Constructing Bachmair-Ganzinger Models*

Christopher Lynch

Department of Mathematics and Computer Science Box 5815,
Clarkson University, Potsdam, NY 13699-5815, USA
clynch@clarkson.edu

Abstract. We give some algorithms for constructing models from sets of clauses saturated by Ordered Resolution (with Selection rules). In the ground case, we give an efficient algorithm for constructing a minimal model. Then we generalize minimal models to preferred models, which may be useful for verification. For the ground case, we also show how to construct all models for a set of clauses saturated by Ordered Resolution, in time polynomial in the number of models. We also generalize our results to nonground models, where we add a restricted splitting rule to our inference rules, and show that for any set of clauses saturated by Ordered Resolution (with Selection), a query about the truth of a particular atom in the model can be decided.

1 Introduction

It is generally believed that a major drawback of Resolution-based theorem proving methods is that a model is not constructed when the set of clauses is satisfiable. If the inference system halts without producing the empty clause, then the set of clauses is determined to be satisfiable. But there is no model constructed. The set of clauses produced by the inference system can be considered to represent a model in some sense. In fact, it is possible to theoretically construct a model in this case. But in the practical sense, there is no known method for determining if a ground atom is true in this theoretically constructed model. This problem has received some attention[9,6,5], and methods have been given in some restricted cases. The main goal of this paper is to determine a more general way to accomplish this.

Interestingly, the method of Bachmair and Ganzinger[1] for proving completeness of the resolution inference system actually constructs a model for a set of clauses saturated by Resolution when the empty clause cannot be produced. But this is only a theoretical construction. It is difficult to use this practically. In the ground case (no variables), it can be done. But not in the nonground case.

In this paper, we first define a notion called a Preferred Model for ground clauses. For each atom, the user defines a preference for that predicate of either true or false. A model of a set of clauses is a Preferred Model if each atom receives the preferred truth value whenever that is consistent with the truth

* This work was supported by NSF grant number CCR-0098270.

A. Voronkov and C. Weidenbach (Eds.): Ganzinger Festschrift, LNCS 7797, pp. 285–301, 2013.

value of all smaller atoms. For example, if each atom is preferred to be false, then the Preferred Model is the Minimal Model. Each set of clauses has a unique Preferred Model. We show that if a set of ground clauses S is saturated by Ordered Resolution, then the Preferred Model of S can be constructed in time $O(|S|lg(|S|))$. Preferred Models could be useful in verification. For example if a program does not meet its specifications, then the programmer would like to see a counterexample. Since all counterexamples might not make sense, it would be useful for the programmer to express some preferences.

We then give an algorithm to show that if a set of clauses S has exactly k models and if S is saturated by Ordered Resolution, then all models of S can be constructed in time $O(|S|lg(|S|) + |S|k)$. In other words, the time needed to construct all the models of S is just the time it takes to write out all the models, plus an initialization time to sort the clauses. In general, a set of clauses may have exponentially many models. But this result shows that if there are only polynomially many models, then they can all be constructed in polynomial time.

We extend our results on Preferred Models to nonground clauses saturated by Ordered Resolution (possibly with selection rules). This is useful, because Ordered Resolution is an inference rule that often halts, the only model construction results of which we are aware which handles clauses saturated by Ordered Resolution is the one of Peltier[9], but that method only handles some sets of saturated clauses. Of course, models of nonground clauses may be infinite. We do not try to schematize all the models. Instead, we are interested in developing an algorithm which will decide if a given atom is true in the Preferred Model. The notion of Preferred Model can be extended to nonground clauses by defining it for the ground instances.

The first result for nonground clauses is related to results for Local Theories. Given an ordering $<$, we can define the order type $ot(n)$ to be the number of atoms smaller than any atom of size n. If S is a set of clauses saturated by Ordered Resolution, and if A is a ground atom, then it can be decided in time polynomial in $ot(|A|)$ whether A is true in the Preferred Model of S. We can extend the result so that if N is a set of ground clauses, then it is decidable in time exponential in $ot(|N| + |A|)$ whether A is true in the Preferred Model of $S \cup N$. The interest of this last result is that rather than just deciding whether atoms are true in a model of S, we are asking whether atoms are true in a model of any set of ground clauses modulo the theory of S.

The above results can only work if the order type is finite, and it is not finite for some orderings. Therefore, we address the problems using a different technique. First, we add a Splitting rule to the Resolution inference system, which is only applicable in a restricted number of cases. Our major result is to show that if S is a set of clauses saturated by Ordered Resolution (plus Selection) with Splitting, then it is decidable whether A is in the Preferred Model of S. The Splitting rule is especially restrictive if the Ordering satisfies some simple conditions, which hold for most standard orderings.