

NOTE

COMPLETENESS OF RESOLUTION REVISITED *

Marc BEZEM

Centre for Mathematics and Computer Science, P.O. Box 4079, 1009 AB Amsterdam, The Netherlands

Communicated by G. Jäger

Received September 1988

Revised July 1989

Abstract. By a novel argument we prove the completeness of (ground) resolution. The argument allows us to give the completeness proofs for various strategies of resolution in a uniform way, thus contributing to the insight into these strategies. For example, our exposition shows how the more efficient strategies can be derived from an analysis of the redundancies in the completeness proofs. Moreover, by using Zorn's Lemma in dealing with infinite sets of ground clauses, we obtain completeness proofs which are completely independent of the cardinality of both the language and the set of clauses. We discuss the set theoretic status of these results.

1. Preliminaries

1.1. Logic

Let $\mathcal{L} = \{p_a \mid a \in A\}$, with A some index set, be a set of propositional atoms. We do not make any assumption about the cardinality of \mathcal{L} . A *literal* is an atom or the negation of an atom. Literals p_a and $\neg p_a$ are called *complementary* (p_a is called *positive*, $\neg p_a$ *negative*). If L is a literal, then its complement is denoted by \bar{L} . The set of all literals will be denoted by \mathcal{Lit} . An *interpretation* is a subset I of \mathcal{L} , corresponding to the truth valuation $\mathcal{V}_I(p_a) = \text{TRUE}$ if $p_a \in I$, and **FALSE** otherwise. A *clause* is a finite set of literals, which should be thought of as the disjunction of these literals.

Truth of a literal L (respectively a clause C) in an interpretation I , denoted by $I \models L$ (respectively $I \models C$), is defined as follows:

$$I \models p_a \text{ iff } p_a \in I,$$

$$I \models \neg p_a \text{ iff } p_a \notin I,$$

$$I \models \{L_1, \dots, L_n\} \ (n \geq 0) \text{ iff } I \models L_i \text{ for some } 1 \leq i \leq n.$$

* The material for this article has been taken from the PRISMA document P346. The research is part of the PRISMA project (PaRallel Inference and Storage MACHine), a joint effort with Philips Research Eindhoven, partially supported by the Dutch "Stimulerings-projectteam Informatica-onderzoek" (SPIN).

Note that the empty clause is false in any interpretation. The *truth set* (respectively *falsity set*) of an interpretation I is defined as $\mathcal{T}(I) = \{L \in \mathcal{Lit} \mid I \models L\}$ (respectively $\mathcal{F}(I) = \{L \in \mathcal{Lit} \mid \text{not } I \models L\} = \mathcal{Lit} - \mathcal{T}(I)$). Note that neither $\mathcal{T}(I)$ nor $\mathcal{F}(I)$ contains complementary literals, and that both contain an occurrence, positive or negative, of every $p_a \in \mathcal{L}$. For a set of clauses S we define $S_{\mathcal{T}}(I) = \{C \in S \mid I \models C\}$ and $S_{\mathcal{F}}(I) = S - S_{\mathcal{T}}(I)$. A set of clauses S is called *satisfiable* if there exists an interpretation I in which every clause from S is true. Such an interpretation I is called a *model* of S (or a *model for* S).

1.2. Transfinite induction

Let $<$ be a well-founded partial ordering on a set S (i.e. $<$ is an irreflexive, transitive relation on S such that every descending sequence is finite). The principle of *transfinite induction* with respect to $<$ states that if some property P of elements of S is progressive, then it holds for every element of S . Here *progressivity* of P means that $P(x)$ is implied by $\forall y < x P(y)$, for all $x \in S$. The validity of this principle of proof, formally requiring an application of the Axiom of Choice, is intuitively obvious. For, assume P is progressive and $\neg P(x_0)$ for some $x_0 \in S$ (towards a contradiction). Then there exists $x_1 < x_0$ with $\neg P(x_1)$. Iteration of this argument yields an infinite descending sequence, which contradicts the well-foundedness of $<$. It should be noted that we do not use at all the full proof theoretic strength of transfinite induction. This strength can be measured by assigning ordinals to elements of S in the usual way: $\|x\| = \sup\{\|y\| + 1 \mid y < x\}$ (with $\sup \emptyset = 0$). We only use transfinite induction with respect to orderings having the property that $\|x\| < \omega$ (the order type of the natural numbers) for all $x \in S$. Whether one prefers transfinite induction or induction on the natural number $\|x\|$ (or even an informal argument in which the induction is not explicit) appears to be a matter of taste. However, for the purpose of unifying completeness proofs of various strategies of resolution and analyzing their differences, transfinite induction suits best.

1.3. Zorn's Lemma

Let \subset be a partial ordering on a set S . A *chain* in S is a totally ordered subset of S (i.e. satisfying the trichotomy axiom). *Zorn's Lemma* states that S contains a maximal (minimal) element, provided that every chain in S has an upper (lower) bound in S . Zorn's Lemma is known to be one of the most practical equivalents of the Axiom of Choice (see [3]). We use it in dealing with infinite sets of clauses.

1.4. Zermelo's Well Ordering Theorem

A *well ordering* of a set S is an ordering such that every non-empty subset of S has a smallest element. The Well Ordering Theorem of Zermelo (see [3]) states that every set can be well ordered. The Well Ordering Theorem is equivalent to the Axiom of Choice; we use it to ensure the possibility of ordered hyperresolution for all \mathcal{L} .

1.5. Resolution

Resolution is the rule according to which a *resolvent* $R = (C - \{L\}) \cup (C' - \{\bar{L}\})$ may be inferred from *parent clauses* C and C' , containing literals L and \bar{L} , respectively, and satisfying the requirements of the strategy. A *strategy* is, intuitively speaking, a prescription telling which clauses may be resolved. For example, we can require that one of the parent clauses consists entirely of positive literals. For some strategies this prescription may extend over more than one inference. For our exposition this informal notion of strategy suffices.

A *derivation* (relative to a strategy) of a clause C from a set of clauses S is a sequence of clauses C_1, \dots, C_n such that $C_n = C$ and, for all $1 \leq k \leq n$, either C_k is in S or C_k is a resolvent of some C_i and C_j with $1 \leq i, j < k$, provided that C_k may be inferred from C_i and C_j according to the strategy. For some strategies the notion of derivation has to be generalized by allowing C_k to be inferred from C_{i_1}, \dots, C_{i_r} , with $1 \leq i_1, \dots, i_r < k, r \geq 2$, instead of just from C_i and C_j .

Resolution can easily be proved *sound*, i.e. for any interpretation I , resolution preserves truth in I . For, at least one of two true parent clauses must contain a true literal different from the two complementary literals which are resolved out. In particular satisfiability is preserved when resolvents are added to a set of clauses; unsatisfiability is of course always preserved when a set of clauses is extended.

Completeness of a resolution strategy is the property that from any unsatisfiable set of clauses the empty clause can be derived. Completeness is usually proved as follows: let S be an unsatisfiable set of clauses, close S under resolution according to the strategy and then applying the following.

Proposition 1.1. *Every set of clauses which is closed under resolution according to the strategy and does not contain the empty clause is satisfiable.*

In the next section we shall prove this proposition for various resolution strategies: binary resolution [7], semantic resolution [6] (where the notion of renaming is introduced), P_1 -resolution as well as SLD-resolution [5], hyperresolution [8] and ordered hyperresolution [10] (where the idea of ordering the atoms is attributed to Reynolds and worked out in the more general setting of semantic resolution). All proofs will be by transfinite induction and have the following general form: first prove for some well-founded ordering $<$ on S and some interpretation J that truth in J of clauses from S is progressive, then conclude that J is a model for S by transfinite induction.

2. Completeness

2.1. Binary resolution

In the case of *binary* resolution, no restrictions are specified and every two clauses containing complementary literals may be resolved. Let S be a set of clauses which is closed under binary resolution and does not contain the empty clause. Note that

we do not assume that S is finite. Fix an arbitrary interpretation I . The interpretation I may informally be seen as a first try for a model of S . If this try fails (i.e. if $S_{\mathcal{F}}(I)$ is not empty), then I has to be “adjusted” on literals occurring in clauses from $S_{\mathcal{F}}(I)$. This adjustment should not affect the truth of clauses from $S_{\mathcal{F}}(I)$. Therefore some “minimal” adjustment is made, yielding a model J of S .

Let X be the set of all literals which occur in clauses from $S_{\mathcal{F}}(I)$ (formally $X = \bigcup S_{\mathcal{F}}(I)$). Note that $X \subseteq \mathcal{F}(I)$ does not contain complementary literals. We say that a subset Y of X covers $S_{\mathcal{F}}(I)$ (or Y is a *covering* of $S_{\mathcal{F}}(I)$) if every clause from $S_{\mathcal{F}}(I)$ contains at least one literal from Y . In particular X itself covers $S_{\mathcal{F}}(I)$, since S does not contain the empty clause. We shall construct a minimal (with respect to set inclusion) subset Y of X that covers $S_{\mathcal{F}}(I)$. If X is finite, then Y is easily obtained from X by deleting elements in such a way that the resulting set still covers $S_{\mathcal{F}}(I)$. If X is infinite, then this process is iterated, intuitively speaking, in a transfinite way until eventually Y is reached:

$$X_0 = X, \dots, X_{i+1} = X_i - \{L\}, \dots, X_\omega = \bigcap_{i < \omega} X_i, X_{\omega+1} = X_\omega - \{L'\}, \dots,$$

with $L \in X_i$ such that X_{i+1} covers $S_{\mathcal{F}}(I)$, and $L' \in X_\omega, \dots$ similarly. Note that we tacitly assumed that, for example, X_ω covers $S_{\mathcal{F}}(I)$. As there are much more ordinals in the universe than elements of X , this process terminates with a minimal set covering $S_{\mathcal{F}}(I)$.

This informal argument can be made rigorous by applying Zorn’s Lemma. Let Z be the set of subsets of X that cover $S_{\mathcal{F}}(I)$. Z is partially ordered by set inclusion. Existence of a minimal Y in Z is guaranteed by Zorn’s Lemma if we prove that every chain in Z has lower bound in Z . Let Z' be a chain in Z . The set $\bigcap Z'$ (with $\bigcap \emptyset = X$) is certainly a lower bound of Z' , so it suffices to prove that $\bigcap Z'$ is in Z , i.e. covers $S_{\mathcal{F}}(I)$. Suppose $\{L_1, \dots, L_n\}$ is a clause in $S_{\mathcal{F}}(I)$ having no literal in common with $\bigcap Z'$ (towards a contradiction). Then there exists for every $1 \leq i \leq n$ an element, say X_i , of Z' which does not contain L_i . Since Z' is a chain, the X_i s are totally ordered. Hence some X_i is a subset of all of them, and hence contains none of the literals L_1, \dots, L_n . This clearly contradicts $X_i \in Z' \subseteq Z$ by the definition of Z .

Given a minimal set Y covering $S_{\mathcal{F}}(I)$ we define J to be the (unique) interpretation such that $\mathcal{F}(I) \cap \mathcal{F}(J) = Y$ (formally $J = \{p_a \in \mathcal{L} \mid p_a \in Y \vee (\neg p_a \notin Y \wedge p_a \in I)\}$). In other words: the interpretation J is such that the truth valuations \mathcal{V}_J and \mathcal{V}_I only differ on the atoms which occur, positively or negatively, in Y . Since Y is a minimal covering of $S_{\mathcal{F}}(I)$ it follows that J is a model of $S_{\mathcal{F}}(I)$ having the property that for every literal $L \in \mathcal{F}(J)$ which occurs in a clause from $S_{\mathcal{F}}(I)$ there exists a clause in $S_{\mathcal{F}}(I)$ in which L is the *only* literal from $\mathcal{F}(J)$; otherwise $Y - \{L\}$ would cover $S_{\mathcal{F}}(I)$. This property of J is crucial and shall be used in the proof of the lemma below.

We now arrive at the point where the ordering $<$ on S is defined. Let $<$ be the transitive closure of the relation $<_1$ on S defined by

$$R <_1 C \text{ iff } R \in S \text{ is the resolvent of } C \in S \text{ and some } C' \in S_{\mathcal{F}}(I).$$

As R contains less literals from $\mathcal{T}(I)$ than C (recall that $\bigcup S_{\bar{x}}(I) \subseteq \mathcal{F}(I)$), it follows that $<$ is a well-founded partial ordering. The lemma below implies that truth in J is progressive. It follows by transfinite induction that J is a model of S , and hence S is satisfiable. This completes the proof of Proposition 1.1 in the case of binary resolution.

Lemma 2.1. *For every C in S we have: if $\forall R <_1 C J \models R$, then $J \models C$.*

Proof. Let C be a clause of S such that $\forall R <_1 C J \models R$. If $C \in S_{\bar{x}}(I)$, then we immediately have $J \models C$ since J is a model of $S_{\bar{x}}(I)$. Now assume $C \in S_{\bar{y}}(I)$ is false in J (towards a contradiction), then C consists entirely of literals from $\mathcal{F}(J)$. Since C is true in I , it follows that C contains a literal $L \in \mathcal{T}(I) \cap \mathcal{F}(J)$, so $\bar{L} \in \mathcal{F}(I) \cap \mathcal{T}(J) = Y$. Now by the crucial property of J stated above there exists a clause $C' \in S_{\bar{x}}(I)$ such that \bar{L} is the only literal of C' which is true in J . Hence $R = (C - \{L\}) \cap (C' - \{\bar{L}\}) <_1 C$ and R consists entirely of literals which are false in J . This clearly contradicts $\forall R <_1 C J \models R$. \square

2.2. Comparison with other completeness proofs

The first completeness proof for ground resolution was given in [7] as a purely combinatorial result for finite sets of clauses. Completeness of quantified resolution (where literals are atoms or negated atoms of predicate logic and clauses are finite sets of literals, thought of as the universal closure of the disjunction of these literals) was obtained by, first, reduction to the finite ground case using Herbrand's Theorem, and then lifting the result back to the quantified level by using the so-called Lifting Lemma [7, 5.15]. It should be noted that the combinatorial argument from [7] immediately generalizes to the countable case, both with respect to the cardinality of the language and of the set of clauses. In fact, with an application of the Well Ordering Theorem, the argument can be generalized to arbitrary cardinalities.

In [8] a completeness proof was presented which was based on a kind of minimality argument such as we used in Section 2.1. In [8], however, minimality was taken in the sense of number, thus limiting the argument to the finite case, whereas we take minimality in the sense of set inclusion.

More recent completeness proofs, such as in [9], use Herbrand map trees, also called semantic trees, and do not appeal to Herbrand's Theorem. However, the Herbrand map tree argument relies on the countability of the language.

In the previous subsection we obtained, by using Zorn's Lemma, a completeness proof for ground resolution which is completely independent of the cardinality of both the language and the set of clauses. The completeness of quantified resolution can now be proved as follows. Let S be an unsatisfiable set of clauses. Then the set $ground(S)$ of all variable-free instances (with respect to the language of S) of clauses from S is an unsatisfiable set of ground clauses, since every model of $ground(S)$ would be a Herbrand model of S . So by the completeness of ground resolution the empty clause can be derived from $ground(S)$. By the Lifting Lemma this derivation

can be lifted to a derivation of the empty clause from S by quantified resolution. Hence we have proved the completeness of quantified resolution independently of the cardinality of both the language and the set of clauses, without appealing to Herbrand's Theorem. Although in practice languages will be finitely generated and hence countable, we think that these features, combined with the uniform approach to various strategies of resolution, indicate the full generality of our proof and have an aesthetic merit.

2.3. Redundancies

If one takes a closer look at the argument developed in Section 2.1, then the following observations can be made:

- the interpretation I on which the argument is based is arbitrary;
- the minimal set Y covering $S_{\bar{x}}(I)$ may not be unique;
- Lemma 2.1 is stronger than progressivity since $\forall R <_1 C J \models R$ is weaker than $\forall R < C J \models R$.

These observations reveal substantial redundancies in the completeness proof, since for any interpretation I , any minimal Y covering $S_{\bar{x}}(I)$, and even with $<_1 = <$ a completeness result can be obtained.

In general, a resolution strategy aims at reducing the costs of finding a derivation of the empty clause from a given unsatisfiable set of clauses S . If a strategy is complete, then we can simply close S under resolution according to the strategy, until eventually the empty clause is derived. The costs of this closing procedure depend on the number of generated resolvents. Thus the importance of reducing the number of generated resolvents becomes evident. To this end various strategies of resolution exploit the redundancies in the completeness proof of Section 2.1 mentioned above: semantic resolution (with P_1 -resolution and SLD-resolution as special cases) fixes I , hyperresolution fixes I and trivializes the ordering (" $<_1 = <$ "), whereas ordered hyperresolution exploits the non-uniqueness of Y as well. We shall discuss these matters in the following subsections.

2.4. Semantic resolution

In the case of *semantic* resolution, an interpretation I is fixed in advance. Given a set of clauses S , resolution is only allowed between a clause from $S_{\bar{x}}(I)$ and one from $S_{\bar{y}}(I)$. This restriction does not at all affect the completeness proof from Section 2.1. Hence semantic resolution is complete.

P_1 -resolution [8] is obtained as a special case of semantic resolution by taking $I = \emptyset$. Then $S_{\bar{x}}(I)$ consists of the clauses from S not containing negated atoms, so-called *positive* clauses.

SLD-resolution [5] is a rule of inference for so-called Horn clauses. A *Horn clause* is a clause with at most one positive literal. Note that the set of all Horn clauses is closed under binary resolution. We distinguish between *program clauses* (or *definite clauses*), which contain exactly one positive literal, and *goal clauses*, which consist

entirely of negated atoms. Thus the empty clause is a goal clause. SLD-resolution uses a *selection rule*, which selects from every goal clause a (negative) literal. Resolution is only allowed between program clauses and goal clauses, and with the restriction that the negation of the positive literal of the program clause is the selected literal of the goal clause. SLD-resolution can be viewed as semantic resolution with $I = \mathcal{L}$: for a set of Horn clauses S , $S_{\mathcal{L}}(\mathcal{L})$ consists of the goal clauses from S , and $S_{\mathcal{L}}(\mathcal{L})$ of the program clauses. With some technical effort (concerning selection rules) the completeness of SLD-resolution can be obtained from the completeness of semantic resolution. We refrain from giving a detailed account on this point.

As done in [10], hyperresolution as well as ordered hyperresolution (and also SLD-resolution) could be treated more generally in the context of semantic resolution. For reasons of simplicity, however, we prefer to specialize to the case $I = \emptyset$. Modulo renaming from [6], we do not lose generality.

2.5. Hyperresolution

In [8] hyperresolution was introduced as a refinement of P_1 -resolution. A hyperresolvent of a set of clauses S is a positive clause which is obtained by successive P_1 -resolutions in a way depicted in Fig. 1. More precisely: a positive clause C_{n+1} is called a *hyperresolvent* of S with parent clause C_1 if $n \geq 1$, $C_1 \in S$, $D_i \in S$ is positive and C_{i+1} is a P_1 -resolvent of C_i and D_i , for all $1 \leq i \leq n$.




Fig. 1.




Fig. 2.

If we assume that S is closed under hyperresolution, then we can define a relation $<_1^h$ on S by

$$R <_1^h C \text{ iff } R \in S \text{ is hyperresolvent of } S \text{ with parent clause } C \in S.$$

Since every hyperresolvent is positive and can as such not act as a parent clause of another hyperresolvent, it follows trivially that $<_1^h$ is a well-founded partial ordering. As $<_1^h$ equals its transitive closure $<^h$ we shall drop the subscript.

In order to compare $<^h$ with $<$ we assume for the time of this paragraph that S is closed under P_1 -resolution, which implies closure under hyperresolution. We then have

$$R <^h C \text{ iff } R < C \text{ and } R \text{ is positive.}$$

Since positive clauses are $<$ -minimal we can view $<^h$ as “cutting short” $<$.

Let S be a set of clauses which is closed under hyperresolution and does not contain the empty clause. The argument that S is satisfiable is again similar to that given in Section 2.1, taking $I = \emptyset$, $<_{(1)}^h$ instead of $<_{(1)}$, reading “hyperresolution” for “resolution”, and so on. Only the proof of the progressivity of truth in the interpretation J needs some more attention.

Lemma 2.2. *Under conditions as above we have for every C in S : if $\forall R <^h C J \models R$, then $J \models C$.*

Proof. Let C be a clause of S such that $\forall R <^h C J \models R$. If $C \in S_{\mathcal{F}}(\emptyset)$ (i.e. C is positive), then we immediately have $J \models C$ since J is a model of $S_{\mathcal{F}}(\emptyset)$. Now assume $C \in S_{\mathcal{T}}(\emptyset)$ is false in J , i.e. C consists entirely of literals which are false in J (towards a contradiction). Let L_1, \dots, L_n ($n > 0$) be the negative literals of C (here we deviate from Section 2.1 if $n > 1$). We have $L_1, \dots, L_n \in \mathcal{T}(\emptyset) \cap \mathcal{F}(J)$, so $\bar{L}_1, \dots, \bar{L}_n \in \mathcal{F}(\emptyset) \cap \mathcal{T}(J) = Y$. It follows by the minimality of Y that there exist $D_1, \dots, D_n \in S_{\mathcal{F}}(\emptyset)$ (i.e. positive clauses) such that \bar{L}_i is the *only* literal of D_i which is true in J ($1 \leq i \leq n$). Hence the hyperresolvent R of S with parent clause C , obtained as in Fig. 2, consists entirely of literals which are false in J . This contradicts $\forall R <^h C J \models R$. \square

2.6. Ordered hyperresolution

In the case of ordered hyperresolution (see [10]), we rely on the Well Ordering Theorem for the existence of a well-ordering on \mathcal{L} . Recall Fig. 1 and assume $C_{i+1} = (C_i - \{\neg p_a\}) \cup (D_i - \{p_a\})$ for $1 \leq i \leq n$. For C_{n+1} to be an *ordered* hyperresolvent we require p_a to be the *maximal* atom of D_i , for all $1 \leq i \leq n$. Note that in a well-ordering every finite set indeed has a maximum.

Ordered hyperresolution is seen to be complete in the same way as hyperresolution. The restriction upon the p_a s has the effect that the deletion process in Section 2.1 should be modified as follows. Instead of starting with a covering set $X = \bigcup S_{\mathcal{F}}(\emptyset)$ we start with $X_0 = \{p_a \mid p_a \text{ is the maximal atom of a clause in } S_{\mathcal{F}}(\emptyset)\}$. Furthermore, the deletion process should be such that the minimal atom $p_a \in X_\alpha$ such that $X_\alpha - \{p_a\}$

covers $S_{\mathcal{F}}(\emptyset)$ is deleted. So we have formally $X_{\alpha+1} = X_{\alpha} - \{\min\{p_a \mid p_a \in X_{\alpha} \text{ and } X_{\alpha} - \{p_a\} \text{ covers } S_{\mathcal{F}}(\emptyset)\}\}$ for all ordinals α for which X_{α} is not a minimal covering. Note that in a well-ordering such minimal atoms do always exist. Thus a minimal set Y covering $S_{\mathcal{F}}(\emptyset)$ is obtained, having the property that for every $p_a \in Y$ there exists a clause in $S_{\mathcal{F}}(\emptyset)$ in which p_a is the maximum and the *only* literal from Y . For, there exist clauses in $S_{\mathcal{F}}(\emptyset)$ in which p_a is the only literal from Y (otherwise Y would not be a minimal covering). Assume all such clauses have a maximum greater than p_a (towards a contradiction). Before deletion of one of these maxima, deletion of p_a would also result in a set which still covers $S_{\mathcal{F}}(\emptyset)$. But then the deletion of the first of these maxima would be contrary to the definition of the deletion process, since p_a is smaller than the deleted literal. By this contradiction we have proved the desired property of the minimal covering Y .

2.7. The role of the Axiom of Choice

One may ask in how far the Axiom of Choice is really necessary for the completeness of resolution. Let us first give simple and intuitive evidence that at least some weak form of the Axiom of Choice is necessary. Consider a collection $\mathcal{C} = \{\{p_a, q_a\} \mid a \in A\}$ of pairwise disjoint sets of two indistinguishable elements. A *choice function* for \mathcal{C} is a function on A such that $f(a) \in \{p_a, q_a\}$ for all $a \in A$. Let S be the set consisting of clauses $\{p_a, q_a\}$ and $\{\bar{p}_a, \bar{q}_a\}$ for every $a \in A$. In fact S consists of the clausal forms of $p_a \leftrightarrow \neg q_a$ ($a \in A$). The closure of S under resolution does not contain the empty clause: it is simply S itself plus tautologies $\{p_a, \bar{p}_a\}, \{q_a, \bar{q}_a\}$ for all $a \in A$. By (the contraposition of) the completeness of resolution it follows that S must be satisfiable, i.e. has a model. But every model of S constitutes a choice function for \mathcal{C} . So we have proved a special case of the Axiom of Choice, which is unprovable in set theory (see [4, Theorem 5.20]). The argument above can easily be extended to a proof of the so-called Axiom of Choice for Finite Sets [4, p. 107]. Note that the argument above requires clauses $\{p_a, q_a\}$ which are not Horn clauses.

For establishing the precise set theoretic status of the completeness of resolution we recall that the Compactness Theorem is a well-known (weak) consequence of the Axiom of Choice, equivalent to, e.g. the Prime Ideal Theorem for Boolean Algebras (see [4, Theorem 2.2]). The completeness of resolution immediately implies (and in fact is equivalent to) the following version of the Compactness Theorem.

Theorem 2.3. *Every unsatisfiable set of clauses has a finite subset which is unsatisfiable.*

By the language restriction to clauses we have to exercise some care in applying [4, Theorem 2.2]: in predicate logic reduction to clausal form involves skolemization, which relies on the Axiom of Choice (see [1, footnote 8]). In the propositional case, however, the reduction simply consists of taking conjunctive normal forms. Inspection of the proof of Theorem 2.2 from [4] tells us that the use of predicate logic there is completely harmless. In fact this proof could be given using propositional

logic only. Note that this proof involves clauses which are not Horn clauses, namely $I(u) \vee I(-u)$, or $\{p_u, p_{-u}\}$ as clauses in our propositional language.

In the previous paragraphs we gave two hints that the restriction to Horn clauses might affect the set theoretic status of the completeness results. Indeed the completeness of resolution for Horn clauses can be proved in the following elementary way. (It does not make sense to consider SLD-resolution in this respect, since the selection function already presupposes a choice function on the set of non-empty goal clauses.) Let S be a set of Horn clauses which is closed under resolution and does not contain the empty clause. Then $S_{\bar{x}}(\emptyset)$ consists entirely of clauses consisting of one single atom. These atoms form a minimal covering of $S_{\bar{x}}(\emptyset)$ and hence a model of S . The set theoretic effect of the restriction to Horn clauses nicely corresponds to the complexity theoretic facts that the satisfiability problem for finite sets of clauses is NP-complete (see [2]), whereas the same problem for Horn clauses can easily be seen to be in the complexity class P.

Another interesting special case that we can consider is the case in which the language is countable. Then the set of all clauses is also countable. The completeness of resolution can then be proved without any appeal to the Axiom of Choice. The easiest way to see this is to inspect the proof of the completeness of ordered hyperresolution in Section 2.6. A countable language can trivially be well-ordered, without appealing to the Well Ordering Theorem, by using the enumeration as ordering. The minimal covering is then constructed in at most ω steps of the deletion process described in Section 2.1, which is made completely deterministic by the modifications from Section 2.6. This observation explains why the Herbrand map tree argument in [9], where the language is countable, does not appeal to the Axiom of Choice, nor to Herbrand's Theorem or the Compactness Theorem.

Acknowledgment

This paper has benefitted much from discussions with Anton Eliëns, Jan Willem Klop, Piet Rodenburg and Professor J.A. Robinson.

References

- [1] M. Davis and H. Putnam, A computing procedure for quantification theory, *J. ACM* 7 (1960) 201–215.
- [2] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-completeness* (Freeman, San Francisco, CA, 1979).
- [3] P.R. Halmos, *Naive Set Theory* (Van Nostrand, Princeton, NJ, 1960).
- [4] T.J. Jech, *The Axiom of Choice* (North-Holland, Amsterdam, 1973).
- [5] R.A. Kowalski, Predicate logic as a programming language, in: J. Rosenfeld, ed., *Information Processing 74*, Stockholm (North-Holland, Amsterdam, 1974) 569–574.
- [6] B. Meltzer, Theorem-proving for computers: some results on resolution and renaming, *Comput. J.* 8 (1966) 341–343.

- [7] J.A. Robinson, A machine oriented logic based on the resolution principle, *J. ACM* **12** (1965) 23–41.
- [8] J.A. Robinson, Automatic deduction with hyper-resolution, *Internat. J. Comput. Math.* **1** (1965) 227–234.
- [9] J.A. Robinson, *Logic: Form and Function* (Edinburgh University Press, Edinburgh, 1979).
- [10] J.R. Slagle, Automatic theorem proving with renamable and semantic resolution, *J. ACM* **14** (1967) 687–697.