR complement tableau) which is not closed is said to be *open.*

Definition 21 *Let $\mathcal{S}$ be a set of range-restricted clauses in implication form and $\mathcal{A}$ a set of ground atoms, disjunctions, and clauses in implication form. $\mathcal{A}$ is said to be saturated with respect to $\mathcal{S}$ for the positive unit hyper-resolution and the complement splitting expansion rules when the following properties hold:*

- *if $(A_1 \wedge ... \wedge A_n \rightarrow E) \in \mathcal{S}$, $B_1 \in \mathcal{A}, ..., B_n \in \mathcal{A}$, and $(A_1 \wedge ... \wedge A_n)$ and $(B_1, ..., B_n)$ are unifiable, then $E\sigma \in \mathcal{A}$ for some most general unifier σ of $(A_1 \wedge ... \wedge A_n)$ and $(B_1, ..., B_n)$.*

- *If* $(E_1 \vee E_2) \in \mathcal{A}$, *then* $\{E_1\} \cup Neg(E_2) \subseteq \mathcal{A}$, *or* $E_2 \in \mathcal{A}$.

Note that if $\mathcal{A}$ is saturated with respect to $\mathcal{S}$ for the positive unit hyper-resolution and the complement splitting expansion rules, then it is also saturated for the positive unit hyper-resolution and the splitting expansion rules.

Model soundness for PUHR complement tableaux follows from Theorem 14.

Lemma 22 *Let* $\mathcal{S}$ *be a set of clauses and* $A_1, ..., A_n (n \geq 1)$ *be atoms.*

1. *If* M *is a minimal Herbrand model of* $\mathcal{S}$ *such that* $M \not\models A_1 \wedge ... \wedge A_n$, *then* M *is a minimal Herbrand model of* $\mathcal{S} \cup \{A_1 \wedge ... \wedge A_n \rightarrow \bot\}$.
2. *If* M *is a minimal Herbrand model of* $\mathcal{S} \cup \{A_1 \wedge ... \wedge A_n \rightarrow \bot\}$, *then* M *is also a minimal Herbrand model of* $\mathcal{S}$.

Proof: 1. Let $H(\mathcal{M})$ be a nonminimal model of $\mathcal{S} \cup \{A_1 \wedge ... \wedge A_n \rightarrow \bot\}$. There exists $\mathcal{M}_1 \subset \mathcal{M}$ such that $H(\mathcal{M}_1)$ is a model of $\mathcal{S} \cup \{A_1 \wedge ... \wedge A_n \rightarrow \bot\}$. Hence, $H(\mathcal{M})$ is not a minimal model of $\mathcal{S}$.

2. Assume that $H(\mathcal{M})$ is a Herbrand model of $\mathcal{S} \cup \{A_1 \wedge ... \wedge A_n \rightarrow \bot\}$ which is not a minimal Herbrand model of $\mathcal{S}$. There is $\mathcal{M}_1 \subset \mathcal{M}$ such that $H(\mathcal{M})$ is a model of $\mathcal{S}$. Since $H(\mathcal{M}) \not\models A_i$ for some $i = 1, ..., n$ and since $\mathcal{M}_1 \subset \mathcal{M}$, $H(\mathcal{M}_1) \not\models A_i$. $H(\mathcal{M}_1)$ is therefore not a minimal Herbrand model of $\mathcal{S} \cup \{A_1 \wedge ... \wedge A_n \rightarrow \bot\}$, and the same holds of $H(\mathcal{M})$. ■

Lemma 23 *Let* $\mathcal{E}$ *be a set of clauses in implication form, ground atoms and disjunctions of ground atoms,* $E_1 \vee E_2 \in \mathcal{E}$ *be a ground clause, and* $\mathcal{M}$ *be a set of ground atoms.* $H(\mathcal{M})$ *is a minimal model of* $\mathcal{E}$ *if and only if*

1. *either it is a minimal model of* $\mathcal{E} \cup \{E_1\} \cup Neg(E_2)$
2. *or it is a minimal model of* $\mathcal{E} \cup \{E_2\}$ *and for all* $\mathcal{M}_1 \subseteq \mathcal{M}$, $H(\mathcal{M}_1)$ *is not a minimal model of* $\mathcal{E} \cup Neg(E_2)$.

Proof: Let $H(\mathcal{M})$ be a minimal model of $\mathcal{E}$. If $H(\mathcal{M})$ does not satisfy E_2, then $H(\mathcal{M})$ is a model of $\mathcal{E} \cup \{E_2 \rightarrow \bot\}$. By Lemma 22, $H(\mathcal{M})$ is a minimal model of $\mathcal{E} \cup Neg(E_2)$. If $H(\mathcal{M})$ satisfies E_2 it is a model of $\mathcal{E} \cup \{E_2\}$. If it is not a minimal model of $\mathcal{E} \cup \{E_2\}$, then there exists $\mathcal{M}_1 \subset \mathcal{M}$ such that $H(\mathcal{M}_1)$ is a model of $\mathcal{E} \cup \{E_2\}$, hence of $\mathcal{E}$, contradicting the hypothesis that $H(\mathcal{M})$ is a minimal model of $\mathcal{E}$. By Lemma 22, if $H(\mathcal{M})$ is a minimal model of $\mathcal{E} \cup Neg(E_2)$, then it is also a minimal model of $\mathcal{E}$. Let $H(\mathcal{M})$ be a minimal model of $\mathcal{E} \cup \{E_2\}$. If $H(\mathcal{M})$ is not a minimal model of $\mathcal{E}$, then there exists $\mathcal{M}_1 \subset \mathcal{M}$ such that $H(\mathcal{M}_1)$ is a minimal model of $\mathcal{E}$. Since $H(\mathcal{M})$ is a minimal model of $\mathcal{E} \cup \{E_2\}$, $H(\mathcal{M}_1)$ does not satisfy E_2. Since $E_1 \vee E_2$ in $\mathcal{E}$, $H(\mathcal{M}_1)$ satisfies E_1 . Therefore, $H(\mathcal{M}_1)$ satisfies $\mathcal{E} \cup \{E_2 \rightarrow \bot\}$, i.e. there exists $\mathcal{M}_2 \subseteq \mathcal{M}_1 \subseteq \mathcal{M}$, such that $H(\mathcal{M}_2)$ is a minimal model of $\mathcal{E} \cup Neg(E_2)$. ■

Theorem 24 (Minimal model completeness for complement tableaux) *Let $\mathcal{S}$ be a satisfiable set of range-restricted clauses in implication form, T be a fair PUHR complement tableau for $\mathcal{S}$, and $\mathcal{M}$ a set of ground atoms. If $H(\mathcal{M})$ is a minimal model of $\mathcal{S}$, then there is a branch $\mathcal{B}$ of T such that $Units(\cup\mathcal{B}) = \mathcal{M}$.*

Proof: Follows from Corollary 16 since by definition every PUHR complement tableau for a set $\mathcal{S}$ can be constructed from a PUHR (noncomplement) tableaux by adding $\perp$ to some of its nodes, and from Lemma 23 which basically states that minimal models are preserved by complement splitting. ∎

The following example shows that complement splitting is not always sufficient to prune all nonminimal models.

Example 4 Let $\mathcal{S} = \{\top \rightarrow P(a), P(x) \rightarrow P(b) \vee P(f(x)), P(a) \rightarrow P(b)\}$. Let T be the PUHR complement tableau for $\mathcal{S}$ by applying first the PUHR rule on $\top \rightarrow P(a)$ and $P(a) \rightarrow P(b)$, and then alternatively the PUHR and splitting rule on $P(x) \rightarrow P(b) \vee P(f(x))$. Although $H(\{P(a), P(b)\})$ is the unique minimal model of $\mathcal{S}$, T also has branches corresponding to the models $H(\{P(a), P(b), P(f(a)), ..., P(f^n(a))\})$ for all $n \in \mathbf{N}$.

Although possibly having branches corresponding to nonminimal models, PUHR complement tableaux never have two distinct branches defining the same model, as established next.

Lemma 25 *Let $\mathcal{S}$ be a satisfiable set of range-restricted clauses in implication form, T be a fair PUHR complement tableau for $\mathcal{S}$, and $\mathcal{B}_L$ and $\mathcal{B}_R$ be two open branches of T. If $\mathcal{B}_L$ appears to the left of $\mathcal{B}_R$ in T, then $Units(\cup\mathcal{B}_R) \not\subseteq Units(\cup\mathcal{B}_L)$.*

Proof: Let A_R be an atom in the first node of $\mathcal{B}_R$ (in a root to leaf traversal) which is not not in $\mathcal{B}_L$. By definition of the complement splitting rule, $(A_R \rightarrow \perp) \in \cup\mathcal{B}_L$. Hence $A_R \notin \cup\mathcal{B}_L$. ∎

Corollary 26 *Let $\mathcal{S}$ be a satisfiable set of range-restricted clauses in implication form, T be a fair PUHR complement tableau for $\mathcal{S}$ and $\mathcal{B}_0, ..., \mathcal{B}_i, ...$ a left-to-right enumeration of the open branches of T.*

1. *$H(Units(\cup\mathcal{B}_0))$ is a minimal model of $\mathcal{S}$.*
2. *If $i \neq j$, then $Units(\cup\mathcal{B}_i) \neq Units(\cup\mathcal{B}_j)$*

Proof: 1. Since $\mathcal{B}_0$ is the leftmost branch of T, by Lemma 25 $H(Units(\mathcal{B}_0))$ is a minimal model of $\mathcal{S}$.

2. Follows directly from Lemma 25. ∎

The following example demonstrates that a PUHR complement tableau can generate nonminimal models.

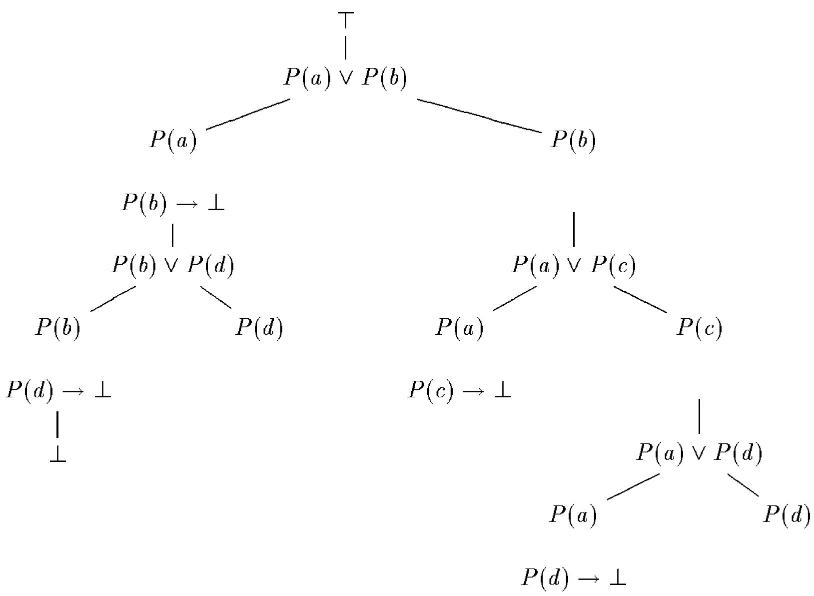

Figure 4: A PUHR complement tableau.

Example 5 Let $\mathcal{S}$ be the set of clauses of Example 2, i.e.:

$$\begin{array}{ll} \top \rightarrow P(a) \vee P(b) & \qquad P(a) \rightarrow P(b) \vee P(d) \\ \top \rightarrow P(a) \vee P(c) & \qquad P(b) \rightarrow P(a) \vee P(d) \end{array}$$

Figure 4 gives a PUHR complement tableau for $\mathcal{S}$. The models generated by this PUHR complement tableau are $H(\{P(a), P(d)\})$, $H(\{P(b), P(c), P(a)\})$, $H(\{P(b), P(a)\})$, and $H(\{P(b), P(c), P(d)\})$. Note that although some are not minimal, the PUHR complement tableau returns no duplicates.

4.3 Implementation of Complement Splitting

Complement splitting can be built into SATCHMO by replacing the procedure `satisfy` by the following procedure `cs_satisfy`, as shown by Figure 5. `cs_component` returns not only the atoms of a disjunction, like `component` does, but also the rest of the disjunction on the right hand side of the returned atom (`false` if this right hand side is empty). This implementation, which we call CS-SATCHMO, departs slightly from Definition 19 since it represents $Neg(A_1 \vee ... \vee A_n)$ as $A_1 \vee ... \vee A_n \rightarrow \perp$ instead of $\{A_1 \rightarrow \perp, ..., A_n \rightarrow \perp\}$. Since the A_i are ground, the two representations are equivalent.

The complete program of CS-SATCHMO is given in Appendix A.

```
cs_satisfy_all([]).
cs_satisfy_all([_B ---> H | Tail]) :-
     H,
     !,
     cs_satisfy_all(Tail).
cs_satisfy_all([_B ---> H | Tail]) :-
     cs_satisfy(H),
     cs_satisfy_all(Tail).

cs_satisfy(E) :-
     cs_component(Atom, Suffix, E),
     not (Atom = false),
     assume(Atom),
     assume_neg(Suffix).

cs_component(A, S, (A ; S)).
cs_component(A, S, (_ ; Rest)) :-
     !,
     cs_component(A, S, Rest).
cs_component(A, false, A).

assume_neg(false) :-
   !.
assume_neg(E) :-
   assume(E ---> false).
```

Figure 5: Complement splitting for SATCHMO

4.4 Constrained Search

By Corollary 26 the first model returned from a depth-first-left-first traversal of a PUHR complement tableau is minimal, and by Lemma 25 no models are $\leq$-larger than subsequently returned models. In order to prune PUHR complement tableaux from nonminimal models, it therefore suffices to constrain any model under construction not to be $\leq$-larger than any previously returned model. This is easily achieved by adding to the set of clauses a constraint $Neg(\{A_1, ..., A_n\}) = \{A_1 \wedge ... \wedge A_n \rightarrow \perp\}$ once a (finite) model $H(\{A_1, ..., A_n\})$ has been constructed.

Definition 27 (Minimal model generation procedure) *Let $\mathcal{S}$ be a set of range restricted clauses in implication form. Applying the minimal model generation procedure to $\mathcal{S}$ consists in a depth-first-left-first construction of a fair PUHR complement tableau for $\mathcal{S}$ such that $\mathcal{S}$ is augmented with $Neg(\mathcal{M})$ after each computation of a model $H(\mathcal{M})$ of $\mathcal{S}$.*

Note that, by Definitions 7 and 19, if $\mathcal{S}_1$ and $\mathcal{S}_2$ are sets of range-restricted clauses in implication form such that $\mathcal{S}_1 \subseteq \mathcal{S}_2$ and all clauses in $\mathcal{S}_2 \setminus \mathcal{S}_1$ are of the form $A_1 \wedge ... \wedge A_n \rightarrow \perp$, then every PUHR complement tableau for $\mathcal{S}_2$ can be obtained from a PUHR complement tableau for $\mathcal{S}_1$ by adding $\perp$ to some nodes. Conversely, every PUHR complement tableau for $\mathcal{S}_1$ can be obtained from a PUHR complement tableau for $\mathcal{S}_2$ by discarding $\perp$ from some nodes.

Recall that a set of clauses is finitary if its minimal Herbrand models are all finitely representable.

Lemma 28 *Let $\mathcal{S}$ be a finitary and finite set of range-restricted clauses in implication form, and T be a PUHR complement tableau for $\mathcal{S}$.*

If t is a node in T, let $\mathcal{B}_0, ..., \mathcal{B}_{n_t}$ be branches of T to the left of t such that $H(Units(\cup\mathcal{B}_0)), ..., H(Units(\cup\mathcal{B}_{n_t}))$ are minimal models of $\mathcal{S}$.

Let T_t be the PUHR complement tableau for $\mathcal{S} \cup Neg(\cup\mathcal{B}_0) \cup ... \cup Neg(\cup\mathcal{B}_{n_t})$ corresponding to T. If $\mathcal{B}$ is a branch of T, let $\mathcal{B}_t$ denote the corresponding branch in T_t and conversely.

$\mathcal{B}_t$ is open in T_t if and only if $\mathcal{B}$ is open in T and $Units(\cup\mathcal{B}_i) \not\subseteq Units(\cup\mathcal{B}_t)$, for all $i = 0, ..., n_t$.

Proof: Assume that $\mathcal{B}$ is an open branch of T and $Units(\cup\mathcal{B}_i) \not\subseteq Units(\cup\mathcal{B})$, for all $i = 0, ..., n_t$. For all $i = 0, ..., n_t$ there exists an atom $A_i \in \cup\mathcal{B}$ such that $A_i \in \cup\mathcal{B} \setminus \cup\mathcal{B}_i$. Therefore, $H(Units(\cup\mathcal{B})) \models Neg(\cup\mathcal{B}_i)$. Hence $\mathcal{B}_t$ is open in T_t.

Assume that $\mathcal{B}_t$ is an open branch of T_t. If $Units(\cup\mathcal{B}_i) \not\subseteq Units(\cup\mathcal{B})$, for all $i = 0, ..., n_t$, then $\perp \notin \cup\mathcal{B}$. Hence $\mathcal{B}$ is open in T. ∎

Theorem 29 (Soundness and completeness of the minimal model generation procedure) *Let $\mathcal{S}$ be a finite set of range-restricted clauses in* *implication form. If $\mathcal{S}$ is finitary, then applied on $\mathcal{S}$, the minimal model generation procedure terminates, returns all minimal models of $\mathcal{S}$ (i.e. it is complete), does not return any nonminimal model of $\mathcal{S}$ (i.e. it is sound), and does not return any minimal model more than once.*

Proof: Let $\mathcal{S}$ be a finitary and finite set of range restricted clauses in implication form.

Soundness: By Corollary 26 the first model returned by the procedure is a minimal model of $\mathcal{S}$. Assume that the first n models $H(\mathcal{M}_0), ..., H(\mathcal{M}_{n-1})$ returned by the procedure are minimal models of $\mathcal{S}$. Let T be the tableau expanded so far. After returning the first n models, the procedure backtracks to a node t of T, such that the branches corresponding to previously returned models are to the left of t. The $(n+1)$-th model returned by the procedure corresponds to the first open branch of a tableau T_t for $\mathcal{S} \cup Neg(\mathcal{M}_0) \cup ... \cup Neg(\mathcal{M}_{n-1})$. By Lemma 28, this model is not $\leq$-larger than any previously returned model. By Corollary 26 it is a minimal model of $\mathcal{S} \cup Neg(\mathcal{M}_0) \cup ... \cup Neg(\mathcal{M}_{n-1})$. Hence, by Lemma 22 it is a minimal model of $\mathcal{S}$ as well. By induction, all models returned are minimal models of $\mathcal{S}$.

Completeness: For any two minimal models $H(\mathcal{M}_1)$ and $H(\mathcal{M}_2)$ of $\mathcal{S}$, $\mathcal{M}_1 \not\subseteq \mathcal{M}_2$ and $\mathcal{M}_2 \not\subseteq \mathcal{M}_1$. Therefore, $H(\mathcal{M}_1) \models Neg(\mathcal{M}_2)$ and $H(\mathcal{M}_2) \models Neg(\mathcal{M}_1)$. Consequently, no branches corresponding to a minimal model $H(\mathcal{M})$ of $\mathcal{S}$ with $\mathcal{M} \notin \{\mathcal{M}_0, ..., \mathcal{M}_n\}$ of a PUHR complement tableau for $\mathcal{S}$ can be closed in a tableau for $\mathcal{S} \cup Neg(\mathcal{M}_0) \cup ... \cup Neg(\mathcal{M}_n)$, for some minimal models $H(\mathcal{M}_0), ..., H(\mathcal{M}_n)$ of $\mathcal{S}$. Since the procedure returns only minimal models, the result follows. From Lemma 28, it follows that no models are generated more than once.

Termination: Since $\mathcal{S}$ is finitary, it has by Theorem 18 finitely many minimal models. Since the procedure returns all and only minimal models of $\mathcal{S}$, and since no minimal models are generated more than once, the procedure terminates. ∎

The following example shows how the minimal model generation procedure generates only minimal models and does not return duplicates.

Example 6 Figure 6 gives the search spaces of the minimal model generation procedure for the set of clauses of Examples 2 and 5, i.e.:

$$\begin{array}{ll} \top \rightarrow P(a) \vee P(b) \qquad & P(a) \rightarrow P(b) \vee P(d) \\ \top \rightarrow P(a) \vee P(c) \qquad & P(b) \rightarrow P(a) \vee P(d) \end{array}$$

Note that all models returned by the procedure are minimal.

It is worth noting that fairness is necessary for the minimal model generation procedure, as the following counter-example shows.

Example 7 Let $\mathcal{S} = \{\top \rightarrow P(a), P(x) \rightarrow P(f(x)) \vee P(b), P(a) \rightarrow P(b)\}$. An unfair PUHR complement tableau for $\mathcal{S}$ with leftmost branch $\{P(a),$

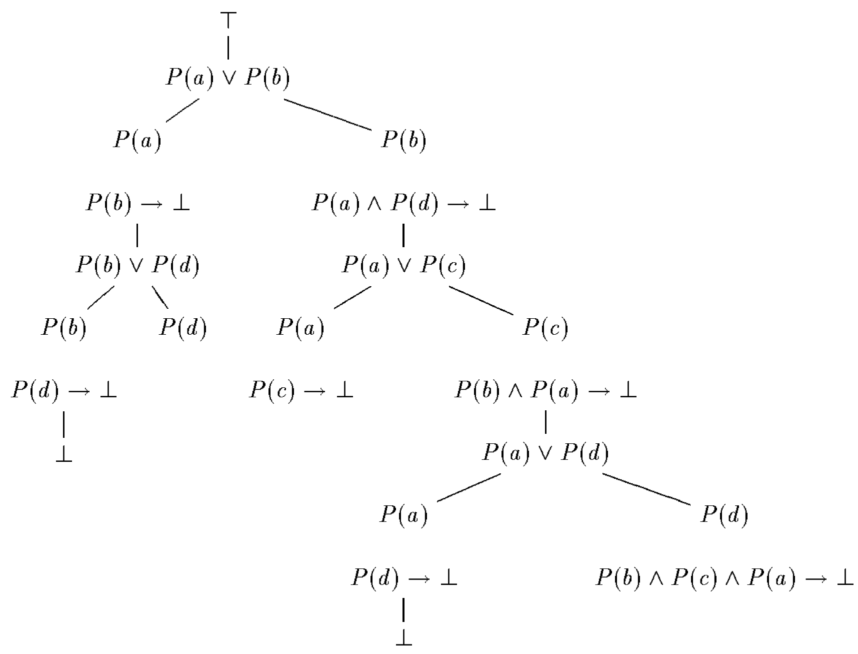

Figure 6: A run of the minimal model generation procedure.

$P(f(a)), \ldots, P(f^n(a)), \ldots\}$ not containing $P(b)$ does not return the minimal model $H(\{P(a), P(b)\})$ and does not give rise to applying the constraint $P(a) \wedge P(b) \rightarrow \perp$ for pruning redundant branches.

4.5 MM-SATCHMO

Figure 7 gives a program, we call MM-SATCHMO, which implements the minimal model generation procedure. It builds upon the implementation of complement splitting described in Section 4.2. A slight modification of `satisfiable` suffices to construct the constraints induced by a (minimal) model.

The argument of the procedure `mm` is the body of the constraint under construction. This data structure is redundant, for the model under construction is also represented in the Prolog database. This redundancy can be easily removed, at the cost of a less readable program. A more serious source of inefficiency lies in the way how violated clauses are detected: the last inserted atoms are not used for an incremental detection. Although quite simple, an incremental evaluation requires longer and more complicated programs. An incremental clause evaluation turns out to be especially beneficial for the constrained search.

The complete program of MM-SATCHMO is given in Appendix B.

```
minimal_model :-
     mm(true).

mm(C1) :-
    findall(Clause, violated_instance(Clause), Set),
    not (Set = []),
    !,
    mm_satisfy_all(Set, C1, C2),
    mm(C2).
mm(C) :-
    asserta(C ---> false).

mm_satisfy_all([], C, C).
mm_satisfy_all([_B ---> H | Tail], C1, C3) :-
     H,
     !,
     mm_satisfy_all(Tail, C1, C3).
mm_satisfy_all([_B ---> H | Tail], C1, C3) :-
     mm_satisfy(H, A),
     and_merge(A, C1, C2),
     mm_satisfy_all(Tail, C2, C3).

mm_satisfy(E, Atom) :-
     cs_component(Atom, Suffix, E),
     not (Atom = false),
     assume(Atom),
     assume_neg(Suffix).

and_merge(Atom, true, Atom) :-
     !.
and_merge(Atom, Conj, (Atom, Conj)).
```

Figure 7: The MM-SATCHMO program.

5 Conclusions and Future Work

This paper presented a positive unit hyper-resolution (PUHR) tableau method for computing the minimal Herbrand models of sets of range restricted clauses. The method was proved to be complete and sound in the sense that it generates all and only minimal models of its input set. A compact implementation of our procedure in the form of a short Prolog program called MM-SATCHMO was also presented.

As a tableau procedure our approach enjoys a good degree of efficiency stemming from its restricted search space, from limiting the applications of expansion rules and the use of matching without occur-check rather than full unification. This is possible because, as a side-effect of a special range-restricted syntactical form, the generated tableaux are ground. Since it makes instantiation necessary, groundness of tableaux might be considered as a source of inefficiency in a refutation procedure. However, since Herbrand models are characterized as sets of ground atoms, this objection does not apply to a model generation procedure.

As a model generation procedure, ours compares well with those reported in the literature, many of which are not sound in the sense that they generate nonminimal models [11, 7]. Compared with approaches based on model generation then testing for minimality [3, 12] our approach avoids nonminimal model generation altogether. The generation of nonminimal models is aborted as soon as possible, in general before they are fully developed. Also, the method we propose is applicable to first-order clauses and not confined to propositional or ground theories as the algorithms reported in [3, 20, 12]. While the applicability of the approach proposed in this article to first-order theories is a major advantage, most of the techniques increasing the efficiency for propositional or ground theories proposed in [20, 12] can be incorporated into a version of our algorithm tailored for that case. Moreover, the approach we propose requires no order to be placed on the sequence in which individual atoms are expanded – although such an order can be incorporated without substantial changes to the algorithm [20]. In [6] the concept of a *ghost tableau* is used to check the minimality of models that may be made nonminimal by the existential instantiation rule (or δ expansion [16]) in the (*primary*) tableau when testing for a "mini-consequence" property. The concept is useful when existential quantifiers are allowed in the theory which is not the case we consider in the present article.

Among the limitations of the procedure are its applicability only to range restricted and so called finitary theories. However, range restriction is not much of a constraint, because a model preserving transformation of general clauses into range restricted ones was given. We believe that much of real-life tasks enjoy the finiteness properties needed for the applicability of our procedure. One of the shortcomings of the procedure as reported here is its lack of incrementality. Further improvements, not discussed in this paper, can also be incorporated into the procedure. Another point is that, in some cases, the large number of constraints corresponding to generated minimal models may overwhelm the process without much positive contribution to discarding nonminimal models. For such cases we developed a second minimal model generation procedure based on a localized test that decides the minimality of the model based on the content of that model alone with no reference to other models. Space considerations prevent us from detailing the approach here.

Initial testing of a prototype of this procedure points to its efficiency both as a model generator, and as a refutation system [5]. Indeed, the restriction to minimal models often dramatically reduces the search space, thus speeding up the closing of a tableaux. The prototype was able to deal with theories with a large number of minimal models with performances comparable to the best reported in the literature [12]. Further testing is needed to better evaluate the gains in performance and compare the minimal model generation procedure with existing systems. We plan also to investigate applying a similar approach for query answering, integrity constraint enforcement, and knowledge assimilation in data and knowledge base applications.

Acknowledgments

We thank Norbert Eisinger, Heribert Schütz and Tim Geisler for the many fruitful discussions on the topic of this paper. This research was done while the second author was visiting at Ludwig-Maximilians-Universität München on an Alexander von Humboldt Research Fellowship. The support of Alexander-von-Humboldt-Stiftung is appreciated. We also thank the three anonymous referees for useful comments on [1].

References

[1] F. Bry and A. Yahya. Minimal model generation with positive unit hyper-resolution tableaux. In *Proc.* 5^{th} *Workshop on Theorem Proving with Analytic Tableaux and Related Methods,* Springer-Verlag, LNCS, 1996.

[2] M. Denecker and D. Schreye. A framework for indeterministic model generation with equality. In *Proc. Conf. on Fifth Generation Computer Systems.*, 1992.

[3] J.A. Fernández and J. Minker. Bottom-up evaluation of Hierarchical Disjunctive Deductive Databases. In *Proc.* 8^{th} *Int. Conf. on Logic Programming*, 660–675. MIT Press, 1991.

[4] M. Fitting. *First-Order Logic and Automated Theorem Proving.* Springer-Verlag, 1987.

[5] H. Schüz and T. Geisler. Efficient model generation through compilation. Tech. Rep. PMS-96-2, Inst. für Informatik, Munich University, 1996. Submitted for publication.

[6] J. Hintikka. Model minimization – an alternative to circumscription. *J. Automated Reasoning*, Vol. 4, 1–13, 1988.

[7] K. Inoue, M. Koshimura, and R. Hasegawa. Embedding negation as failure into a model generation theorem prover. In *Proc.* 11^{th} *Int. Conf. on Automated Deduction*, 1992.

[8] R. Letz, K. Mayr, and C. Goller. Controlled integration of the cut rule into connection tableau calculi. *J. Automated Reasoning*, Vol. 13 No. 3, 297–338, 1994.

[9] D. Loveland, D. Reed, and D. Wilson. SATCHMORE: SATCHMO with RElevancy. *J. Automated Reasoning*, Vol. 14, 325–351, 1995.

[10] R. Manthey and F. Bry. A hyperresolution-based proof procedure and its implementation in prolog. In *Proc.* 11^{th} *German Workshop on Artificial Intelligence*, 456–459, 1987.

[11] R. Manthey and F. Bry. Satchmo: a theorem prover implemented in Prolog. In *Proc.* 9^{th} *Conf. on Automated Deduction*, 456–459, 1988.

[12] I. Niemelä. A tableau calculus for minimal model reasoning. In *Proc.* 5^{th} *Workshop on Theorem Proving with Analytic Tableaux and Related Methods,* 1996. Springer-Verlag.

[13] N. Olivetti. Tableaux and sequent calculus for minimal entailment. *J. Automated Reasoning*, Vol. 9, 99–139, 1992.

[14] D. Prawitz. A new improved proof procedure *Theoria*, Vol. 26, 102–139, 1960.

[15] J.A. Robinson. Automatic deduction with hyper-resoultion. *Int. J. Computational Mathematics*, Vol. 1, 227–234, 1965.

[16] R. Smullyan. *First-Order Logic.* Springer-Verlag, 1968.

[17] M. Suchenek. First-order syntactic characterizations of minimal entailment, domain minimal entailment and herbrand entailment. *J. Automated Reasoning*, Vol. 10, 237–236, 1993.

[18] G. Wrightson (Editor). Special issue on automated reasoning with analytic tableaux, Part I. *J. Automated Reasoning*, Vol. 13 No. 2, 173–281, 1994.

[19] G. Wrightson (Editor). Special issue on automated reasoning with analytic tableaux, Part II. *J. Automated Reasoning*, Vol. 13 No. 3, 283–421, 1994.

[20] A. Yahya, J.A. Fernandez, and J. Minker. Ordered model trees: A normal form for disjunctive deductive databases. *J. Automated Reasoning*, Vol. 13 No. 1, 117–144, 1994.

Appendix A: CS-SATCHMO

```
cs_satisfiable :-
     findall(Clause, violated_instance(Clause), Set),
     not (Set = []),
     !,
     cs_satisfy_all(Set),
     cs_satisfiable.
cs_satisfiable.

violated_instance(Body ---> Head) :-
     (Body ---> Head),
     Body,
     not Head.

cs_satisfy_all([]).
cs_satisfy_all([_B ---> H | Tail]) :-
     H,
     !,
     cs_satisfy_all(Tail).
cs_satisfy_all([_B ---> H | Tail]) :-
     cs_satisfy(H),
     cs_satisfy_all(Tail).

cs_satisfy(E) :-
     cs_component(Atom, Suffix, E),
     not (Atom = false),
     assume(Atom),
     assume_neg(Suffix).

cs_component(Atom, Suffix, (Atom ; Suffix)).
cs_component(Atom, Suffix, (_Atom ; Rest)) :-
     !,
     cs_component(Atom, Suffix, Rest).
cs_component(Atom, false, Atom).

assume(Atom) :-
     asserta(Atom).
assume(Atom) :-
     once(retract(Atom)),
     fail.

assume_neg(false) :-
     !.
assume_neg(E) :-
     assume(E ---> false).
```

Appendix B: MM-SATCHMO

```
minimal_model :-
     mm(true).


mm(C1) :-
     findall(Clause, violated_instance(Clause), Set),
     not (Set = []),
     !,
     mm_satisfy_all(Set, C1, C2),
     mm(C2).
mm(C) :-
     asserta(C ---> false).

violated_instance(Body ---> Head) :-
     (Body ---> Head),
     Body,
     not Head.

mm_satisfy_all([], C, C).
mm_satisfy_all([_B ---> H | Tail], C1, C3) :-
     H,
     !,
     mm_satisfy_all(Tail, C1, C3).
mm_satisfy_all([_B ---> H | Tail], C1, C3) :-
     mm_satisfy(H, A),
     and_merge(A, C1, C2),
     mm_satisfy_all(Tail, C2, C3).

mm_satisfy(E, Atom) :-
    cs_component(Atom, Suffix, E),
    not (Atom = false),
    assume(Atom),
    assume_neg(Suffix).

and_merge(Atom, true, Atom) :-
    !.
and_merge(Atom, Conj, (Atom, Conj)).

cs_component(Atom, Suffix, (Atom ; Suffix)).
cs_component(Atom, Suffix, (_Atom ; Rest)) :-
     !,
     cs_component(Atom, Suffix, Rest).
cs_component(Atom, false, Atom).

assume(Atom) :-
     asserta(Atom).
assume(Atom) :-
     once(retract(Atom)),
     fail.
```

```
assume_neg(false) :-
     !.
assume_neg(E) :-
     assume(E ---> false).
```