Extensions of the Knuth-Bendix Ordering with LPO-like Properties

Michel Ludwig*

*Department of Computer Science
University of Liverpool, UK
michel@csc.liv.ac.uk

Abstract

The Knuth-Bendix ordering is usually preferred over the lexicographic path ordering in successful implementations of resolution and superposition calculi. However, it is incompatible with certain requirements of hierarchic superposition calculi, and it also does not allow non-linear definition equations to be oriented in a natural way. We present two extensions of the Knuth-Bendix ordering that make it possible to overcome these restrictions.

1 Introduction

In theorem proving calculi like Knuth-Bendix completion, resolution, or superposition, reduction orderings like the Knuth-Bendix ordering (KBO) or the lexicographic path ordering (LPO) are crucial to reduce the search space. Among these orderings, the Knuth-Bendix ordering is usually preferred in state-of-the-art implementations of theorem provers. There are several reasons for this: it can be efficiently implemented – the most efficient known algorithm only needs linear time – and it correlates well with the sizes of terms; so, reductions with respect to a KBO usually lead to smaller terms. In comparison, computing the lexicographic path ordering requires at least quadratic time and reductions with respect to an LPO may result in arbitrarily larger terms.

On the other hand, it is exactly this correlation between a KBO and term sizes that renders the KBO incompatible with the particular requirements occurring in certain applications. One example is hierarchic theorem proving (Bachmair et al., 1994; Ganzinger et al., 2006; Prevosto and Waldmann, 2006), where one considers two signatures $\Sigma \supseteq \Sigma_0$ and needs an ordering in which every ground term involving a symbol from $\Sigma \setminus \Sigma_0$ is larger than every ground term over Σ_0. With an LPO, this property is easy to establish, with a KBO it is usually impossible.

A second example are definitions of the form $f(t_1, \ldots, t_n) \approx t_0$ where f does not occur in t_0 as part of a larger specification. Such definitions can easily be ordered from left to right using an LPO where f is larger in the precedence than every symbol occurring in t_0. With a KBO, however, we have the additional requirement that no variable occurs more often in t_0 than in $f(t_1, \ldots, t_n)$; thus, non-linear definitions cannot be handled adequately by the KBO.

A generalisation of the second example involves the orientation of a system of definitions. Suppose that we have a sequence of signatures Σ_i $(0 \leq i \leq n)$ where $\Sigma_i = \{f_i\} \cup \Sigma_{i-1}$ for $i \geq 1$, and that we have a set of non-recursive definition equations of the form $f_i(s_{i1}, \ldots, s_{ik}) \approx t_i$, with terms $t_i, s_{ij} \in \mathrm{T}_{\Sigma_{i-1}}(X)$ and $\mathrm{Var}(t_i) \subseteq \mathrm{Var}(f_i(s_{i1}, \ldots, s_{ik}))$ (where the s_{ij} are often, but not necessarily, variables). If we use a lexicographic path ordering with a precedence $f_n > \cdots > f_2 > f_1 > \ldots$, then every term t with a top symbol f_i is larger than every term in $\mathrm{T}_{\Sigma_{i-1}}(\mathrm{Var}(t))$; so, all these equations can be oriented from left to right (and can hopefully be used to eliminate all occurrences of the f_i in the remainder of the specification completely). If we try to get a similar effect with a KBO, there are two problems: a KBO correlates with term sizes. So, in general, a term cannot be larger than *every* term over some subsignature, and moreover a term cannot be larger than another term in which some variable occurs more often.

Our goal is to find a computable (total) simplification ordering that behaves exactly like the KBO on terms over Σ_0 but in which every term t with a top symbol f_i is larger than every term in $\mathrm{T}_{\Sigma_{i-1}}(\mathrm{Var}(t))$. To this end, we have developed two extensions of the KBO, namely the *KBO with Pairs* and the *Transfinite KBO*. Due to lack of space, we just illustrate the basic ideas of the two extensions. For concrete definitions and proofs, the reader is referred to (Ludwig, 2006).

2 Extensions

Generally speaking, the KBO uses three criteria to compare two terms, which are based on weights assigned to signature symbols and a precedence ordering on these symbols. At first, the weights assigned to signature symbols are used to compute and compare the weight of terms. Then, if both terms should have the same weight, the precedence on signature symbols is considered to compare the top symbols of the terms. And finally only, if both terms should also have the same top symbols, the KBO compares their subterms. It is clear that in implementations of a KBO, the term weights allow for performance improvements over the implementations of a LPO as the weights alone suffice in many practical cases to compare two terms.

The two extensions that we will introduce now therefore keep the general structure of the KBO, but they differ in the objects that they assign as weights to signature symbols.

2.1 KBO with Pairs

The Knuth-Bendix Ordering with Pairs (PKBO) can be used to solve the aforementioned problem in which one considers two signatures $\Sigma \supseteq \Sigma_0$ and one needs an ordering in which every ground term involving a symbol from $\Sigma \setminus \Sigma_0$ is larger than every ground term over Σ_0. The key concept of the PKBO is the use of pairs of real numbers as weights for term signature symbols, whereas the original KBO uses simple reals. The pairs of real numbers will be compared by using the lexicographic extension of the regular greater-than relation on reals. For example, by assigning the pair $(1, 0)$ as weight to a symbol f occurring in $\Sigma \setminus \Sigma_0$ and by giving to every symbol in the signature Σ_0 a pair of reals whose first component is equal to 0, we obtain that every ground term which contains at least one occurrence of the symbol f is bigger with respect to the PKBO than every other ground term which only contains symbols from Σ_0.

Moreover, the conditions on variables imposed by the KBO (i.e. every variable must occur more often in the bigger term with respect to the KBO than in the smaller term) can be relaxed for the PKBO; a slight modification of the notion of “simplification ordering” is necessary however.

2.2 Transfinite KBO

Similar to the previous extension, the Transfinite Knuth-Bendix Ordering (TKBO), which can be used to handle iterated, non-recursive term definitions, does not assign real numbers as symbol weights, but it uses *ordinal numbers* instead (more specifically, ordinals that are smaller than ϵ_0). In that way, we can make use of the property that some specific ordinals are bigger, with respect to the comparison relation on ordinal numbers, than infinitely many other ordinals.

Furthermore, by introducing the concept of *position coefficients*, the conditions on variable occurrences can be relaxed somewhat. We can allow some variables to occur more often in the smaller than in the bigger term of a TKBO comparison. Position coefficients introduce an additional layer on the structure of terms, i.e. not only do we assign an ordinal as weight to every signature symbol, but additionally an ordinal different from 0, which we will call its *subterm coefficient*. The coefficient of a position in a term will be defined as the (commutative) product of all the subterm coefficients on the path from the root symbol to the considered position. Then, the condition on variables of the regular KBO will be replaced for the TKBO by the condition that for every variable x in the smaller term, the sum of all the position coefficients for occurrences of x must be smaller than the sum of all the position coefficients for occurrences of the variable x in the bigger term (with respect to the TKBO).

If we now reconsider the example from the introduction, i.e. we are given a sequence of signatures Σ_i $(0 \leq i \leq n)$ where $\Sigma_i = \{f_i\} \cup \Sigma_{i-1}$ for $i \geq 1$, and a set of non-recursive definition equations of the form $f_i(s_{i1}, \ldots, s_{ik}) \approx t_i$, with $t_i, s_{ij} \in \mathrm{T}_{\Sigma_{i-1}}(X)$ and $\mathrm{Var}(t_i) \subseteq \mathrm{Var}(f_i(s_{i1}, \ldots, s_{ik}))$, then we assign the ordinal ω^{ω^i} both as symbol weight and subterm coefficient to the signature symbol f_i, and smaller ordinals to the other symbols. It follows then from the greater-than relation on ordinals that for all terms $t \in \mathrm{T}_{\Sigma_{i-1}}(\mathrm{Var}(f_i(s_{i1}, \ldots, s_{ik})))$, we obtain $f_i(s_{i1}, \ldots, s_{ik}) \succ_{\mathrm{T}} t$.

It remains to be noted that in many practical cases, it is not necessary for the previous example to evaluate the full ordinal weights of terms (which cannot be done in constant time). It is often sufficient to check whether one of the terms contains a definition symbol; for more details, see (Ludwig and Waldmann, 2007).

References

Leo Bachmair, Harald Ganzinger, and Uwe Waldmann. Refutational theorem proving for hierarchic first-order theories. *Applicable Algebra in Engineering, Communication and Computing (AAECC)*, 5(3/4):193–212, April 1994.

Harald Ganzinger, Viorica Sofronie-Stokkermans, and Uwe Waldmann. Modular proof systems for partial functions with Evans equality. *Information and Computation*, 204:1453–1492, 2006.

Michel Ludwig. Extensions of the Knuth-Bendix ordering with LPO-like properties. Diploma thesis, Universität des Saarlandes, Saarbrücken, Germany, July 2006.

Michel Ludwig and Uwe Waldmann. Extensions of the Knuth-Bendix ordering with LPO-like properties. Submitted to CADE 2007, February 2007.

Virgile Prevosto and Uwe Waldmann. SPASS+T. In *Proc. ESCoR'06*, volume 192 of *CEUR Workshop Proceedings*, pages 18–33, August 21, 2006.